Early Detection of Digital Dermatitis in Cattle using Computer Vision

By

SRIKANTH ARAVAMUTHAN

A dissertation submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

(Comparative Biomedical Sciences)

at the

UNIVERSITY OF WISCONSIN – MADISON

2023

Date of final oral examination: 12/7/2023

The dissertation is approved by the following members of the Final Oral Committee:

Dörte Döpfer, Professor, Medical Sciences
Brian Yandell, Professor, Statistics
Guilherme Rosa, Professor, Animal and Dairy Sciences
João Dórea, Assistant Professor, Animal and Dairy Sciences
Paul Merkatoris, Assistant Teaching Professor, Surgical Sciences

ACKNOWLEDGEMENTS

I would like to express my gratitude to everyone who has contributed to the completion of this thesis.

First and foremost, I would like to express my appreciation to my thesis advisor, Professor Dörte Döpfer, for their support and guidance. Their mentorship has been one of the most challenging, but also one of the most meaningful experiences.

I am indebted to the members of my thesis committee, Professor Brian Yandell, Professor Guilherme Rosa, Professor João Dórea, and Professor Paul Merkatoris, for their insightful feedback and invaluable suggestions. Their expertise in the field has contributed to the depth and breadth of this research.

I would like to extend my thanks to my parents for their continued support, encouragement, and understanding. Their belief in my abilities when I was not as sure has been the motivation behind the completion of this degree.

I would like extend my appreciation to my friends and colleagues, namely Emil Walleser, who provided both intellectual and emotional support. Their friendships and assistance made the challenges more manageable and the successes more meaningful.

I am also indebted to the staff and faculty in the Comparative Biomedical Sciences Graduate Program, namely Susan Thideman and Professor Lyric Bartholomay, for their administrative support and technical assistance that facilitated the navigation of degree requirements and the progress of my research.

My graduate studies would not have been possible without funding support from US Department of Agriculture through the National Institute for Food and Agriculture - Animal Health Grant (WIS03082).

This thesis reflects the collaborative effort of many, and I am truly grateful for the support and guidance I have received.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF ABBREVIATIONS

AI: Artificial intelligence

ANN: Artificial neural network

ANOVA: Analysis of Variance

AP: Average precision

API: Application programming interface

ASIC: Application specific integrated circuit

AUC: Area under the curve

Conv-BN: Convolution batch normalization

CI: Confidence interval

CNN: Convolutional neural network

CSV: Comma-separated value

CV: Computer vision

DC: DairyComp

DD: Digital dermatitis

DIM: Days in milk

E-ELAN: Extended efficient layer aggregation network

FPGA: Field programmable gate array

FPS: Frames per second

FSTBF: First-test butterfat percent

GPS: Global positioning system

GPU: Graphic processing unit

H: Hyperkeratosis

HSV: Hue, saturation, and value

HTTP: Hypertext transfer protocol

ID: Identification

IDE: Integrated development environment

IoT: Internet of Things

IoU: Intersection over Union

IP: Internet protocol

IRT: Infrared thermography

KET: Ketosis

LACT: Lactation number

LCTGP: Lactation group

mAP: Mean average precision

ML: Machine learning

MSE: Mean squared error

NMS: Non-maximum suppression

NPU: Neural processing unit

OR: Odds Ratio

OREPA: Online convolutional re-parameterization

P: Proliferation

R-CNN: Region-based convolutional neural network

ReLU: Rectified linear unit

RFID: Radio-frequency identification

RGB: Red, green, and blue

ROC: receiver operating characteristic

RP: Retained placenta

RPN: Region proposal network

RTMP: Real-time messaging protocol

RTSP: Real-time streaming protocol

SBC: Single-board computer

SGD: Stochastic gradient descent

SSD: Single shot multibox detector

TOPS: Tera operation per second

TPU: Tensor processing unit

TWIN: Twins

UAV: Unmanned aerial vehicle

VPU: Vision processing unit

YOLO: You Only Look Once

ABSTRACT

Digital dermatitis (DD) is an infectious bovine claw disease that presents as painful, ulcerative lesions on the coronary band at the skin horn border of the hoof. The disease is the main cause of lameness in both dairy and beef cattle, leading to a decrease in production and fertility and an increase in culling, resulting in diminished economic well-being and animal welfare. Currently, visual inspection of the lifted hoof is the gold standard for DD diagnosis. However, this method is time-consuming, labor-intensive, subjective, inaccurate, and unable to perform early detection of DD. Tools for early detection and prompt treatment of DD are critical for good practice. Early detection of DD reduces the use of topical treatments and disinfecting foot bath chemicals such as the environmentally problematic copper sulfate and the carcinogenic formalin.

Computer vision (CV) can be used to perform object detection and calculate the associated class probabilities from a series of images or videos. Object detection locates the presence of an object with a bounding box, class label, and class probability. Thus, CV provides a unique opportunity to improve early detection of DD and optimize treatment plans for cattle. Such methods have been used for animal detection to assess the health and well-being of livestock in cattle and small ruminants. However, applications for the classification of health events are still rare in veterinary medicine and food production. The purpose of this project is to implement a CV model for the real-time detection of DD in dairy and beef cattle. The motivation is to minimize the effects of DD-associated lameness in all cattle by means of early detection, prompt treatments, and prevention measures.

Chapter 2 aims to train and compare applications for the real-time detection of DD in dairy cows. Nine CV models were trained for detection and scoring, compared using performance metrics and inference time, and the best model was automated for real-time detection using images and

video. YOLOv4, Tiny YOLOv4, and YOLOv5s outperformed all other models with almost perfect precision, perfect recall, and a high mean average precision (mAP). Tiny YOLOv4 outperformed all other models with respect to inference time at 333 frames per second (FPS). Chapter 3 aims to train, embed, and benchmark DD detection models on edge devices. The Tiny YOLOv4 model deployed on a CV specific integrated camera module connected to a single board computer achieved high mAP, high overall prediction accuracy, and substantial agreement between the computer vision model and the trained investigator. The model reached a final inference speed of 40 FPS. Chapter 4 aims to develop, deploy, and evaluate DD detection models on cloud platforms. The TensorFlow.js application outperformed all other deployments with respect to agreement. All deployments exceeded the minimum threshold for image processing at approximately 10 FPS. Chapter 5 extend the current workflow from bovine DD to canine pododermatitis and neoplasia of the paw. The Pawgnosis tool is a novel object detection model deployed on a microcomputer with a camera for the rapid detection of canine pododermatitis and neoplasia. The proposed tool achieved a high precision, recall, and mAP with an inference speed of 20 FPS such that the novel object detection model has the potential for application in the field of veterinary dermatology.

Chapter 6 aims to provide a singular, automated workflow for extracting health and event information from DairyComp to DairyCoPilot and compile and analyze the data using graphical user interface. DairyCoPilot provides statistical tools for both categorical data analysis and regression analysis. The program provides an elegant and powerful web framework where the user interface is dynamic to conditionally generate input controls and uses reactive programming to automatically update outputs when inputs change. The web application can clean data for downstream data analysis, create interactive graphics for exploratory and expository visualization, and export high quality figures using interactive views for documents, reports, and presentations. Lastly, DairyCoPilot is mobile-friendly, desktop progressive web application for all device platforms.

Chapter 7 aims to update and improve the DD Check app for individuals with limited statistical training or experience. The DDCheckPlus app is designed to include a DD detection module where cows can be scored for M-stages using a custom object detection model. Additionally, the app is developed to streamline data analysis for automated prediction of current DD trends and forecasting of future DD proportions. All plots are interactive with details on demand and all tables are interactive with filtering, pagination, and sorting. The DDCheckPlus app standardizes M-stage data recording, automates comprehensive data analysis including trends over time, calculates predictions, and assigns Cow Types (I – III) based on the presence or absence of active DD lesions.

Chapter 8 serves as a placeholder for research in progress. It describes a comparative analysis of different mobile apps for real-time detection of DD on both Android and iOS devices. Chapter 9 discusses the subsequent workflow for object detection to train and deploy a model that can accurately detect and localize objects of interest within images or videos. This approach can save significant training time and resources while adapting the model to new tasks. Chapter 10 provides a discussion extending the research to further analysis and future studies. The focus of the chapter describes the major opportunities and challenges for deployment of computer vision algorithms and object detection models in precision farming and veterinary medicine.

CHAPTER 1 INTRODUCTION

Definition of Digital Dermatitis

Digital dermatitis (DD) is a polybacterial disease primarily affecting the skin horn border on cattle heels. The infection causes inflammation and skin damage, resulting in pain and discomfort (Laven and Proven, 2000; Plummer and Krull, 2017). The disease is a leading cause of lameness in both dairy cows and beef cattle, posing a significant problem for the dairy and beef industry in numerous countries (Bruijnis et al., 2012; Laven and Logue, 2006; Plummer and Krull, 2017; Sullivan et al., 2013; USDA, 2007). This condition leads to reduced animal welfare and economic loss (Laven, 2001). Recently, DD has emerged as a growing concern in beef cattle (Sullivan et al., 2013). The bacteria is causally associated with DD have also been identified in similar lesions in sheep (Duncan et al., 2014), dairy goats (Sullivan et al., 2015), and even wild North American Elk (Clegg et al., 2014). Furthermore, these DD-associated bacteria have been detected in three severe bovine foot lesions that have been studied during the past 15 years: toe necrosis, non-healing white line disease, and non-healing sole ulcer (Evans et al., 2011). These developments underscore the increasing importance of DD for both domestic and wild animals, as well as for farmers and veterinarians.

Despite the significant economic and welfare implications of this disease, numerous uncertainties persist regarding its cause, transmission, prevention, and treatment. Several factors contribute to the challenges associated with addressing this disease. The infection appears to involve multiple types of bacteria, including hundreds of species, namely of the genus *Treponema*, which have been isolated from lesions (Choi et al., 1997; Cruz et al., 2005; Dhawi et al., 2005; Evans et al., 2008; Krull et al., 2014; Walker et al., 1995; Zinicola et al., 2015). Additionally, the initial cultivation of these bacteria proved to be problematic (Evans et al., 2008), and the development of experimental infection models has been challenging (Gomez et al., 2012; Read and Walker, 1998), leaving the mechanisms of disease transmission largely unknown. However, recent advancements resulted in significant progress in identifying the pathogenic bacteria and detecting their presence in affected animals on the farm (Evans et al., 2012; Klitgaard et al., 2014; Krull et al., 2014; Zinicola et al., 2015).

Previous research has been conducted on endemically infected farms to study both farm and herd level risk factors and animal level risk factors associated with the occurrence of DD (Refaai et al., 2013; Rodríguez Lainz et al., 1996; Somers et al., 2005). Both risk factors, at the farm and individual animal levels, provide valuable insights for minimizing DD infection levels and understanding infection risk. The time between diagnosis and treatment, tie-stall housing, and the amount of concentrate fed are associated with the occurrence of DD (Weber et al., 2023). Hoof knives post-trimming were frequently contaminated with DD-associated treponeme DNA (Gillespie et al., 2020). Treponemes were identified on hoof trimming knives, and suggested a mode of transmission of DD in ruminants. The cleaning and disinfection of hoof trimming equipment between animals and between farms was emphasized to be particularly important, and may control the spread of DD (Sullivan et al., 2014). Another intriguing aspect, albeit less explored, is the presence of cattle variation in susceptibility to the disease. Studies have observed repeated infections in individual cattle within a herd, while others of the same breed, parity, and under identical conditions remained uninfected (Capion et al., 2012; Gomez et al., 2015; Laven, 1999). Identifying the reasons behind this individual variation in susceptibility could enhance our understanding of the disease and contribute to the development of effective prevention and treatment strategies.

Symptoms of Digital Dermatitis

Digital Dermatitis was initially documented in Italy in 1974 (Cheli and Mortellaro, 1974) and has since been reported in the majority of countries with dairy and beef industries (Demirkan et al., 2000; Holzhauer et al., 2006). The infection leads to ulcerative lesions along the coronary band (Cheli and Mortellaro, 1974) or the skin adjacent to the interdigital cleft (Holzhauer et al., 2008), resulting in pain and discomfort (Laven and Proven, 2000). Heifers with DD exhibited less time ruminating and more time inactive compared to heifers without lesions (Thomas et al., 2021). Döpfer et al. described the disease progression in four stages (Döpfer et al., 1997). The initial stage, known as M1, is characterized by a limited granulomatous area that is typically smaller than 2 cm diameter in size (Döpfer et al., 1997) and generally not painful (Holzhauer et al., 2008). The lesion then progresses to the M2 stage, characterized by classic ulceration (Döpfer et al., 1997). During this stage, the lesion is larger than 2 cm diameter and causes pain upon palpation (Holzhauer et al., 2008). As the lesion begins to heal, typically after topical treatment, a scab forms over the ulcerated area, resulting in the M3 stage (Döpfer et al., 1997). In certain cases, the lesion progresses to the M4 stage, known as the chronic stage (Döpfer et al., 1997), characterized by surface proliferation or hyperkeratosis. This stage is typically non-painful but infectious and can revert back to an active M2 lesion (Berry et al., 2012; Döpfer et al., 2012). Berry et al. introduced the M4.1 lesion type, which describes a chronic lesion (M4) with a small area displaying an active M1 lesion (Berry et al., 2012). The differences in rumination time and inactivity are associated with DD and different M-stages in feedlot cattle (Thomas et al., 2021). It is common to find cattle with lesions on both hind feet simultaneously such that in one herd, 51% of diameter have lesions on both feet, 22% on the left foot only, and 27% on the right foot only (Laven, 1999). The reason why some animals develop DD on only one hind foot, despite both feet being exposed to the same risk factors, remains unknown. Investigating the factors influencing foot-specific infection could contribute to understanding of DD and variation in susceptibility to the disease.

Causes of Digital Dermatitis

Spirochaetes belonging to the genus *Treponema* have consistently been found to be associated with DD lesions (Choi et al., 1997; Demirkan et al., 1999; Dhawi et al., 2005; Evans et al., 2009; Krull et al., 2014; Walker et al., 1995; Zinicola et al., 2015) and the bacteria are believed to be the primary microbial agents associated with DD pathogenesis (Evans et al., 2014). The presence of Spirochaetes deep within the lesions (Blowey et al., 1994) indicates that the bacteria are invasive to the epidermis and dermis rather than merely colonizing damaged tissue (Edwards et al., 2003). Treponemes that have been isolated from disease lesions show that cause tissue destruction in various hosts, including humans (Edwards et al., 2003). While treponemes are the most frequently detected bacteria in DD lesions, other types of bacteria have also been identified, including *Borrelia burgdorferi* (Blowey et al., 1994; Collighan and Woodward, 1997), *Bacteroides* and *Mycoplasma* species (Collighan and Woodward, 1997), *Campylobacter* species (Döpfer et al., 1997), and *Candidatus Amoebophilus asiaticus* (Zinicola et al., 2015).

In recent studies, efforts have been made to examine the microbiome of healthy skin and DD lesions at various stages of infection (Ariza et al., 2022; Bay et al., 2023; Caddey and De Buck, 2021; Krull et al., 2014; Zinicola et al., 2015). Bay et al. showed evidence of dysbiosis and differences in taxonomy and functional profiles in the bovine foot skin microbiome of clinically healthy animals that subsequently develop DD lesions (Bay et al., 2023). Krull et al. investigated the microbiota of lesions at seven different developmental stages and discovered a significant increase in the proportion of treponemes detected in lesion biopsies as the lesions progressed

(Krull et al., 2014). They also observed changes in the proportions of different treponeme species throughout lesion development. Similarly, Zinicola et al., who discovered distinct differences in the microbiomes of skin with active DD lesions (M1, M2, and M4.1) compared to those with inactive lesions (M3 and M4), as well as healthy skin (Zinicola et al., 2015). The microbiome of active, ulcerative DD lesions is predominantly composed of six groups of treponemes: *Treponema denticola*, *Treponema maltophilum*, *Treponema medium*, *Treponema putidum*, *Treponema phagedenis*, and *Treponema paraluiscuniculi* (Zinicola et al., 2015). Moreira et al. identified eleven different *Treponema* strains belonging to the six major phylotypes of *Treponema* in DD lesions. Furthermore, Moreira et al. identify *Dichelobacter nodosus* in DD lesions in almost half of biopsy specimens in areas with mild epithelial damage and together with *Treponema*. Ariza et al. confirmed the skin microbiota associated with DD lesions, dominated by *Treponema* spp., is different from the microbiota of healthy skin (Ariza et al., 2022). The diversity and structure of the microbiota in DD lesions did not change based on the footbath disinfectant or the individual topical antibiotic treatments. Microbiotas from proliferative lesions showed a different structure and diversity compared to non-proliferative lesions. Ariza et al. confirmed the role of *Treponema* spp. and emphasized the role of *Mycoplasmopsis* spp. in the onset of DD lesion (Ariza et al., 2022). Members of *Treponema*, *Mycoplasma*, *Porphyromonas*, and *Fusobacterium* were consistently identified in the majority of DD lesions, were the best genera at differentiating DD lesions from normal skin, and may have significant roles in DD pathogenesis (Caddey and De Buck, 2021). Early-stage lesions were associated with *T. medium*, *T. phagedenis*, and *P. levii* (Caddey et al., 2021). Caddey et al. suggested a core DD microbial group consisting of species of *Treponema*, *Fusobacterium*, *Porphyromonas*, and *Bacteroides* may be closely tied with the etiopathogenesis of DD (Caddey et al., 2021). These studies suggest that the variation in bacterial species observed in previous studies of DD lesions may be attributed to sampling at different stages of lesion development.

Environmental transmission of DD-associated treponemes has been hypothesized for years. Initially, these bacteria were challenging to detect in the environment. Evans et al. found no evidence of treponemes in dairy cow feces or environmental slurry samples (Evans et al., 2012). Technological advancements have enabled the detection of treponemes in both dairy cow feces and environmental slurry samples on infected farms. This was achieved by utilizing targeted deep-sequencing methods to detect trace amounts of bacterial RNA (Klitgaard et al., 2014).

Direct skin-to-skin transmission from infected to uninfected feet has been proposed as a potential route of infection (Evans et al., 2012). Another suggested mode of transmission is through hoof trimming tools (Sullivan et al., 2014). The gut has been suggested as a possible reservoir of DD bacteria. Shibahara et al. found similar spirochaetes responsible for simultaneous infections of bovine dysentery and DD in two Japanese dairy cows, indicating a potential connection (Shibahara et al., 2002). Evans et al. detected DD-related treponemes in both oral and rectal tissues of dairy cows on DD-affected farms (Evans et al., 2012). These findings were corroborated by Zinicola et al., who identified DD-associated treponemes in the rumen and fecal microbiomes (Zinicola et al., 2015). These results suggest that the gastrointestinal tract serves as a reservoir for DD-related treponemes (Zinicola et al., 2015), with feces and slurry acting as a mode of transmission between the reservoir and the site of infection (Evans et al., 2012).

Treatment of Digital Dermatitis

Treatment options for DD include systemic and topical antibiotics (Laven and Proven, 2000; Refaai et al., 2013; Walker et al., 1995). Currently, non-antimicrobial topical treatment agents are preferred (Jacobs et al., 2018; Laven and Logue, 2006; Moore et al., 2001; Paudyal et al., 2020). In herds with a high prevalence of DD, individual treatment is time-consuming and labor-intensive, prompting many farmers to opt for footbaths to prevent DD the entire herd (Laven and Proven, 2000).

Footbathing is widely employed as a herd-level prevention method for DD. There is significant variation in the types of products used, their concentrations, and application frequencies (Jacobs et al., 2019). A recent study examined 141 freestall farms and found that 87% of them utilized footbaths. Interestingly, each farm employed a range of 1 to 4 different products (Solano et al., 2015). The most commonly used products were copper sulfate ($CuSO_4$) and formalin. Unfortunately, the use of $CuSO_4$ raises environmental concerns due to soil accumulation (Epperson and Midla, 2007; Flemming and Trevors, 1989; Hoff et al., 1998), while formalin is known to be carcinogenic (Doane and Sarenbo, 2014). These drawbacks underscore the necessity to explore alternative options for footbath strategies.

Unfortunately, complete eradication of DD is rarely achieved, necessitating the repeated application of treatments to prevent the recurrence of infection (Laven and Logue, 2006). Dairy cows develop an antibody response against treponemes when infected with DD, yet this response does not appear sufficient to prevent subsequent infections, as some animals experience repeated infections (Demirkan et al., 1999; Gomez et al., 2014). One possible explanation for the challenges encountered in treating DD and preventing its recurrence is the presence of treponemes in the epithelium that have been found to exist in both encysted and spiral forms (Döpfer et al., 2012). It is conceivable that these encysted bacterial forms could persist deep within the lesions and lead to a recurrence of clinical symptoms at a later stage. However, further research is needed to elucidate the significance of the encysted form of bacteria and its response to DD treatments (Döpfer et al., 2012).

Costs of Digital Dermatitis

The majority of active DD lesions are reported to be painful (Laven and Proven, 2000; Manske et al., 2002; Somers et al., 2005). If left untreated, infected animals may experience lameness for a prolonged period of time (Frankena et al., 2009). Cows that suffer from sufficient pain to develop lameness also exhibit behavioral changes compared to healthy cows. These changes include increased lying time (Hassall et al., 1993; Juarez et al., 2003; Margerison et al., 2002; Walker et al., 2008) and reduced total feeding time (Almeida et al., 2008; Hassall et al., 1993; Juarez et al., 2003; Palmer et al., 2012). These behavior changes are expected to have an impact on productivity.

Digital dermatitis is the leading cause of infectious lameness in dairy cattle, resulting in ulcerative skin lesions that severely affect animal production and welfare. Lameness ranks as the second most significant health issue in dairy cattle in terms of production losses, preceded by mastitis and followed by Bovine viral diarrhea, and the most prominent with respect to welfare concerns (Bennett and IJpelaar, 2005). A comprehensive review examining lameness prevalence studies worldwide from 1993 to 2014 reported that, on average, 25-55% of dairy cattle are clinically lame (Cook, 2017; Solano et al., 2017; Underwood et al., 2015). Additionally, it was found that approximately 15% to 22% of cows have one or more DD lesions, and up to 94% of herds were affected by DD (Cook, 2017; Solano et al., 2017; Underwood et al., 2015).

Various studies examining the impact of DD on milk yield have yielded different findings. While DD is a major contributor to lameness, its association with a significant reduction in milk yield is not consistently observed. Amory et al. conducted a study in England and Wales and found that cows with DD did not exhibit decreased milk yield during infection (Amory et al., 2008). Instead, a slight increase in milk yield was observed after treatment (Amory et al., 2008). In two other studies, cows with DD showed slightly lower milk production, but the differences were not statistically significant (Argaez-Rodriguez and Hird, 1997; Hernandez et al., 2002). Warnick et al. examined two US dairy farms and found that cows with DD experienced a reduction in milk yield (Warnick et al., 2001). This reduction was less pronounced compared to cows affected by other causes of lameness (Warnick et al., 2001). In a study using a large dataset from Holstein cows on French farms, Relun et al. identified a small but statistically significant decrease in milk yield (<1 kg per day) associated with DD (Relun et al., 2013). Pavlenko et al. studied DD-affected Swedish Red and Swedish Holstein cows and found that they had significantly lower milk yield (5.5 kg energy-corrected milk per day) compared to healthy control cows (Pavlenko et al., 2011). Additionally, Gomez et al. discovered that the DD infection history of heifers influenced their milk yield during the first lactation (Gomez et al., 2015). Cows that experienced one or more DD infections before their first calving exhibited a reduction in milk yield of 199 kg or 335 kg over 305 days, respectively, compared to cows without a prior DD infection before calving (Gomez et al., 2015).

While it is established that lameness, in general, can have a negative impact on fertility (Alawneh et al., 2011; Melendez et al., 2003; Morris et al., 2011), there is limited research specifically investigating the relationship between DD infection and fertility changes. A study conducted on Mexican Holstein-Friesian cows revealed an extended calving to conception interval and increased days open in cows with DD compared to healthy cows (Argaez-Rodriguez and Hird, 1997). Similarly, Gomez et al. found that heifers with repeated DD infections prior to their first calving had a lower conception rate at first service and an increase in the number of days open compared to cows without any prior DD infection (Gomez et al., 2015). Zinicola et al. estimated that DD results in an annual economic loss of US $1.1 billion for dairy cows in the United States and European Union, assuming a DD prevalence rate of 25% (Zinicola et al., 2015).

Estimates of the cost per case of non-specific lameness ranged from $76 to $533, depending on the location of the study (Dolecheck and Bewley, 2018). The cost of DD over all combinations of parity group and incidence timing, regardless of incidence likelihood, was $64 ± 24 for DD using a stochastic simulation model (Dolecheck et al., 2019). The cost of a lame cow and a DD-affected cow was $345.09 ± 9.43 and $439.70 ± 11.22 per year on average, respectively using a bioeconomic model (Robcis et al., 2023a). Charfeddine and Pérez-Cabal (2017) noted that it costs $53 to $402 per cow affected with DD (Charfeddine and Pérez-Cabal, 2017), while Cha et al. noted that it costs $132.96 per DD case (Cha et al., 2010). The disease surpasses all other causes and incurs an additional cost of approximately $100 (Robcis et al., 2023a). For each subsequent week that a cow remained lame, the financial burden on the farmer increased by $13.26 per week (Robcis et al., 2023a). The cost of a mild lesion was $53 per affected cow and year, whereas the cost of a severe lesion ranged from $402 per affected cow and year for DD (Charfeddine and Pérez-Cabal, 2017). It was recommended that 95.5% of DD cases be treated and the main contributor to the total cost per case is treatment cost for DD (42%) (Cha et al., 2010). The current estimates of cost for DD do not account for the cost of transmission.

Digital dermatitis is a contagious disease, and disregarding the costs associated with additional cases resulting from an initial case overlooks the complete economic impact of the disorder (Dolecheck et al., 2019). Döpfer et al. determined the reproductive ratio of DD to range from 0.5

to 3.3, depending on the preventive measures implemented (Döpfer et al., 2012). This cost category could constitute a significant portion of the overall cost per DD case, contingent upon the herd-level prevention strategies in place. Due to the complexity of modeling contagious diseases, this cost category is typically omitted from cost estimates (Dolecheck et al., 2019).

The impact of DD varies depending on the stage of the disease. During acute ulcerative phases, cows experience pain that affects their behavior, leading to reduced milk yield and fertility. Chronically affected animals can perpetuate the disease within the herd, acting as reservoirs for pathogens and contributing to the establishment of an endemic state of infection. This can ultimately lead to premature culling of cows and additional costs for disease control and eradication (Biemans et al., 2018; Döpfer et al., 2012). Consequently, determining the disease stage is crucial for informing prevention and control strategies and accurately assessing the economic implications of DD (Gomez et al., 2015).

Detection of Digital Dermatitis

Currently, various methods are employed to gather data related to DD. Herd mobility scoring serves as a widely used screening tool to identify lame animals, followed by clinical investigation of feet to determine the underlying cause of lameness and implement the appropriate treatment. Mobility scoring is subjective and susceptible to biases from both intra- and inter-observer variations (Archer et al., 2010; Whay, 2002). The presence of DD lesions does not always correlate with lameness, potentially resulting in underreporting (Krull et al., 2016). This approach to detecting DD is labor-intensive, expensive, and causes stress to the animals, as it requires thorough cleaning and examination (Relun et al., 2011; Stokes et al., 2012a). Locomotion scoring may overlook cases with less obvious lesions (Solano et al., 2017). The subjective approach of locomotion scoring is often employed to evaluate lameness in dairy cattle, but it is time-consuming and may not be sufficiently sensitive to detect hoof lesions (Flower and Weary, 2006; Sprecher et al., 1997; Tadich et al., 2010). The standard method for detecting and classifying DD involves clinical observation of hooves lifted in a foot trimming chute. Researchers have proposed alternative approaches for monitoring DD, including visual inspections during milking routines in the parlor (Relun et al., 2011; Stokes et al., 2012a), as well as exploring the use of infrared thermography for lesion identification (Stokes et al., 2012b). Unfortunately, these alternatives generally exhibit lower diagnostic accuracy compared to the standard method of foot inspection (Orsel et al., 2018).

Benefits of Early Detection

Currently, visual inspection of lifted hooves is considered the gold standard for DD detection in dairy cattle (Afonso et al., 2021). However, this method is labor-intensive and subjective in nature. The diagnosis of DD is commonly performed during routine claw trimming by lifting the foot in a claw-trimming chute (Holzhauer et al., 2006; Manske et al., 2002; Thomsen et al., 2008b). Alternative methods have been explored to directly inspect feet affected by DD while the cow is standing in the milking parlor. These methods can involve using a swiveling mirror or a rigid borescope for assessment (Relun et al., 2011; Rodriguez-Lainz et al., 1998; Thomsen et al., 2008a).

As previously stated, DD can be present even without noticeable lameness, and severe lameness may not manifest until the DD lesion has progressed (Laven and Proven, 2000; Stokes et al., 2009). Hence, early detection of DD holds significant value in preventing further advancement and enabling timely and effective treatment (Döpfer et al., 2012; Shearer and Hernandez, 2000). There is a need for reliable, practical, and non-invasive screening methods that can rapidly and frequently assess the presence of DD in real time at both the foot and cow level.

Early detection allows for prompt intervention and treatment, minimizing the progression and severity of the disease (Main et al., 2010; Robcis et al., 2023b). This increases the chances of successful treatment outcomes and faster recovery for affected animals. Treating DD in its early stages can be more effective and less costly compared to treating advanced chronic cases. Early detection can help avoid potential complications and more extensive treatments, resulting in cost savings for farmers (Charfeddine and Pérez-Cabal, 2017; Dolecheck and Bewley, 2018). Early detection of DD can help mitigate negative effects and alleviate pain, leading to improved milk yield and overall performance of the animals (Charfeddine and Pérez-Cabal, 2017; Dolecheck and Bewley, 2018; Dolecheck et al., 2019; Galligan, 2006; Laven et al., 2008; Liang et al., 2017; Robcis et al., 2023b; Whay and Shearer, 2017). It early detection of DD enables farmers to implement appropriate biosecurity measures to control the spread of DD (Brennan and Christley, 2012; Oliveira et al., 2017).

Overview of Computer Vision

Deep learning algorithms draw inspiration from the remarkable capabilities of the human brain, which utilizes an extensive network of interconnected neurons to perform complex tasks such as speaking, moving, thinking, and seeing (Goodfellow et al., 2016). In the realm of artificial intelligence, deep learning architectures typically comprise multiple layers of artificial neurons, forming what is known as "deep" neural networks. The fundamental building block of these networks is the neuron (Goodfellow et al., 2016). Neurons are typically organized into layers, where each neuron has a specific function while learning distinct parameters. Through a so-called “feed-forward” process, data input is continually transformed as it passes through the layers, ultimately mapping the data transformations to a desired output.

To optimize the performance of the network, the parameters of neurons and the weights of connections between neurons are updated through a learning process known as backpropagation (Goodfellow et al., 2016). During backpropagation, the network computes the error, which represents the difference between the observed and predicted outcome. This error is then propagated backward through the network using parameter gradients. These gradients are employed to update the parameters of individual neurons and the weights of connections between neurons, aiming to minimize the observed error in the output. The network learns the optimal parameters and weights for each neuron to accurately predict the desired outcome. Through the iterative error minimization process, the network's parameters tend to converge within the network architecture. This convergence results in the network making precise and accurate predictions when presented with new data points or images.

Deep neural networks are constructed using various types of layers, each serving a distinct purpose (Oliveira et al., 2021). Commonly employed layer types include fully connected (or dense), convolutional, deconvolutional, pooling, recurrent, and others. Fully connected (or dense) layers consist of neurons that apply a single activation function to transform the summed weighted input from a node to an output value.

Convolutional layers utilize so-called “kernels” to perform convolutions on input images, where each node convolves its kernel with the input image, generating a convolved image as output (Oliveira et al., 2021). These layers can also modify image scales using strides, resulting in smaller output images, or employ transpose convolutions, often referred to as deconvolutional layers, to generate larger output images (Oliveira et al., 2021).

Pooling and upsampling layers serve similar functions of processing images but with opposite effects. Pooling layers aggregate input image values into smaller images, while upsampling layers interpolate values from smaller images to create larger images (Oliveira et al., 2021). These layers are often used in conjunction with convolutional layers to build encoders and decoders, such as for image segmentation (Oliveira et al., 2021).

Deep learning algorithms, a subset of machine learning techniques, have emerged as a powerful tool in various computer vision (CV) applications. These algorithms extend traditional artificial neural networks (ANN) and demonstrate their ability to learn more abstract representations of input data by constructing a hierarchical structure of nested concepts (Goodfellow et al., 2016). These nested concepts, also known as hidden layers, contribute to the creation of highly complex models with numerous trainable parameters. The training of such models became feasible due to significant advancements in the field of deep learning including the availability of extensive datasets, the utilization of data augmentation techniques, and the progress made in ANN, such as the development of stochastic gradient descent for learning optimization, the introduction of new activation functions such as the rectified linear unit (ReLU), the implementation of regularization techniques, and the efficient utilization of graphics processing units (GPU) (Goodfellow et al., 2016; LeCun et al., 2015).

Object detection is a prominent research area within CV and has generated a substantial body of literature. Deep learning techniques have consistently demonstrated their effectiveness in object detection, and they can be broadly classified into two categories: region proposal-based methods and regression-based methods (Li et al., 2020).

Region proposal-based methods involve proposing potential object regions in an image and classifying them into one or more categories. On the other hand, regression-based methods approach object detection as a regression problem, determining object coordinates within images. One of the early successful region proposal-based methods was the region-based

Convolutional Neural Network (CNN) (Girshick et al., 2016). This algorithm generates a multitude of region proposals by mean of “selective search” (Uijlings et al., 2013). It then extracts deep convolutional features from these regions using a CNN and employs a support vector machine to label object candidates with PASCAL VOC classes. Deep learning approaches in object detection have significantly advanced the field, and these region proposal-based methods have played a crucial role in achieving state-of-the-art object predictions.

The initial concept of region-based CNN has undergone significant enhancements, resulting in the development of new approaches. One such improvement is Fast R-CNN, which introduced the concept of feeding the input image directly into the CNN to generate a convolutional feature maps for identifying region proposals (Girshick, 2015). This differs from the previous approach of feeding region proposals into the CNN. In the case of Faster R-CNN, a separate network was used to predict the region proposals instead of a selective search algorithm applied to the feature map (Ren et al., 2015). This modification improved efficiency and accuracy in generating region proposals. Cascaded R-CNN introduced a sequential framework comprising cascaded object detectors employing R-CNN-like networks (Cai and Vasconcelos, 2018). Each detector in the sequential framework exhibits increasing accuracy and is trained with progressively higher Intersection over Union (IoU) thresholds to evaluate the overlap between a predicted bounding box and a ground truth bounding box. This cascaded approach enhances selectivity against close false positives. These innovations have contributed to the refinement of region-based CNNs, offering improved performance, increased efficiency, and enhanced selectivity in object detection tasks.

The concept of regression-based networks for object detection was initially introduced in OverFeat (Sermanet et al., 2014), and subsequently gained popularity with the emergence of

YOLO (Redmon et al., 2016) and RetinaNet (Lin et al., 2018). YOLO is a regression network that employs a single CNN backbone to predict bounding boxes with label probabilities in a unified evaluation process (Redmon et al., 2016). This architecture employs a grid structure to divide the input image space and locate objects at the center of each grid cell. Multiple bounding boxes and class probabilities are generated for each cell, and the network training involves maximizing these values for every cell.

In YOLOv2, several modifications were proposed to enhance precision compared to the original YOLO architecture (Redmon and Farhadi, 2017). These include the replacement of fully connected layers with anchor boxes that are more robust and have fewer parameters, for bounding box prediction. This approach is similar to what was proposed in the Single Shot Multibox detector (SSD) (Liu et al., 2016). YOLOv3 further improved the use of anchor boxes by employing dimension clusters to predict bounding boxes and logistic classifiers to generate class probabilities for each bounding box (Redmon and Farhadi, 2018).

RetinaNet, on the other hand, integrates ResNet (He et al., 2015) and Feature Pyramid Networks (Lin et al., 2017) as backbone networks for feature extraction. These features are concatenated with two task-specific subnetworks dedicated to classification and bounding box regression. The training schema involves utilizing focal loss, resulting in a single-stage regression approach that outperforms Faster R-CNN.

For scenarios where reduced computational resources are available, Tiny-YOLO (Redmon et al., 2016) was introduced as a smaller version of YOLO. This network architecture requires less time for training and testing but generally exhibits lower accuracy compared to the original models. These advancements in regression-based networks have significantly contributed to the field of object detection, offering improved performance and varying trade-offs between accuracy and computational efficiency.

Applications of Computer Vision

In the field of animal and veterinary sciences, sensors have been widely employed for various purposes, although most applications have been limited to experimental settings, with only a few developed for commercial farm use. These applications encompass a wide range of functions, including assessing the composition of beef cuts, identifying and tracking live animals, monitoring behavior, and measuring relevant phenotypes such as body weight, condition score, and gait (Fernandes et al., 2020). This technology holds great potential for precision livestock farming and high-throughput phenotyping applications. The continuous measurement of phenotypic traits through CV has the potential to reduce management costs, enhance decision-making in livestock operations, and create new opportunities for selective breeding (Fernandes et al., 2020). While applications of CV are currently a burgeoning research area, there are already commercial products available. The remaining challenges necessitate further research to achieve the successful development of autonomous solutions capable of delivering critical information.

Object detection tasks have been employed in various livestock animal studies for multiple purposes. As previously mentioned, the primary objective of these tasks is to identify and locate one or more objects within an image. In animal science research, object detection algorithms have predominantly been applied to animal detection, with a particular focus on swine (Cowton et al., 2019; Lee et al., 2019; Psota et al., 2019; Seo et al., 2020). However, other applications have been explored, including the detection of lameness (Kang et al., 2020) and DD (Cernek et al., 2020) in dairy cattle.

Lee et al. proposed a hybrid approach that combines image processing techniques with a deep learning algorithm (Lee et al., 2019). They utilized a Gaussian Mixture Model to detect the moving frame within a 24-hour timeframe. Tiny YOLOv3 was employed to identify individual pigs in each selected frame. Finally, they applied Otsus' method for pig body segmentation, which utilizes automatic image thresholding to calculate pig size. Instead of relying on end-to-end deep learning methods, the authors introduced these hybrid models and demonstrated that they offer faster analysis speed and satisfactory accuracy when implemented on single-board GPUs, such as the Jeston TX2. This approach was compared to deep learning strategies like Mask R-CNN.

In animal science, YOLO and Faster R-CNN emerged as the main algorithms used for object detection. The specific deep learning architectures employed varied across the different studies. Popular architectures for deep learning included ResNet, Xception, VGG16, Inception, and Darknet. Barbedo et al. utilized an Unmanned Aerial Vehicle (UAV) to capture aerial images of beef farms and evaluated 15 CNN architectures for animal detection (Barbedo et al., 2017). The authors concluded that many CNN architectures exhibited robustness in detecting cattle on farms, with NASNet and Xception networks demonstrating particularly impressive performance. Geffen et al. employed Faster R-CNN with a ResNet-101 backbone network to detect and count hens per cage, achieving a detection accuracy of 89.6% (Geffen et al., 2020). Ye et al. utilized R-CNN and reported an accuracy of 98.6% in predicting the stunned condition of broilers (Ye et al., 2020).

Cernek et al. implemented YOLOv2 using RGB images and achieved an accuracy of 88% in detecting cows with DD (Cernek et al., 2020). These promising results demonstrated the significant potential of CV in accurately identifying animals affected by DD, thereby reducing its prevalence and improving animal welfare. Kang et al. developed a lameness scoring system for dairy cows using the Receptive Field Block Net Single Shot Detector deep learning network (Kang et al., 2020). They achieved a mean average precision of 87.0% in accurately locating cow hooves within video footage. The located legs were subsequently utilized as input in a proposed algorithm called the "supporting phase", which involved calculating the difference between the hoof lifting time and the hoof load time.

CHAPTER SUMMARY

Chapter 2, 3, and 4, discuss real-time detection of DD in dairy cows using computer vision algorithms. Chapter 2 performs a comparative analysis of different computer vision algorithms for the real-time detection of DD in dairy cows. Chapter 3 applies the prior information in Chapter 2 and performs benchmarking analysis of the best computer vision algorithm on an edge device. Chapter 4 extends the benchmarking analysis on an edge device in Chapter 3 to various cloud computing platforms and updates the computer vision algorithm to the latest version. The chapters demonstrate the feasibility of implementing a portable solution for early detection of DD in precision farming.

Chapter 5 expands on real-time detection of DD in dairy cows to pododermatitis in canines. Additionally, the study examines the differences in performance using single or multiple labelers with various levels of instruction. The custom models, edge devices, and cloud deployments are a significant step towards the integration of CV algorithms in veterinary medicine and signifies forward progress in the real-time detection of health outcomes in agriculture.

Chapter 6 and Chapter 7 describe data analysis tools for dairy records. Chapter 6 demonstrates the extraction of health and production data from DairyComp, then data cleaning and analysis using a menu-driven point-and-click approach via the DairyCoPilot app. Chapter 7 demonstrates real-time detection of DD on a mobile app, data analysis including trends of M-stages over time, and predictions and assignments of Cow Types via the DairyCheckPlus app. Chapter 8 serves as a placeholder for research in progress. It describes a comparative analysis of different mobile apps for real-time detection of DD on both Android and iOS devices.

The tools are the first step in developing advanced freely available tools for rapid analysis of records from dairy farms that are translatable to all other animal species and healthcare settings. Additionally, the proposed tools are user-friendly and can be download to a mobile device or can be accessed via the web. The applications are an important step in converting farm records into data assets for more customized decision-making processes by informed consultants in the life sciences.

Chapter 9 discusses the subsequent workflow for object detection to train and deploy a model that can accurately detect and localize objects of interest within images or videos. This approach can save significant training time and resources while adapting the model to new tasks. Chapter 10 provides a discussion extending the research to further analysis and future studies. The focus of the chapter describes the major opportunities and challenges for deployment of computer vision algorithms and object detection models in precision farming and veterinary medicine.

REFERENCES

Afonso, J.S., Oikonomou, G., Carter, S., Clough, H.E., Griffiths, B.E., Rushton, J., 2021. Diagnosis of Bovine Digital Dermatitis: Exploring the Usefulness of Indirect ELISA. Frontiers in Veterinary Science 8.

Alawneh, J.I., Laven, R.A., Stevenson, M.A., 2011. The effect of lameness on the fertility of dairy cattle in a seasonally breeding pasture-based system. Journal of Dairy Science 94, 5487–5493. https://doi.org/10.3168/jds.2011-4395

Almeida, P.E., Weber, P.S.D., Burton, J.L., Zanella, A.J., 2008. Depressed DHEA and increased sickness response behaviors in lame dairy cows with inflammatory foot lesions. Domestic Animal Endocrinology 34, 89–99. https://doi.org/10.1016/j.domaniend.2006.11.006

Amory, J.R., Barker, Z.E., Wright, J.L., Mason, S.A., Blowey, R.W., Green, L.E., 2008. Associations between sole ulcer, white line disease and digital dermatitis and the milk yield of 1824 dairy cows on 30 dairy cow farms in England and Wales from February 2003–November 2004. Preventive Veterinary Medicine 83, 381–391. https://doi.org/10.1016/j.prevetmed.2007.09.007

Archer, S., Bell, N., Huxley, J., 2010. Lameness in UK dairy cows: a review of the current status. In Practice 32, 492–504.

Argaez-Rodriguez, F., Hird, D., 1997. deAnda J., Read D., RodriguezLainz A. Papillomatous digital dermatitis on a commercial dairy farm in Mexicali, Mexico: Incidence and effect on reproduction and milk production. Prev. Vet. Med 32, 275–286.

Ariza, J.M., Döpfer, D., Oberle, K., Anklam, K., Labrut, S., Bareille, N., Relun, A., Guatteo, R., 2022. Skin microbiota dynamics of dairy cows affected by digital dermatitis. https://doi.org/10.1101/2019.12.20.882787

Barbedo, J.G.A., Gomes, C.C.G., Cardoso, F.F., Domingues, R., Ramos, J.V., McManus, C.M., 2017. The use of infrared images to detect ticks in cattle and proposal of an algorithm for quantifying the infestation. Veterinary Parasitology 235, 106–112. https://doi.org/10.1016/j.vetpar.2017.01.020

Bay, V., Gillespie, A., Ganda, E., Evans, N.J., Carter, S.D., Lenzi, L., Lucaci, A., Haldenby, S., Barden, M., Griffiths, B.E., Sánchez-Molano, E., Bicalho, R., Banos, G., Darby, A., Oikonomou, G., 2023. The bovine foot skin microbiota is associated with host genotype and the development of infectious digital dermatitis lesions. Microbiome 11, 4. https://doi.org/10.1186/s40168-022-01440-7

Bennett, R., IJpelaar, J., 2005. Updated Estimates of the Costs Associated with Thirty Four Endemic Livestock Diseases in Great Britain: A Note. Journal of Agricultural Economics 56, 135–144. https://doi.org/10.1111/j.1477-9552.2005.tb00126.x

Berry, S.L., Read, D.H., Famula, T.R., Mongini, A., Döpfer, D., 2012. Long-term observations on the dynamics of bovine digital dermatitis lesions on a California dairy after topical treatment with lincomycin HCl. The Veterinary Journal, Special Issue:Bovine Lameness 193, 654–658. https://doi.org/10.1016/j.tvjl.2012.06.048

Biemans, F., Bijma, P., Boots, N.M., de Jong, M.C.M., 2018. Digital Dermatitis in dairy cattle: The contribution of different disease classes to transmission. Epidemics 23, 76–84. https://doi.org/10.1016/j.epidem.2017.12.007

Blowey, R.W., Carter, S.D., White, A.G., Barnes, A., 1994. Borrelia burgdorferi infections in UK cattle: a possible association with digital dermatitis. Veterinary Record 135, 577–578.

Brennan, M.L., Christley, R.M., 2012. Biosecurity on cattle farms: a study in north-west England. PloS one 7, e28139.

Bruijnis, M., Beerda, B., Hogeveen, H., Stassen, E., 2012. Assessing the welfare impact of foot disorders in dairy cattle by a modeling approach. Animal 6, 962–970.

Caddey, B., De Buck, J., 2021. Meta-Analysis of Bovine Digital Dermatitis Microbiota Reveals Distinct Microbial Community Structures Associated With Lesions. Frontiers in Cellular and Infection Microbiology 11.

Caddey, B., Orsel, K., Naushad, S., Derakhshani, H., De Buck, J., 2021. Identification and Quantification of Bovine Digital Dermatitis-Associated Microbiota across Lesion Stages in Feedlot Beef Cattle. mSystems 6, 10.1128/msystems.00708-21. https://doi.org/10.1128/msystems.00708-21

Cai, Z., Vasconcelos, N., 2018. Cascade R-CNN: Delving Into High Quality Object Detection. Presented at the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 6154–6162.

Capion, N., Boye, M., Ekstrøm, C.T., Jensen, T.K., 2012. Infection dynamics of digital dermatitis in first-lactation Holstein cows in an infected herd. Journal of Dairy Science 95, 6457–6464. https://doi.org/10.3168/jds.2012-5335

Cernek, P., Bollig, N., Anklam, K., Döpfer, D., 2020. Hot topic: Detecting digital dermatitis with computer vision. Journal of Dairy Science 103, 9110–9115. https://doi.org/10.3168/jds.2019-17478

Cha, E., Hertl, J.A., Bar, D., Gröhn, Y.T., 2010. The cost of different types of lameness in dairy cows calculated by dynamic programming. Preventive Veterinary Medicine 97, 1–8. https://doi.org/10.1016/j.prevetmed.2010.07.011

Charfeddine, N., Pérez-Cabal, M.A., 2017. Effect of claw disorders on milk production, fertility, and longevity, and their economic impact in Spanish Holstein cows. Journal of Dairy Science 100, 653–665. https://doi.org/10.3168/jds.2016-11434

Cheli, R., Mortellaro, C., 1974. Digital dermatitis in cattle. Proc 8th Int Meet Dis Cattle 8, 208–213.

Choi, B.-K., Nattermann, H., Grund, S., Haider, W., GöBel, U.B., 1997. Spirochetes from Digital Dermatitis Lesions in Cattle Are Closely Related to Treponemes Associated with Human Periodontitis. International Journal of Systematic and Evolutionary Microbiology 47, 175–181. https://doi.org/10.1099/00207713-47-1-175

Clegg, S.R., Mansfield, K.G., Newbrook, K., Sullivan, L.E., Blowey, R.W., Carter, S.D., Evans, N.J., 2014. Isolation of Digital Dermatitis Treponemes from Hoof Lesions in Wild North American Elk (Cervus elaphus) in Washington State, USA. Journal of Clinical Microbiology 53, 88–94. https://doi.org/10.1128/jcm.02276-14

Collighan, R.J., Woodward, M.J., 1997. Spirochaetes and other bacterial species associated with bovine digital dermatitis. FEMS Microbiology Letters 156, 37–41. https://doi.org/10.1111/j.1574-6968.1997.tb12702.x

Cook, N.B., 2017. A Life Cycle, Lesion Oriented Approach to Lameness Control.

Cowton, J., Kyriazakis, I., Bacardit, J., 2019. Automated Individual Pig Localisation, Tracking and Behaviour Metric Extraction Using Deep Learning. IEEE Access 7, 108049–108060. https://doi.org/10.1109/ACCESS.2019.2933060

Cruz, C.E.F., Pescador, C.A., Nakajima, Y., Driemeier, D., 2005. Immunopathological investigations on bovine digital epidermitis. Veterinary Record 157, 834–840. https://doi.org/10.1136/vr.157.26.834

Demirkan, I., Murray, R., Carter, S., 2000. Skin diseases of the bovine digit associated with lameness. Veterinary Bulletin 70, 149–171.

Demirkan, I., Walker, R.L., Murray, R.D., Blowey, R.W., Carter, S.D., 1999. Serological Evidence of Spirochaetal Infections Associated with Digital Dermatitis in Dairy Cattle. The Veterinary Journal 157, 69–77. https://doi.org/10.1053/tvjl.1998.0276

Dhawi, A., Hart, C.A., Demirkan, I., Davies, I.H., Carter, S.D., 2005. Bovine digital dermatitis and severe virulent ovine foot rot: A common spirochaetal pathogenesis. The Veterinary Journal 169, 232–241. https://doi.org/10.1016/j.tvjl.2004.01.029

Doane, M., Sarenbo, S., 2014. Exposure of farm laborers and dairy cattle to formaldehyde from footbath use at a dairy farm in New York State. Science of The Total Environment 487, 65–71. https://doi.org/10.1016/j.scitotenv.2014.04.007

Dolecheck, K., Bewley, J., 2018. Animal board invited review: Dairy cow lameness expenditures, losses and total cost. Animal 12, 1462–1474. https://doi.org/10.1017/S1751731118000575

Dolecheck, K.A., Overton, M.W., Mark, T.B., Bewley, J.M., 2019. Use of a stochastic simulation model to estimate the cost per case of digital dermatitis, sole ulcer, and white line disease by parity group and incidence timing. Journal of Dairy Science 102, 715–730. https://doi.org/10.3168/jds.2018-14901

Döpfer, D., Anklam, K., Mikheil, D., Ladell, P., 2012. Growth curves and morphology of three Treponema subtypes isolated from digital dermatitis in cattle. The Veterinary Journal, Special Issue:Bovine Lameness 193, 685–693. https://doi.org/10.1016/j.tvjl.2012.06.054

Döpfer, Dörte, Holzhauer, M., Boven, M. van, 2012. The dynamics of digital dermatitis in populations of dairy cattle: Model-based estimates of transition rates and implications for control. The Veterinary Journal, Special Issue:Bovine Lameness 193, 648–653. https://doi.org/10.1016/j.tvjl.2012.06.047

Döpfer, D., Huurne, A.A.H.M. ter, Cornelisse, J.L., Asten, A.J.A.M. van, Koopmans, A., Meijer, F.A., Schukken, Y.H., Szakáll, I., Klee, W., Bosma, R.B., 1997. Histological and bacteriological evaluation of digital dermatitis in cattle, with special reference to spirochaetes and Campylobacter faecalis. Veterinary Record 140, 620–623. https://doi.org/10.1136/vr.140.24.620

Duncan, J.S., Angell, J.W., Carter, S.D., Evans, N.J., Sullivan, L.E., Grove-White, D.H., 2014. Contagious ovine digital dermatitis: An emerging disease. The Veterinary Journal 201, 265–268. https://doi.org/10.1016/j.tvjl.2014.06.007

Edwards, A.M., Dymock, D., Jenkinson, H.F., 2003. From tooth to hoof: treponemes in tissue-destructive diseases. Journal of Applied Microbiology 94, 767–780. https://doi.org/10.1046/j.1365-2672.2003.01901.x

Epperson, B., Midla, L., 2007. Copper Sulfate for Footbaths - Issues and Alternatives.

Evans, N.J., Blowey, R.W., Timofte, D., Isherwood, D.R., Brown, J.M., Murray, R., Paton, R.J., Carter, S.D., 2011. Association between bovine digital dermatitis treponemes and a range of 'non-healing' bovine hoof disorders. Veterinary Record 168, 214–214. https://doi.org/10.1136/vr.c5487

Evans, N.J., Brown, J.M., Demirkan, I., Murray, R.D., Vink, W.D., Blowey, R.W., Hart, C.A., Carter, S.D., 2008. Three unique groups of spirochetes isolated from digital dermatitis lesions in UK cattle. Veterinary Microbiology 130, 141–150. https://doi.org/10.1016/j.vetmic.2007.12.019

Evans, N.J., Brown, J.M., Demirkan, I., Singh, P., Getty, B., Timofte, D., Vink, W.D., Murray, R.D., Blowey, R.W., Birtles, R.J., Hart, C.A., Carter, S.D., 2009. Association of Unique, Isolated Treponemes with Bovine Digital Dermatitis Lesions. Journal of Clinical Microbiology 47, 689–696. https://doi.org/10.1128/JCM.01914-08

Evans, N.J., Brown, J.M., Scholey, R., Murray, R.D., Birtles, R.J., Hart, C.A., Carter, S.D., 2014. Differential inflammatory responses of bovine foot skin fibroblasts and keratinocytes to digital dermatitis treponemes. Veterinary Immunology and Immunopathology 161, 12–20. https://doi.org/10.1016/j.vetimm.2014.05.005

Evans, N.J., Timofte, D., Isherwood, D.R., Brown, J.M., Williams, J.M., Sherlock, K., Lehane, M.J., Murray, R.D., Birtles, R.J., Anthony Hart, C., Carter, S.D., 2012. Host and environmental reservoirs of infection for bovine digital dermatitis treponemes. Veterinary Microbiology 156, 102–109. https://doi.org/10.1016/j.vetmic.2011.09.029

Fernandes, A.F.A., Dórea, J.R.R., Rosa, G.J. de M., 2020. Image Analysis and Computer Vision Applications in Animal Sciences: An Overview. Frontiers in Veterinary Science 7.

Flemming, C.A., Trevors, J.T., 1989. Copper toxicity and chemistry in the environment: a review. Water Air Soil Pollut 44, 143–158. https://doi.org/10.1007/BF00228784

Flower, F.C., Weary, D.M., 2006. Effect of Hoof Pathologies on Subjective Assessments of Dairy Cow Gait. Journal of Dairy Science 89, 139–146. https://doi.org/10.3168/jds.S0022-0302(06)72077-X

Frankena, K., Somers, J.G.C.J., Schouten, W.G.P., van Stek, J.V., Metz, J.H.M., Stassen, E.N., Graat, E.A.M., 2009. The effect of digital lesions and floor type on locomotion score in Dutch dairy cows. Preventive Veterinary Medicine, Special Section: Schwabe Symposium 2007 88, 150–157. https://doi.org/10.1016/j.prevetmed.2008.08.004

Galligan, D., 2006. Economic Assessment of Animal Health Performance. Veterinary Clinics: Food Animal Practice 22, 207–227. https://doi.org/10.1016/j.cvfa.2005.11.007

Geffen, O., Yitzhaky, Y., Barchilon, N., Druyan, S., Halachmi, I., 2020. A machine vision system to detect and count laying hens in battery cages. Animal 14, 2628–2634. https://doi.org/10.1017/S1751731120001676

Gillespie, A.V., Carter, S.D., Blowey, R.W., Staton, G.J., Evans, N.J., 2020. Removal of bovine digital dermatitis-associated treponemes from hoof knives after foot-trimming: a disinfection field study. BMC Veterinary Research 16, 330. https://doi.org/10.1186/s12917-020-02552-8

Girshick, R., 2015. Fast R-CNN, in: 2015 IEEE International Conference on Computer Vision (ICCV). Presented at the 2015 IEEE International Conference on Computer Vision (ICCV), pp. 1440–1448. https://doi.org/10.1109/ICCV.2015.169

Girshick, R., Donahue, J., Darrell, T., Malik, J., 2016. Region-Based Convolutional Networks for Accurate Object Detection and Segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence 38, 142–158. https://doi.org/10.1109/TPAMI.2015.2437384

Gomez, A., Anklam, K.S., Cook, N.B., Rieman, J., Dunbar, K.A., Cooley, K.E., Socha, M.T., Döpfer, D., 2014. Immune response against Treponema spp. and ELISA detection of digital dermatitis. Journal of Dairy Science 97, 4864–4875. https://doi.org/10.3168/jds.2013-7616

Gomez, A., Cook, N.B., Bernardoni, N.D., Rieman, J., Dusick, A.F., Hartshorn, R., Socha, M.T., Read, D.H., Döpfer, D., 2012. An experimental infection model to induce digital dermatitis infection in cattle. Journal of Dairy Science 95, 1821–1830. https://doi.org/10.3168/jds.2011-4754

Gomez, A., Cook, N.B., Socha, M.T., Döpfer, D., 2015. First-lactation performance in cows affected by digital dermatitis during the rearing period. Journal of Dairy Science 98, 4487–4498. https://doi.org/10.3168/jds.2014-9041

Goodfellow, I., Bengio, Y., Courville, A., 2016. Deep Learning, Illustrated edition. ed. The MIT Press, Cambridge, Massachusetts.

Hassall, S., Ward, W., Murray, R., 1993. Effects of lameness on the behaviour of cows during the summer. The veterinary record 132, 578–580.

He, K., Zhang, X., Ren, S., Sun, J., 2015. Deep Residual Learning for Image Recognition. https://doi.org/10.48550/arXiv.1512.03385

Hernandez, J., Shearer, J.K., Webb, D.W., 2002. Effect of lameness on milk yield in dairy cows. Journal of the American Veterinary Medical Association 220, 640–644. https://doi.org/10.2460/javma.2002.220.640

Hoff, B., Boermans, H.J., Baird, J.D., 1998. Retrospective study of toxic metal analyses requested at a veterinary diagnostic toxicology laboratory in Ontario (1990-1995). Can Vet J 39, 39–43.

Holzhauer, M., Bartels, C.J.M., Döpfer, D., van Schaik, G., 2008. Clinical course of digital dermatitis lesions in an endemically infected herd without preventive herd strategies. The Veterinary Journal 177, 222–230. https://doi.org/10.1016/j.tvjl.2007.05.004

Holzhauer, M., Hardenberg, C., Bartels, C.J.M., Frankena, K., 2006. Herd- and Cow-Level Prevalence of Digital Dermatitis in The Netherlands and Associated Risk Factors. Journal of Dairy Science 89, 580–588. https://doi.org/10.3168/jds.S0022-0302(06)72121-X

Jacobs, C., Beninger, C., Hazlewood, G.S., Orsel, K., Barkema, H.W., 2019. Effect of footbath protocols for prevention and treatment of digital dermatitis in dairy cattle: A systematic review and network meta-analysis. Preventive Veterinary Medicine 164, 56–71. https://doi.org/10.1016/j.prevetmed.2019.01.011

Jacobs, C., Orsel, K., Mason, S., Barkema, H.W., 2018. Comparison of effects of routine topical treatments in the milking parlor on digital dermatitis lesions. Journal of Dairy Science 101, 5255–5266. https://doi.org/10.3168/jds.2017-13984

Juarez, S.T., Robinson, P.H., DePeters, E.J., Price, E.O., 2003. Impact of lameness on behavior and productivity of lactating Holstein cows. Applied Animal Behaviour Science 83, 1–14. https://doi.org/10.1016/S0168-1591(03)00107-2

Kang, X., Zhang, X.D., Liu, G., 2020. Accurate detection of lameness in dairy cattle with computer vision: A new and individualized detection strategy based on the analysis of the supporting phase. Journal of Dairy Science 103, 10628–10638. https://doi.org/10.3168/jds.2020-18288

Klitgaard, K., Nielsen, M.W., Ingerslev, H.-C., Boye, M., Jensen, T.K., 2014. Discovery of Bovine Digital Dermatitis-Associated Treponema spp. in the Dairy Herd Environment by a Targeted Deep-Sequencing Approach. Applied and Environmental Microbiology 80, 4427–4432. https://doi.org/10.1128/AEM.00873-14

Krull, A.C., Shearer, J.K., Gorden, P.J., Cooper, V.L., Phillips, G.J., Plummer, P.J., 2014. Deep Sequencing Analysis Reveals Temporal Microbiota Changes Associated with Development of

Bovine Digital Dermatitis. Infection and Immunity 82, 3359–3373. https://doi.org/10.1128/iai.02077-14

Krull, A.C., Shearer, J.K., Gorden, P.J., Scott, H.M., Plummer, P.J., 2016. Digital dermatitis: Natural lesion progression and regression in Holstein dairy cattle over 3 years. Journal of Dairy Science 99, 3718–3731. https://doi.org/10.3168/jds.2015-10535

Laven, R., 2001. Control of digital dermatitis in cattle. In Practice 23, 336–341. https://doi.org/10.1136/inpract.23.6.336

Laven, R., 1999. The environment and digital dermatitis. Cattle Practice 7, 349–354.

Laven, R., Lawrence, K., Weston, J., Dowson, K., Stafford, K., 2008. Assessment of the duration of the pain response associated with lameness in dairy cows, and the influence of treatment. New Zealand Veterinary Journal 56, 210–217. https://doi.org/10.1080/00480169.2008.36835

Laven, R.A., Logue, D.N., 2006. Treatment strategies for digital dermatitis for the UK. The Veterinary Journal 171, 79–88. https://doi.org/10.1016/j.tvjl.2004.08.009

Laven, R.A., Proven, M.J., 2000. Use of an antibiotic footbath in the treatment of bovine digital dermatitis. Veterinary Record 147, 503–506. https://doi.org/10.1136/vr.147.18.503

LeCun, Y., Bengio, Y., Hinton, G., 2015. Deep learning. Nature 521, 436–444. https://doi.org/10.1038/nature14539

Lee, S., Ahn, H., Seo, J., Chung, Y., Park, D., Pan, S., 2019. Practical Monitoring of Undergrown Pigs for IoT-Based Large-Scale Smart Farm. IEEE Access 7, 173796–173810. https://doi.org/10.1109/ACCESS.2019.2955761

Li, K., Wan, G., Cheng, G., Meng, L., Han, J., 2020. Object detection in optical remote sensing images: A survey and a new benchmark. ISPRS Journal of Photogrammetry and Remote Sensing 159, 296–307. https://doi.org/10.1016/j.isprsjprs.2019.11.023

Liang, D., Arnold, L.M., Stowe, C.J., Harmon, R.J., Bewley, J.M., 2017. Estimating US dairy clinical disease costs with a stochastic simulation model. Journal of Dairy Science 100, 1472–1486. https://doi.org/10.3168/jds.2016-11565

Lin, T.-Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S., 2017. Feature Pyramid Networks for Object Detection. https://doi.org/10.48550/arXiv.1612.03144

Lin, T.-Y., Goyal, P., Girshick, R., He, K., Dollár, P., 2018. Focal Loss for Dense Object Detection. https://doi.org/10.48550/arXiv.1708.02002

Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.-Y., Berg, A.C., 2016. SSD: Single Shot MultiBox Detector, in: Leibe, B., Matas, J., Sebe, N., Welling, M. (Eds.), Computer Vision – ECCV 2016, Lecture Notes in Computer Science. Springer International Publishing, Cham, pp. 21–37. https://doi.org/10.1007/978-3-319-46448-0_2

Main, D.C.J., Barker, Z.E., Leach, K.A., Bell, N.J., Whay, H.R., Browne, W.J., 2010. Sampling strategies for monitoring lameness in dairy cattle. Journal of Dairy Science 93, 1970–1978. https://doi.org/10.3168/jds.2009-2500

Manske, T., Hultgren, J., Bergsten, C., 2002. Prevalence and interrelationships of hoof lesions and lameness in Swedish dairy cows. Preventive Veterinary Medicine 54, 247–263. https://doi.org/10.1016/S0167-5877(02)00018-1

Margerison, J.K., Winkler, B., Stephens, G., 2002. The effect locomotion score and lameness and on dry matter intake and behaviour in dairy cattle. Proceedings of the British Society of Animal Science 2002, 199–199. https://doi.org/10.1017/S1752756200008553

Melendez, P., Bartolome, J., Archbald, L.F., Donovan, A., 2003. The association between lameness, ovarian cysts and fertility in lactating dairy cows. Theriogenology 59, 927–937. https://doi.org/10.1016/S0093-691X(02)01152-4

Moore, D.A., Berry, S.L., Truscott, M.L., Koziy, V., 2001. Efficacy of a nonantimicrobial cream administered topically for treatment of digital dermatitis in dairy cattle. Journal of the American Veterinary Medical Association 219, 1435–1438. https://doi.org/10.2460/javma.2001.219.1435

Morris, M.J., Kaneko, K., Walker, S.L., Jones, D.N., Routly, J.E., Smith, R.F., Dobson, H., 2011. Influence of lameness on follicular growth, ovulation, reproductive hormone concentrations and estrus behavior in dairy cows. Theriogenology 76, 658–668. https://doi.org/10.1016/j.theriogenology.2011.03.019

Oliveira, D.A.B., Pereira, L.G.R., Bresolin, T., Ferreira, R.E.P., Dorea, J.R.R., 2021. A review of deep learning algorithms for computer vision systems in livestock. Livestock Science 253, 104700. https://doi.org/10.1016/j.livsci.2021.104700

Oliveira, V.H.S., Sørensen, J.T., Thomsen, P.T., 2017. Associations between biosecurity practices and bovine digital dermatitis in Danish dairy herds. Journal of Dairy Science 100, 8398–8408. https://doi.org/10.3168/jds.2017-12815

Orsel, K., Plummer, P., Shearer, J., De Buck, J., Carter, S.D., Guatteo, R., Barkema, H.W., 2018. Missing pieces of the puzzle to effectively control digital dermatitis. Transboundary and Emerging Diseases 65, 186–198. https://doi.org/10.1111/tbed.12729

Palmer, M.A., Law, R., O'Connell, N.E., 2012. Relationships between lameness and feeding behaviour in cubicle-housed Holstein–Friesian dairy cows. Applied Animal Behaviour Science 140, 121–127. https://doi.org/10.1016/j.applanim.2012.06.005

Paudyal, S., Manriquez, D., Velasquez, A., Shearer, J.K., Plummer, P.J., Melendez, P., Callan, R.J., Sorge, U.S., Bothe, H., Velez, J., Pinedo, P.J., 2020. Efficacy of non-antibiotic treatment options for digital dermatitis on an organic dairy farm. The Veterinary Journal 255, 105417. https://doi.org/10.1016/j.tvjl.2019.105417

Pavlenko, A., Bergsten, C., Ekesbo, I., Kaart, T., Aland, A., Lidfors, L., 2011. Influence of digital dermatitis and sole ulcer on dairy cow behaviour and milk production. Animal 5, 1259–1269.

Plummer, P.J., Krull, A., 2017. Clinical Perspectives of Digital Dermatitis in Dairy and Beef Cattle. Veterinary Clinics of North America: Food Animal Practice 33, 165–181. https://doi.org/10.1016/j.cvfa.2017.02.002

Psota, E.T., Mittek, M., Pérez, L.C., Schmidt, T., Mote, B., 2019. Multi-Pig Part Detection and Association with a Fully-Convolutional Network. Sensors 19. https://doi.org/10.3390/s19040852

Read, D., Walker, R., 1998. Experimental transmission of papillomatous digital dermatitis (footwarts) in dairy cattle, in: Proceedings of the 10th International Symposium on Lameness in Ruminants, Casino Lucerne, Switzerland. pp. 7–10.

Redmon, J., Divvala, S., Girshick, R., Farhadi, A., 2016. You Only Look Once: Unified, Real-Time Object Detection. arXiv:1506.02640 [cs].

Redmon, J., Farhadi, A., 2018. YOLOv3: An Incremental Improvement. arXiv:1804.02767 [cs].

Redmon, J., Farhadi, A., 2017. YOLO9000: Better, Faster, Stronger. Presented at the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 7263–7271.

Refaai, W., Van Aert, M., Abd El-Aal, A.M., Behery, A.E., Opsomer, G., 2013. Infectious diseases causing lameness in cattle with a main emphasis on digital dermatitis (Mortellaro disease). Livestock Science, Lameness and claw lesions in sows, cows and small ruminants 156, 53–63. https://doi.org/10.1016/j.livsci.2013.06.004

Relun, A., Guatteo, R., Roussel, P., Bareille, N., 2011. A simple method to score digital dermatitis in dairy cows in the milking parlor. Journal of Dairy Science 94, 5424–5434. https://doi.org/10.3168/jds.2010-4054

Relun, A., Lehebel, A., Chesnin, A., Guatteo, R., Bareille, N., 2013. Association between digital dermatitis lesions and test-day milk yield of Holstein cows from 41 French dairy farms. Journal of Dairy Science 96, 2190–2200. https://doi.org/10.3168/jds.2012-5934

Ren, S., He, K., Girshick, R., Sun, J., 2015. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks, in: Advances in Neural Information Processing Systems. Curran Associates, Inc.

Robcis, R., Ferchiou, A., Berrada, M., Ndiaye, Y., Herman, N., Lhermie, G., Raboisson, D., 2023a. Cost of lameness in dairy herds: An integrated bioeconomic modeling approach. Journal of Dairy Science 106, 2519–2534. https://doi.org/10.3168/jds.2022-22446

Robcis, R., Ferchiou, A., Berrada, M., Raboisson, D., 2023b. Management of Digital Dermatitis in Dairy Herds: Optimization and Time Allocation. Animals (Basel) 13, 1988. https://doi.org/10.3390/ani13121988

Rodríguez Lainz, A., Hird, D.W., Walker, R.L., Read, D.H., 1996. Papillomatous digital dermatitis in 458 dairies. Journal of the American Veterinary Medical Association 209, 1464–1467.

Rodriguez-Lainz, A., Melendez-Retamal, P., Hird, D.W., Read, D.H., 1998. Papillomatous digital dermatitis in Chilean dairies and evaluation of a screening method. Preventive Veterinary Medicine 37, 197–207. https://doi.org/10.1016/S0167-5877(98)00091-9

Seo, J., Ahn, H., Kim, D., Lee, S., Chung, Y., Park, D., 2020. EmbeddedPigDet—Fast and Accurate Pig Detection for Embedded Board Implementations. Applied Sciences 10. https://doi.org/10.3390/app10082878

Sermanet, P., Eigen, D., Zhang, X., Mathieu, M., Fergus, R., LeCun, Y., 2014. OverFeat: Integrated Recognition, Localization and Detection using Convolutional Networks. https://doi.org/10.48550/arXiv.1312.6229

Shearer, J.K., Hernandez, J., 2000. Efficacy of Two Modified Nonantibiotic Formulations (Victory) for Treatment of Papillomatous Digital Dermatitis in Dairy Cows. Journal of Dairy Science 83, 741–745. https://doi.org/10.3168/jds.S0022-0302(00)74936-8

Shibahara, T., Ohya, T., Ishii, R., Ogihara, Y., Maeda, T., Ishikawa, Y., Kadota, K., 2002. Concurrent spirochaetal infections of the feet and colon of cattle in Japan. Australian veterinary journal 80, 497–502.

Solano, L., Barkema, H.W., Jacobs, C., Orsel, K., 2017. Validation of the M-stage scoring system for digital dermatitis on dairy cows in the milking parlor. Journal of Dairy Science 100, 1592–1603. https://doi.org/10.3168/jds.2016-11365

Solano, L., Barkema, H.W., Pajor, E.A., Mason, S., LeBlanc, S.J., Zaffino Heyerhoff, J.C., Nash, C.G.R., Haley, D.B., Vasseur, E., Pellerin, D., Rushen, J., de Passillé, A.M., Orsel, K., 2015. Prevalence of lameness and associated risk factors in Canadian Holstein-Friesian cows housed in freestall barns. Journal of Dairy Science 98, 6978–6991. https://doi.org/10.3168/jds.2015-9652

Somers, J.G.C.J., Frankena, K., Noordhuizen-Stassen, E.N., Metz, J.H.M., 2005. Risk factors for digital dermatitis in dairy cows kept in cubicle houses in The Netherlands. Preventive Veterinary Medicine 71, 11–21. https://doi.org/10.1016/j.prevetmed.2005.05.002

Sprecher, D.J., Hostetler, D.E., Kaneene, J.B., 1997. A lameness scoring system that uses posture and gait to predict dairy cattle reproductive performance. Theriogenology 47, 1179–1187. https://doi.org/10.1016/S0093-691X(97)00098-8

Stokes, J., Leach, K., Main, D., Whay, H., 2009. Can locomotion scoring be used to identify cows with digital dermatitis? Cattle Lameness Conference, Sutton Bonnington, Nottingham, 25 March.

Stokes, J.E., Leach, K.A., Main, D.C.J., Whay, H.R., 2012a. The reliability of detecting digital dermatitis in the milking parlour. The Veterinary Journal, Special Issue:Bovine Lameness 193, 679–684. https://doi.org/10.1016/j.tvjl.2012.06.053

Stokes, J.E., Leach, K.A., Main, D.C.J., Whay, H.R., 2012b. An investigation into the use of infrared thermography (IRT) as a rapid diagnostic tool for foot lesions in dairy cattle. The Veterinary Journal, Special Issue:Bovine Lameness 193, 674–678. https://doi.org/10.1016/j.tvjl.2012.06.052

Sullivan, L.E., Blowey, R.W., Carter, S.D., Duncan, J.S., Grove-White, D.H., Page, P., Iveson, T., Angell, J.W., Evans, N.J., 2014. Presence of digital dermatitis treponemes on cattle and sheep hoof trimming equipment. Veterinary Record 175, 201–201. https://doi.org/10.1136/vr.102269

Sullivan, L.E., Carter, S.D., Blowey, R., Duncan, J.S., Grove-White, D., Evans, N.J., 2013. Digital dermatitis in beef cattle. Veterinary Record 173, 582–582. https://doi.org/10.1136/vr.101802

Sullivan, L.E., Evans, N.J., Clegg, S.R., Carter, S.D., Horsfield, J.E., Grove- White, D., Duncan, J.S., 2015. Digital dermatitis treponemes associated with a severe foot disease in dairy goats. Veterinary Record 176, 283–283. https://doi.org/10.1136/vr.102858

Tadich, N., Flor, E., Green, L., 2010. Associations between hoof lesions and locomotion score in 1098 unsound dairy cows. The Veterinary Journal 184, 60–65. https://doi.org/10.1016/j.tvjl.2009.01.005

Thomas, A.D., Orsel, K., Cortés, J.A., Pajor, E.A., 2021. Impact of digital dermatitis on feedlot cattle behaviour. Applied Animal Behaviour Science 244, 105468. https://doi.org/10.1016/j.applanim.2021.105468

Thomsen, P.T., Klaas, I.C., Bach, K., 2008a. Short Communication: Scoring of Digital Dermatitis During Milking as an Alternative to Scoring in a Hoof Trimming Chute. Journal of Dairy Science 91, 4679–4682. https://doi.org/10.3168/jds.2008-1342

Thomsen, P.T., Sørensen, J.T., Ersbøll, A.K., 2008b. Evaluation of Three Commercial Hoof-Care Products Used in Footbaths in Danish Dairy Herds. Journal of Dairy Science 91, 1361–1365. https://doi.org/10.3168/jds.2007-0820

Uijlings, J.R.R., van de Sande, K.E.A., Gevers, T., Smeulders, A.W.M., 2013. Selective Search for Object Recognition. Int J Comput Vis 104, 154–171. https://doi.org/10.1007/s11263-013-0620-5

Underwood, W.J., Blauwiekel, R., Delano, M.L., Gillesby, R., Mischler, S.A., Schoell, A., 2015. Biology and Diseases of Ruminants (Sheep, Goats, and Cattle). Laboratory Animal Medicine 623–694. https://doi.org/10.1016/B978-0-12-409527-4.00015-8

USDA, 2007. Dairy 2007, Part I: Reference of Dairy Cattle Health and Management Practices in the United States, 2007.

Walker, R.L., Read, D.H., Loretz, K.J., Nordhausen, R.W., 1995. Spirochetes isolated from dairy cattle with papillomatous digital dermatitis and interdigital dermatitis. Veterinary Microbiology 47, 343–355. https://doi.org/10.1016/0378-1135(95)00114-X

Walker, S.L., Smith, R.F., Jones, D.N., Routly, J.E., Dobson, H., 2008. Chronic stress, hormone profiles and estrus intensity in dairy cattle. Hormones and Behavior 53, 493–501. https://doi.org/10.1016/j.yhbeh.2007.12.003

Warnick, L.D., Janssen, D., Guard, C.L., Gröhn, Y.T., 2001. The Effect of Lameness on Milk Production in Dairy Cows. Journal of Dairy Science 84, 1988–1997. https://doi.org/10.3168/jds.S0022-0302(01)74642-5

Weber, J., Becker, J., Syring, C., Ruiters, M.W., Locher, I., Bayer, M., Schüpbach-Regula, G., Steiner, A., 2023. Farm-level risk factors for digital dermatitis in dairy cows in mountainous regions. Journal of Dairy Science 106, 1341–1350. https://doi.org/10.3168/jds.2022-22243

Whay, H., 2002. Locomotion scoring and lameness detection in dairy cattle. In Practice 24, 444–449. https://doi.org/10.1136/inpract.24.8.444

Whay, H.R., Shearer, J.K., 2017. The Impact of Lameness on Welfare of the Dairy Cow. Veterinary Clinics: Food Animal Practice 33, 153–164. https://doi.org/10.1016/j.cvfa.2017.02.008

Ye, C., Yousaf, K., Qi, C., Liu, C., Chen, K., 2020. Broiler stunned state detection based on an improved fast region-based convolutional neural network algorithm. Poultry Science 99, 637–646. https://doi.org/10.3382/ps/pez564

Zinicola, M., Lima, F., Lima, S., Machado, V., Gomez, M., Döpfer, D., Guard, C., Bicalho, R., 2015. Altered Microbiomes in Bovine Digital Dermatitis Lesions, and the Gut as a Pathogen Reservoir. PLOS ONE 10, e0120504. https://doi.org/10.1371/journal.pone.0120504

CHAPTER 2 COMPARATIVE ANALYSIS OF COMPUTER VISION ALGORITHMS FOR THE REAL-TIME DETECTION OF DIGITAL DERMATITIS IN DAIRY COWS

ABSTRACT

Digital dermatitis (DD) is a bovine claw disease responsible for ulcerative lesions on the coronary band of the hoof. DD is associated with massive herd outbreaks of lameness and influences cattle welfare and production. Early detection of DD can lead to prompt treatment and decrease lameness. Computer vision (CV) provides a unique opportunity to improve early detection. The study aims to train and compare applications for the real-time detection of DD in dairy cows. Nine CV models were trained for detection and scoring, compared using performance metrics and inference time, and the best model was automated for real-time detection using images and video.

Images were collected from commercial dairy farms while facing the interdigital space on the plantar surface of the foot. Images were scored for M-stages of DD by a trained investigator using the M-stage DD classification system. Two sets of images were compiled: the first dataset (Dataset 1) containing 1,177 M0/M4H and 1,050 M2/M2P images and the second dataset (Dataset 2) containing 240 M0, 17 M2, 51 M2P, 114 M4H, and 108 M4P images. Models were trained to detect and score DD lesions and compared for precision, recall, and mean average precision (mAP) in addition to inference time in frame per second (FPS).

Seven of the nine CV models performed well compared to the ground truth of labeled images using Dataset 1. The six models, Faster R-CNN, Cascade R-CNN, YOLOv3, Tiny YOLOv3, YOLOv4, Tiny YOLOv4, and YOLOv5s achieved an mAP between 0.964 to 0.998, whereas the other two models, SSD and SSD Lite, yielded an mAP of 0.371 and 0.387 respectively. Overall,

YOLOv4, Tiny YOLOv4, and YOLOv5s outperformed all other models with almost perfect precision, perfect recall, and a higher mAP. Tiny YOLOv4 outperformed all other models with respect to inference time at 333 FPS, followed by YOLOv5s at 133 FPS and YOLOv4 at 65 FPS. YOLOv4 and Tiny YOLOv4 performed better than YOLOv5s compared to the ground truth using Dataset 2. YOLOv4 and Tiny YOLOv4 yielded a similar mAP of 0.896 and 0.895, respectively. However, Tiny YOLOv4 achieved both higher precision and recall compared to YOLOv4.

Finally, Tiny YOLOv4 was able to detect DD lesions on a commercial dairy farm with high performance and speed. The proposed CV tool can be used for early detection and prompt treatment of DD in dairy cows. This result is a step towards applying CV algorithms to veterinary medicine and implementing real-time DD detection on dairy farms.

INTRODUCTION

Digital dermatitis (DD) is the most prevalent bovine infectious claw disease worldwide (Refaai et al., 2013; Logue et al., 2005; Teixeira et al., 2010). The disease is responsible for painful circumscribed ulcerative lesions of the coronary band of the hoof in both dairy and beef cattle (Cheli and Mortellaro, 1974; Döpfer et al., 1997). Massive herd outbreaks of lameness are caused by DD, influencing cattle welfare and production (Cramer et al., 2009; Krull et al., 2016; Jacobs et al., 2019). In addition, the disease results in major losses for the cattle industry due to decreased milk production, decreased fertility rate, and increased premature culling. (Vanhoudt et al., 2019). Early detection can lead to prompt treatment and decrease lameness within the herd (Schulz et al., 2016). Lesions are detected through visual inspection and scored using the M-stage classification system by trained investigators (Döpfer et al., 1997; Berry et al., 2012). Early identification and prompt intervention requires extensive employee training in an industry with a high employee turnover rate (Döpfer et al., 2012; Shearer and Hernandez, 2000).

Recent advances in computational power and state-of-the-art algorithms have made object detection and classification possible by applying machine learning techniques (Vaidya and Paunwala, 2019; Wu et al., 2020). Such methods for animal detection have been used to assess the health and well-being of livestock in cattle and small ruminants (Alsaaod et al., 2014; Byrne et al., 2018, 2017; Gomes et al., 2016; Hovinen et al., 2008; Martins et al., 2013; Metzner et al., 2015; Salau et al., 2017; Scoley et al., 2018; Yang et al., 2018; Zaninelli et al., 2018). However, applications for the classification of health events are still rare in veterinary medicine and food production (Vinicki et al., 2018). There are many unexplored applications in veterinary medicine and an untapped potential for machine learning algorithms to predict relevant biological outcomes (Liakos et al., 2018). Thus, machine learning provides a unique opportunity to improve early detection of DD and optimized treatment plans for cattle.

Computer vision (CV) can be used to perform object detection and calculate the associated class probabilities from a series of images or videos (Szeliski, 2022). Object detection locates the presence of an object with a bounding box, class label, and class probability (Zhao et al., 2019). The input is an image with one or more objects and the output is an image or video with zero or more class detections. The bounding box is defined by x- and y- coordinates of the center, width, height, and a class label for each object. The conditional class probability is the probability that the detected object belongs to a particular class (Zhao et al., 2019; Wu et al., 2020; Sharma and Mir, 2020).

Object detectors are mainly divided into two-stage object detectors and one-stage object detectors. The first stage of two-stage networks identifies region proposals, or subsets of the image that might contain an object. The second stage classifies the objects within the region proposals (Wu et al., 2020; Sharma and Mir, 2020). One-stage networks produce predictions for regions across the entire image using anchor boxes, and the predictions are decoded to generate the final bounding boxes for the objects. The approach skips the region proposal stage and runs detection directly over a dense sampling of possible locations (Wu et al., 2020; Sharma and Mir, 2020). Single-stage networks can be much faster than two-stage networks, but they may not reach the same level of accuracy, especially for images containing small objects.

One-stage object detectors such as You Only Look Once (YOLO) and Single Shot MultiBox Detector (SSD) in addition to two-stage object detectors such as Faster and Cascade region-based convolutional neural networks (R-CNNs) are commonly used, because of higher speed and higher accuracy compared to other computational approaches (Soviany and Ionescu, 2018). Such models are used for real-time detection of health outcomes to generate bounding boxes and class probabilities from labeled objects in images (Ünver and Ayan, 2019; Cao et al., 2017; Al-antari et al., 2020). The current study compares and contrasts five versions of YOLO: YOLOv3, Tiny YOLOv3, YOLOv4, Tiny YOLOv4, and YOLOv5s for differences in prediction performance in addition to the SSD, SSD Lite, Faster R-CNN, and Cascade R-CNN (Redmon and Farhadi, 2018; Adarsh et al., 2020; Bochkovskiy et al., 2020; Liu et al., 2016; Ren et al., 2016; Cai and Vasconcelos, 2017; ultralytics, 2022).

YOLO models have several advantages over other object detectors. First, YOLO can recognize multiple objects in a single frame (Redmon et al., 2016). It is extremely fast, because it uses a single network evaluation to simultaneously predict multiple bounding boxes and class probabilities, where other systems require thousands of network evaluations for a single image (Redmon et al., 2016). YOLO looks at the entire image during training and test time to implicitly encode contextual information about classes including their appearance and generalize representations of the objects (Redmon et al., 2016).

YOLO has achieved high performance on a high-end graphics processing unit (GPU), but does not perform fast enough on portable devices, because of the high memory requirements (Redmon, 2018; Redmon and Farhadi, 2018). Consequently, lightweight network architectures are used for constrained environments such as portable devices resulting in reduced model size and computation cost (Redmon, 2018). Tiny YOLO is the compressed version of YOLO with a simpler network structure and a reduced number of corresponding parameters (Adarsh et al., 2020). For real-time object detection, Tiny YOLO is the better option compared to YOLO, where speed is prioritized over precision or accuracy (Adarsh et al., 2020). The speed of Tiny YOLOv4 is approximately eight times faster than YOLOv4 (Alexey, 2021; Bochkovskiy et al., 2020; Techzizou, 2021). Furthermore, the speed of YOLOv5s is approximately two-and-a-half times faster than YOLOv4 (Kin-Yiu, 2020). However, the accuracy for YOLOv4 is 33% better than Tiny YOLOv4 tested on the MS COCO dataset (Alexey, 2021; Bochkovskiy et al., 2020; Techzizou, 2021). Moreover, the accuracy for YOLOv4 is better than 11% better than YOLOv5s on the MS COCO dataset using the same image and batch size (Kin-Yiu, 2020).

A YOLOv2 model was previously trained on a database of more than 3,500 DD lesion images to detect M-stages of DD (Cernek et al., 2020). The resulting predicting accuracy was 71% and the agreement between the human investigator and CV model was quantified as “moderate” by Cohen’s kappa during internal validation (Cernek et al., 2020). In the external validation, the YOLOv2 model detected DD with an accuracy of 88% and agreement between a human investigator and CV model was quantified as “fair” by Cohen’s kappa (Cernek et al., 2020).

The YOLOv3 model builds on the YOLOv2 framework by making detections at three different scales and predicts ten times more bounding boxes to improve performance on smaller objects (Redmon and Farhadi, 2018). YOLOv4 improves on the average precision (AP) by 10% and frames per second (FPS) by 12% compared to YOLOv3 on the COCO dataset (Bochkovskiy et al., 2020). For the purpose of our use case, different state-of-the-art CV models are trained using the same datasets to determine the best performing approach for fast and accurate object detection. Model performance is characterized by common benchmarking measures for speed and accuracy.

The study aims to design, develop, and implement an application for the real-time detection of DD in dairy cattle with the intention of minimizing the adverse effects of DD and lameness in all cattle by means of early detection and prompt treatments. Nine state-of-the-art CV models are trained for detection and scoring of DD, compared for performance, and implemented for real-time detection. Such tools have not yet been implemented and will help dairy farmers and producers to improve DD prevention strategies for early intervention as well as helping increase cattle welfare and production. We hope the proposed tool can be employed in combination with current practice for prevention of DD in dairy cattle. Automated detection of DD on streaming video can be used to generate treatment lists in real-time for dairy and beef cattle. Ultimately, the workflow can be deployed as a web-based application for the real-time detection of DD lesions. Such approaches can help identify high risk cattle for DD. Monitoring herds with endemic DD for changes in lesion prevalence or severity and classifying cattle based on lesion stages can improve on-farm decision-making processes (Plummer and Krull, 2017).

MATERIALS AND METHODS

Data Collection and Image Labeling

We used two different datasets for the study: the first dataset (Dataset 1) of images was used to compare the performance of all eight CV models and a second dataset (Dataset 2) of images was used to evaluate the performance of the best models for real-world application. Dataset 1 contained a collection of 2,227 images of cattle feet from midwestern United States commercial dairy farms and images were labeled for M-stages of DD (Cernek et al., 2020). Images were generated using GoPro Hero 5 Black cameras to collect MP4 video recordings of the plantar aspect of standing feet of dairy cows. Additionally, traditional and cell phone cameras were used to obtain JPG images from other studies in Wisconsin; Manitoba, Canada; and Utrecht, The Netherlands. All images were collected between June 2018 and March 2019. A rich, diverse library of images was compiled from various settings and scenarios including images of feet lifted in stand-up restraining chutes and cattle standing in the cattle housing area, in automated milking systems, or in rotary milking parlors. Dataset 2 contained a collection of 409 images of cattle feet from milking parlors in Wisconsin and images were labeled for M-stages of DD. As previously described, GoPro Hero 5 Black cameras were used to take MP4 video recordings of the backside of the hind feet at claw level. Images were taken during May 2021 of hooves lifted in the cattle housing area, automated milking system, or exiting a disinfecting foot bath.

Images were scored for M-stages of DD by a trained investigator using the M-stage DD classification system. The classification system describes various clinical stages over the course of the disease based on morphological observations between healthy (M0), active (M2), and chronic (M4) with hyperkeratosis (H) or proliferation (P) stages (Döpfer et al., 1997; Berry et al., 2012). For Dataset 1, the images were combined and reclassified into M0/M4H lesions versus active ulcerative M2/M2P lesions. For Dataset 2, the images were classified into the different stages of DD as follows: M0, M2, M2P, M4H, and M4P.

Overall, the first library (Dataset 1) of a single foot per image and two class labels included 1,177 M0/M4H class labels (936 JPG images from MP4 video recordings and 241 JPG images from camera images) and 1,050 M2 class labels (660 screenshot JPG images from MP4 video recordings and 390 JPG images from camera images). The second library (Dataset 2) of multiple feet per image and five labels consisted of 240 M0, 17 M2, 51 M2P, 114 M4H, and 108 M4P class labels (409 screenshot JPG images from MP4 video recordings). The associated annotations for the bounding boxes and class labels were generated in LabelImg version 1.8.1 using Python (Lin, 2018). Labels were converted from Pascal VOC format for R-CNN and SSD models to YOLO format for YOLO models (Tashiev, 2022).

Model Training and Performance Evaluation

The datasets were split with 90% of the images placed into a training set and 10% of the images placed into a testing set at random. For Dataset 1, the testing set contained 222 images including 117 M0/M4H and 105 M2 class labels. For Dataset 2, the testing set contained 40 images consisting of 20 M0, 4 M2, 4 M2P, 14 M4H, and 10 M4P class labels. All models were trained and tested using Google Colab and a 12GB NVIDIA Tesla K80 GPU with a batch size of 64 and subdivision size of 16 ("Google Colaboratory," 2022). The training weights were initialized using convolutional weights that were pre-trained on Imagenet ("ImageNet," 2020; Krizhevsky et al., 2012). The eight different computer vision models were trained until the model converged or the maximum number of batches was reached.

The R-CNNs were performed using ResNet-50 as the backbone feature extractor and input images were resized to 640 by 640 pixels. Faster R-CNN improved on Fast R-CNN by utilizing a Region Proposal Network to generate high-quality region proposals used by Fast R-CNN for detection (Ren et al., 2016). Cascade R-CNN extended Faster R-CNN where a sequence of detectors was trained stage by stage to leverage the output of a previous detector to train the next detector (Cai and Vasconcelos, 2017).

SSD uses MobileNet V1 as the backbone feature extractor and SSDLite network uses MobileNet V2. The SSD network accepts images that are resized to an input resolution of 300 by 300 pixels (Liu et al., 2016). SSDLite is similar to SSD, but implements depthwise-separable convolutions rather than regular convolution layers, which increases the speed compared to the regular SSD making it perfectly suited for use on mobile devices (Fan et al., 2018).

The YOLOv3 network uses a variant of the custom deep architecture Darknet, originally a 53-layer network trained on Imagenet with a 53-layer stack of neurons added for the task of detection to result in a total of a 106-layer fully convolutional underlying architecture (Redmon et al., 2016; Redmon and Farhadi, 2016, 2018). The YOLOv4 and YOLOv5 network uses CSPDarknet53 as the backbone, a spatial pyramid pooling additional module, PANet path-aggregation neck, and a YOLOv3 head (Bochkovskiy et al., 2020; ultralytics, 2022). The YOLOv4 network size height and width is 416 by 416 pixels with a learning rate of 0.001 for a maximum number of 4,000 batches. The YOLOv5 network size height and width is also 416 by 416 pixels with an initial learning rate of 0.001 for a maximum number of 1,500 epochs. The best training weights were selected for model evaluation.

For speed, frames per second (FPS) were used to evaluate the time to predict the bounding boxes and class labels of the located objects in an image. For detection and identification, precision, recall, and mean average precision (mAP) at intersection over union (IOU) of 0.5 were used as performance measures to compare between the predictions made by the CV models, and labels made by a trained investigator (the so-called 'ground truth'). Lastly, the best CV model was validated for detection of each M-stage on video files from dataset 2.

RESULTS

Tiny YOLOv4 outperformed all other models with a speed of 333 FPS (Figure 1). The closest model was YOLOv5s at 133 FPS, then SSD and SSD Lite at approximately 100 FPS, followed by YOLOv4 at 65 FPS. Cascade R-CNN and Faster R-CNN performed worse than the seven other models with a speed less than 10 FPS. For the purpose of real-time detection on a portable, stand-alone device, Tiny YOLOv4 was the best model for the use case, based on the size and speed of the implementation.

Seven of the nine CV models performed well compared to the ground truth using the first dataset with a single object per image and two class labels for object detection (Figure 2 - 4). In addition, the seven models outperformed SSD and SSD Lite. The six models achieved an mAP between 0.964 to 0.998, whereas SSD and SSD Lite yielded an mAP of 0.371 and 0.387 respectively. Cascade R-CNN significantly outperformed Faster R-CNN for all three performance measures including precision, recall, and mAP. YOLOv4, Tiny YOLOv4, and YOLOv5s outperformed YOLOv3 and Tiny YOLOv3 for all three performance measures of precision, recall, and mAP. In addition, YOLOv3 had a lower precision but higher recall compared to Tiny-YOLOv3, but still had a slightly higher mAP. YOLOv4, Tiny YOLOv4, and YOLOv5s achieved similar precision and recall, however Tiny YOLOv4 and YOLOv5s achieved the highest mAP of 0.998 and 0.995, respectively. Overall, YOLOv4, Tiny YOLOv4, and YOLOv5s outperformed all other models with almost perfect precision, perfect recall, and a higher mAP.

Tiny YOLOv4 was an extremely fast and accurate model for real-time object detection (333 FPS and mAP of 0.998) and the best model for object detection on a portable, stand-alone device, followed by YOLOv5s (133 FPS and mAP of 0.995) and YOLOv4 (65 FPS and mAP of 0.974). Consequently, YOLOv4, Tiny YOLOv4, and YOLOv5s were evaluated for real-world application using the second dataset of multiple objects per image and more class labels for object detection (Figure 5 - 6). YOLOv4 and Tiny YOLOv4 performed better than YOLOv5s compared to the ground truth yielding a similar mAP of 0.896 and 0.895, respectively. However, Tiny YOLOv4 achieved both higher precision and recall. Tiny YOLOv4 was able to detect all five class labels well (Figure 7 - 8). The model was able to detect M2P, M4H, and M4P lesions with a higher average precision compared to M2 lesions. Overall, Tiny YOLOv4 was able to detect all five M-stages of DD on video files.

DISCUSSION

Seven of the nine CV models were able to detect DD lesions in dairy cattle from an image. For speed, Tiny YOLOv4 clearly outperformed all other models with the highest FPS. For detection, YOLOv4, Tiny YOLOv4, and YOLOv5 outperformed all other models with high precision, high recall, and the highest mAP. The next closest model was Cascade R-CNN followed by Faster R-CNN, YOLOv3, and Tiny YOLOv3. YOLOv4, Tiny YOLOv4, and YOLOv5 achieved similar precision and recall using the first dataset, but Tiny YOLOv4 yielded higher mAP. YOLOv4 and Tiny YOLOv4 achieved similar mAP using the Dataset 2, however Tiny YOLOv4 yielded higher precision and recall. Tiny YOLOv4 was the best CV model for the use case of real-time detection on streaming video. Additionally, the YOLOv3, YOLOv4, and YOLOv5 models had better performance than the YOLOv2 model using the same input data (Cernek et al., 2020). As a proof-of-concept, the Tiny YOLOv4 model was able to detect the different M-stages of DD on video files.

For future application, the comparison of CV models for a given task can help identify which approaches perform best for a given situation (Cao et al., 2019; Jia et al., 2020; Liu et al., 2020; Xu et al., 2020). For instance, the R-CNN models, including Faster, Mask, and Cascade R-CNN, provide better detection performance when hooves were in close proximity, such as in densely populated dairy barns or feedlots. For applications where hooves were clearly separated from each other, such as a milking parlor, the YOLO models were a superior approach for the detection of DD lesions.

For Dataset 2, Tiny YOLOv4 was able to detect all five class labels well. However, the model was able to detect M2P, M4H, and M4P lesions better than M2 lesions with a higher average precision. This was probably caused by unbalanced classes of data in the second dataset with a relatively small number of M2 in the training set. Unbalanced datasets where the number of instances for class labels, such as M2, is less than the other class labels can create problems with regards to overall prediction performance of the CV model (Oksuz et al., 2021). Adding images to the dataset for infrequent classes will improve the model performance. The difference in distance between the camera and the object can create inaccurate object detection performance when using YOLO models (Redmon and Farhadi, 2018).

The CV models applied during this project were able to identify and classify DD lesions on commercial dairy farms with high accuracy, in real-time, and at high speed. This result is a step in applying CV algorithms to veterinary medicine and implementing real-time DD detection on dairy farms. The proposed CV tool can be used for early detection and prompt treatment of DD in dairy cattle and has been extended to beef cattle (unpublished data). The trained model may be applied to different cattle breeds and various locations. The CV models for DD detection can be expanded to include more images from more farms for implementation on different farms. The YOLO framework can classify multiple claw diseases simultaneously in the future after training the models using more images of claw lesions. For increased generalization, the model requires additional images originating from additional cattle breeds, animal species, and geographical regions for robust predictions. CV has the potential to be used in precision livestock farming and high-throughput phenotyping applications. Such applications include animal identification and tracking, behavior monitoring, and phenotype classification and recognition (Fernandes et al., 2020). Applications of CV in animal production systems and veterinary sciences are a growing research area (Fernandes et al., 2020; Li et al., 2021; Wurtz et al., 2019).

The TensorFlow 1.x framework used in the study can be expanded to other open-source machine learning frameworks for deployment and overall ease of use. TensorFlow 2 offers better usability and high performance (NVIDIA Corporation, 2022; tensorflow, 2017). TensorFlow models can be converted to TensorFlow Lite models reducing model size and power consumption and subsequently, they can be deployed on edge devices like mobiles or microcontrollers (tensorflow, 2021). PyTorch framework is quite fast and extremely efficient compared to some of the alternatives (Paszke et al., 2019; pytorch, 2021). YOLOv3 has been implemented in PyTorch by Ultralytics. YOLOv5 is the PyTorch implementation of YOLOv4 and achieves similar speed and accuracy, but reduced training time (ultralytics, 2022, 2021). The PyTorch implementation of YOLOv3 or different size models of YOLOv5 can be evaluated to see if there is an improvement in performance or speed. Lastly, YOLOv4 and YOLOv5 can be updated to the most recent version of each implementation, namely YOLOv7 and YOLOv8, providing an increase in accuracy and efficiency (Jocher, 2023; Jocher et al., 2023; Wang et al., 2022; Wong, 2023).

Accessibility to CV tools can promote the growth of such techniques in veterinary medicine and reduce adoption thresholds for the models. The CV models can be extended to portable platforms, including iOS or Android application for mobile use, as well as using a docker container for deployment of CV models as cloud-based applications (Docker, 2021). The implementation of such CV models in hand-held devices and in collaboration with cattle professionals will generate a rich, diverse library of images for the optimization and validation of CV models. Ultimately, the proposed CV tool will improve animal welfare and increase production for large-scale cattle facilities as an example for other applications in veterinary medicine and agriculture.

ACKNOWLEDGEMENTS

Funding support for the research was provided by the US Department of Agriculture through the National Institute for Food and Agriculture - Animal Health Grant (WIS03082). We would also like to thank Jamie Sullivan (Carman, MB, Canada), Chris Bauer (Mondovi, WI), and all other dairy producers for contributing to the images of hoof lesions. The authors have no financial and personal relationships with other people or organizations that could inappropriately influence (bias) their work.

REFERENCES

Adarsh, P., Rathi, P., Kumar, M., 2020. YOLO v3-Tiny: Object Detection and Recognition using one stage improved model, in: 2020 6th International Conference on Advanced Computing and Communication Systems (ICACCS). Presented at the 2020 6th International Conference on Advanced Computing and Communication Systems (ICACCS), pp. 687–694. https://doi.org/10.1109/ICACCS48705.2020.9074315

Al-antari, M.A., Al-masni, M.A., Kim, T.-S., 2020. Deep Learning Computer-Aided Diagnosis for Breast Lesion in Digital Mammogram, in: Lee, G., Fujita, H. (Eds.), Deep Learning in Medical Image Analysis : Challenges and Applications, Advances in Experimental Medicine and Biology. Springer International Publishing, Cham, pp. 59–72. https://doi.org/10.1007/978-3-030-33128-3_4

Alexey, 2021. Yolo v4, v3 and v2 for Windows and Linux.

Alsaaod, M., Syring, C., Dietrich, J., Doherr, M.G., Gujan, T., Steiner, A., 2014. A field trial of infrared thermography as a non-invasive diagnostic tool for early detection of digital dermatitis in dairy cows. The Veterinary Journal 199, 281–285. https://doi.org/10.1016/j.tvjl.2013.11.028

Berry, S.L., Read, D.H., Famula, T.R., Mongini, A., Döpfer, D., 2012. Long-term observations on the dynamics of bovine digital dermatitis lesions on a California dairy after topical treatment with lincomycin HCl. The Veterinary Journal, Special Issue:Bovine Lameness 193, 654–658. https://doi.org/10.1016/j.tvjl.2012.06.048

Bochkovskiy, A., Wang, C.-Y., Liao, H.-Y.M., 2020. YOLOv4: Optimal Speed and Accuracy of Object Detection. arXiv:2004.10934 [cs, eess].

Byrne, D.T., Berry, D.P., Esmonde, H., McGovern, F., Creighton, P., McHugh, N., 2018. Infrared thermography as a tool to detect hoof lesions in sheep. Transl Anim Sci 3, 577–588. https://doi.org/10.1093/tas/txy132

Byrne, D.T., Berry, D.P., Esmonde, H., McHugh, N., 2017. Temporal, spatial, inter-, and intra-cow repeatability of thermal imaging1. Journal of Animal Science 95, 970–979. https://doi.org/10.2527/jas.2016.1005

Cai, Z., Vasconcelos, N., 2017. Cascade R-CNN: Delving into High Quality Object Detection. arXiv:1712.00726 [cs].

Cao, C., Wang, B., Zhang, W., Zeng, X., Yan, X., Feng, Z., Liu, Y., Wu, Z., 2019. An Improved Faster R-CNN for Small Object Detection. IEEE Access 7, 106838–106846. https://doi.org/10.1109/ACCESS.2019.2932731

Cao, Z., Duan, L., Yang, G., Yue, T., Chen, Q., Fu, H., Xu, Y., 2017. Breast Tumor Detection in Ultrasound Images Using Deep Learning, in: Wu, G., Munsell, B.C., Zhan, Y., Bai, W., Sanroma, G., Coupé, P. (Eds.), Patch-Based Techniques in Medical Imaging, Lecture Notes in

Computer Science. Springer International Publishing, Cham, pp. 121–128. https://doi.org/10.1007/978-3-319-67434-6_14

Cernek, P., Bollig, N., Anklam, K., Döpfer, D., 2020. Hot topic: Detecting digital dermatitis with computer vision. Journal of Dairy Science 103, 9110–9115. https://doi.org/10.3168/jds.2019-17478

Cheli, R., Mortellaro, C., 1974. La dermatite digitale del bovino, in: Proceedings of the 8th International Conference on Diseases of Cattle. Piacenza, pp. 208–213.

Cramer, G., Lissemore, K.D., Guard, C.L., Leslie, K.E., Kelton, D.F., 2009. Herd-level risk factors for seven different foot lesions in Ontario Holstein cattle housed in tie stalls or free stalls. Journal of Dairy Science 92, 1404–1411. https://doi.org/10.3168/jds.2008-1134

Docker, 2021. Docker Documentation [WWW Document]. Docker Documentation. URL https://docs.docker.com/

Döpfer, D., Holzhauer, M., Boven, M. van, 2012. The dynamics of digital dermatitis in populations of dairy cattle: Model-based estimates of transition rates and implications for control. The Veterinary Journal, Special Issue:Bovine Lameness 193, 648–653. https://doi.org/10.1016/j.tvjl.2012.06.047

Döpfer, D., Huurne, A.A.H.M. ter, Cornelisse, J.L., Asten, A.J.A.M. van, Koopmans, A., Meijer, F.A., Schukken, Y.H., Szakáll, I., Klee, W., Bosma, R.B., 1997. Histological and bacteriological evaluation of digital dermatitis in cattle, with special reference to spirochaetes and Campylobacter faecalis. Veterinary Record 140, 620–623. https://doi.org/10.1136/vr.140.24.620

Fan, H., Liu, Shuanglong, Ferianc, M., Ng, H.-C., Que, Z., Liu, Shen, Niu, X., Luk, W., 2018. A Real-Time Object Detection Accelerator with Compressed SSDLite on FPGA, in: 2018 International Conference on Field-Programmable Technology (FPT). Presented at the 2018 International Conference on Field-Programmable Technology (FPT), IEEE, Naha, Okinawa, Japan, pp. 14–21. https://doi.org/10.1109/FPT.2018.00014

Fernandes, A.F.A., Dórea, J.R.R., Rosa, G.J. de M., 2020. Image Analysis and Computer Vision Applications in Animal Sciences: An Overview. Frontiers in Veterinary Science 7.

Gomes, R.A., Monteiro, G.R., Assis, G.J.F., Busato, K.C., Ladeira, M.M., Chizzotti, M.L., 2016. Technical note: Estimating body weight and body composition of beef cattle trough digital image analysis. J Anim Sci 94, 5414–5422. https://doi.org/10.2527/jas.2016-0797

Google Colaboratory [WWW Document], 2022. URL https://colab.research.google.com/?utm_source=scs-index (accessed 2.24.22).

Hovinen, M., Siivonen, J., Taponen, S., Hänninen, L., Pastell, M., Aisla, A.-M., Pyörälä, S., 2008. Detection of Clinical Mastitis with the Help of a Thermal Camera. Journal of Dairy Science 91, 4592–4598. https://doi.org/10.3168/jds.2008-1218

ImageNet [WWW Document], 2020. URL https://www.image-net.org/ (accessed 2.24.22).

Jacobs, C., Beninger, C., Hazlewood, G.S., Orsel, K., Barkema, H.W., 2019. Effect of footbath protocols for prevention and treatment of digital dermatitis in dairy cattle: A systematic review and network meta-analysis. Preventive Veterinary Medicine 164, 56–71. https://doi.org/10.1016/j.prevetmed.2019.01.011

Jia, W., Tian, Y., Luo, R., Zhang, Z., Lian, J., Zheng, Y., 2020. Detection and segmentation of overlapped fruits based on optimized mask R-CNN application in apple harvesting robot. Computers and Electronics in Agriculture 172, 105380. https://doi.org/10.1016/j.compag.2020.105380

Jocher, G., 2023. Home - Ultralytics YOLOv8 Docs [WWW Document]. URL https://docs.ultralytics.com/ (accessed 9.9.23).

Jocher, G., Chaurasia, A., Qiu, J., 2023. YOLO by Ultralytics.

Kin-Yiu, W., 2020. some comparison · Issue #32 · WongKinYiu/CrossStagePartialNetworks [WWW Document]. URL https://github.com/WongKinYiu/CrossStagePartialNetworks/issues/32#issuecomment-638064640 (accessed 9.9.23).

Krizhevsky, A., Sutskever, I., Hinton, G.E., 2012. ImageNet Classification with Deep Convolutional Neural Networks, in: Advances in Neural Information Processing Systems. Curran Associates, Inc.

Krull, A.C., Shearer, J.K., Gorden, P.J., Scott, H.M., Plummer, P.J., 2016. Digital dermatitis: Natural lesion progression and regression in Holstein dairy cattle over 3 years. Journal of Dairy Science 99, 3718–3731. https://doi.org/10.3168/jds.2015-10535

Li, G., Huang, Y., Chen, Z., Chesser, G.D., Purswell, J.L., Linhoss, J., Zhao, Y., 2021. Practices and Applications of Convolutional Neural Network-Based Computer Vision Systems in Animal Farming: A Review. Sensors 21, 1492. https://doi.org/10.3390/s21041492

Liakos, K.G., Busato, P., Moshou, D., Pearson, S., Bochtis, D., 2018. Machine Learning in Agriculture: A Review. Sensors 18, 2674. https://doi.org/10.3390/s18082674

Lin, T., 2018. LabelImg.

Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.-Y., Berg, A.C., 2016. SSD: Single Shot MultiBox Detector. arXiv:1512.02325 [cs] 9905, 21–37. https://doi.org/10.1007/978-3-319-46448-0_2

Liu, X., Hu, C., Li, P., 2020. Automatic segmentation of overlapped poplar seedling leaves combining Mask R-CNN and DBSCAN. Computers and Electronics in Agriculture 178, 105753. https://doi.org/10.1016/j.compag.2020.105753

Logue, D.N., Offer, J.E., Laven, R.A., Ellis, W.A., 2005. Digital dermatitis – The aetiological soup. The Veterinary Journal 170, 12–13. https://doi.org/10.1016/j.tvjl.2005.04.030

Martins, R.F.S., do Prado Paim, T., de Abreu Cardoso, C., Stéfano Lima Dallago, B., de Melo, C.B., Louvandini, H., McManus, C., 2013. Mastitis detection in sheep by infrared thermography. Research in Veterinary Science 94, 722–724. https://doi.org/10.1016/j.rvsc.2012.10.021

Metzner, M., Sauter-Louis, C., Seemueller, A., Petzl, W., Zerbe, H., 2015. Infrared thermography of the udder after experimentally induced Escherichia coli mastitis in cows. The Veterinary Journal 204, 360–362. https://doi.org/10.1016/j.tvjl.2015.04.013

NVIDIA Corporation, 2022. TensorRT Open Source Software. NVIDIA Corporation.

Oksuz, K., Cam, B.C., Kalkan, S., Akbas, E., 2021. Imbalance Problems in Object Detection: A Review. IEEE Transactions on Pattern Analysis and Machine Intelligence 43, 3388–3415. https://doi.org/10.1109/TPAMI.2020.2981890

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S., 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library, in: Advances in Neural Information Processing Systems. Curran Associates, Inc.

Plummer, P.J., Krull, A., 2017. Clinical Perspectives of Digital Dermatitis in Dairy and Beef Cattle. Veterinary Clinics of North America: Food Animal Practice 33, 165–181. https://doi.org/10.1016/j.cvfa.2017.02.002

pytorch, 2021. pytorch/pytorch. pytorch.

Redmon, J., 2018. YOLO: Real-Time Object Detection [WWW Document]. URL https://pjreddie.com/darknet/yolo/ (accessed 2.24.22).

Redmon, J., Divvala, S., Girshick, R., Farhadi, A., 2016. You Only Look Once: Unified, Real-Time Object Detection. arXiv:1506.02640 [cs].

Redmon, J., Farhadi, A., 2018. YOLOv3: An Incremental Improvement. arXiv:1804.02767 [cs].

Redmon, J., Farhadi, A., 2016. YOLO9000: Better, Faster, Stronger. arXiv:1612.08242 [cs].

Refaai, W., Van Aert, M., Abd El-Aal, A.M., Behery, A.E., Opsomer, G., 2013. Infectious diseases causing lameness in cattle with a main emphasis on digital dermatitis (Mortellaro disease). Livestock Science, Lameness and claw lesions in sows, cows and small ruminants 156, 53–63. https://doi.org/10.1016/j.livsci.2013.06.004

Ren, S., He, K., Girshick, R., Sun, J., 2016. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. arXiv:1506.01497 [cs].

Salau, J., Haas, J.H., Junge, W., Thaller, G., 2017. Automated calculation of udder depth and rear leg angle in Holstein-Friesian cows using a multi-Kinect cow scanning system. Biosystems Engineering 160, 154–169. https://doi.org/10.1016/j.biosystemseng.2017.06.006

Schulz, T., Gundelach, Y., Feldmann, M., Hoedemaker, M., 2016. Early detection and treatment of lame cows. Tierarztl Prax Ausg G Grosstiere Nutztiere 44, 5–11. https://doi.org/10.15653/TPG-150534

Scoley, G.E., Gordon, A.W., Morrison, S.J., 2018. Use of thermal imaging in dairy calves: exploring the repeatability and accuracy of measures taken from different anatomical regions1. Transl Anim Sci 3, 564–576. https://doi.org/10.1093/tas/txy126

Sharma, V., Mir, R.N., 2020. A comprehensive and systematic look up into deep learning based object detection techniques: A review. Computer Science Review 38, 100301. https://doi.org/10.1016/j.cosrev.2020.100301

Shearer, J.K., Hernandez, J., 2000. Efficacy of Two Modified Nonantibiotic Formulations (Victory) for Treatment of Papillomatous Digital Dermatitis in Dairy Cows. Journal of Dairy Science 83, 741–745. https://doi.org/10.3168/jds.S0022-0302(00)74936-8

Soviany, P., Ionescu, R.T., 2018. Optimizing the Trade-Off between Single-Stage and Two-Stage Deep Object Detectors using Image Difficulty Prediction, in: 2018 20th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC). Presented at the 2018 20th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC), pp. 209–214. https://doi.org/10.1109/SYNASC.2018.00041

Szeliski, R., 2022. Computer Vision: Algorithms and Applications, 2nd ed. 2022 edition. ed. Springer, Cham.

Tashiev, I., 2022. XmlToTxt.

Techzizou, 2021. YOLOv4 VS YOLOv4-tiny. Analytics Vidhya. URL https://medium.com/analytics-vidhya/yolov4-vs-yolov4-tiny-97932b6ec8ec (accessed 2.24.22).

Teixeira, A.G.V., Machado, V.S., Caixeta, L.S., Pereira, R.V., Bicalho, R.C., 2010. Efficacy of formalin, copper sulfate, and a commercial footbath product in the control of digital dermatitis. Journal of Dairy Science 93, 3628–3634. https://doi.org/10.3168/jds.2010-3246

tensorflow, 2021. TensorFlow Lite for Microcontrollers. tensorflow.

tensorflow, 2017. Documentation for TensorRT in TensorFlow (TF-TRT). tensorflow.

ultralytics, 2022. ultralytics/yolov5. Ultralytics.

ultralytics, 2021. ultralytics/yolov3. Ultralytics.

Ünver, H.M., Ayan, E., 2019. Skin Lesion Segmentation in Dermoscopic Images with Combination of YOLO and GrabCut Algorithm. Diagnostics 9, 72. https://doi.org/10.3390/diagnostics9030072

Vaidya, B., Paunwala, C., 2019. Deep Learning Architectures for Object Detection and Classification, in: Mishra, M.K., Mishra, B.S.P., Patel, Y.S., Misra, R. (Eds.), Smart Techniques for a Smarter Planet: Towards Smarter Algorithms, Studies in Fuzziness and Soft Computing.

Springer International Publishing, Cham, pp. 53–79. https://doi.org/10.1007/978-3-030-03131-2_4

Vanhoudt, A., Yang, D.A., Armstrong, T., Huxley, J.N., Laven, R.A., Manning, A.D., Newsome, R.F., Nielen, M., Werven, T. van, Bell, N.J., 2019. Interobserver agreement of digital dermatitis M-scores for photographs of the hind feet of standing dairy cattle. Journal of Dairy Science 102, 5466–5474. https://doi.org/10.3168/jds.2018-15644

Vinicki, K., Ferrari, P., Belic, M., Turk, R., 2018. Using Convolutional Neural Networks for Determining Reticulocyte Percentage in Cats. arXiv:1803.04873 [cs].

Wang, C.-Y., Bochkovskiy, A., Liao, H.-Y.M., 2022. YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. https://doi.org/10.48550/arXiv.2207.02696

Wong, K.-Y., 2023. Official YOLOv7.

Wu, X., Sahoo, D., Hoi, S.C.H., 2020. Recent advances in deep learning for object detection. Neurocomputing 396, 39–64. https://doi.org/10.1016/j.neucom.2020.01.085

Wurtz, K., Camerlink, I., D'Eath, R.B., Fernández, A.P., Norton, T., Steibel, J., Siegford, J., 2019. Recording behaviour of indoor-housed farm animals automatically using machine vision technology: A systematic review. PLOS ONE 14, e0226669. https://doi.org/10.1371/journal.pone.0226669

Xu, B., Wang, W., Falzon, G., Kwan, P., Guo, L., Chen, G., Tait, A., Schneider, D., 2020. Automated cattle counting using Mask R-CNN in quadcopter vision system. Computers and Electronics in Agriculture 171, 105300. https://doi.org/10.1016/j.compag.2020.105300

Yang, C., Li, G., Zhang, X., Gu, X., 2018. Udder skin surface temperature variation pre- and post- milking in dairy cows as determined by infrared thermography. Journal of Dairy Research 85, 201–203. https://doi.org/10.1017/S0022029918000213

Zaninelli, M., Redaelli, V., Luzi, F., Bronzo, V., Mitchell, M., Dell'Orto, V., Bontempo, V., Cattaneo, D., Savoini, G., 2018. First Evaluation of Infrared Thermography as a Tool for the Monitoring of Udder Health Status in Farms of Dairy Cows. Sensors 18, 862. https://doi.org/10.3390/s18030862

Zhao, Z.-Q., Zheng, P., Xu, S.-T., Wu, X., 2019. Object Detection With Deep Learning: A Review. IEEE Transactions on Neural Networks and Learning Systems 30, 3212–3232. https://doi.org/10.1109/TNNLS.2018.2876865

TABLES AND FIGURES

Figure 2.1 DairyCoPilot Inference time of the nine CV algorithms for the testing set in Dataset 1 (n= 2,227). Inference time was measured in frames per second (FPS) and the nine CV models include the five latest versions of You Only Look Once (YOLO) models: YOLOv3, Tiny YOLOv3, YOLOv4, Tiny YOLOv4, and YOLOv5s in addition to Single-Shot Multibox Detector (SSD), SSD Lite, Faster Region-Based Convolutional Neural Networks (R-CNN), and Cascade R-CNN.

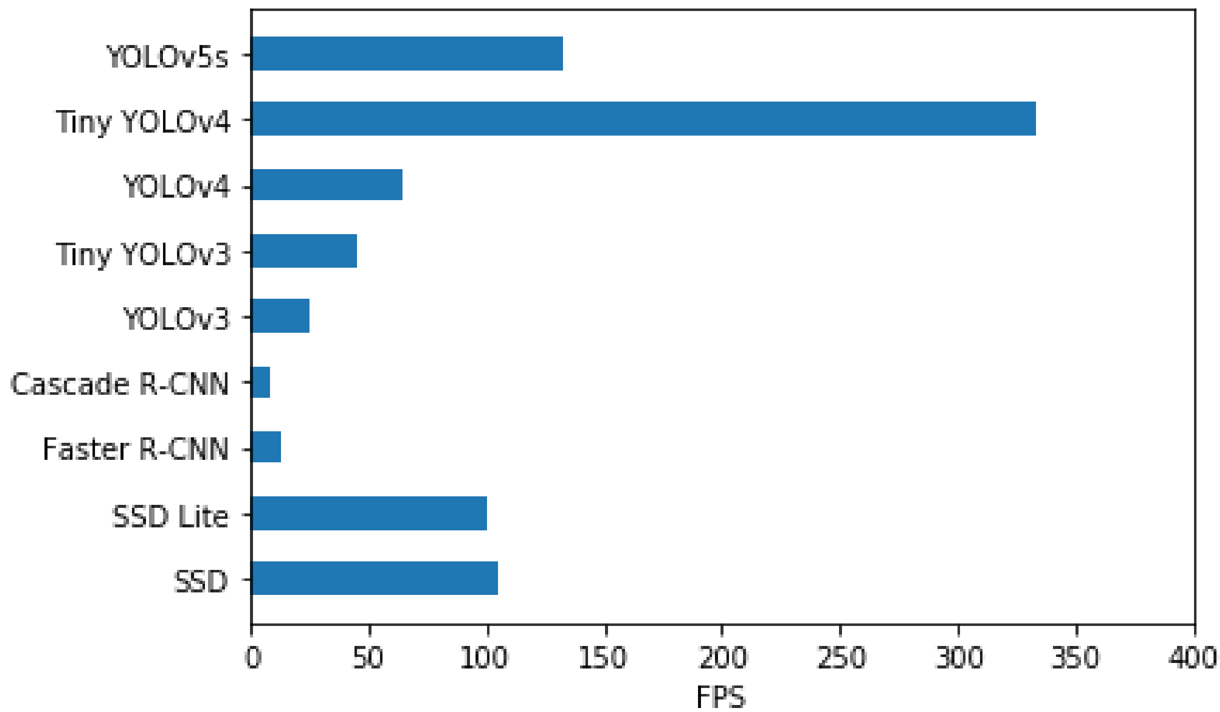

Figure 2.2 Performance of the nine CV algorithms for the testing set in Dataset 1 (n= 2,227). Performance was measured in mean average precision (mAP) and the nine CV models include the five of the latest versions of You Only Look Once (YOLO) models: YOLOv3, Tiny YOLOv3, YOLOv4, Tiny YOLOv4, and YOLOv5s in addition to Single-Shot Multibox Detector (SSD), SSD Lite, Faster Region-Based Convolutional Neural Networks (R-CNN), and Cascade R-CNN.

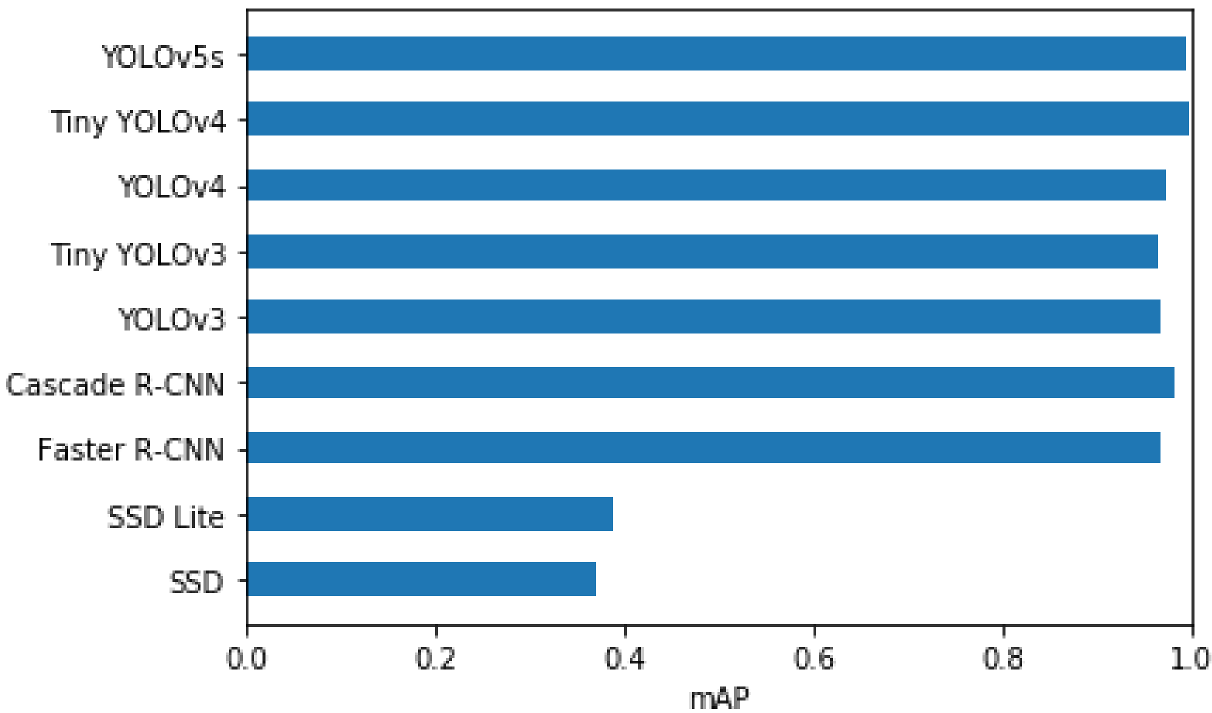

Figure 2.3 Precision and recall of the nine CV algorithms (bottom-left) with inset of the top seven CV model (top-right) for the testing set in Dataset 1 (n= 2,227). Performance was measured in precision and recall and the nine CV models include the five of the latest versions of You Only Look Once (YOLO) models: YOLOv3, Tiny YOLOv3, YOLOv4, Tiny YOLOv4, and YOLOv5s in addition to Single-Shot Multibox Detector (SSD), SSD Lite, Faster Region-Based Convolutional Neural Networks (R-CNN), and Cascade R-CNN.

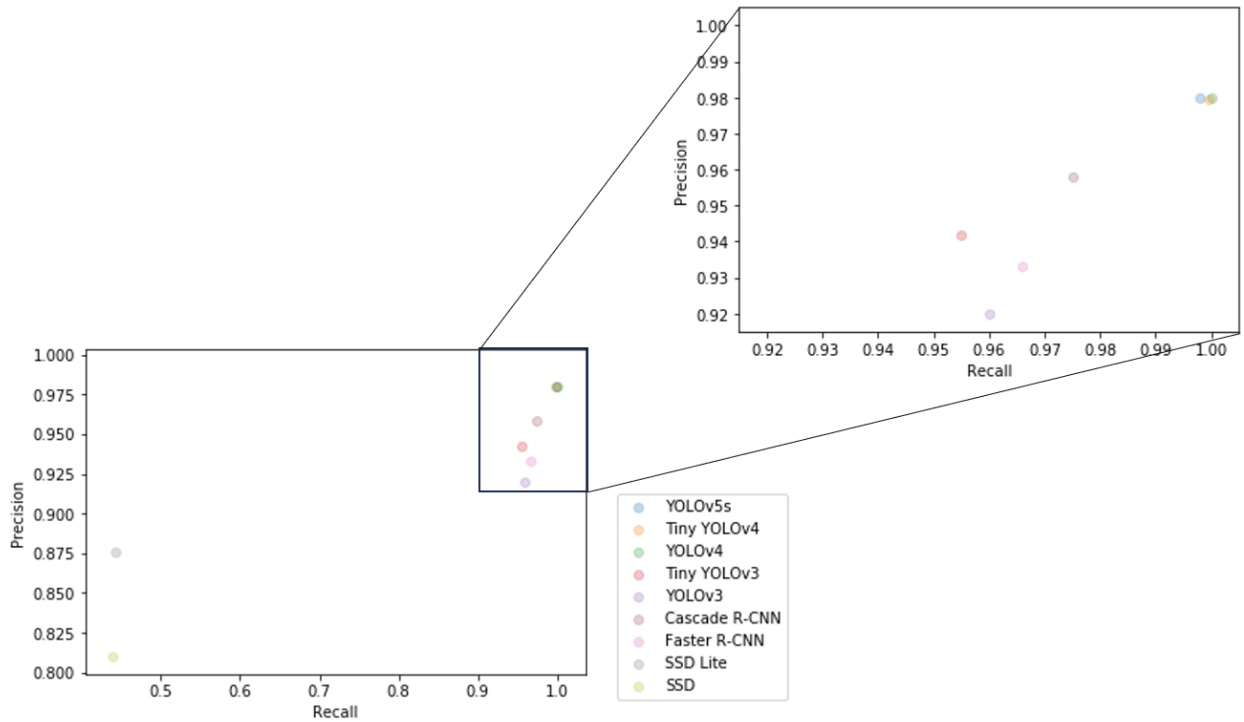

Figure 2.4 Bounding box predictions of M0/M4 and M2 by the nine CV algorithms for the testing set in Dataset 1 (n= 2,227). The nine CV models include the five latest versions of You Only Look Once (YOLO) models: YOLOv3, Tiny YOLOv3, YOLOv4, Tiny YOLOv4, YOLOv5s in addition to Single-Shot Multibox Detector (SSD), SSD Lite, Faster Region-Based Convolutional Neural Networks (R-CNN), and Cascade R-CNN compared to labels by a trained investigator (‘Ground Truth’).

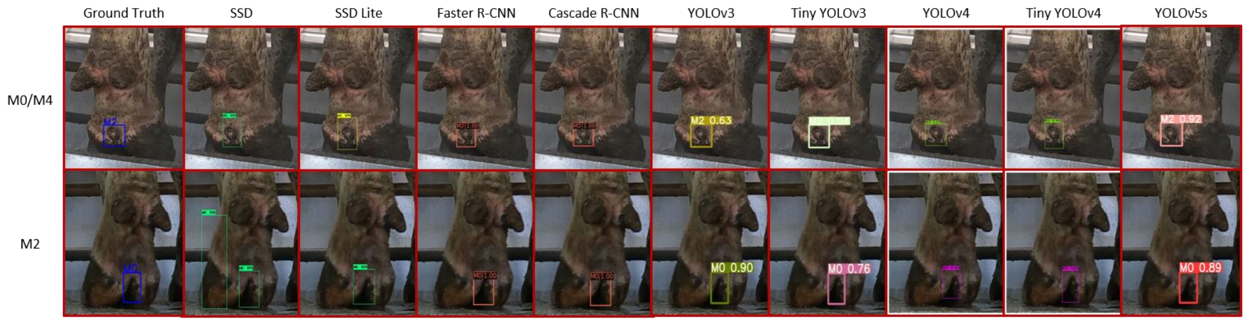

Figure 2.5 Performance of the top three CV algorithms for the testing set of Dataset 2 (n= 409). Performance was measured in mean average precision (mAP) and the CV models include the three of the latest versions of You Only Look Once (YOLO) models: YOLOv4, Tiny YOLOv4, and YOLOv5s.

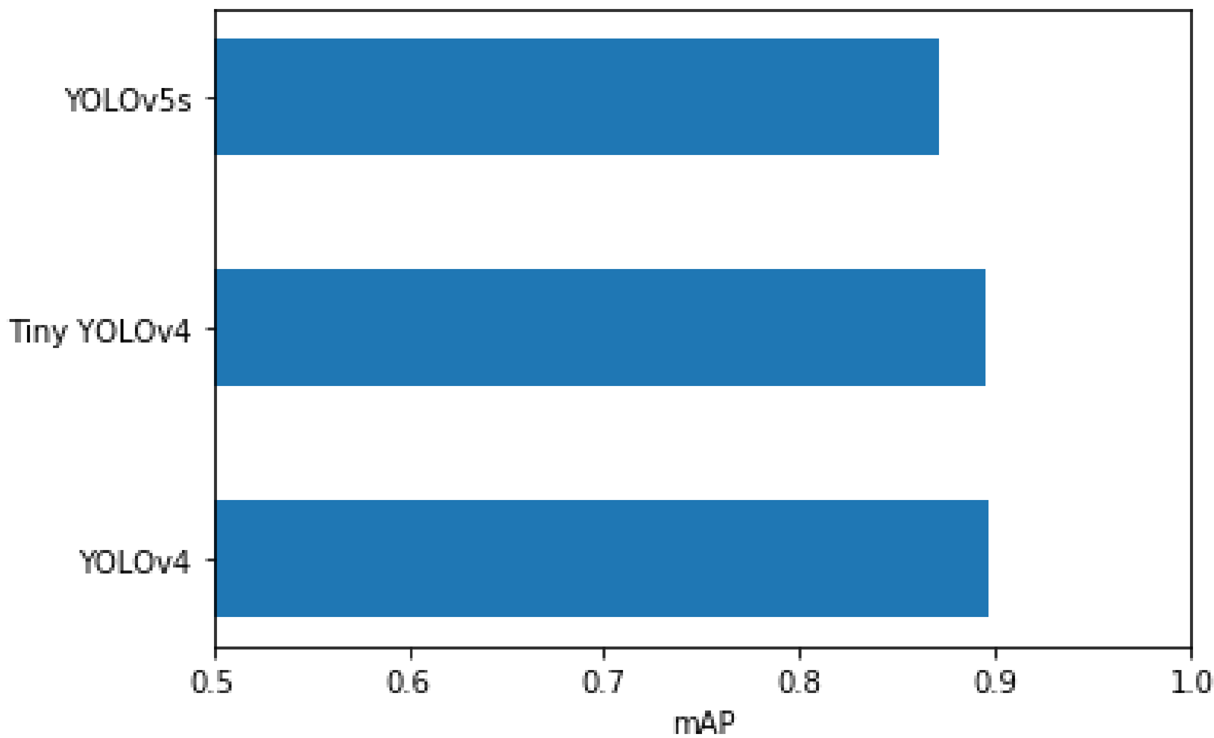

Figure 2.6 Precision and recall of the top three CV algorithms for the testing set of Dataset 2 (n= 409). Performance was measured in precision and recall and the CV models include the three of the latest versions of You Only Look Once (YOLO) models: YOLOv4, Tiny YOLOv4, and YOLOv5s.

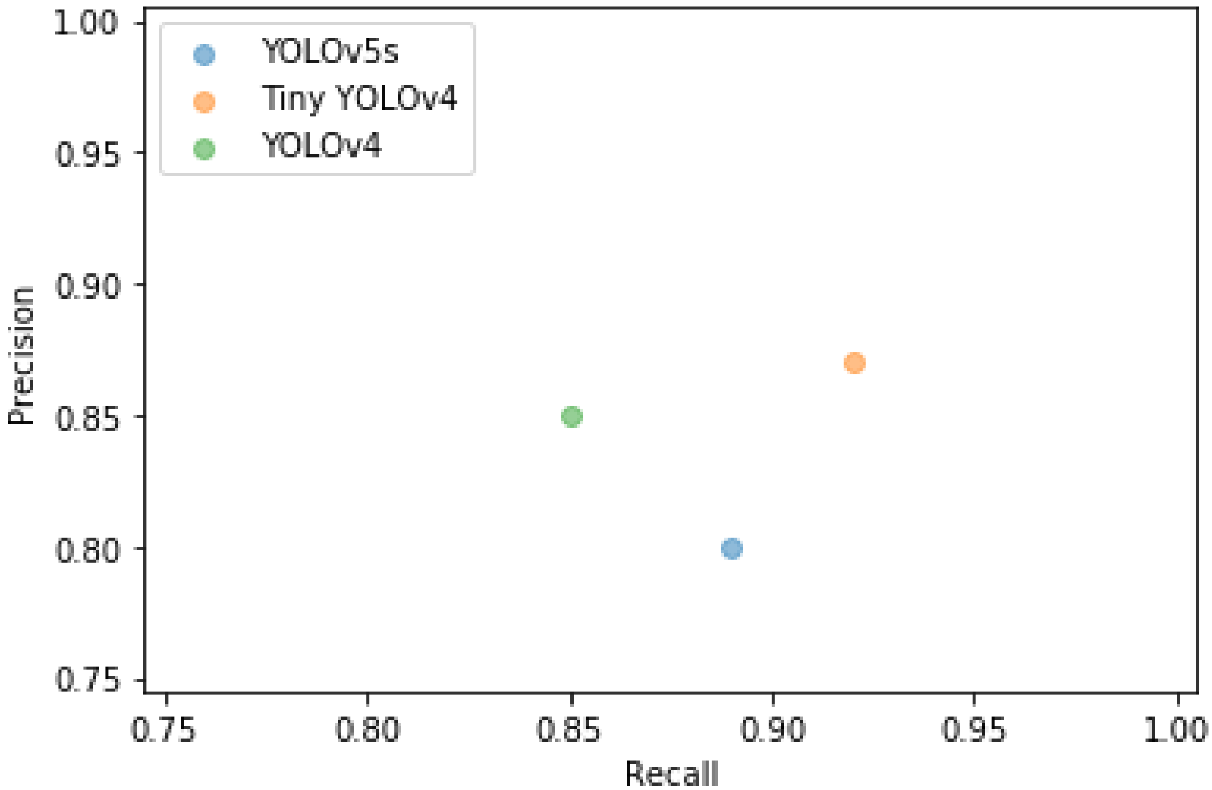

Figure 2.7 Bounding box predictions of M-stages by Tiny YOLOv4 for the testing set in Dataset 2. Predictions were made for each of the M-stages: M0 (top-center; magenta), M2, M2P (bottom-left; cyan), M4H, and M4P (bottom-right; green).

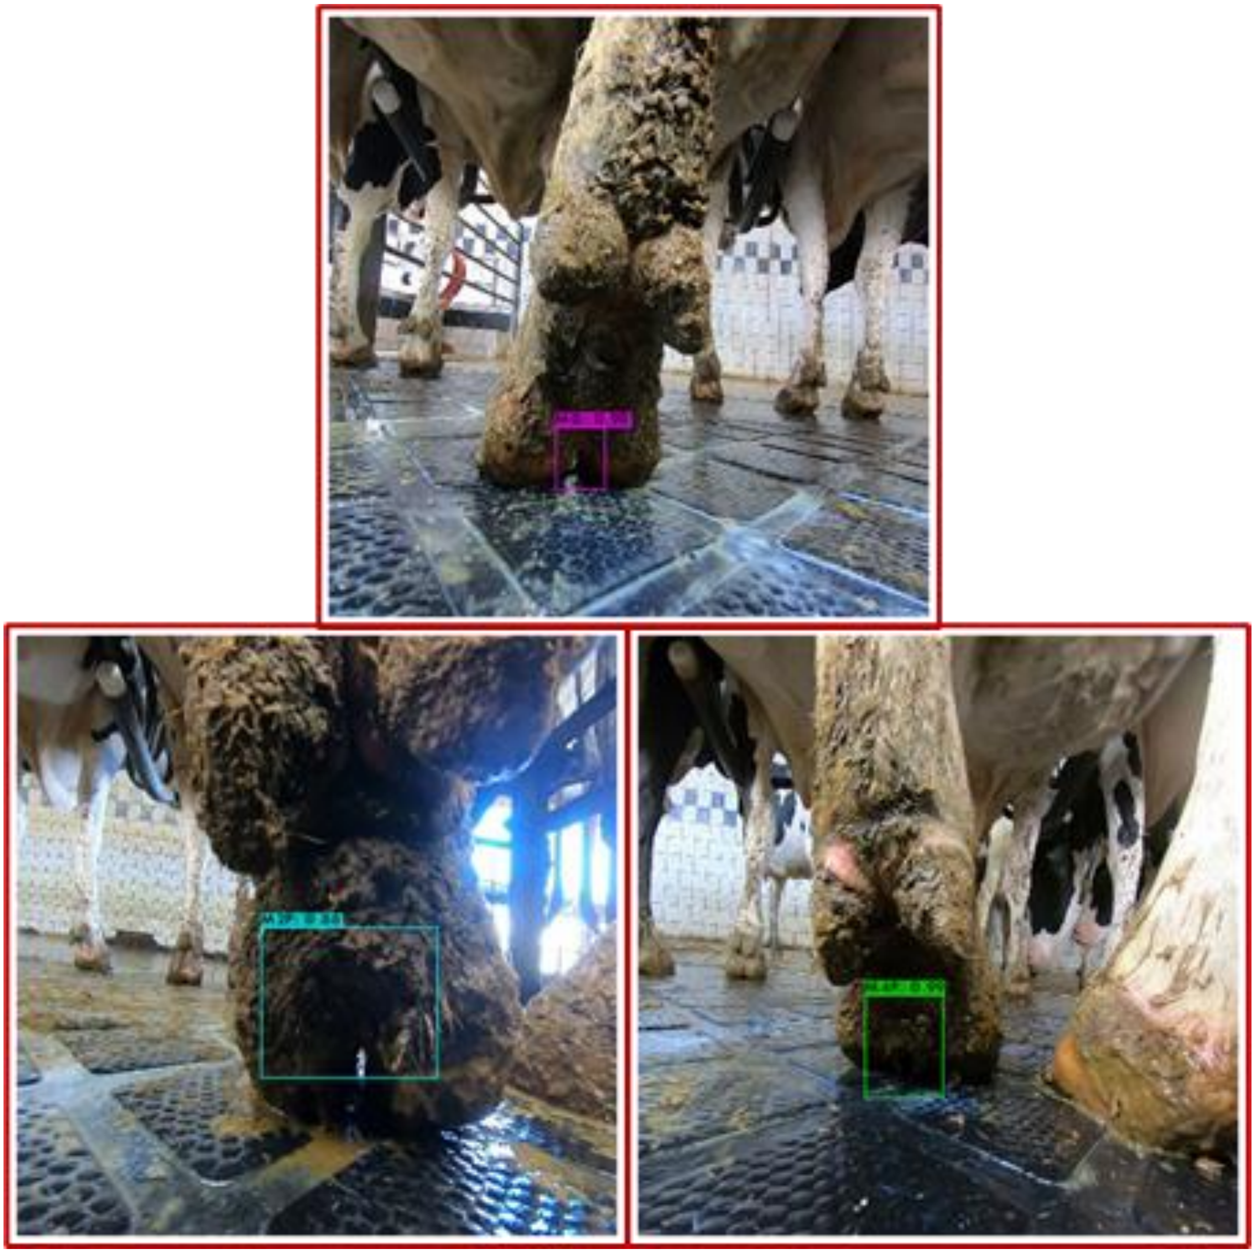

Figure 2.8 Average precision of the five M-stages by Tiny YOLOv4 for the testing set in Dataset 2 (n= 409). Performance was measured in average precision and calculated for each of the five M-stages: M0, M2, M2P, M4H, and M4P.

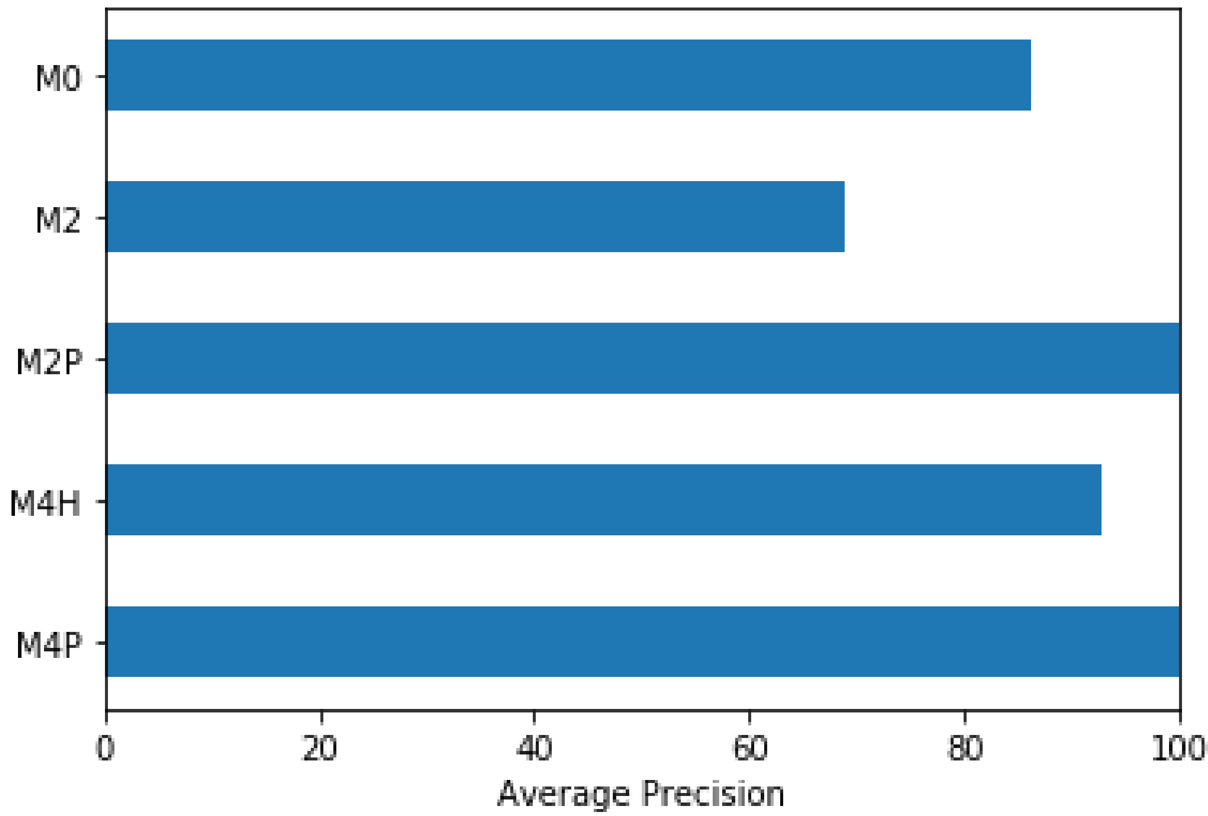

CHAPTER 3 BENCHMARKING ANALYSIS OF COMPUTER VISION ALGORITHMS ON EDGE DEVICES FOR THE REAL-TIME DETECTION OF DIGITAL DERMATITIS IN DAIRY COWS

ABSTRACT

Digital dermatitis (DD) is a bovine claw disease responsible for ulcerative lesions on the coronary band of the foot. It causes significant animal welfare and economic losses to the cattle industry. Early detection of DD can lead to prompt treatment and decrease lameness. Current detection and staging methods require a trained individual to evaluate the interdigital space on each foot for clinical signs of DD.

Computer vision (CV), a type of artificial intelligence for image analysis, has demonstrated promising results on object detection tasks. However, farms require robust solutions that can be deployed in harsh conditions including dust, debris, humidity, precipitation, other equipment issues. The study aims to train, embed, and benchmark DD detection models on edge devices.

Images were collected from commercial dairy farms with the camera facing the interdigital space on the plantar surface of the foot. Images were scored for M-stages of DD by a trained investigator using the M-stage DD classification system. Models were trained to detect and score DD lesions and embedded on an edge device.

The Tiny YOLOv4 model deployed on a CV specific integrated camera module connected to a single board computer achieved a mean average precision (mAP) of 0.895, an overall prediction accuracy of 0.873, and a Cohen’s kappa of 0.830 for agreement between the computer vision model and the trained investigator. The model reached a final inference speed of 40 frames per second (FPS) and was stable without any interruptions.

The CV model was able to detect DD lesions on an edge device with high performance and speed. The CV tool can be used for early detection and prompt treatment of DD in dairy cows. Real-time detection of DD on edge device will improve health outcomes, while simultaneously decreasing labor costs. We demonstrate that the deployed model can be a low-power and portable solution for real-time detection of DD on dairy farms. This result is a step towards applying CV algorithms to veterinary medicine and implementing real-time detection of health outcomes in precision farming.

INTRODUCTION

Bovine digital dermatitis (DD) is an infectious disease affecting the coronary band of cattle feet (Cheli and Mortellaro, 1974; Döpfer et al., 1997). The disease is characterized by circumscribed ulcero-proliferative lesions generally situated on the plantar aspect of the hoof and spread through direct contact between animals or contact with contaminated surfaces or objects (Blowey and Sharp, 1988; Read and Walker, 1998). Approximately 70 percent of US dairies and 95 percent of dairies with more than 500 cows have reported to have DD (Gomez, 2018; Bruijnis et al., 2012; Solano et al., 2016). Up to 80 percent of lameness is caused by DD in the US and world cattle industry, resulting in production losses, decreased fertility, and increased mortality (Gomez, 2018; USDA, 2009). Lost productivity, labor, treatment or preventive measures, and other indirect economic costs attributable to DD cost the US cattle industry hundreds of millions of dollars annually (Losinger, 2006; de Jesús Argáez-Rodríguez et al., 1997; Hernandez et al., 2001; Whay et al., 1997; Cha et al., 2010; Evans et al., 2016; Solano et al., 2017; Evans et al., 2009; Bruijnis et al., 2010; Gomez et al., 2015; Ettema et al., 2010; Relun et al., 2013). Prevention and proper treatment of DD are essential to reduce the economic burden of the disease on the US cattle industry (Orsel et al., 2018; Evans et al., 2016; Branine et al., 2014). After detection, treating DD lesions involves a combination of topical applications of antimicrobials or disinfecting agents (Shearer et al., 2015; Solano et al., 2015; Döpfer et al., 2011; Berry et al., 2010, 2012; Nishikawa and Taguchi, 2008; Hernandez et al., 1999; el-Ghoul and Shaheed, 2001; Read and Walker, 1998). Long-term management and prevention of DD often requires disinfecting footbaths and environmental remediation to control for risk factors (Shearer et al., 2015; Solano et al., 2015; Laven and Proven, 2000; Laven and Hunt, 2002; Silva et al., 2005; Teixeira et al., 2010; Thomsen et al., 2008; Speijers et al., 2012). Regular inspection of cattle for signs of DD and early detection of DD can help to reduce the risk of environmental contamination and spread of the disease (Orsel et al., 2018; Refaai et al., 2013).

The M-stage scoring system is a highly effective way to classify and monitor DD lesions and signs of chronicity (Döpfer et al., 2012). Recording M-stages, particularly at regular time intervals, enables herdsmen to improve claw health and to quantify economic and welfare losses (Döpfer et al., 2012). The M-stage scoring system classifies DD lesions into one of five stages: M0, M1, M2, M3, M4, and M4.1 (Döpfer et al., 1997; Berry et al., 2012). The combination of M-stages with signs of chronicity results in 5 recognizable stages while the foot is standing on the ground: M0, M2, M2P, M4H, and M4P (Figure 1). The first stage, M0, is characterized by normal digital skin upon visual inspection. The M2 lesions are characterized by an acute, active ulcerative or granulomatous digital skin lesion, less than 2.0 cm in diameter and M2P lesions are characterized by M2 lesions with proliferative growth. Chronic stages of DD have a thickened epithelium (hyperkeratosis) for M4H lesions or proliferative growth of the epithelium in the case of M4P lesions. This classification system helps farmers, hoof trimmers, agricultural consultants, and veterinarians gauge the transition between stages of DD and to quantify the effectiveness of current treatment and prevention protocols (Döpfer et al., 2012). The management of DD, whether endemic or during outbreak scenarios, is an ongoing challenge and early detection of DD lesions is critical in the eventual treatment of the disease (Döpfer et al., 2012).

Computer vision (CV) is a subfield of artificial intelligence (AI) using machine learning algorithms to identify objects in digital images and videos (Szeliski, 2022). This technology is capable of learning the shape, size, and color of the target object and distinguishing the target object from other objects and background (Z.-Q. Zhao et al., 2019; Wu et al., 2020; Sharma and Mir, 2020). Agricultural CV applications continue to develop at a rapid pace (Arthur F A

Fernandes et al., 2020; Arthur Francisco Araújo Fernandes et al., 2020). Object detection can be used to precisely monitor animal health and accurately diagnose a variety of medical conditions (Arthur Francisco Araújo Fernandes et al., 2020; Ryu and Tai, 2021). This technique assists in identifying illnesses and infections by means of early detection, reducing the need for invasive diagnostic tests and aiding the development of optimized treatment plans. By automating visual observation tasks, object detection is able to reduce the time needed for diagnosis using standardized clinical signs that are otherwise difficult to see and may go untreated. Early detection and automatization are particularly important for underserved and underprivileged areas where veterinary services for food animal production is increasingly scarce.

You Only Look Once (YOLO) is a CV algorithm for real-time object detection and classification. YOLO splits the input image into a grid of cells (Redmon et al., 2016). Each cell runs a classifier to determine the type of object using a single neural network (Redmon et al., 2016). The model output is a bounding box around each object to detect the specific location of each object, a class label to classify the type of object, and a class probability to indicate the confidence scores for each box (Redmon et al., 2016).

YOLO models have several advantages over other object detection algorithms which results in high accuracy and speed (Redmon et al., 2016). YOLO looks at the entire image at once during training and testing of the detection model, resulting in prediction of bounding boxes and class probabilities simultaneously (Redmon et al., 2016). YOLO is able to leverage pre-trained weights on ImageNet for transfer learning, saving resources and improving efficiency to train a new model (Redmon et al., 2016). YOLO is energy efficient in terms of computational costs, resulting in implementation on low-power, portable devices (Redmon et al., 2016).

YOLO achieves best performance on a high-end graphics processing unit (GPU), but is limited by high memory requirements, specifically on portable devices (Redmon, 2018; Redmon and Farhadi, 2018). Lightweight network architectures are needed for constrained environments, such as portable devices, resulting in reduced model size and computation cost (Redmon, 2018). Tiny YOLO is the compressed version of YOLO with a simpler network structure and a reduced number of corresponding parameters (Adarsh et al., 2020). For real-time object detection, Tiny YOLO is the better option compared to YOLO, where speed is prioritized over precision or accuracy (Adarsh et al., 2020). The speed of Tiny YOLOv4 is approximately eight times faster than YOLOv4 (Alexey, 2021; Bochkovskiy et al., 2020; Techzizou, 2021). However, the accuracy for YOLOv4 is 33% better than Tiny YOLOv4 when tested on the MS COCO dataset (Alexey, 2021; Bochkovskiy et al., 2020; Techzizou, 2021).

A YOLOv2 model was previously trained on a database of more than 3,500 DD lesion images to detect M-stages of DD (Cernek et al., 2020). The resulting predicting accuracy was 71% and the agreement between the human investigator and CV model was quantified as “moderate” by Cohen’s kappa during internal validation (Cernek et al., 2020). In the external validation, the YOLOv2 model detected DD with an accuracy of 88% and agreement between a human investigator and CV model was quantified as “fair” by Cohen’s kappa (Cernek et al., 2020).

There have been multiple subsequent versions of YOLO models. The YOLOv3 model builds on the YOLOv2 framework by making detections at three different scales and predicts ten times more bounding boxes to improve performance on smaller objects (Redmon and Farhadi, 2018). YOLOv4 improves on the average precision (AP) by 10% and on the frames per second (FPS) by 12% compared to YOLOv3 on the COCO dataset (Bochkovskiy et al., 2020). For the purpose of our use case, Tiny YOLOv4 and Single Shot MultiBox Detector (SSD) models are trained to detect M-stages of DD, because of higher speed and lower size compared to other computer vision algorithms for use on portable, edge devices (Redmon and Farhadi, 2018; Bochkovskiy et al., 2020; Liu et al., 2016).

Edge devices are used to collect, process, analyze, and report data as close as possible to the source data for data-driven decisions. Edge devices are essential for real-time detection or extremely rapid results on location where data transmission from the camera or sensor to the cloud takes time (Cao et al., 2020; Chen and Ran, 2019; Murshed et al., 2022). Edge lets users process data faster and allows them to make split-second decisions and complete tasks efficiently (Cao et al., 2020; Chen and Ran, 2019; Murshed et al., 2022). They are typically deployed on the edge of network infrastructures, such as mobile networks and Internet of Things (IoT) devices, to better support distributed data processing (Cao et al., 2020; Chen and Ran, 2019; Pan and McElhannon, 2018; Murshed et al., 2022). Edge devices are characterized by their low power consumption, high portability, and ability to perform real-time data processing (Cao et al., 2020; Chen and Ran, 2019; Pan and McElhannon, 2018; Murshed et al., 2022).

The popularity of edge devices is due to the various advantages for a wide-array of use cases. For example, edge devices that do not require internet connectivity make the technology ideal for use in remote, rural locations such as cattle feedyards and dairy farms (Cao et al., 2020; Chen and Ran, 2019; Pan and McElhannon, 2018; Murshed et al., 2022). Edge devices provide increased scalability and flexibility, allowing users to quickly adjust to changing demands (Cao et al., 2020; Chen and Ran, 2019; Pan and McElhannon, 2018). Edge devices can process and store data locally, such as real-time video and audio that are otherwise unable to be sent over wide area networks (Chen and Ran, 2019; Pan and McElhannon, 2018; Murshed et al., 2022). By processing and analyzing data at the source, edge devices can reduce latency, eliminating the need to transfer large data sets over long distances and helping improve overall response times (Cao et al., 2020; Chen and Ran, 2019; Pan and McElhannon, 2018; Murshed et al., 2022).

An example of edge devices are single-board computers (SBCs) used in IoT scenarios because of low power consumption, tiny volume, and high performance. SBCs are combined with various edge accelerators to accelerate machine learning algorithms on edge devices. There are three ways to combine SBCs and edge accelerators: (1) GPU-based SBCs, that are suitable for general-purpose applications because of their good compatibility, such as NVIDIA Jetson Nano and NVIDIA Xavier NX; (2) Application Specific Integrated Circuit (ASIC)-based SBCs, that specialize in deep neural network applications, such as Intel's vision processing unit (VPU) and Google's tensor processing unit (TPU); and (3) Field Programmable Gate Array (FPGA)-based SBCs, that are more energy efficient (Feng et al., 2021; Nair et al., 2020; Intel, 2022, 2023; OpenVINO, 2023a, 2023b; Intel, 2021a, 2021b; Coral, 2020a, 2020b, 2020c; Attaran et al., 2018). This work will be carried out on GPU-based SBCs (NVIDIA Jetson Nano and NVIDIA Xavier NX) and an ASIC-based SBC (OAK-1 and OAK-D-Lite with Intel Neural Compute Stick2).

Previous edge detection device evaluation has been completed using the NVIDIA Jetson family and Raspberry Pi platform (Magalhães et al., 2023). Object detection models on Raspberry Pi platform were too slow for real-time detection (Suzen and Duman, 2020). NVIDIA Jetson TX2 and NVIDIA GTX Titan X was benchmarked with similar accuracy for both (H. Zhao et al., 2019). Inference speed of the NVIDIA Jetson TX2, running at 18 FPS, was ten times slower, but consumed 20 times less power. Additional research to benchmark machine learning models on NVIDIA Jetson boards, also showed an inverse relationship between inference speed and power consumption where the NVIDIA Jetson AGX Xavier had the highest inference speed and the

NVIDIA Jetson Nano with the lower power consumption (Suzen and Duman, 2020; Chiu et al., 2020; Rahmaniar and Hernawan, 2021; Panero Martinez et al., 2021). Therefore, there exist a tradeoff between performance and other costs on edge devices to achieve adequate accuracy at the requisite speed for the use case.

The most benchmarked object detection models include SSD MobileNet networks family and YOLO family. YOLOLACT model achieved 66 FPS on an NVIDIA Jetson AGX Xavier and 16 FPS on an NVIDIA Jetson TX2 (Panero Martinez et al., 2021). SSD model benchmarking have achieved best performance on the NVIDIA Jetson TX2 (Chiu et al., 2020; Rahmaniar and Hernawan, 2021). Additionally, there exist implementations of object detection models on edge devices in medicine and agriculture. A skin lesion image classification model was deployed on a Raspberry Pi to provide a handheld diagnostic support tool (Sahu et al., 2018). A person-recognition system to improve safety of autonomous tractors was implemented using a YOLOv3 model on a Jetson AGX Xavier with faster speed and higher precision compared to a general YOLO model trained on the COCO dataset (Jung et al., 2020).

The study aims to train lightweight CV models for constrained environments and compare edge devices for the real-time detection of DD in dairy cows. CV models are trained for detection and scoring of DD, compared using performance metrics and inference time, and automated for real-time detection using images and video streams on portable devices. Edge devices offer advantages in terms of scalability, flexibility, energy efficiency, and privacy. Furthermore, edge devices can reduce latency and provide businesses with the ability to process highly localized data streams. Object detection in precision farming and veterinary medicine can improve the accuracy of diagnoses for health outcomes. It can help producers, herdsmen, and veterinarians make more informed decisions regarding the diagnosis that would otherwise go untreated and better customize treatment plans according to severity and dynamics of the disease. We hope the proposed tool can be employed in combination with current practice for prevention of DD in dairy cattle. Ultimately, object detection on the edge can help identify high risk cattle for DD, monitoring herds with endemic DD, and improve cattle welfare via on-farm decision-making processes (Plummer and Krull, 2017).

MATERIALS AND METHODS

Data Collection and Image Labeling

We used a collection of 409 images of cattle feet for the study from farms in Wisconsin and images were labeled for M-stages of DD. GoPro Hero 5 Black cameras were used to take MP4 video recordings of the backside of the hind feet at claw level, while standing on the ground. Videos were taken during May 2021 with feet located in cattle pens, standing in an automated milking system, or while exiting a disinfecting foot bath.

Images were scored for M-stages of DD by a trained investigator using the M-stage DD classification system. The classification system describes various clinical stages over the course of the disease, based on morphological observations between healthy (M0), active (M2), and chronic (M4) with hyperkeratosis (H) or proliferation (P) stages (Döpfer et al., 1997; Berry et al., 2012). The images were classified into the different stages of DD as follows: M0, M2, M2P, M4H, and M4P.

Overall, the library of multiple feet per image and five labels consisted of 240 M0, 17 M2, 51 M2P, 114 M4H, and 108 M4P class labels (409 screenshot JPG images from MP4 video recordings). The associated annotations for the bounding boxes and class labels were generated using LabelImg version 1.8.1 in Python version 3.8.15 (Lin, 2018). Labels were converted from

Pascal VOC format for R-CNN and SSD models to YOLO format for YOLO models (Tashiev, 2022).

Model Training

The datasets were split into 90% of the images placed into a training set and 10% of the images placed into a testing set at random. The testing set contained 40 images consisting of 20 M0, 4 M2, 4 M2P, 14 M4H, and 10 M4P class labels. Models were trained and tested using Google Colab and a 12GB NVIDIA Tesla K80 GPU with a batch size of 64 and subdivision size of 16 ("Google Colaboratory," 2022). The training weights were initialized using convolutional weights that were pre-trained on Imagenet ("ImageNet," 2020; Krizhevsky et al., 2012). The framework used for training is TensorFlow 1.15.2.

SSD uses MobileNet V2 as the backbone feature extractor from the Tensorflow Object Detection API model zoo. The SSD network accepts images that are resized to an input resolution of 300 by 300 pixels (Liu et al., 2016). SSDLite is similar to SSD, but implements depthwise-separable convolutions rather than regular convolution layers, which increases the speed compared to the regular SSD making it perfectly suited for use on mobile devices (Fan et al., 2018).

The YOLOv3 network uses a variant of the custom deep architecture Darknet, originally a 53-layer network trained on Imagenet with a 53-layer stack of neurons added for the task of detection to result in a total of a 106-layer fully convolutional underlying architecture (Redmon et al., 2016; Redmon and Farhadi, 2016, 2018). The YOLOv4 network uses CSPDarknet53 as the backbone, a spatial pyramid pooling additional module, PANet path-aggregation neck, and a YOLOv3 head (Bochkovskiy et al., 2020). The YOLO network size height and width is 416 by 416 pixels with a learning rate of 0.001 for a maximum number of 4,000 batches.

Model Embedding

The Tiny YOLOv4 best weights file from the training step was converted to a TensorFlow frozen model using OpenVINO-YOLOV4. The TensorFlow frozen model was converted to an OpenVINO 21.03 IRv10 model and FP16 data type using OpenVINO 2021.3 for optimizing and deploying AI inference. Model optimizer converts other model formats to OpenVINO's IR format, which produces xml and bin files. This format of the model can be deployed across multiple Intel devices including VPU. The TinyIR model was compiled to an IR model to a MyriadX blob with FP16 precision and 5 shave cores for use on DepthAI modules/platform. DepthAI is a Spatial AI platform built around Movidius VPU to running custom trained models on OAK cameras. Models were converted using Google Colab and a 2.30GHz 12GB Intel Xeon Dual-Core vCPU ("Google Colaboratory," 2022).

RESULTS

Edge Device

The build of the portable edge device for deployment and demonstration contains four main components: Jetson Xavier NX SBC, OAK-1 camera, an LCD screen, and a portable 12V/5mA power source (Figure 2). The four constituent parts are enclosed in a hard plastic container to protect it from the elements (Figure 3).

Jetson Xavier NX module delivers server-class performance from 21 Tera Operations Per Second (TOPS) at 15W or 20W or up to 14 TOPS at 10W. It is smaller than a credit card (70 x 45mm), extremely energy-efficient, and can run multiple modern neural networks in parallel and process data from multiple high-resolution sensors (NVIDIA, 2023). Jetson Xavier NX is production-ready and supports popular AI frameworks. This opens new possibilities for embedded edge-computing devices that demand increased performance to support AI workloads but are constrained by size, weight, power budget, or cost (NVIDIA, 2023).

The OAK-1 adds USB3 Type-C device power and connectivity, and a single 12MP RGB camera module (Luxonis, 2023a). The integrated camera module communicates over an on-board interface directly with the Robotics Vision Core 2, that processes the data and performs object detection, returning the results over USB (Luxonis, 2023a). Such a data path relieves the host processor from all of this work. In the common use case of object detection from an image, the host processes the data of the objects and the location in the image, instead of the entire video.

The LCD screen is used to display the results in a window. The set-up can run without dedicated LCD screen as well, where the detection results are collected on a remotely connected PC. The device can be plugged into a wall outlet or a portable power bank using an AC power adaptor or USB-C cable respectively. DepthAI is installed and a Python script is prompted at the command line to open a preview of the RGB camera from the device. While the Python script is running, the detection results were displayed as an overlay on the video stream. The Tiny YOLOv4 object detector trained on the M-stages of DD was deployed on the OAK-1 camera.

Benchmarking

The network performance and efficiency on the edge device are evaluated using accuracy and inference speed. For assessing model training accuracy, mean average precision (mAP) at intersection over union (IOU) of 0.5 are evaluated. Alternatively, for assessing model deployment accuracy, Cohen’s kappa is evaluated. Both stages are measured for inference speed using frames per second (FPS).

Both Tiny YOLOv4 and SSD are first tested using images on Google Colab. Tiny YOLOv4 outperforms SSD with a speed of 333 FPS compared to approximately 100 FPS, validating previous findings. In addition, Tiny YOLOv4 outperforms SSD where Tiny YOLOv4 achieves an mAP of 0.895 whereas SSD and SSD Lite yields an mAP of 0.451, similarly validating previous findings. For the purpose of real-time detection on a portable, stand-alone device, Tiny YOLOv4 is the best model for the use case, based on the size and speed of the implementation.

Tiny YOLOv4 is then tested on the edge device using live streaming video (Figure 4). Images are at a sufficient distance from the camera such that the object in the image is focused and does not present glare or darkness. Video is scored for M-stages of DD and compared to the corresponding labels made by a trained investigator (the so-called ‘ground truth’). Tiny YOLOv4 is extremely fast and accurate for real-time object detection. The Cohen’s kappa is used to measure the agreement of two raters (Viera and Garrett, 2005). The model performs well compared to the ground truth where the Cohen’s kappa is determined to be 0.830 and interpreted as strong to almost perfect agreement between the two raters ($n = 55$; $z = 11.4$; $p < 0.001$). In addition, the prediction accuracy is 0.873 with perfect detection for M2P, M4H, and M4P (Figure 5). Overall, Tiny YOLOv4 on the edge device is able to detect all five M-stages of DD on streaming video.

The initial inference speed is 25 FPS and it climbs to the final inference speed of 40 FPS over the course of 10 minutes. The model on the edge device runs stably for 8 hours and longer without any interruptions. The power usage for OAK-1 cameras ranges between 1.94 W (standby) and 4.56 W (max consumption). The OAK-1 camera also resolves issues regarding heat dissipation by running the CV model on the OAK-1 camera instead of the NVIDIA Jetson Xavier NX. The Robotics Vision Core 2 on the OAK-1 camera is rated for industrial use and has an operating temperature range of -40°C to 105°C, while other components have a higher temperature range (Luxonis, 2023b). In general, for the use case at an ambient temperature of 25°C, the chip running at full power (i.e. worst-case heat generation) is approximately 70°C or a thermal margin of 45°C. Theoretically, the maximum ambient temperature is the difference between the maximum operating temperature and the thermal margin or 60°C (Luxonis, 2023b).

DISCUSSION

The DD detection device developed in this study was able to identify DD lesions and classify M-stages in real-time with high accuracy and speed. For speed, Tiny YOLOv4 outperformed the threshold for real-time detection with a high inference speed of approximately 40 FPS. For detection, Tiny YOLOv4 was able to perform real-time detection of DD with high mAP on Google Colab and high Cohen's kappa on the edge device. The Tiny YOLOv4 model on the edge device had better performance than the YOLOv2 model for additional M-stages of DD on more complex dataset analogous to practical scenarios (Cernek et al., 2020). As a proof-of-concept, the Tiny YOLOv4 model was able to detect the different M-stages of DD on video files and streaming video.

The current model was implemented using Tiny YOLOv4 and the TensorFlow 1.0 framework with high speed and performance. TensorFlow 2.0 framework is now available and offers improved usability and performance (NVIDIA Corporation, 2022; tensorflow, 2017). TensorFlow models can be converted to TensorFlow Lite models reducing model size and power consumption making them easier to deploy to edge devices (tensorflow, 2021). Alternatively, the model architecture can be updated to the PyTorch framework using YOLOv5 or YOLOv7 and likely improving speed and performance (Paszke et al., 2019; pytorch, 2021).

Tiny YOLOv4 was able to detect all five class labels well. However, the model was able to detect M2P, M4H, and M4P lesions with higher average precision than M2. This was probably caused by the imbalanced number of class labels in the dataset with relatively few M2 training examples. Model performance can suffer when the number of training examples available for each training class is unequal (Oksuz et al., 2021). Remedial measures for improving model performance include adding images for minority classes, artificial dataset balancing, image augmentation, or adjusting the loss function for training.

Object detection on edge devices is subject to a number of constraints. Edge devices are generally less costly, but also less powerful than cloud-based systems. This means that the performance for object detection tasks on edge devices is less than on cloud-based systems. Many CV algorithms used for object detection are computationally expensive, and more sophisticated algorithms require more data and training data. Edge devices may not have access to a consistent power source, resulting in the device to restart frequently or to lose work progress. Another challenge for edge devices is that the environmental conditions can affect object detection on an edge device, such as recognizing or classifying an object in low light, shadows, and other optical occlusions. Edge devices may require advanced sensors that are necessary for differentiating between objects in complex environments. Data collection and storage on edge devices may have slower rates and larger latency due to connectivity of advanced sensors. This means that the amount of data that can be processed at once on an edge device is limited.

A wide variety of precision agriculture tools are available, primarily in agricultural applications (Arthur Francisco Araújo Fernandes et al., 2020). Improving edge device implementation of animal health CV models enhances their adoption within animal agriculture and veterinary medicine. The CV models developed in this work can be extended to other portable platforms and distributions, including lighter-weight framework and smaller size of enclosure. Increasing accessibility to these models can be accomplished through mobile applications or containers for a lightweight, standalone package (Docker, 2021). Increasing engagement for these models will also help generate a rich, diverse library of images for the optimization and validation of CV models. Ultimately, these CV tools will improve access and enhance development speed for animal agriculture applications.

As currently implemented, this portable CV tool is used for automated detection of DD in dairy cattle and has been extended to beef cattle (unpublished data). Model development can be extended to include model training on a wide variety of breeds, locations, and other sources of data variability. Ear tag identification paired with DD detection improves the ability to create automated treatment lists and monitoring data. The scope of research can be expanded to include additional M-stages of DD, white line disease, foot rot, and other hoof diseases (Lely, 2016).

The development and deployment of CV algorithms for automatic monitoring and measuring traits of interest in animals is a widely investigated topic (Arthur Francisco Araújo Fernandes et al., 2020; Li et al., 2021; Wurtz et al., 2019). Studies indicate that different imaging technologies can improve model performance. Infrared thermography has shown acceptable performance for DD detection, while animal identification using 3D cameras often outperforms standard cameras. As technology continues to advance, these tools will be more widely implemented.

There are still several challenges for the successful development and deployment of practical CV solutions. Current challenges include the implementation of CV algorithms for true on-farm deployment. There are still few studies that evaluated CV algorithms using validation datasets for reliability and robustness across multiple farms (Arthur Francisco Araújo Fernandes et al., 2020; Li et al., 2021; Wurtz et al., 2019). Individual animal identification and tracking is still prone to error. Combining multiple devices into a single application is also necessary for true on-farm deployment of many CV models (Arthur Francisco Araújo Fernandes et al., 2020; Li et al., 2021; Wurtz et al., 2019). This may enable the implementation of more sophisticated predictive algorithms based on multiple inputs and multiple outputs (joint prediction of multiple traits). The final challenge is delivering CV model predictions to farmers in a convenient and actionable manner (Arthur Francisco Araújo Fernandes et al., 2020; Li et al., 2021; Wurtz et al., 2019). Long-term improvements in each of these areas will drive adoption of these technologies and development of new applications.

ACKNOWLEDGEMENTS

Funding support for the research was provided by the US Department of Agriculture through the National Institute for Food and Agriculture - Animal Health Grant (WIS03082). We would also like to thank Jamie Sullivan (Carman, MB, Canada), Chris Bauer (Mondovi, WI), and all other dairy producers for contributing to the images of hoof lesions. The authors have no financial and personal relationships with other people or organizations that could inappropriately influence (bias) their work.

REFERENCES

Adarsh, P., Rathi, P., Kumar, M., 2020. YOLO v3-Tiny: Object Detection and Recognition using one stage improved model, in: 2020 6th International Conference on Advanced Computing and Communication Systems (ICACCS). Presented at the 2020 6th International Conference on Advanced Computing and Communication Systems (ICACCS), pp. 687–694. https://doi.org/10.1109/ICACCS48705.2020.9074315

Alexey, 2021. Yolo v4, v3 and v2 for Windows and Linux.

Attaran, N., Puranik, A., Brooks, J., Mohsenin, T., 2018. Embedded Low-Power Processor for Personalized Stress Detection. IEEE Transactions on Circuits and Systems II: Express Briefs 65, 2032–2036. https://doi.org/10.1109/TCSII.2018.2799821

Berry, S.L., Read, D.H., Famula, T.R., Mongini, A., Döpfer, D., 2012. Long-term observations on the dynamics of bovine digital dermatitis lesions on a California dairy after topical treatment with lincomycin HCl. The Veterinary Journal, Special Issue:Bovine Lameness 193, 654–658. https://doi.org/10.1016/j.tvjl.2012.06.048

Berry, S.L., Read, D.H., Walker, R.L., Famula, T.R., 2010. Clinical, histologic, and bacteriologic findings in dairy cows with digital dermatitis (footwarts) one month after topical treatment with lincomycin hydrochloride or oxytetracycline hydrochloride. Journal of the American Veterinary Medical Association 237, 555–560. https://doi.org/10.2460/javma.237.5.555

Blowey, R.W., Sharp, M.W., 1988. Digital dermatitis in dairy cattle. Vet Rec 122, 505–508. https://doi.org/10.1136/vr.122.21.505

Bochkovskiy, A., Wang, C.-Y., Liao, H.-Y.M., 2020. YOLOv4: Optimal Speed and Accuracy of Object Detection. arXiv:2004.10934 [cs, eess].

Branine, M., Döpfer, D., Edwards, T., Larson, C., Mülling, C., Tomlinson, D., 2014. Cattle lameness: identification, prevention and control of claw lesions.

Bruijnis, M.R.N., Beerda, B., Hogeveen, H., Stassen, E.N., 2012. Assessing the welfare impact of foot disorders in dairy cattle by a modeling approach. animal 6, 962–970. https://doi.org/10.1017/S1751731111002606

Bruijnis, M.R.N., Hogeveen, H., Stassen, E.N., 2010. Assessing economic consequences of foot disorders in dairy cattle using a dynamic stochastic simulation model. Journal of Dairy Science 93, 2419–2432. https://doi.org/10.3168/jds.2009-2721

Cao, K., Liu, Y., Meng, G., Sun, Q., 2020. An Overview on Edge Computing Research. IEEE Access 8, 85714–85728. https://doi.org/10.1109/ACCESS.2020.2991734

Cernek, P., Bollig, N., Anklam, K., Döpfer, D., 2020. Hot topic: Detecting digital dermatitis with computer vision. Journal of Dairy Science 103, 9110–9115. https://doi.org/10.3168/jds.2019-17478

Cha, E., Hertl, J.A., Bar, D., Gröhn, Y.T., 2010. The cost of different types of lameness in dairy cows calculated by dynamic programming. Preventive Veterinary Medicine 97, 1–8. https://doi.org/10.1016/j.prevetmed.2010.07.011

Cheli, R., Mortellaro, C., 1974. La dermatite digitale del bovino, in: Proceedings of the 8th International Conference on Diseases of Cattle. Piacenza, pp. 208–213.

Chen, J., Ran, X., 2019. Deep Learning With Edge Computing: A Review. Proceedings of the IEEE 107, 1655–1674. https://doi.org/10.1109/JPROC.2019.2921977

Chiu, Y.-C., Tsai, C.-Y., Ruan, M.-D., Shen, G.-Y., Lee, T.-T., 2020. Mobilenet-SSDv2: An Improved Object Detection Model for Embedded Systems, in: 2020 International Conference on System Science and Engineering (ICSSE). Presented at the 2020 International Conference on System Science and Engineering (ICSSE), pp. 1–5. https://doi.org/10.1109/ICSSE50014.2020.9219319

Coral, 2020a. USB Accelerator [WWW Document]. Coral. URL https://coral.ai/products/accelerator/ (accessed 4.8.23).

Coral, 2020b. Get started with the USB Accelerator [WWW Document]. Coral. URL https://coral.ai/docs/accelerator/get-started/ (accessed 4.8.23).

Coral, 2020c. USB Accelerator datasheet [WWW Document]. Coral. URL https://coral.ai/docs/accelerator/datasheet/ (accessed 4.8.23).

de Jesús Argáez-Rodríguez, F., Hird, D.W., de Anda, J.H., Read, D.H., Rodríguez-Lainz, A., 1997. Papillomatous digital dermatitis on a commercial dairy farm in Mexicali, Mexico: Incidence and effect on reproduction and milk production. Preventive Veterinary Medicine 32, 275–286. https://doi.org/10.1016/S0167-5877(97)00031-7

Docker, 2021. Docker Documentation [WWW Document]. Docker Documentation. URL https://docs.docker.com/

Döpfer, D., Holzhauer, M., Boven, M. van, 2012. The dynamics of digital dermatitis in populations of dairy cattle: Model-based estimates of transition rates and implications for control. The Veterinary Journal, Special Issue:Bovine Lameness 193, 648–653. https://doi.org/10.1016/j.tvjl.2012.06.047

Döpfer, D., Huurne, A.A.H.M. ter, Cornelisse, J.L., Asten, A.J.A.M. van, Koopmans, A., Meijer, F.A., Schukken, Y.H., Szakáll, I., Klee, W., Bosma, R.B., 1997. Histological and bacteriological evaluation of digital dermatitis in cattle, with special reference to spirochaetes and Campylobacter faecalis. Veterinary Record 140, 620–623. https://doi.org/10.1136/vr.140.24.620

Döpfer, D., Lopez-Benavides, M., Buchalova, M., Mateus, C., Cook, N., Dusick, A., Hemling, T., Socha, M., Read, D., Gomez, A., 2011. Clinical, Histopathological, and microbiological cure ofacute digital dermatitis lesions and the influence of topical treatments unpon cure, in: Proceedings. 16th Lameness in Ruminants Symposium.

el-Ghoul, W., Shaheed, B.I., 2001. Ulcerative and papillomatous digital dermatitis of the pastern region in dairy cattle: clinical and histopathological studies. Dtsch Tierarztl Wochenschr 108, 216–222.

Ettema, J., Østergaard, S., Kristensen, A.R., 2010. Modelling the economic impact of three lameness causing diseases using herd and cow level evidence. Preventive Veterinary Medicine 95, 64–73. https://doi.org/10.1016/j.prevetmed.2010.03.001

Evans, N.J., Brown, J.M., Demirkan, I., Singh, P., Getty, B., Timofte, D., Vink, W.D., Murray, R.D., Blowey, R.W., Birtles, R.J., Hart, C.A., Carter, S.D., 2009. Association of Unique, Isolated Treponemes with Bovine Digital Dermatitis Lesions. Journal of Clinical Microbiology 47, 689–696. https://doi.org/10.1128/JCM.01914-08

Evans, N.J., Murray, R.D., Carter, S.D., 2016. Bovine digital dermatitis: Current concepts from laboratory to farm. The Veterinary Journal 211, 3–13. https://doi.org/10.1016/j.tvjl.2015.10.028

Fan, H., Liu, Shuanglong, Ferianc, M., Ng, H.-C., Que, Z., Liu, Shen, Niu, X., Luk, W., 2018. A Real-Time Object Detection Accelerator with Compressed SSDLite on FPGA, in: 2018 International Conference on Field-Programmable Technology (FPT). Presented at the 2018 International Conference on Field-Programmable Technology (FPT), IEEE, Naha, Okinawa, Japan, pp. 14–21. https://doi.org/10.1109/FPT.2018.00014

Feng, H., Mu, G., Zhong, S., Zhang, P., Yuan, T., 2021. Benchmark Analysis of YOLO Performance on Edge Intelligence Devices, in: 2021 Cross Strait Radio Science and Wireless Technology Conference (CSRSWTC). Presented at the 2021 Cross Strait Radio Science and Wireless Technology Conference (CSRSWTC), pp. 319–321. https://doi.org/10.1109/CSRSWTC52801.2021.9631594

Fernandes, Arthur Francisco Araújo, Dórea, J.R.R., Rosa, G.J. de M., 2020. Image Analysis and Computer Vision Applications in Animal Sciences: An Overview. Frontiers in Veterinary Science 7.

Fernandes, Arthur F A, Dórea, J.R.R., Valente, B.D., Fitzgerald, R., Herring, W., Rosa, G.J.M., 2020. Comparison of data analytics strategies in computer vision systems to predict pig body composition traits from 3D images. Journal of Animal Science 98, skaa250. https://doi.org/10.1093/jas/skaa250

Gomez, A., 2018. What We Know and Don’t Know About Digital Dermatitis in Cattle. Zinpro®. URL https://www.zinpro.com/what-we-know-and-dont-know-about-digital-dermatitis-in-cattle/ (accessed 4.4.23).

Gomez, A., Cook, N.B., Socha, M.T., Döpfer, D., 2015. First-lactation performance in cows affected by digital dermatitis during the rearing period. Journal of Dairy Science 98, 4487–4498. https://doi.org/10.3168/jds.2014-9041

Google Colaboratory [WWW Document], 2022. URL https://colab.research.google.com/?utm_source=scs-index (accessed 2.24.22).

Hernandez, J., Shearer, J.K., Elliott, J.B., 1999. Comparison of topical application of oxytetracycline and four nonantibiotic solutions for treatment of papillomatous digital dermatitis in dairy cows. J Am Vet Med Assoc 214, 688–690.

Hernandez, J., Shearer, J.K., Webb, D.W., 2001. Effect of lameness on the calving-to-conception interval in dairy cows. Journal of the American Veterinary Medical Association 218, 1611–1614. https://doi.org/10.2460/javma.2001.218.1611

ImageNet [WWW Document], 2020. URL https://www.image-net.org/ (accessed 2.24.22).

Intel, 2023. Intel® Movidius™ Vision Processing Units (VPUs) [WWW Document]. Intel. URL https://www.intel.com/content/www/us/en/products/details/processors/movidius-vpu.html (accessed 4.8.23).

Intel, 2022. Intel® Vision Accelerator Design with Intel® Movidius™ VPU [WWW Document]. Intel. URL https://www.intel.com/content/www/us/en/developer/topic-technology/edge-5g/hardware/vision-accelerator-movidius-vpu.html (accessed 4.5.23).

Intel, 2021a. Intel® Movidius™ Myriad™ X VPU Product Brief [WWW Document]. Intel. URL https://www.intel.com/content/www/us/en/products/docs/processors/movidius-vpu/myriad-x-product-brief.html (accessed 4.8.23).

Intel, 2021b. Intel® Myriad™ X VPU Product Brief [WWW Document]. Intel. URL https://www.intel.com/content/www/us/en/content-details/760604/intel-myriad-x-vpu-product-brief.html (accessed 4.8.23).

Jung, T.-H., Cates, B., Choi, I.-K., Lee, S.-H., Choi, J.-M., 2020. Multi-Camera-Based Person Recognition System for Autonomous Tractors. Designs 4, 54. https://doi.org/10.3390/designs4040054

Krizhevsky, A., Sutskever, I., Hinton, G.E., 2012. ImageNet Classification with Deep Convolutional Neural Networks, in: Advances in Neural Information Processing Systems. Curran Associates, Inc.

Laven, R.A., Hunt, H., 2002. Evaluation of copper sulphate, formalin and peracetic acid in footbaths for the treatment of digital dermatitis in cattle. Veterinary Record 151, 144–146. https://doi.org/10.1136/vr.151.5.144

Laven, R.A., Proven, M.J., 2000. Use of an antibiotic footbath in the treatment of bovine digital dermatitis. Veterinary Record 147, 503–506. https://doi.org/10.1136/vr.147.18.503

Lely, 2016. Claw diseases - Lely [WWW Document]. URL https://www.lely.com/us/farming-insights/claw-diseases/ (accessed 4.8.23).

Li, G., Huang, Y., Chen, Z., Chesser, G.D., Purswell, J.L., Linhoss, J., Zhao, Y., 2021. Practices and Applications of Convolutional Neural Network-Based Computer Vision Systems in Animal Farming: A Review. Sensors 21, 1492. https://doi.org/10.3390/s21041492

Lin, T., 2018. LabelImg.

Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.-Y., Berg, A.C., 2016. SSD: Single Shot MultiBox Detector. arXiv:1512.02325 [cs] 9905, 21–37. https://doi.org/10.1007/978-3-319-46448-0_2

Losinger, W.C., 2006. Economic impacts of reduced milk production associated with papillomatous digital dermatitis in dairy cows in the USA. Journal of Dairy Research 73, 244–256. https://doi.org/10.1017/S0022029906001798

Luxonis, 2023a. OAK-1 — DepthAI Hardware Documentation 1.0.0 documentation [WWW Document]. URL https://docs.luxonis.com/projects/hardware/en/latest/pages/BW1093.html (accessed 4.8.23).

Luxonis, 2023b. Operative temperature range — DepthAI Hardware Documentation 1.0.0 documentation [WWW Document]. URL https://docs.luxonis.com/projects/hardware/en/latest/pages/articles/operative_temperature_range.html (accessed 4.8.23).

Magalhães, S.C., dos Santos, F.N., Machado, P., Moreira, A.P., Dias, J., 2023. Benchmarking edge computing devices for grape bunches and trunks detection using accelerated object detection single shot multibox deep learning models. Engineering Applications of Artificial Intelligence 117, 105604. https://doi.org/10.1016/j.engappai.2022.105604

Murshed, M.G.S., Murphy, C., Hou, D., Khan, N., Ananthanarayanan, G., Hussain, F., 2022. Machine Learning at the Network Edge: A Survey. ACM Comput. Surv. 54, 1–37. https://doi.org/10.1145/3469029

Nair, D., Pakdaman, A., Plöger, P.G., 2020. Performance Evaluation of Low-Cost Machine Vision Cameras for Image-Based Grasp Verification. https://doi.org/10.48550/arXiv.2003.10167

Nishikawa, A., Taguchi, K., 2008. Healing of digital dermatitis after a single treatment with topical oxytetracycline in 89 dairy cows. Veterinary Record 163, 574–576. https://doi.org/10.1136/vr.163.19.574

NVIDIA, 2023. The World's Smallest AI Supercomputer [WWW Document]. NVIDIA. URL https://www.nvidia.com/en-us/autonomous-machines/embedded-systems/jetson-xavier-nx/ (accessed 4.8.23).

NVIDIA Corporation, 2022. TensorRT Open Source Software. NVIDIA Corporation.

Oksuz, K., Cam, B.C., Kalkan, S., Akbas, E., 2021. Imbalance Problems in Object Detection: A Review. IEEE Transactions on Pattern Analysis and Machine Intelligence 43, 3388–3415. https://doi.org/10.1109/TPAMI.2020.2981890

OpenVINO, 2023a. VPU Devices — OpenVINO™ documentation [WWW Document]. URL https://docs.openvino.ai/latest/openvino_docs_OV_UG_supported_plugins_VPU.html?wapkw=vpu (accessed 4.8.23).

OpenVINO, 2023b. How to Implement Custom Layers for VPU (Intel® Neural Compute Stick 2) — OpenVINO™ documentation [WWW Document]. URL https://docs.openvino.ai/latest/openvino_docs_Extensibility_UG_VPU_Kernel.html?wapkw=vpu (accessed 4.8.23).

Orsel, K., Plummer, P., Shearer, J., De Buck, J., Carter, S.D., Guatteo, R., Barkema, H.W., 2018. Missing pieces of the puzzle to effectively control digital dermatitis. Transboundary and Emerging Diseases 65, 186–198. https://doi.org/10.1111/tbed.12729

Pan, J., McElhannon, J., 2018. Future Edge Cloud and Edge Computing for Internet of Things Applications. IEEE Internet of Things Journal 5, 439–449. https://doi.org/10.1109/JIOT.2017.2767608

Panero Martinez, R., Schiopu, I., Cornelis, B., Munteanu, A., 2021. Real-Time Instance Segmentation of Traffic Videos for Embedded Devices. Sensors 21, 275. https://doi.org/10.3390/s21010275

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S., 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library, in: Advances in Neural Information Processing Systems. Curran Associates, Inc.

Plummer, P.J., Krull, A., 2017. Clinical Perspectives of Digital Dermatitis in Dairy and Beef Cattle. Veterinary Clinics of North America: Food Animal Practice 33, 165–181. https://doi.org/10.1016/j.cvfa.2017.02.002

pytorch, 2021. pytorch/pytorch. pytorch.

Rahmaniar, W., Hernawan, A., 2021. Real-Time Human Detection Using Deep Learning on Embedded Platforms: A Review. Journal of Robotics and Control (JRC) 2, 462–468. https://doi.org/10.18196/jrc.26123

Read, D.H., Walker, R.L., 1998. Papillomatous Digital Dermatitis (Footwarts) in California Dairy Cattle: Clinical and Gross Pathologic Findings. J VET Diagn Invest 10, 67–76. https://doi.org/10.1177/104063879801000112

Redmon, J., 2018. YOLO: Real-Time Object Detection [WWW Document]. URL https://pjreddie.com/darknet/yolo/ (accessed 2.24.22).

Redmon, J., Divvala, S., Girshick, R., Farhadi, A., 2016. You Only Look Once: Unified, Real-Time Object Detection. arXiv:1506.02640 [cs].

Redmon, J., Farhadi, A., 2018. YOLOv3: An Incremental Improvement. arXiv:1804.02767 [cs].

Redmon, J., Farhadi, A., 2016. YOLO9000: Better, Faster, Stronger. arXiv:1612.08242 [cs].

Refaai, W., Van Aert, M., Abd El-Aal, A.M., Behery, A.E., Opsomer, G., 2013. Infectious diseases causing lameness in cattle with a main emphasis on digital dermatitis (Mortellaro disease). Livestock Science, Lameness and claw lesions in sows, cows and small ruminants 156, 53–63. https://doi.org/10.1016/j.livsci.2013.06.004

Relun, A., Lehebel, A., Chesnin, A., Guatteo, R., Bareille, N., 2013. Association between digital dermatitis lesions and test-day milk yield of Holstein cows from 41 French dairy farms. Journal of Dairy Science 96, 2190–2200. https://doi.org/10.3168/jds.2012-5934

Ryu, H.W., Tai, J.H., 2021. Object detection and tracking using a high-performance artificial intelligence-based 3D depth camera: towards early detection of African swine fever. J Vet Sci 23, e17. https://doi.org/10.4142/jvs.21252

Sahu, P., Yu, D., Qin, H., 2018. Apply lightweight deep learning on internet of things for low-cost and easy-to-access skin cancer detection, in: Medical Imaging 2018: Imaging Informatics for Healthcare, Research, and Applications. Presented at the Medical Imaging 2018: Imaging Informatics for Healthcare, Research, and Applications, SPIE, pp. 254–262. https://doi.org/10.1117/12.2293350

Sharma, V., Mir, R.N., 2020. A comprehensive and systematic look up into deep learning based object detection techniques: A review. Computer Science Review 38, 100301. https://doi.org/10.1016/j.cosrev.2020.100301

Shearer, J.K., Plummer, P.J., Schleining, J.A., 2015. Perspectives on the treatment of claw lesions in cattle. VMRR 6, 273–292. https://doi.org/10.2147/VMRR.S62071

Silva, L.A.F., Silva, C.A., Borges, J.R.J., Fioravanti, M.C.S., Borges, G.T., Atayde, I.B., 2005. A clinical trial to assess the use of sodium hypochlorite and oxytetracycline on the healing of digital dermatitis lesions in cattle. Can Vet J 46, 345–348.

Solano, L., Barkema, H.W., Jacobs, C., Orsel, K., 2017. Validation of the M-stage scoring system for digital dermatitis on dairy cows in the milking parlor. Journal of Dairy Science 100, 1592–1603. https://doi.org/10.3168/jds.2016-11365

Solano, L., Barkema, H.W., Mason, S., Pajor, E.A., LeBlanc, S.J., Orsel, K., 2016. Prevalence and distribution of foot lesions in dairy cattle in Alberta, Canada. Journal of Dairy Science 99, 6828–6841. https://doi.org/10.3168/jds.2016-10941

Solano, L., Barkema, H.W., Pajor, E.A., Mason, S., LeBlanc, S.J., Zaffino Heyerhoff, J.C., Nash, C.G.R., Haley, D.B., Vasseur, E., Pellerin, D., Rushen, J., de Passillé, A.M., Orsel, K., 2015. Prevalence of lameness and associated risk factors in Canadian Holstein-Friesian cows housed in freestall barns. Journal of Dairy Science 98, 6978–6991. https://doi.org/10.3168/jds.2015-9652

Speijers, M.H.M., Finney, G.A., McBride, J., Watson, S., Logue, D.N., O'Connell, N.E., 2012. Effectiveness of different footbathing frequencies using copper sulfate in the control of digital dermatitis in dairy cows. Journal of Dairy Science 95, 2955–2964. https://doi.org/10.3168/jds.2011-4683

Suzen, A., Duman, B., 2020. Benchmark analysis of jetson TX2, jetson nano and raspberry PI using deep-CNN. International Congress on Human-Computer Interaction, Optimization and Robotic Applications (HORA) 10, 1–5.

Szeliski, R., 2022. Computer Vision: Algorithms and Applications, 2nd ed. 2022 edition. ed. Springer, Cham.

Tashiev, I., 2022. XmlToTxt.

Techzizou, 2021. YOLOv4 VS YOLOv4-tiny. Analytics Vidhya. URL https://medium.com/analytics-vidhya/yolov4-vs-yolov4-tiny-97932b6ec8ec (accessed 2.24.22).

Teixeira, A.G.V., Machado, V.S., Caixeta, L.S., Pereira, R.V., Bicalho, R.C., 2010. Efficacy of formalin, copper sulfate, and a commercial footbath product in the control of digital dermatitis. Journal of Dairy Science 93, 3628–3634. https://doi.org/10.3168/jds.2010-3246

tensorflow, 2021. TensorFlow Lite for Microcontrollers. tensorflow.

tensorflow, 2017. Documentation for TensorRT in TensorFlow (TF-TRT). tensorflow.

Thomsen, P.T., Sørensen, J.T., Ersbøll, A.K., 2008. Evaluation of Three Commercial Hoof-Care Products Used in Footbaths in Danish Dairy Herds. Journal of Dairy Science 91, 1361–1365. https://doi.org/10.3168/jds.2007-0820

USDA, 2009. NAHMS dairy 2007 part IV: reference of dairy cattle health and management practices in the United States.

Viera, A.J., Garrett, J.M., 2005. Understanding interobserver agreement: the kappa statistic. Fam Med 37, 360–363.

Whay, H.R., Waterman, A.E., Webster, A.J.F., 1997. Associations between locomotion, claw lesions andnociceptive threshold in dairy heifers during the peri-partum period. The Veterinary Journal 154, 155–161. https://doi.org/10.1016/S1090-0233(97)80053-6

Wu, X., Sahoo, D., Hoi, S.C.H., 2020. Recent advances in deep learning for object detection. Neurocomputing 396, 39–64. https://doi.org/10.1016/j.neucom.2020.01.085

Wurtz, K., Camerlink, I., D'Eath, R.B., Fernández, A.P., Norton, T., Steibel, J., Siegford, J., 2019. Recording behaviour of indoor-housed farm animals automatically using machine vision technology: A systematic review. PLOS ONE 14, e0226669. https://doi.org/10.1371/journal.pone.0226669

Zhao, H., Zhang, W., Sun, H., Xue, B., 2019. Embedded Deep Learning for Ship Detection and Recognition. Future Internet 11, 53. https://doi.org/10.3390/fi11020053

Zhao, Z.-Q., Zheng, P., Xu, S.-T., Wu, X., 2019. Object Detection With Deep Learning: A Review. IEEE Transactions on Neural Networks and Learning Systems 30, 3212–3232. https://doi.org/10.1109/TNNLS.2018.2876865

TABLES AND FIGURES

Figure 3.1 M-stage scoring system of digital dermatitis (DD) with signs of chronicity. Images are presented for each of the M-stages: M0 (top-left), M2 (top-center), M2P (top-left), M4H (bottom-left), and M4P (bottom-right).

Figure 3.2 Schematic representation of edge device for deployment.

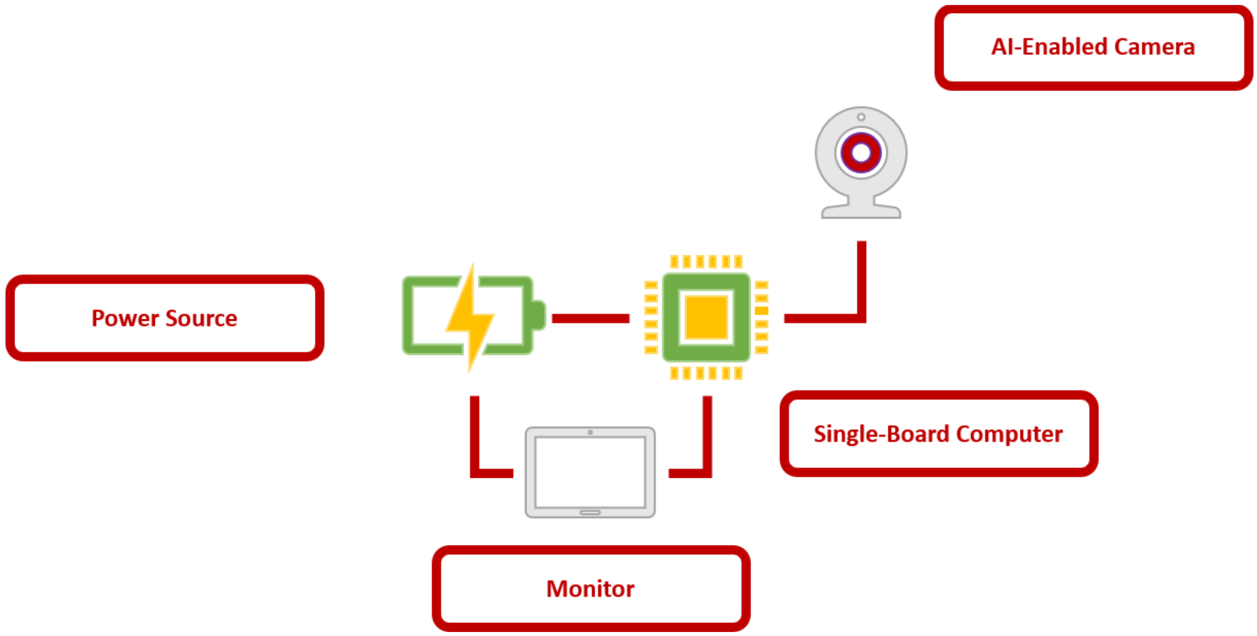

Figure 3.3 Real-time detection of DD on a portable, self-contained edge device in action.

Figure 3.4 Bounding box predictions of M-stages by Tiny YOLOv4 on the edge device. Predictions are made for each of the M-stages: M0 (top-center; magenta), M2, M2P (bottom-left; yellow), M4H, and M4P (bottom-right; green).

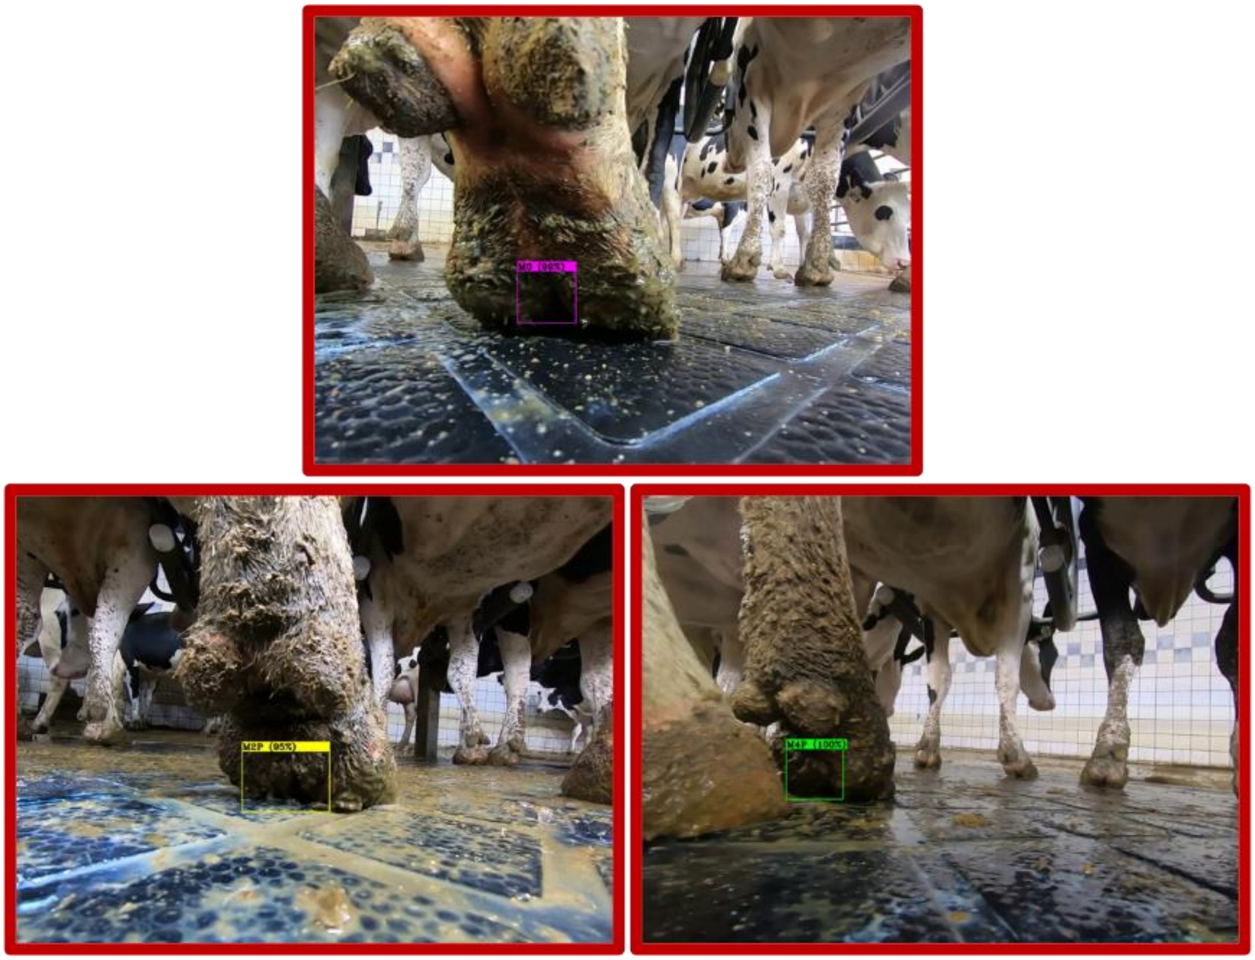

Figure 3.5 Accuracy of the five M-stages by Tiny YOLOv4 on the edge device. Accuracy is measured as the number of correct predictions and calculated for each of the five M-stages: M0, M2, M2P, M4H, and M4P.

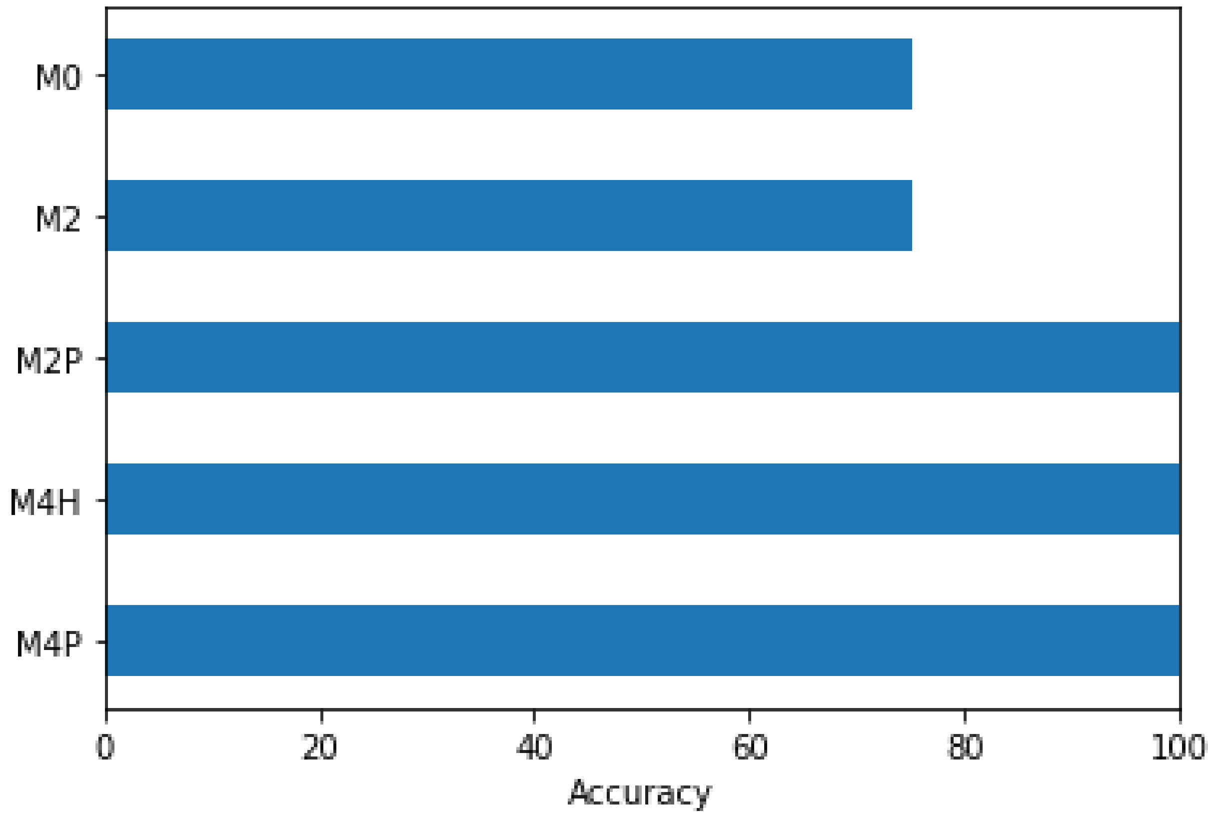

CHAPTER 4 BENCHMARKING ANALYSIS OF COMPUTER VISION ALGORITHMS ON CLOUD PLATFORMS FOR THE EARLY DETECTION OF DIGITAL DERMATITIS IN DAIRY COWS

ABSTRACT

Digital dermatitis (DD) is an infectious disease in cattle that leads to ulcerative lesions on the coronary band of the hoof, resulting in diminished animal welfare and economic well-being. DD detection at an early stage can facilitate timely treatment and reduce lameness. Computer vision (CV) models including YOLOv5 can identify hoof lesions and classify M-stages of DD. However, precision farming demands portable solutions for constrained environments and equipment challenges. Therefore, the objective of this study is to develop, deploy, and evaluate DD detection models on cloud platforms.

Data collection included images from the interdigital space on the plantar surface of the foot. The images were then assessed for M-stages of DD by a trained investigator. A YOLOv5s model was trained to identify and classify DD lesions. The final model was deployed on various cloud platforms for testing. Implementations included Google Colab, Docker using an IP camera via HTTP, Docker using an IP camera via RTSP, and TensorFlow.js.

YOLOv5s achieved an mAP of 0.946 with high AP for all five M-stages during model training. The TensorFlow.js application outperformed all other deployments with respect to agreement. The Cohen's kappa for TensorFlow.js was 0.763 and interpreted as substantial agreement. The prediction accuracy for TensorFlow.js was 0.842 with high detection for M0, M4H, and M4P and perfect detection for M2 and M2P. All deployments exceeded the minimum threshold for image processing at approximately 10 FPS.

The CV model was able to effectively identify and classify DD lesions on the TensorFlow.js application with high performance and sufficient speed. The study demonstrates the feasibility of implementing a cloud-based, portable solution for real-time detection of DD for precision farming. This achievement is a significant step towards the integration of CV algorithms in veterinary medicine and signifies forward progress in the real-time detection of health outcomes in agriculture.

INTRODUCTION

Digital dermatitis (DD) is an infectious claw disease causing significant lameness in cattle worldwide across different production systems (Bruijnis et al., 2012; Holzhauer et al., 2006). This disease is characterized by distinct ulceroproliferative lesions primarily found on the plantar aspect of the hoof on the coronary band, leading to discomfort or severe pain (Blowey and Sharp, 1988; Cheli and Mortellaro, 1974; Döpfer et al., 1997; Read and Walker, 1998). The consequences of DD include reduced animal welfare and economic losses from decreased milk production, compromised reproductive performance, and premature culling, resulting in an estimated cost of approximately US$133 per case (Bruijnis et al., 2010; Cha et al., 2010; de Jesús Argáez-Rodríguez et al., 1997; Ettema et al., 2010; Garbarino et al., 2004; Hernandez et al., 2001; Losinger, 2006; Whay et al., 1997). The reported prevalence rates of DD at the cow level range from 21.2% to 29.2%, accounting for 61.8% of lameness cases in bred heifers and 49.1% in adult cows (Brown et al., 2000; Holzhauer et al., 2006; Solano et al., 2016; USDA, 2009). Consequently, the combined annual economic loss from DD in the United States (approximately 9 million dairy cows) and the European Union (around 24.5 million dairy cows) surpasses US$1.1 billion based on an incidence rate of only 25% (Zinicola et al., 2015).

Cows affected by DD exhibit altered gait or posture to reduce discomfort, decreased mobility, lifting or shaking of the affected leg, and adopting a toe-down posture to minimize contact with the ground (Rodriguez-Lainz et al., 1998; Shearer et al., 2005). The disease is a multifactorial and polymicrobial, with consistent isolation of treponemes from DD lesions (Döpfer et al., 2012; Gomez et al., 2012; Krull et al., 2016). The pathogen can persist endemically within a herd, and infected cattle may experience active or chronic stages of infection. Chronically infected cattle serve as a reservoir for DD and pose a potential risk for herd outbreaks (Döpfer, 2009; Döpfer et al., 2012). Affected cows may develop changes in the skin of the heel area, promoting the persistence of DD, emergence of dermal treponeme reservoirs, and occurrence of heel horn erosion (Gomez et al., 2015a). Consequently, DD presents a significant animal welfare concern because of prolonged and recurrent painful episodes (Bruijnis et al., 2012; Dörte Döpfer et al., 2012; Gomez et al., 2015b).

Early detection and treatment of DD by means of regular monitoring for signs of DD plays a crucial role in effectively managing the disease (Döpfer et al., 2012; Döpfer and Morlán, 2008). Various classification systems have been developed for the different clinical stages of DD (Krull et al., 2014; Laven, 1999; Manske et al., 2002; Vink, 2006). The M-stage scoring system, initially introduced by Döpfer et al. (1997) and updated by Berry et al. (2012), categorizes different clinical stages of DD throughout the disease progression, allowing for the observation of transitions between active, chronic, and healed stages (Berry et al., 2012; Döpfer et al., 1997). This scoring system serves as a valuable tool for researchers, farmers, and hoof trimmers to monitor the effectiveness of DD control programs at both the individual-animal and herd levels (Döpfer et al., 2012; Kofler et al., 2019).

The dynamics of DD in cattle groups are characterized by different clinical stages and the transitions between these stages (Berry et al., 2012; Döpfer et al., 1997). The M-stage system for macroscopic scoring is based on lesion development to facilitate clinical inspections of the bovine foot (Berry et al., 2012; Döpfer et al., 1997). This classification system defines five M-stages (Figure 1):

M0: Normal skin appearance

M1: Small focal damage of the epithelium at the skin horn border (less than 2.0 cm in diameter)

M2: Circumscribed ulcerative skin defect with a red or greyish surface, possibly with a white epithelial margin and overlong hair (greater than 2.0 cm in diameter) or proliferative growth of the epithelium in the case of M2P lesions

M3: Healing stage of DD, characterized by the lesion being covered with a scab

M4: Chronic stage of DD, featuring a thickened epithelium (hyperkeratosis) for M4H lesions or proliferative growth of the epithelium in the case of M4P lesions

The scoring of DD lesions can be used to identify animals in need of treatment (Jorritsma et al., 2017; Schultz and Capion, 2013). The M-stage scoring system is widely recognized as the most accurate and comprehensive method for macroscopic evaluation of DD lesions (Egger-Danner et al., 2020; Greenough et al., 2008). However, visual inspection can introduce misclassification bias (Relun et al., 2011). While lifting the foot for inspection of DD lesions in the trimming chute remains the gold standard for DD detection, it presents several drawbacks such as high costs, increased labor, and elevated stress for the cows (Relun et al., 2011; Stokes et al., 2012; Thomsen et al., 2008). Relying solely on trimming chute inspections is impractical for assessing disease prevalence on a regular basis in medium to large scale herds or for early detection and treatment of DD.

Although M-stages is commonly used by hoof-care professionals, the training methods for scorers are rarely discussed in the published literature. Some publications mention the use of "experienced or trained scorers" to generate M-stages, while in other cases, scorers undergo a comprehensive training program that involves recognizing M-stages from color photographs and sometimes scoring live animals (Alsaaod et al., 2014; Cutler et al., 2013; Kulow et al., 2017; Logue et al., 2012; Solano et al., 2017; Yang et al., 2017). The absence of standardized training programs for DD scoring poses challenges to the reliability and repeatability of the scoring process (Vanhoudt et al., 2019).

Computer vision (CV) is a branch of artificial intelligence (AI) that focuses on enabling computers to process, analyze, and interpret visual data (Szeliski, 2022). By leveraging machine learning models, computer vision can identify and classify objects within images and videos, allowing computers to make informed decisions based on visual input. Object detection is a CV technique that focuses on the identification and localization of objects within images or videos. The primary objective of object detection is to accurately determine the position and boundaries of objects present in an image, while also assigning the appropriate categories. Object detection algorithms has shifted towards deep learning techniques, driven by advancements in high-performance graphics processing units (GPUs) (Sanders and Kandrot, 2010), the availability of large-scale datasets (e.g., DOTA, ImageNet, COCO) (Deng et al., 2009; Lin et al., 2014; Xia et al., 2018), and the proliferation of convolutional neural networks (CNNs) (LeCun et al., 2015).

Object detection algorithms can be categorized into two groups: one-stage and two-stage detectors (Wu et al., 2020; Sharma and Mir, 2020). Two-shot object detection algorithms use two passes of the input image to predict the presence and location of objects. The first pass generates a set of proposals or potential object locations. The second pass refines these proposals and produces final predictions. This approach is more accurate compared to other methods, but computationally expensive and resource intensive (Wu et al., 2020; Sharma and Mir, 2020). One-stage object detection algorithms use a single pass of the input image, making predictions about the presence and location of objects. By processing the entire image in a single pass, one-stage detectors are computationally efficient, but can be less accurate compared to other methods and less effective for detecting small objects (Wu et al., 2020; Sharma and Mir, 2020). Such algorithms are particularly useful for real-time detection in resource-constrained environments.

The You Only Look Once (YOLO) algorithm is a CV technique used for real-time object detection and classification. The YOLO models combine bounding box prediction and object classification within a single end-to-end differentiable network. The algorithm divides the input image into a grid of cells, and within each cell, a single neural network is employed to detect and classify the objects (Redmon et al., 2016). The output of the model includes bounding boxes for the location of the detected objects, the class labels that identify the type of object, and the corresponding class probabilities that represent the confidence scores for each bounding box (Redmon et al., 2016).

The YOLO algorithm provides several advantages over other object detection algorithms, leading to a combination of high accuracy and speed for inference (Redmon et al., 2016). As a one-stage object detector, YOLO looks at the entire image at once such that the model predicts both bounding boxes and class probabilities simultaneously (Redmon et al., 2016). Additionally, YOLO is able to leverage pre-trained weights, namely ImageNet for transfer learning. Starting from pre-trained weights not only conserves resources, but also improves the efficiency of training a new model (Redmon et al., 2016). Since YOLO is computationally efficient, it is ideal for deployment (Redmon et al., 2016).

The YOLO algorithm exhibits optimal performance on high-end GPUs, but its memory requirements can be limiting, particularly for deployment (Redmon, 2018; Redmon and Farhadi, 2018). There is a need for lightweight network architectures in object detection that can operate efficiently in resource-constrained environments, resulting in reduced model size and computational cost (Redmon, 2018). Tiny YOLO is a compressed version of YOLO, featuring a simpler network structure and a reduced number of parameters (Adarsh et al., 2020). Tiny YOLO is preferred compared to YOLO when real-time object detection is prioritized over accuracy (Adarsh et al., 2020). The speed of Tiny YOLOv4 is approximately eight times faster than YOLOv4 (Alexey, 2021; Bochkovskiy et al., 2020; Techzizou, 2021). However, the more complex YOLOv4 achieves 33% better accuracy compared to the less complex Tiny YOLOv4 when evaluated on the MS COCO dataset (Alexey, 2021; Bochkovskiy et al., 2020; Techzizou, 2021).

YOLOv5 is a recent release of the YOLO family of models by Ultralytics (Jocher, 2020). The previous versions of YOLO were implemented and are currently maintained in the Darknet framework (Bochkovskiy et al., 2020; Redmon, 2018; Redmon et al., 2016; Redmon and Farhadi, 2018). The YOLOv5 algorithm is the first model in the YOLO family to be developed using the PyTorch framework, resulting in a more lightweight and user-friendly (Jocher, 2020). YOLOv5 uses mosaic augmentation to combine four images into four tiles of random ratio, helping the model learn small objects (Jocher, 2020). Anchor boxes are learned based on the distribution of bounding boxes in the custom dataset with K-means and genetic learning algorithms (Jocher, 2020). The PyTorch framework is able to half the floating-point precision in training and inference from 32-bit to 16-bit precision, significantly speeds up the inference time of the models (Jocher, 2020). However, no major architectural changes are introduced in YOLOv5 compared to YOLOv4. Moreover, YOLOv5 does not surpass YOLOv4 in terms of performance on the commonly used COCO dataset benchmark. The YOLOv5 family of models is available in five variants: YOLOv5n, YOLOv5s, YOLOv5m, YOLOv5l, and YOLOv5x where YOLOv5s is analogous to Tiny YOLOv4 and YOLOv5m is analogous to YOLOv4

(Jocher, 2020). Similar to Tiny YOLOv4 and YOLOv4, YOLOv5s is 2.29 times faster than YOLOv5m, whereas YOLOv5m achieves 13% better accuracy than YOLOv5s on the MS COCO dataset (Jocher, 2020).

A YOLOv2 model was previously trained on a dataset comprising over 3,500 images of DD lesions to detect M-stages of DD (Cernek et al., 2020). The model achieved an accuracy of 71%, and the agreement between the human investigator and the CV model was determined to be "moderate" level according to Cohen's kappa during internal validation (Cernek et al., 2020). The YOLOv2 model achieved an accuracy of 88% in detecting DD, and the agreement between the human investigator and the CV model was determined to be "fair" based on Cohen's kappa during external validation (Cernek et al., 2020).

The CV tools and techniques used in agriculture and veterinary medicine are experiencing rapid growth for both research and commercial applications (Arthur F A Fernandes et al., 2020; Arthur Francisco Araújo Fernandes et al., 2020). Object detection is used to accurately monitor animal health and diagnose various medical conditions (Arthur Francisco Araújo Fernandes et al., 2020; Ryu and Tai, 2021). This technique aids in the identification of illnesses and infections, thereby reducing the reliance on invasive diagnostic tests and facilitating the development of optimized treatment plans. The application of object detection model on live-streaming video for inference requires significant storage and computational resources. The more complex AI models developed for video processing are commonly deployed on cloud servers with sufficient storage capacity and computing power compared to the resource-limited field environments where cattle are raised and crops are grown. Similarly, cloud computing leads to several other challenges, including increased latency, network congestion, and high computing costs.

Object detection has been used in video surveillance and monitoring (Oliveira et al., 2021). Cloud computing services such as Amazon Web Services, Microsoft Azure, and Google Cloud can be used to transfer large amounts of data from on-site devices and servers to the cloud. Cloud platforms can be used to share and maintain data and other computing resources, thereby ensuring availability and security (Vashisht and Kumar, 2020). The reliance of cloud platforms can result in network latency where the processing time depends on the availability of computational resources and the current workload on the system and network connection. Cloud computing services generally provide near real-time performance, but network latency can be too high for real-time detection. Therefore, onboard processing on the edge may be necessary for critical tasks and decisions to ensure timely and efficient execution (Lee et al., 2017).

For animal science and veterinary medicine, object detection models are mostly used for animal and disease detection (Cowton et al., 2019; Lee et al., 2019; Psota et al., 2019; Seo et al., 2020). Kang et al. trained a lameness scoring model for dairy cows utilizing the Receptive Field Block Net Single Shot Detector deep learning network (Kang et al., 2020). The network achieved a mean average precision of 87.0% and number of frames per second of 83.3 for locating cow hooves on video. The hooves are used as the input for an algorithm to compute the supporting phase of a cow's hoof when walking, which represents the time difference between the hoof lifting from the ground and the hoof load is in contact with the ground. Both YOLO and Faster R-CNN models are the main algorithms used for object detection in animal science studies (Oliveira et al., 2021). Lee et al. introduce a hybrid approach for the practical monitoring of undergrown pigs consisting of two steps: an image processing stage followed by a deep learning algorithm (Lee et al., 2019). They first use a Gaussian Mixture Model to identify the moving frames during a 24-hour period. Next, they employ Tiny YOLOv3 to detect individual pigs within each of the selected frames.

Additionally, Wang et al. proposes a smart surface inspection system using Faster R-CNN algorithm in the cloud-edge computing environment (Wang et al., 2020b). The cloud-edge computing method significantly increases the computation speed by more than 10 times compared to an embedded system where a microprocessor is dedicated to performing a specific task (Wang et al., 2020b). The cloud platform provided strong computation power, but the image transmission depended on the internet connection status, resulting in extra time and energy (Wang et al., 2020b). Wang et al. propose a CNN-based visual sorting system supported by a cloud-edge computing environment developed for fast computation and continuous service maintenance and upgrading (Wang et al., 2020a). The proposed R-CNN model is the most robust and accurate with high precision, but high computation time per image as well.

In the current study, lightweight CV models are trained for cloud deployment and evaluated for real-time detection of DD in dairy cows. The CV models are trained to detect and score DD lesions, and compared using common performance metrics and inference time. The process is automated for real-time detection using images and video streams on cloud platforms. Cloud deployment provides several advantages, including scalability, flexibility, improved efficiency, and enhanced privacy (Carroll et al., 2011; Subramanian and Jeyaraj, 2018; Zissis and Lekkas, 2012). Object detection for precision farming and veterinary medicine has the potential to improve accuracy of diagnoses for health outcomes, enabling producers, herdsmen, and veterinarians to make more informed decisions in real-time at the cow-level. This technology allows for timely detection and customized treatment plans based on the severity and progression of a disease including infectious claw conditions such as DD. The proposed tool is intended to complement existing practices in preventing DD in dairy cattle. Object detection on a cloud platform has the potential to identify cattle at high risk for DD, monitor herds affected by endemic DD, and enhance cattle welfare by enabling informed decision-making on the farm. This technology empowers farmers to take proactive measures to prevent and manage DD, leading to improved health outcomes for the cattle and overall herd welfare (Plummer and Krull, 2017).

MATERIALS AND METHODS

Data Collection and Image Labeling

For the current study, a dataset of 409 images of cattle feet was used. The images were collected from multiple farms in Wisconsin during May 2021. The JPG images were captured using GoPro Hero 5 Black cameras, which recorded MP4 videos of the hind feet from behind, while the cattle were standing on the ground. The videos were taken in multiple settings, including cattle pens, automated milking systems, and while exiting disinfecting foot baths. The videos were converted into images using FFmpeg with a frame rate of one frame per second (FFmpeg, 2023). All resulting images in the dataset were labeled for the M-stages of DD.

A trained investigator scored the images using the M-stage DD classification system. This system categorizes the different clinical stages of the disease based on morphological characteristics, including healthy (M0), active (M2), and chronic (M4) stages, with additional indications of hyperkeratosis (H) or proliferation (P) (Berry et al., 2012; Döpfer et al., 1997). The images were classified into the following stages: M0, M2, M2P, M4H, and M4P.

The dataset consisted of multiple feet per image, resulting in a total of 240 M0, 17 M2, 51 M2P, 114 M4H, and 108 M4P class labels. These labels were associated with bounding box annotations to precisely identify the location of the DD lesions in the images. The annotations and class labels were generated using LabelImg version 1.8.1 in Python version 3.8.15 (Lin, 2018). For compatibility with the object detection models used, the labels were converted from Pascal VOC format for R-CNN and SSD models to YOLO format for YOLO models (Tashiev, 2022).

Model Training

The datasets were randomly divided, with 90% of the images allocated to the training set and the remaining 10% to the testing set. The testing set consisted of 41 images, comprising 22 M0, 2 M2, 3 M2P, 11 M4H, and 16 M4P class labels. The models were trained and tested using Google Colab and an NVIDIA Tesla T4 GPU, with a batch size of 32 and image size of 416 ("Google Colaboratory," 2022). The training weights were initialized using convolutional weights pre-trained on ImageNet ("ImageNet," 2020; Krizhevsky et al., 2012). The PyTorch 1.7.0 framework was used for training in Python version 3.10.12 (Paszke et al., 2019).

YOLOv5 passes training data through a data loader making augmentations including scaling, color space adjustments, and mosaic augmentation (Jocher, 2020). The model configuration for YOLOv5 is defined in the yaml file compared to the model configuration for YOLOv4 is specified in a cfg file (Jocher, 2020). A stochastic gradient descent (SGD) optimizer was used with an initial learning rate of 0.001 and momentum of 0.937 for a maximum number of 1,500 epochs with early stoppage, taking the best set of trained weights (Jocher, 2020). Results are logged during model training using TensorBoard 2.4.1 (TensorBoard, 2023). The custom YOLOv5s model was deployed on four implementations using live streaming video for benchmarking: Google Colab, Docker using an IP camera via HTTP, Docker using an IP camera via RTSP, and TensorFlow.js.

Implementations

Google Colaboratory abbreviated as Colab is a hosted Jupyter Notebook service that provides free access to computing resources, including GPUs and TPUs runtime to significantly speed up the training and inference process ("colab.google," 2022; "Google Colaboratory," 2022). Google Colab is particularly appropriate for machine learning and deep learning tasks. It is a cloud-based web Integrated Development Environment (IDE) primarily designed for Python, enabling the user to run code directly in the browser without the requirement of setting up a local environment. The YOLOv5 repository was cloned from GitHub and the best weights for the custom model were loaded (Jocher, 2020). The detector was run on a live webcam using JavaScript to create the live video stream using the webcam as input.

Docker is an open-source platform that allows developers to automate the deployment, scaling, and management of applications inside lightweight, portable containers (Docker, 2022a; *Docker*, 2022b). Containers are images that package an application and its dependencies with runtime, ensuring consistent behavior across different environments, such as development, testing, and production. This isolated environment maintains that the application works consistently across various environments and eliminates portability issues. It can run on any platform that supports Docker, no matter if it is a local machine, a staging server, or a cloud-based production environment. This enables seamless deployment and scaling across different environments. Docker containers are lightweight and share the host OS kernel, which means they use fewer resources compared to traditional virtual machines. This leads to faster startup times and better resource utilization. The YOLOv5 docker image was pulled from DockerHub and the Docker container was run with local file access and GPU access (Ultralytics, 2023a; Ultralytics, 2023b).

The detector was run within the running Docker container on a live webcam using a wireless IP camera via Hypertext Transfer Protocol (HTTP) or Real Time Streaming Protocol (RTSP).

TensorFlow.js is an open-source JavaScript library developed by Google that allows developers to build and train machine learning models directly in the browser or on Node.js servers (TensorFlow.js, 2023a; TensorFlow.js, 2023b). It is an extension of TensorFlow, a popular open-source machine learning framework, and provides a way to run TensorFlow models in JavaScript environments. The TensorFlow.js library supports training machine learning models in a local web browser using client-side data. Developers can create, train, and fine-tune models for tasks such as speech recognition, natural language processing, and object detection. TensorFlow.js leverages WebGL, a browser-based graphics library, to accelerate numerical computations and optimize performance. This enables running deep learning models efficiently on GPUs, if available. The TensorFlow.js library provides support for running models on mobile devices and Internet of Things (IoT) devices, allowing developers to deploy machine learning applications on a wide range of platforms. Models trained in Python using TensorFlow can be converted to TensorFlow.js format, enabling seamless deployment of models from server to browser.

Node.js is an open-source, server-side runtime environment built on Chrome's V8 JavaScript engine (Node.js, 2023). It allows developers to execute JavaScript code outside of a web browser, enabling server-side scripting and development of server applications. Node.js provides an event-driven, non-blocking I/O model, making it efficient and well-suited for building scalable and real-time applications. It is designed to handle asynchronous I/O operations efficiently, which allows it to handle many concurrent connections without blocking the execution of other tasks, making it suitable for building scalable applications that need to handle high traffic. Node.js is commonly used for building server-side applications, such as web servers, APIs, real-time applications, and networking tools. Node.js is designed to be cross-platform across various operating systems including Windows, macOS, and Linux. It is lightweight and fast, making it ideal for applications that require low latency and quick response times. Node.js is widely used for building real-time applications because of its ability to handle simultaneous connections efficiently.

The custom YOLOv5 model was deployed in Google Chrome using TensorFlow.js with WebGL backend. The trained model was converted from PyTorch to TensorFlow.js layers format by means of a TensorFlow frozen model. The exported YOLOv5 web model contains a json file and a set of sharded weights files in a binary format. The web model is dropped into an open-sourced implementation of YOLOv5 using Tensorflow.js. A local server is started and the application is deployed directly in the web browser and offline. The implementation can perform object detection on images, videos, and streaming webcam at a minimum class threshold of 0.25 (Figure 2).

Performance Evaluation

The evaluation of network performance and efficiency of the implementations included two main components: accuracy and inference speed. For assessing model training, the mean average precision (mAP) at intersection over union (IoU) of 0.5 was used to evaluate accuracy. Additionally, precision, recall, and average precision (AP) were used as performance measures to compare between the predictions made by the CV models, and labels made by a trained investigator (the so-called ‘ground truth’) (Ultralytics, 2023c). For assessing model deployment, Cohen's kappa was used to evaluate the agreement between the model predictions and the ground truth labels, providing valuable insights into the model's ability to make accurate predictions in real-world scenarios (Landis and Koch, 1977; Viera and Garrett, 2005).

In addition to accuracy, the inference speed of the model is also significant, especially in cloud computing environments. The inference speed is measured in frames per second (FPS), calculating the number of images that the model has processed and inferred within one second. A higher FPS value implies a faster and more efficient model, wanted in real-time applications or scenarios with limited computing resources. By considering both accuracy and inference speed, we can comprehensively evaluate the performance of the model deployment. This evaluation process helps to select the best model for specific use cases, ensuring optimal accuracy and efficiency in real-world applications.

RESULTS

Benchmarking

A YOLOv5s model was trained and tested using 409 images and 530 instances in Google Colab. Best results were observed at epoch 107 where the precision was 0.858 and the recall was 0.907. Overall, the mAP of the best model was 0.946 with high AP for all five M-stages of DD including extremely high AP for M4P and near perfect AP for M2 and M2P (Figure 3).

The custom YOLOv5s model was deployed on four implementations using live streaming video: Google Colab, Docker using an IP camera via HTTP, Docker using an IP camera via RTSP, and TensorFlow.js. Images were at a sufficient distance from the camera such that the object in the image is focused and the lighting does not present glare or darkness on the image. Videos were scored for M-stages of DD and compared to the corresponding labels made by a trained investigator.

Cohen’s kappa was used to measure the agreement between two raters (Landis and Koch, 1977; Viera and Garrett, 2005). The Cohen’s kappa for Google Colab compared to the trained investigator was determined to be 0.706 and interpreted as substantial agreement between the two raters ($z = 11.4$; $p < 0.001$). The Cohen's kappa for Docker using an IP camera via HTTP compared to the trained investigator was determined to be 0.691 and interpreted as substantial agreement between the two raters ($z = 9.9$; $p < 0.001$). The Cohen's kappa for Docker using an IP camera via RTSP compared to the trained investigator was determined to be 0.568 and interpreted as moderate agreement between the two raters ($z = 9.24$; $p < 0.001$). The Cohen's kappa for TensorFlow.js compared to the trained investigator was determined to be 0.763 and interpreted as substantial agreement between the two raters ($z = 9.83$; $p < 0.001$). The model performed well compared to the ground truth of the trained investigator for all four implementations where the TensorFlow.js application performed the best (Figure 4). The TensorFlow.js application was able to detect all five M-stages of DD on live streaming video. The prediction accuracy for TensorFlow.js was 0.842 with high detection for M0, M4H, and M4P and perfect detection for M2 and M2P (Figure 5).

The FPS metric was used to measure the inference time of the implementations using live streaming video (Figure 6). The Docker container using an IP camera via RTSP stream outperformed all other implementations at a maximum of 25 FPS followed closely by the Docker container using an IP camera via HTTP stream at a maximum of 22 FPS. On the other hand, both the Google Colab and TensorFlow.js implementations ran at a maximum of 12 FPS and 10 FPS, or half the speed of the Docker container using an IP camera. Regardless, all four implementations were deployed serving YOLOv5s for live object detection.

DISCUSSION

The real-time detection applications in this study demonstrated proficiency in identifying DD lesions and classifying M-stages with high accuracy speed while deployed in the cloud. For speed, the Docker container using an IP camera via RTSP displayed the highest inference speed at 25 FPS, followed closely by the Docker container using an IP camera via HTTP at FPS. Notably, all deployments exceeded the minimum threshold for image processing by a human visual system at approximately 10 FPS. All four implementations underperformed with respect to inference speed compared to a DD detector run on the edge at 40 FPS. This is to be expected because model deployment includes the additional step of network connectivity and associated latency.

For detection, the Tensorflow.js application was able to classify the M-stages of DD achieving the highest Cohen's kappa value of 0.763 and an overall accuracy of 0.842. Comparatively, the YOLOv5s model in the browser outperformed the YOLOv2 model with additional M-stages of DD on a more complex dataset, representative of practical scenarios (Cernek et al., 2020). As a proof-of-concept, the YOLOv5s model, deployed in the web browser using TensorFlow.js, was able to effectively detect the various M-stages of DD on streaming video and has the ability to detect M-stages of DD across other media types, including images and video files.

The current YOLOv5 model was developed using the PyTorch framework. PyTorch models can also be converted to TensorFlow Lite models, instead of TensorFlow.js Layer models, providing the additional advantage of reduced model size and latency (tensorflow, 2021). This facilitates deployment to cloud platforms and mobile devices, improving accessibility and versatility (tensorflow, 2021). The YOLOv5 model architecture can be updated to leverage the current framework with YOLOv7 or YOLOv8 for further technological advancements and potential performance gains (Paszke et al., 2019; pytorch, 2021). The integration of YOLOv7 or YOLOv8 would increase speed and performance for real-time detection where the model architecture and the format of the exported model presents the opportunity for further optimization and expand the number of possible deployment options (Paszke et al., 2019; pytorch, 2021).

The custom YOLOv5s model deployed on the TensorFlow.js live detection application demonstrated high performance with respect to detection for all five M-stages of DD. However, it displayed higher accuracy in detecting M2 and M2P (Fig. 5). This difference in performance can be caused by the presence of data imbalanced where the number of class labels within the dataset where M2 and M2P had relatively fewer training examples compared to M0, M4H, and M4P.

The impact of class label imbalance on model performance is a persistent challenge in machine learning tasks and can adversely affect the ability of a model to accurately detect and classify underrepresented classes (Johnson and Khoshgoftaar, 2019). Increasing the number of training examples for minor classes using data augmentation techniques can help improve model performance (Johnson and Khoshgoftaar, 2019). Augmentation involves applying various transformations to existing images, such as rotation, flipping, zooming, and color adjustments, to generate new examples that retain the same class label (Albumentations, 2021, 2020; albumentations-team, 2018; Buslaev et al., 2020). As previously mentioned, the YOLOv5 model training uses mosaic augmentation to combine four images into four tiles of random ratio to help the model learn small objects (Jocher, 2020). Additionally, the model training uses image reflection, image translation, image scaling, and color space adjustments via Hue, Saturation, and Value (HSV) augmentation to mitigate data imbalance (Jocher, 2020). However, additional augmentation using blurring, random cropping, and random brightness and contrast can prevent overfitting and help build better models (Albumentations, 2021, 2020; albumentations-team, 2018; Buslaev et al., 2020).

Artificially balancing the dataset involves duplicating or replicating instances of minor classes to match the number of training examples in the major classes. Similarly, balancing the dataset can include deleting or removing instances of major classes to match the number of training examples in the minor classes. This guarantees that all classes have equal representation and reduces the impact of class imbalance on model performance, but also decreases the training data for model learning. Other remedial measures for data imbalance include adjusting the loss function during model training can also address the class imbalance issue (Johnson and Khoshgoftaar, 2019) in addition to techniques such as weighted loss functions or focal loss can be employed to assign greater importance to the minor classes, imposing the model to focus on improving the detection accuracy (Lin et al., 2018; Phan and Yamamoto, 2020). Expanding the dataset by collecting more data specifically for minor classes can help resolve the class imbalance issue. By obtaining additional instances for underrepresented classes, the model can better learn their unique features and characteristics.

By implementing these remedial measures, model performance can be significantly enhanced, ensuring accuracy and reliability for detection of all class labels, including M2 and M2P lesions. Additionally, continuous evaluation and optimization of the model can lead to further improvements and make it more robust for a variety of real-world scenarios and real-time applications.

Object detection on cloud platforms is faced with various constraints. Cloud-based object detection requires a stable and reliable internet connection. If there are network disruptions or slow connections, it can affect the speed and reliability of the detection process. In areas with poor or limited internet connectivity, accessing cloud services may be difficult or not feasible, making offline processing necessary. Cloud-based object detection can be subject to latency, as it involves transmitting data to and from the cloud server. The speed and reliability of detection may be affected by the network connection, leading to delays in real-time applications. Transmitting large amounts of data between edge devices and the cloud can result in significant data transfer overhead, consuming bandwidth and potentially affecting the overall performance. Lastly, cloud services for object detection can be costly, especially for large-scale or continuous data processing where cost is based on data usage, computation time, and storage. Scalability may be an issue as the volume of data increases, potentially leading to higher expenses. Cloud platforms may experience downtime or performance issues, affecting the availability and scalability of object detection services.

A wide range of precision agriculture tools are readily available, mainly in agricultural applications (Arthur Francisco Araújo Fernandes et al., 2020). Expanding the deployment of CV models on cloud computing services will contribute to their widespread adoption in agriculture and veterinary medicine. The CV models developed in this study can be further extended to various portable platforms and distributions, including different lightweight frameworks, compact enclosures, and mobile devices.

Improving accessibility to these models can be achieved through the development of mobile applications such as TensorFlow Lite, Ultralytics Hub Android App, or Docker containers, creating a lightweight, standalone package (Docker, 2022a; *Docker*, 2022b; PyTorch, 2023; techzizou, 2021; Ultralytics, 2023d, 2023e). By increasing user acceptance of these models, we can generate a rich and diverse library of images for optimization and validation of CV models. Ultimately, the availability and deployment of these CV tools will greatly benefit agriculture applications, providing enhanced access to valuable information and accelerating advancements in the field.

The current implementation of this portable CV tool focuses on automating DD detection in dairy cattle, and preliminary results show its potential applicability to beef cattle as well. Further development of the model can be extended to include training on a diverse range of cattle breeds, geographical locations, and other sources of data variability, increasing its robustness. By integrating ear tag identification with DD detection, the tool can be used for automated treatment lists and monitoring data, streamlining management for farmers and veterinarians. Moreover, the research can be broadened to encompass additional M-stages of DD, white line disease, foot rot, and other hoof diseases (Lely, 2016), making it a comprehensive and versatile tool for addressing various challenges in the livestock industry. By extending the capabilities of the CV tool and exploring a broader range of applications, this technology can play a widespread role in supporting livestock management, promoting animal welfare, and modernizing agricultural practices.

Automatic monitoring and measurement of animal traits using CV algorithms is a highly researched area (Arthur Francisco Araújo Fernandes et al., 2020; Li et al., 2021; Wurtz et al., 2019). Researchers have explored various imaging technologies to enhance the performance of these models. Previous studies have demonstrated that infrared thermography can achieve acceptable results for detecting diseases like DD. Additionally, the use of 3D cameras for animal identification has proven to outperform standard cameras in accuracy and reliability. As technology continues to advance, we can expect these CV tools to be more widely adopted and integrated into various agricultural practices. The ongoing development and improvement of imaging technologies will likely further enhance the performance and applicability of CV algorithms for animal monitoring and management, paving the way for more efficient and data-driven agricultural practices.

The successful development and deployment of practical CV solutions still faces several challenges. One prominent challenge is adapting CV algorithms for true on-farm deployment. Limited studies have thoroughly evaluated these algorithms using validation datasets to assess their reliability and robustness across multiple farms (Arthur Francisco Araújo Fernandes et al., 2020; Li et al., 2021; Wurtz et al., 2019). This validation process is essential for the confirmation of these algorithms to effectively perform in diverse real-world farm settings.

Another significant challenge is individual animal identification and tracking, as it remains susceptible to errors. Achieving accurate and consistent identification of animals is essential for the success of various CV applications in precision farming and veterinary medicine. Additionally, merging multiple devices into a single application is imperative for true on-farm deployment of various CV models (Arthur Francisco Araújo Fernandes et al., 2020; Li et al., 2021; Wurtz et al., 2019). This integration allows the implementation of more sophisticated predictive algorithms that can handle multiple inputs and provide joint predictions for multiple traits.

Delivering CV model predictions to farmers in a convenient and actionable manner poses a significant challenge. Farmers need easy access to the results and insights provided by the CV models to make informed decisions (Arthur Francisco Araújo Fernandes et al., 2020; Li et al., 2021; Wurtz et al., 2019). The user interface and presentation of the CV outputs play a fundamental role in facilitating adoption and practical use of these technologies. Addressing these challenges and making long-term improvements in each of these areas will direct the adoption of CV technologies in agriculture and foster the development of new applications to benefit the industry.

ACKNOWLEDGEMENTS

Funding support for the research was provided by the US Department of Agriculture through the National Institute for Food and Agriculture - Animal Health Grant (WIS03082). We would also like to thank Jamie Sullivan (Carman, MB, Canada), Chris Bauer (Mondovi, WI), and all other dairy producers for contributing to the images of hoof lesions. The authors have no financial and personal relationships with other people or organizations that could inappropriately influence (bias) their work.

REFERENCES

Adarsh, P., Rathi, P., Kumar, M., 2020. YOLO v3-Tiny: Object Detection and Recognition using one stage improved model, in: 2020 6th International Conference on Advanced Computing and Communication Systems (ICACCS). Presented at the 2020 6th International Conference on Advanced Computing and Communication Systems (ICACCS), pp. 687–694. https://doi.org/10.1109/ICACCS48705.2020.9074315

Albumentations, 2021. Albumentations Documentation [WWW Document]. Albumentations Documentation. URL https://albumentations.ai/docs/ (accessed 8.21.23).

Albumentations, 2020. Bounding boxes augmentation for object detection [WWW Document]. Albumentations Documentation. URL https://albumentations.ai/docs/getting_started/bounding_boxes_augmentation/ (accessed 8.21.23).

albumentations-team, 2018. albumentations-team/albumentations: Fast image augmentation library and an easy-to-use wrapper around other libraries. https://doi.org/10.3390/info11020125

Alexey, 2021. Yolo v4, v3 and v2 for Windows and Linux.

Alsaaod, M., Syring, C., Dietrich, J., Doherr, M.G., Gujan, T., Steiner, A., 2014. A field trial of infrared thermography as a non-invasive diagnostic tool for early detection of digital dermatitis in dairy cows. The Veterinary Journal 199, 281–285. https://doi.org/10.1016/j.tvjl.2013.11.028

Berry, S.L., Read, D.H., Famula, T.R., Mongini, A., Döpfer, D., 2012. Long-term observations on the dynamics of bovine digital dermatitis lesions on a California dairy after topical treatment with lincomycin HCl. The Veterinary Journal, Special Issue:Bovine Lameness 193, 654–658. https://doi.org/10.1016/j.tvjl.2012.06.048

Blowey, R.W., Sharp, M.W., 1988. Digital dermatitis in dairy cattle. Vet Rec 122, 505–508. https://doi.org/10.1136/vr.122.21.505

Bochkovskiy, A., Wang, C.-Y., Liao, H.-Y.M., 2020. YOLOv4: Optimal Speed and Accuracy of Object Detection. arXiv:2004.10934 [cs, eess].

Brown, C.C., Kilgo, P.D., Jacobsen, K.L., 2000. Prevalence of papillomatous digital dermatitis among culled adult cattle in the southeastern United States. American Journal of Veterinary Research 61, 928–930. https://doi.org/10.2460/ajvr.2000.61.928

Bruijnis, M.R.N., Beerda, B., Hogeveen, H., Stassen, E.N., 2012. Assessing the welfare impact of foot disorders in dairy cattle by a modeling approach. animal 6, 962–970. https://doi.org/10.1017/S1751731111002606

Bruijnis, M.R.N., Hogeveen, H., Stassen, E.N., 2010. Assessing economic consequences of foot disorders in dairy cattle using a dynamic stochastic simulation model. Journal of Dairy Science 93, 2419–2432. https://doi.org/10.3168/jds.2009-2721

Buslaev, A., Iglovikov, V.I., Khvedchenya, E., Parinov, A., Druzhinin, M., Kalinin, A.A., 2020. Albumentations: Fast and Flexible Image Augmentations. Information 11, 125. https://doi.org/10.3390/info11020125

Carroll, M., van der Merwe, A., Kotzé, P., 2011. Secure cloud computing: Benefits, risks and controls, in: 2011 Information Security for South Africa. Presented at the 2011 Information Security for South Africa, pp. 1–9. https://doi.org/10.1109/ISSA.2011.6027519

Cernek, P., Bollig, N., Anklam, K., Döpfer, D., 2020. Hot topic: Detecting digital dermatitis with computer vision. Journal of Dairy Science 103, 9110–9115. https://doi.org/10.3168/jds.2019-17478

Cha, E., Hertl, J.A., Bar, D., Gröhn, Y.T., 2010. The cost of different types of lameness in dairy cows calculated by dynamic programming. Preventive Veterinary Medicine 97, 1–8. https://doi.org/10.1016/j.prevetmed.2010.07.011

Cheli, R., Mortellaro, C., 1974. La dermatite digitale del bovino, in: Proceedings of the 8th International Conference on Diseases of Cattle. Piacenza, pp. 208–213.

colab.google [WWW Document], 2022. . colab.google. URL http://0.0.0.0:8080/ (accessed 7.23.23).

Cowton, J., Kyriazakis, I., Bacardit, J., 2019. Automated Individual Pig Localisation, Tracking and Behaviour Metric Extraction Using Deep Learning. IEEE Access 7, 108049–108060. https://doi.org/10.1109/ACCESS.2019.2933060

Cutler, J.H.H., Cramer, G., Walter, J.J., Millman, S.T., Kelton, D.F., 2013. Randomized clinical trial of tetracycline hydrochloride bandage and paste treatments for resolution of lesions and pain associated with digital dermatitis in dairy cattle. Journal of Dairy Science 96, 7550–7557. https://doi.org/10.3168/jds.2012-6384

de Jesús Argáez-Rodríguez, F., Hird, D.W., de Anda, J.H., Read, D.H., Rodríguez-Lainz, A., 1997. Papillomatous digital dermatitis on a commercial dairy farm in Mexicali, Mexico: Incidence and effect on reproduction and milk production. Preventive Veterinary Medicine 32, 275–286. https://doi.org/10.1016/S0167-5877(97)00031-7

Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., Fei-Fei, L., 2009. ImageNet: A large-scale hierarchical image database, in: 2009 IEEE Conference on Computer Vision and Pattern Recognition. Presented at the 2009 IEEE Conference on Computer Vision and Pattern Recognition, pp. 248–255. https://doi.org/10.1109/CVPR.2009.5206848

Docker, 2022a. Docker: Accelerated, Containerized Application Development [WWW Document]. URL https://www.docker.com/ (accessed 7.23.23).

Docker, 2022b.

Döpfer, D., 2009. The dynamics of digital dermatitis in dairy cattle and the manageable state of disease, in: Proceedings of the CanWest Veterinary Conference. pp. 17–20.

Döpfer, D., Anklam, K., Mikheil, D., Ladell, P., 2012. Growth curves and morphology of three Treponema subtypes isolated from digital dermatitis in cattle. The Veterinary Journal, Special Issue:Bovine Lameness 193, 685–693. https://doi.org/10.1016/j.tvjl.2012.06.054

Döpfer, Dörte, Holzhauer, M., Boven, M. van, 2012. The dynamics of digital dermatitis in populations of dairy cattle: Model-based estimates of transition rates and implications for control. The Veterinary Journal, Special Issue:Bovine Lameness 193, 648–653. https://doi.org/10.1016/j.tvjl.2012.06.047

Döpfer, D., Huurne, A.A.H.M. ter, Cornelisse, J.L., Asten, A.J.A.M. van, Koopmans, A., Meijer, F.A., Schukken, Y.H., Szakáll, I., Klee, W., Bosma, R.B., 1997. Histological and bacteriological evaluation of digital dermatitis in cattle, with special reference to spirochaetes and Campylobacter faecalis. Veterinary Record 140, 620–623. https://doi.org/10.1136/vr.140.24.620

Döpfer, D., Morlán, J.B., 2008. The paradox of modern animal husbandry and lameness. Veterinary journal 175, 153–154. https://doi.org/10.1016/j.tvjl.2007.02.013

Egger-Danner, C., Nielen, P., Fielder, A., Müller, K., Fjeldaas, T., Döpfer, D., Daniel, V., Bergsten, C., Cramer, G., Christen, A.-M., Stock, K.F., Thomas, G., Holzhauer, M., Steiner, A., Clarke, J., Capion, N., Charfeddine, N., Pryce, J.E., Oakes, E., Burgstaller, J., Heringstad, B., Ødegård, C., Kofler, J., 2020. ICAR Claw Health Atlas. ICAR Technical Series.

Ettema, J., Østergaard, S., Kristensen, A.R., 2010. Modelling the economic impact of three lameness causing diseases using herd and cow level evidence. Preventive Veterinary Medicine 95, 64–73. https://doi.org/10.1016/j.prevetmed.2010.03.001

Fernandes, Arthur Francisco Araújo, Dórea, J.R.R., Rosa, G.J. de M., 2020. Image Analysis and Computer Vision Applications in Animal Sciences: An Overview. Frontiers in Veterinary Science 7.

Fernandes, Arthur F A, Dórea, J.R.R., Valente, B.D., Fitzgerald, R., Herring, W., Rosa, G.J.M., 2020. Comparison of data analytics strategies in computer vision systems to predict pig body composition traits from 3D images. Journal of Animal Science 98, skaa250. https://doi.org/10.1093/jas/skaa250

FFmpeg, 2023. ffmpeg Documentation [WWW Document]. FFmpeg. URL https://ffmpeg.org/ffmpeg.html (accessed 9.26.23).

Garbarino, E.J., Hernandez, J.A., Shearer, J.K., Risco, C.A., Thatcher, W.W., 2004. Effect of Lameness on Ovarian Activity in Postpartum Holstein Cows*. Journal of Dairy Science 87, 4123–4131. https://doi.org/10.3168/jds.S0022-0302(04)73555-9

Gomez, A., Cook, N.B., Bernardoni, N.D., Rieman, J., Dusick, A.F., Hartshorn, R., Socha, M.T., Read, D.H., Döpfer, D., 2012. An experimental infection model to induce digital dermatitis infection in cattle. Journal of Dairy Science 95, 1821–1830. https://doi.org/10.3168/jds.2011-4754

Gomez, A., Cook, N.B., Rieman, J., Dunbar, K.A., Cooley, K.E., Socha, M.T., Döpfer, D., 2015a. The effect of digital dermatitis on hoof conformation. Journal of Dairy Science 98, 927–936. https://doi.org/10.3168/jds.2014-8483

Gomez, A., Cook, N.B., Socha, M.T., Döpfer, D., 2015b. First-lactation performance in cows affected by digital dermatitis during the rearing period. Journal of Dairy Science 98, 4487–4498. https://doi.org/10.3168/jds.2014-9041

Google Colaboratory [WWW Document], 2022. URL https://colab.research.google.com/?utm_source=scs-index (accessed 2.24.22).

Greenough, P., Muelling, C., Döpfer, D., Tomlinson, D., 2008. International atlas of lesions of cattle feet. Nomenclature and atlas update, in: Proc. 15th Int. Symp. 7th Conf. Lameness Ruminants, Kuopio, Finland. Savonia University of Applied Sciences, Kuopio, Finland, p. 40.

Hernandez, J., Shearer, J.K., Webb, D.W., 2001. Effect of lameness on the calving-to-conception interval in dairy cows. Journal of the American Veterinary Medical Association 218, 1611–1614. https://doi.org/10.2460/javma.2001.218.1611

Holzhauer, M., Hardenberg, C., Bartels, C.J.M., Frankena, K., 2006. Herd- and Cow-Level Prevalence of Digital Dermatitis in The Netherlands and Associated Risk Factors. Journal of Dairy Science 89, 580–588. https://doi.org/10.3168/jds.S0022-0302(06)72121-X

Jocher, G., 2020. YOLOv5 by Ultralytics. https://doi.org/10.5281/zenodo.3908559

Johnson, J.M., Khoshgoftaar, T.M., 2019. Survey on deep learning with class imbalance. Journal of Big Data 6, 27. https://doi.org/10.1186/s40537-019-0192-5

Jorritsma, R., Nielen, M., Dotinga, A., 2017. A Randomised Non-inferiority Trial on the Effect of an Antibiotic or Non-antibiotic Topical Treatment Protocol for Digital Dermatitis in Dairy Cattle: a knowledge summary. Veterinary Evidence 2. https://doi.org/10.18849/ve.v2i4.111

Kang, X., Zhang, X.D., Liu, G., 2020. Accurate detection of lameness in dairy cattle with computer vision: A new and individualized detection strategy based on the analysis of the supporting phase. Journal of Dairy Science 103, 10628–10638. https://doi.org/10.3168/jds.2020-18288

Kofler, J., Fielder, A., Charfeddine, N., Capion, N., Fjeldaas, T., Cramer, G., Bell, N.J., Müller, K.E., Christen, A.-M., Thomas, G., Heringstad, B., Stock, K.F., Holzhauer, M., Nieto, J.M., Egger-Danner, C., Döpfer, D., 2019. ICAR Claw Health Atlas – Appendix 1: Digital Dermatitis Stages (M-stages). ICAR Technical Series.

Krull, A.C., Cooper, V.L., Coatney, J.W., Shearer, J.K., Gorden, P.J., Plummer, P.J., 2016. A Highly Effective Protocol for the Rapid and Consistent Induction of Digital Dermatitis in Holstein Calves. PLOS ONE 11, e0154481. https://doi.org/10.1371/journal.pone.0154481

Krull, A.C., Shearer, J.K., Gorden, P.J., Cooper, V.L., Phillips, G.J., Plummer, P.J., 2014. Deep Sequencing Analysis Reveals Temporal Microbiota Changes Associated with Development of

Bovine Digital Dermatitis. Infection and Immunity 82, 3359–3373. https://doi.org/10.1128/iai.02077-14

Kulow, M., Merkatoris, P., Anklam, K.S., Rieman, J., Larson, C., Branine, M., Döpfer, D., 2017. Evaluation of the prevalence of digital dermatitis and the effects on performance in beef feedlot cattle under organic trace mineral supplementation1. Journal of Animal Science 95, 3435–3444. https://doi.org/10.2527/jas.2017.1512

Landis, J.R., Koch, G.G., 1977. The Measurement of Observer Agreement for Categorical Data. Biometrics 33, 159–174. https://doi.org/10.2307/2529310

Laven, R., 1999. The environment and digital dermatitis. Cattle Practice 7, 349–354.

LeCun, Y., Bengio, Y., Hinton, G., 2015. Deep learning. Nature 521, 436–444. https://doi.org/10.1038/nature14539

Lee, J., Wang, J., Crandall, D., Šabanović, S., Fox, G., 2017. Real-Time, Cloud-Based Object Detection for Unmanned Aerial Vehicles, in: 2017 First IEEE International Conference on Robotic Computing (IRC). Presented at the 2017 First IEEE International Conference on Robotic Computing (IRC), pp. 36–43. https://doi.org/10.1109/IRC.2017.77

Lee, S., Ahn, H., Seo, J., Chung, Y., Park, D., Pan, S., 2019. Practical Monitoring of Undergrown Pigs for IoT-Based Large-Scale Smart Farm. IEEE Access 7, 173796–173810. https://doi.org/10.1109/ACCESS.2019.2955761

Lely, 2016. Claw diseases - Lely [WWW Document]. URL https://www.lely.com/us/farming-insights/claw-diseases/ (accessed 4.8.23).

Li, G., Huang, Y., Chen, Z., Chesser, G.D., Purswell, J.L., Linhoss, J., Zhao, Y., 2021. Practices and Applications of Convolutional Neural Network-Based Computer Vision Systems in Animal Farming: A Review. Sensors 21, 1492. https://doi.org/10.3390/s21041492

Lin, T., 2018. LabelImg.

Lin, T.-Y., Goyal, P., Girshick, R., He, K., Dollár, P., 2018. Focal Loss for Dense Object Detection. https://doi.org/10.48550/arXiv.1708.02002

Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L., 2014. Microsoft COCO: Common Objects in Context, in: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (Eds.), Computer Vision – ECCV 2014, Lecture Notes in Computer Science. Springer International Publishing, Cham, pp. 740–755. https://doi.org/10.1007/978-3-319-10602-1_48

Logue, D.N., Gibert, T., Parkin, T., Thomson, S., Taylor, D.J., 2012. A field evaluation of a footbathing solution for the control of digital dermatitis in cattle. The Veterinary Journal, Special Issue:Bovine Lameness 193, 664–668. https://doi.org/10.1016/j.tvjl.2012.06.050

Losinger, W.C., 2006. Economic impacts of reduced milk production associated with papillomatous digital dermatitis in dairy cows in the USA. Journal of Dairy Research 73, 244–256. https://doi.org/10.1017/S0022029906001798

Manske, T., Hultgren, J., Bergsten, C., 2002. Topical treatment of digital dermatitis associated with severe heel-horn erosion in a Swedish dairy herd. Preventive Veterinary Medicine 53, 215–231. https://doi.org/10.1016/S0167-5877(01)00268-9

Node.js, 2023. Node.js [WWW Document]. Node.js. URL https://nodejs.org/en (accessed 7.23.23).

Oliveira, D.A.B., Pereira, L.G.R., Bresolin, T., Ferreira, R.E.P., Dorea, J.R.R., 2021. A review of deep learning algorithms for computer vision systems in livestock. Livestock Science 253, 104700. https://doi.org/10.1016/j.livsci.2021.104700

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S., 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library, in: Advances in Neural Information Processing Systems. Curran Associates, Inc.

Phan, T.H., Yamamoto, K., 2020. Resolving Class Imbalance in Object Detection with Weighted Cross Entropy Losses.

Plummer, P.J., Krull, A., 2017. Clinical Perspectives of Digital Dermatitis in Dairy and Beef Cattle. Veterinary Clinics of North America: Food Animal Practice 33, 165–181. https://doi.org/10.1016/j.cvfa.2017.02.002

Psota, E.T., Mittek, M., Pérez, L.C., Schmidt, T., Mote, B., 2019. Multi-Pig Part Detection and Association with a Fully-Convolutional Network. Sensors 19. https://doi.org/10.3390/s19040852

PyTorch, 2023. android-demo-app/ObjectDetection at master · pytorch/android-demo-app [WWW Document]. GitHub. URL https://github.com/pytorch/android-demo-app/tree/master/ObjectDetection (accessed 7.24.23).

pytorch, 2021. pytorch/pytorch. pytorch.

Read, D.H., Walker, R.L., 1998. Papillomatous Digital Dermatitis (Footwarts) in California Dairy Cattle: Clinical and Gross Pathologic Findings. J VET Diagn Invest 10, 67–76. https://doi.org/10.1177/104063879801000112

Redmon, J., 2018. YOLO: Real-Time Object Detection [WWW Document]. URL https://pjreddie.com/darknet/yolo/ (accessed 2.24.22).

Redmon, J., Divvala, S., Girshick, R., Farhadi, A., 2016. You Only Look Once: Unified, Real-Time Object Detection. arXiv:1506.02640 [cs].

Redmon, J., Farhadi, A., 2018. YOLOv3: An Incremental Improvement. arXiv:1804.02767 [cs].

Relun, A., Guatteo, R., Roussel, P., Bareille, N., 2011. A simple method to score digital dermatitis in dairy cows in the milking parlor. Journal of Dairy Science 94, 5424–5434. https://doi.org/10.3168/jds.2010-4054

Rodriguez-Lainz, A., Melendez-Retamal, P., Hird, D.W., Read, D.H., 1998. Papillomatous digital dermatitis in Chilean dairies and evaluation of a screening method. Preventive Veterinary Medicine 37, 197–207. https://doi.org/10.1016/S0167-5877(98)00091-9

Ryu, H.W., Tai, J.H., 2021. Object detection and tracking using a high-performance artificial intelligence-based 3D depth camera: towards early detection of African swine fever. J Vet Sci 23, e17. https://doi.org/10.4142/jvs.21252

Sanders, J., Kandrot, E., 2010. CUDA by example: an introduction to general-purpose GPU programming. Addison-Wesley Professional.

Schultz, N., Capion, N., 2013. Efficacy of salicylic acid in the treatment of digital dermatitis in dairy cattle. The Veterinary Journal 198, 518–523. https://doi.org/10.1016/j.tvjl.2013.09.002

Seo, J., Ahn, H., Kim, D., Lee, S., Chung, Y., Park, D., 2020. EmbeddedPigDet—Fast and Accurate Pig Detection for Embedded Board Implementations. Applied Sciences 10. https://doi.org/10.3390/app10082878

Sharma, V., Mir, R.N., 2020. A comprehensive and systematic look up into deep learning based object detection techniques: A review. Computer Science Review 38, 100301. https://doi.org/10.1016/j.cosrev.2020.100301

Shearer, J.K., Van Amstel, S.R., Gonzalez, A., 2005. Manual of foot care in cattle. Hoard's Dairyman Books.

Solano, L., Barkema, H.W., Jacobs, C., Orsel, K., 2017. Validation of the M-stage scoring system for digital dermatitis on dairy cows in the milking parlor. Journal of Dairy Science 100, 1592–1603. https://doi.org/10.3168/jds.2016-11365

Solano, L., Barkema, H.W., Mason, S., Pajor, E.A., LeBlanc, S.J., Orsel, K., 2016. Prevalence and distribution of foot lesions in dairy cattle in Alberta, Canada. Journal of Dairy Science 99, 6828–6841. https://doi.org/10.3168/jds.2016-10941

Stokes, J.E., Leach, K.A., Main, D.C.J., Whay, H.R., 2012. The reliability of detecting digital dermatitis in the milking parlour. The Veterinary Journal, Special Issue:Bovine Lameness 193, 679–684. https://doi.org/10.1016/j.tvjl.2012.06.053

Subramanian, N., Jeyaraj, A., 2018. Recent security challenges in cloud computing. Computers & Electrical Engineering 71, 28–42. https://doi.org/10.1016/j.compeleceng.2018.06.006

Szeliski, R., 2022. Computer Vision: Algorithms and Applications, 2nd ed. 2022 edition. ed. Springer, Cham.

Tashiev, I., 2022. XmlToTxt.

Techzizou, 2021. YOLOv4 VS YOLOv4-tiny. Analytics Vidhya. URL https://medium.com/analytics-vidhya/yolov4-vs-yolov4-tiny-97932b6ec8ec (accessed 2.24.22).

techzizou, 2021. Build Android app for custom object detection (TF 2.x). TECHZIZOU. URL https://techzizou.com/build-android-app-for-custom-object-detection-using-tf2/ (accessed 7.24.23).

TensorBoard, 2023. . tensorflow.

tensorflow, 2021. TensorFlow Lite for Microcontrollers. tensorflow.

TensorFlow.js, 2023a. . tensorflow.

TensorFlow.js, 2023b. Machine Learning for JavaScript Developers [WWW Document]. TensorFlow. URL https://www.tensorflow.org/js (accessed 7.22.23).

Thomsen, P.T., Klaas, I.C., Bach, K., 2008. Short Communication: Scoring of Digital Dermatitis During Milking as an Alternative to Scoring in a Hoof Trimming Chute. Journal of Dairy Science 91, 4679–4682. https://doi.org/10.3168/jds.2008-1342

Ultralytics, 2023a. Docker Quickstart [WWW Document]. GitHub. URL https://github.com/ultralytics/yolov5/wiki/Docker-Quickstart (accessed 7.23.23).

Ultralytics, 2023b. ultralytics/yolov5 - Docker Image | Docker Hub [WWW Document]. URL https://hub.docker.com/r/ultralytics/yolov5 (accessed 7.23.23).

Ultralytics, 2023c. metrics [WWW Document]. URL https://docs.ultralytics.com/reference/utils/metrics (accessed 9.26.23).

Ultralytics, 2023d. Android [WWW Document]. URL https://docs.ultralytics.com/hub/app/android (accessed 7.24.23).

Ultralytics, 2023e. HUB App [WWW Document]. Ultralytics. URL https://ultralytics.com/app (accessed 7.24.23).

USDA, 2009. NAHMS dairy 2007 part IV: reference of dairy cattle health and management practices in the United States.

Vanhoudt, A., Yang, D.A., Armstrong, T., Huxley, J.N., Laven, R.A., Manning, A.D., Newsome, R.F., Nielen, M., Werven, T. van, Bell, N.J., 2019. Interobserver agreement of digital dermatitis M-scores for photographs of the hind feet of standing dairy cattle. Journal of Dairy Science 102, 5466–5474. https://doi.org/10.3168/jds.2018-15644

Vashisht, M., Kumar, B., 2020. A Survey Paper on Object Detection Methods in Image Processing, in: 2020 International Conference on Computer Science, Engineering and Applications (ICCSEA). Presented at the 2020 International Conference on Computer Science, Engineering and Applications (ICCSEA), pp. 1–4. https://doi.org/10.1109/ICCSEA49143.2020.9132871

Viera, A.J., Garrett, J.M., 2005. Understanding interobserver agreement: the kappa statistic. Fam Med 37, 360–363.

Vink, W.D., 2006. Investigating the epidemiology of Bovine Digital Dermatitis: Causality, transmission and infection dynamics (PhD Thesis). University of Liverpool.

Wang, Y., Hong, K., Zou, J., Peng, T., Yang, H., 2020a. A CNN-Based Visual Sorting System With Cloud-Edge Computing for Flexible Manufacturing Systems. IEEE Transactions on Industrial Informatics 16, 4726–4735. https://doi.org/10.1109/TII.2019.2947539

Wang, Y., Liu, M., Zheng, P., Yang, H., Zou, J., 2020b. A smart surface inspection system using faster R-CNN in cloud-edge computing environment. Advanced Engineering Informatics 43, 101037. https://doi.org/10.1016/j.aei.2020.101037

Whay, H.R., Waterman, A.E., Webster, A.J.F., 1997. Associations between locomotion, claw lesions andnociceptive threshold in dairy heifers during the peri-partum period. The Veterinary Journal 154, 155–161. https://doi.org/10.1016/S1090-0233(97)80053-6

Wu, X., Sahoo, D., Hoi, S.C.H., 2020. Recent advances in deep learning for object detection. Neurocomputing 396, 39–64. https://doi.org/10.1016/j.neucom.2020.01.085

Wurtz, K., Camerlink, I., D’Eath, R.B., Fernández, A.P., Norton, T., Steibel, J., Siegford, J., 2019. Recording behaviour of indoor-housed farm animals automatically using machine vision technology: A systematic review. PLOS ONE 14, e0226669. https://doi.org/10.1371/journal.pone.0226669

Xia, G.-S., Bai, X., Ding, J., Zhu, Z., Belongie, S., Luo, J., Datcu, M., Pelillo, M., Zhang, L., 2018. DOTA: A Large-Scale Dataset for Object Detection in Aerial Images. Presented at the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 3974–3983.

Yang, D., Heuer, C., Laven, R., Vink, W., Chesterton, R., 2017. Farm and cow-level prevalence of bovine digital dermatitis on dairy farms in Taranaki, New Zealand. New Zealand Veterinary Journal 65, 252–256. https://doi.org/10.1080/00480169.2017.1344587

Zinicola, M., Lima, F., Lima, S., Machado, V., Gomez, M., Döpfer, D., Guard, C., Bicalho, R., 2015. Altered Microbiomes in Bovine Digital Dermatitis Lesions, and the Gut as a Pathogen Reservoir. PLOS ONE 10, e0120504. https://doi.org/10.1371/journal.pone.0120504

Zissis, D., Lekkas, D., 2012. Addressing cloud computing security issues. Future Generation Computer Systems 28, 583–592. https://doi.org/10.1016/j.future.2010.12.006

TABLES AND FIGURES

Figure 4.1 M-stage scoring system of digital dermatitis (DD) with signs of chronicity. Images are presented for each of the M-stages: M0 (top-left), M2 (top-center), M2P (top-left), M4H (bottom-left), and M4P (bottom-right).

Figure 4.2 YOLOv5s live detection in browser using TensorFlow.js on an image. Prediction is made using a bounding box for an M2 lesion with a confidence score of 72.6%.

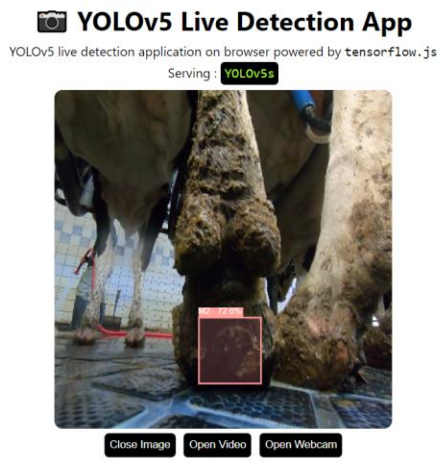

Figure 4.3 Average precision of the five M-stages by YOLOv5s for the testing set. Performance was measured in average precision and calculated for each of the five M-stages: M0, M2, M2P, M4H, and M4P.

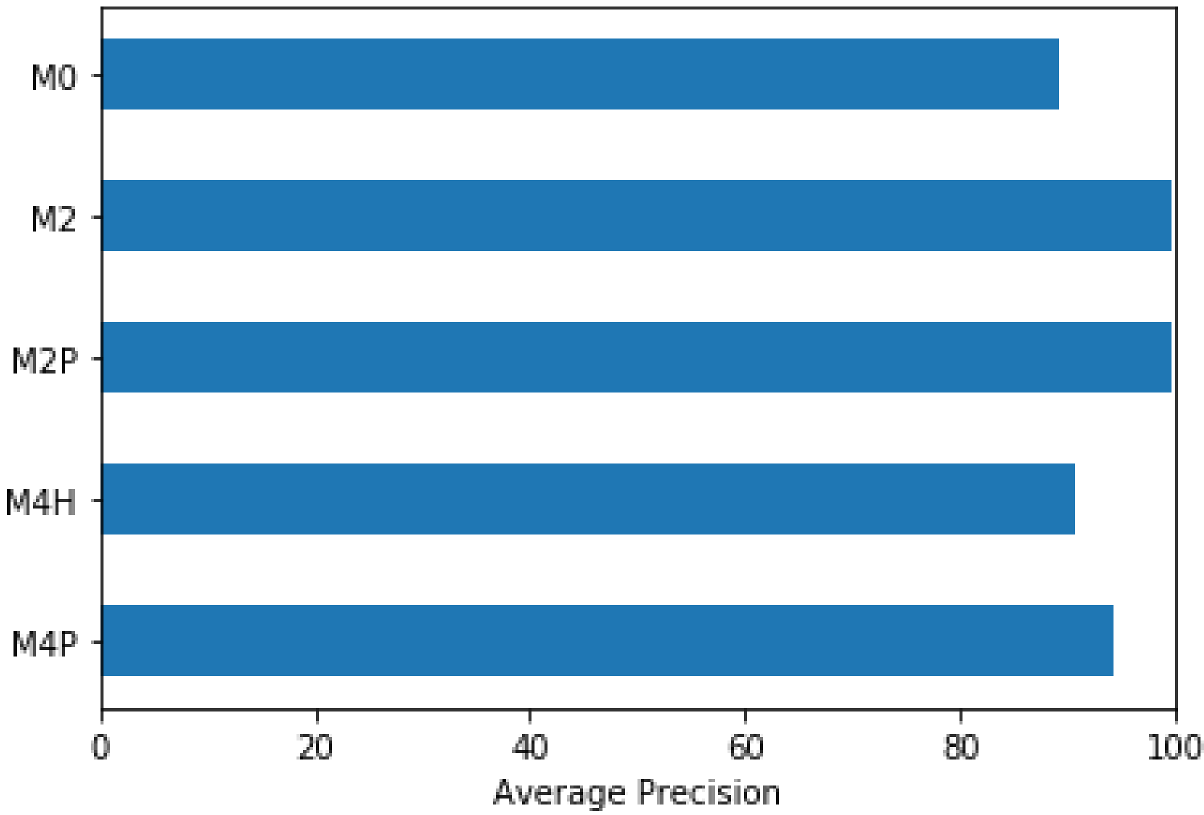

Figure 4.4 Bounding box predictions of M-stages by YOLOv5s on TensorFlow.js. Predictions are made for each of the M-stages: M0 (top-left; red), M2 (top-center; pink), M2P (top-right; orange), M4H (bottom-left; light orange), and M4P (bottom-right; light green).

Figure 4.5 Accuracy of the five M-stages by YOLOv5s on TensorFlow.js. Accuracy is measured as the number of correct predictions and calculated for each of the five M-stages: M0, M2, M2P, M4H, and M4P.

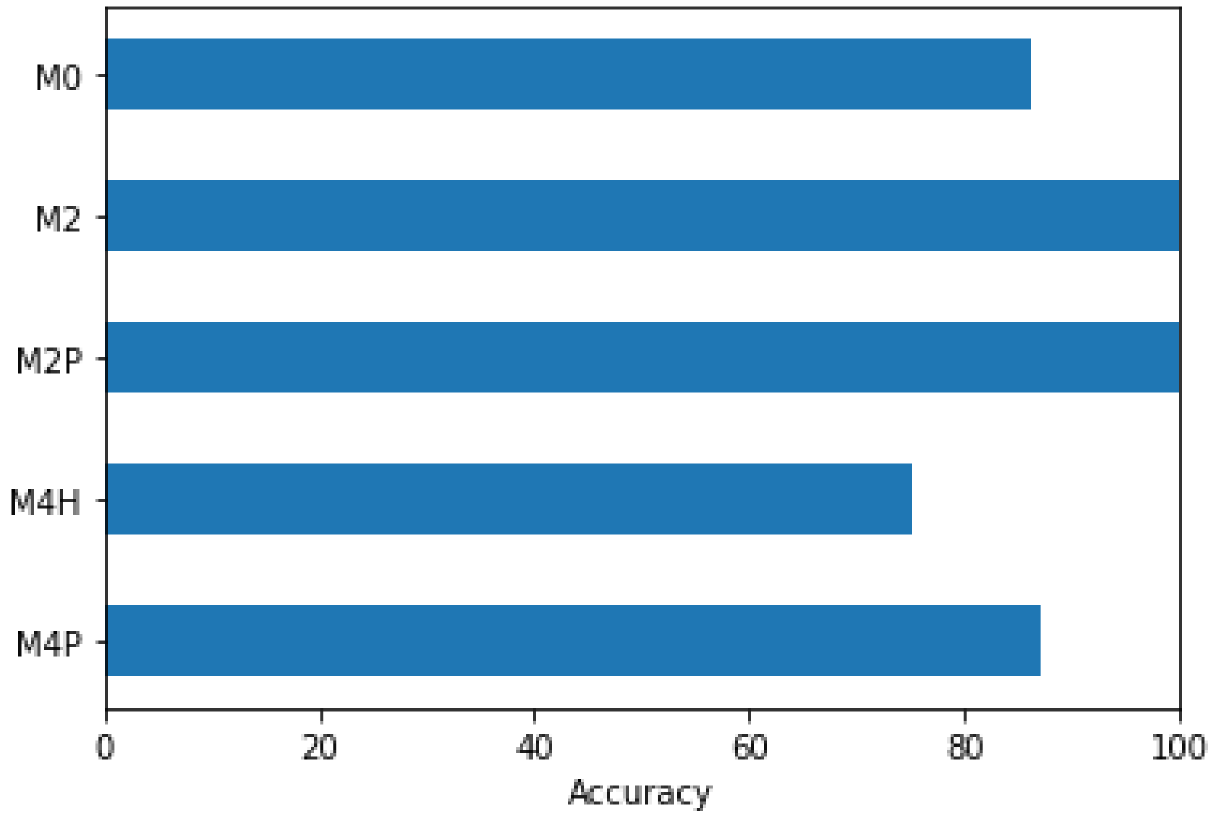

Figure 4.6 Inference time of the four implementations. Inference time was measured in frames per second (FPS) and the fours implementations Google Colab, Docker using an IP camera via HTTP, Docker using an IP camera via RTSP, and TensorFlow.js.

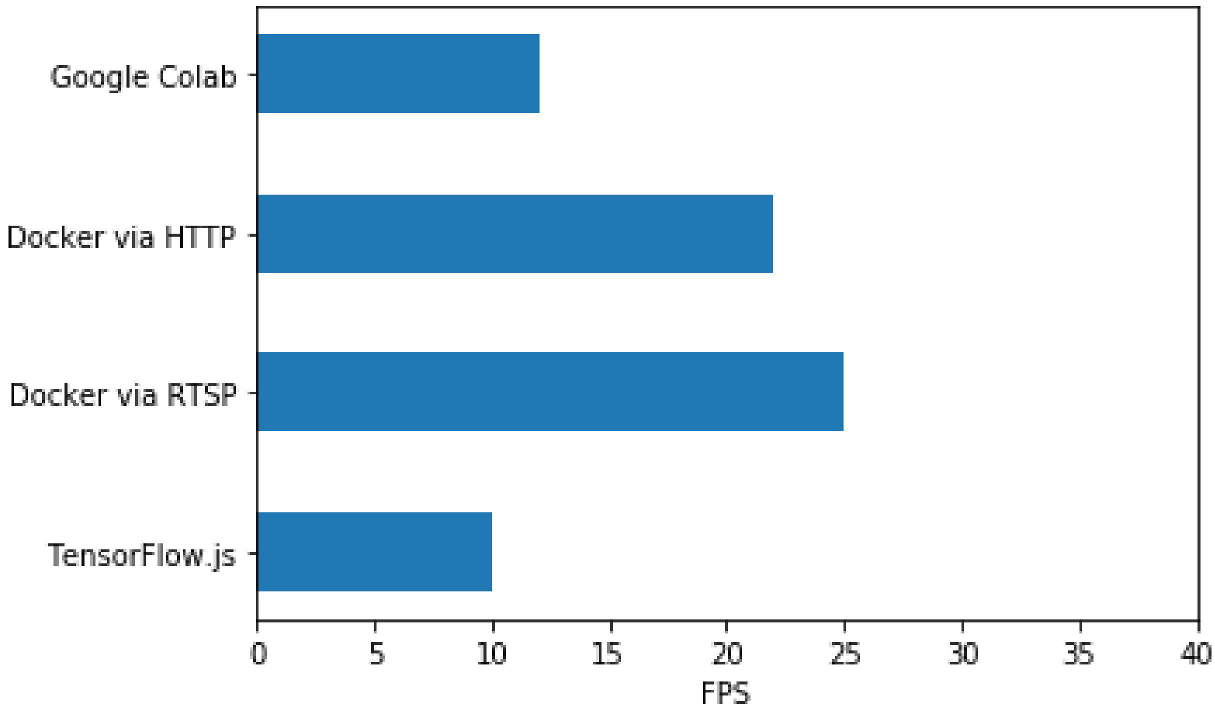

CHAPTER 5 PAWGNOSIS - COMPUTER VISION MODELS FOR THE DETECTION OF CANINE PODODERMATITIS AND NEOPLASIA OF THE PAW

ABSTRACT

Artificial intelligence (AI) has been used successfully in human dermatology. AI utilizes convolutional neural networks (CNN) to accomplish tasks such as image classification, object detection, and segmentation, facilitating early diagnosis. Computer vision (CV), a field of AI, has shown great results in detecting signs of human skin diseases. Canine paw skin diseases are a common problem in general veterinary practice, and computer vision tools could facilitate the detection and monitoring of disease processes. Currently, no such tool is available for veterinary dermatology. Digital images of paws from healthy dogs as well as paws with pododermatitis or neoplasia were used. We tested a novel object detection model - Pawgnosis, a Tiny YOLOv4 image analysis model deployed on a microcomputer with a camera for the rapid detection of canine pododermatitis and neoplasia. The prediction performance metrics used to evaluate the models included mean average precision (mAP), precision, recall, average precision (AP) for accuracy, and frames per second (FPS) for speed. A large dataset labeled by a single individual (Dataset A) used to train a Tiny YOLOv4 model provided the best results with a mean mAP of 0.95, precision of 0.86, recall of 0.93, and 20 FPS. The novel object detection model has the potential for application in the field of veterinary dermatology.

INTRODUCTION

Artificial intelligence (AI) can provide valuable tools for veterinary medicine to alleviate diagnosis difficulties and long-term disease management problems, similar to their applications in human medicine. Computer vision (CV) is a field of AI that enables computers to derive meaningful predictions of various aspects of diseases, by interpreting digital images, videos, and other visual sources (Vaidya and Paunwala, 2019; Wu et al., 2020). In human dermatology, CV approaches to detecting signs of skin diseases have been demonstrated to perform equally to those of board-certified specialists (Pham et al., 2021). Object detection is a CV approach and is the process of identifying and locating specific objects in images or videos using computer vision algorithms. It can facilitate a fast, non-invasive diagnosis, that does not require additional staff or manual labor. Diagnostic support for skin lesions or disease detection would be a significant asset to veterinarians worldwide.

Convolutional neural networks (CNN) are types of neural networks frequently utilized for medical image analysis (Yu et al., 2021). These networks can accomplish tasks such as image classification, object detection, and segmentation. Large datasets of thousands to millions of images have been utilized to train CNN models to classify and detect a wide array of objects (Deng et al., 2009; Weiss et al., 2016).

One limiting factor for creating such models is obtaining considerable and comprehensive enough dataset to train a CNN structure. The use of transfer learning, a method in which pre-existing models are pretrained on extensive image datasets, represents a feasible approach for training CV models for detection of signs of disease (Deng et al., 2009). After transfer learning, the CV models can be tailored to the specific task of interest (Weiss et al., 2016). In machine learning and CV, object detection models analyze an image and identify the presence of specified object classes within it. Upon detection, bounding boxes demarcate these objects. The model identifies the class by shape, color, and texture, differentiating from the background and calculates a class probability (Zhao et al., 2019).

You Only Look Once (YOLO) is a CNN for performing real-time object detection deep learning that has become a standard method in the field of CV (Redmon et al., 2016). The first version of YOLO, was released in 2015 (Redmon et al., 2016). YOLOv4 is an architecture originally implemented in a TensorFlow framework detecting objects in one step (Bochkovskiy et al., 2020). Tiny YOLOv4 is a compressed version of YOLOv4 designed for smaller devices with limited computing power (Jiang et al., 2020). Tiny YOLOv4 models have a simpler network structure and fewer parameters to train compared to YOLOv4 (Adarsh et al., 2020). Higher speed is traded-off to lower accuracy by Tiny YOLOv4 compared to YOLOv4 (Alexey, 2023; Bochkovskiy et al., 2020; Techzizou, 2021).

The paws are a discrete body region with well-recognized disease conditions for exploring the application of object detection, particularly for research purposes. One of the most common diseases is canine pododermatitis (Bajwa, 2016). This study aims to train and deploy a CV pododermatitis model named Pawgnosis, which can detect healthy paws, canine pododermatitis, and neoplasia cases based on images, videos, or live camera capture in real-time.

MATERIALS AND METHODS

Image Collection and Definition of Classes

The canine paw images used in this study for Datasets A, B, and C described below were collected by board-certified dermatologists with owner permission from dermatology patients of the last author, friends, and family. Written informed consent for photographs or video imaging was obtained from the owners. The criteria for image collection of canine paws included being healthy, having pododermatitis, or showing signs of neoplasia. Images used in this study were taken in ventral, dorsal, and lateral position of the dog's paw, with various backgrounds and lighting. Images were taken with the interdigital spaces open and closed, as well as of the palmar and plantar aspects between the footpads.

The three object classes used in this study are defined as follows. "Healthy" paws were canine paws free from any clinically observable disease with no signs of inflammation, abrasions, or masses. "Pododermatitis" paws were defined as canine paws exhibiting clinical signs of inflammation due to any underlying diseases without signs of neoplasia. Signs of inflammation on the feet include single or multiple lesions that are dry or crusted, edematous, erythematous, nodular, ulcerated, exudative, and may have focal areas of alopecia. "Neoplasia" paws were defined as paws that had a mass and, after additional diagnostic testing, were diagnosed as squamous cell carcinoma, melanoma, osteosarcoma, mast cell tumors, or malignant soft tissue sarcomas. Lymphoma cases were not included, due to lack of images.

Image Labeling and Definition of Datasets

After collection, the images were manually labeled using the three previously defined classes (Healthy, Pododermatitis, and Neoplasia) via LabelImg (Lin, 2018). YOLO formatted bounding boxes and annotation files were created for each image of the canine paws for each of the three classes. These boxes were then labeled "healthy", "pododermatitis", or "neoplasia".

Three different datasets were generated to compare the predictive performances of the resulting object detection models. Mean average precision (mAP), precision, and recall metrics for all classes and average precision (AP) for each class were computed (Chollet, 2021). For each model, the inference time to predict bounding boxes and corresponding class labels in an image was measured using frames per second (FPS). Dataset A contained 575 images labeled by one person drawing relatively wide boxes around the entire feet area during labeling. Dataset B contained the same 575 images labeled by three people. Labeling was divided equally with a mix of narrow boxes around lesions or wide boxes around the area of the affected feet during labeling resulting in multiple boxes per image. Dataset C included 301 images labeled by two people drawing a mix of relatively narrow boxes around the affected feet area, resulting in multiple nonoverlapping bounding boxes, or wide boxes around the area of the affected feet. For Dataset B and C, images were randomly split between the labelers, labeled by single labeler, and reviewed by the other labelers for consistency of box size and consensus of classification.

Model Building

All data processing and model training were performed in Python 3.8 (Van Rossum et al., 1995). The three data sets were split at random into 90% train and 10% validation image sets without data leakage where images of the same paw are deleted and are not part of the train and validation data sets simultaneously.

For the purpose of transfer learning, Tiny YOLOv4 model was initialized using classification weights pre-trained on ImageNet (Redmon and Farhadi, 2018). Models were trained using an input size of 416 x 416 pixels for a maximum number of batches of 6000 with a batch size of 64 and a learning rate of 0.00261. The models were trained in triplicates for each Dataset A, B, and C and prediction performance metrics were averaged over the three training runs.

The weights of the best model over the three training runs of Dataset A were used for deployment of the Tiny YOLOv4 model on an edge device. Using the same training and validation data sets from the best of three Tiny YOLOv4 training runs, YOLO5s, and Tiny YOLOv7 models were trained. The YOLOv5 and YOLOv7 implementations are based on Pytorch and were used as a comparison for speed of detections on an edge device (Jocher, 2020; Paszke et al., 2019; Wang et al., 2023). The edge device used for deployment of the Tiny YOLOV4, YOLOV5s, and Tiny YOLOv7 models was an NVIDIA Jetson Xavier NX connected to a Luxonis OAK-1 camera that applies the DepthAI framework for inference (Luxonis, n.d.; NVIDIA, n.d.). This set-up of the edge device prototype was independent from web access and battery-powered for 8 hours of continuous detection. For deployment of the Tiny YOLOv4 model in our veterinary dermatology clinic, a Google Colab notebook was adapted to run a real-time built-in detection camera on a smartphone (Google, n.d.). This deployment required web access and a free Google Colab account.

RESULTS

The training duration for each model was under two hours. The mAP values for all three datasets were stable after an iteration number of approximately 3000. The mAP for Dataset A was higher than for both Datasets B and C (Figure 1). The mAP and recall for Dataset A were higher than for both Dataset B and C, while the precision was similar for all three datasets (Figure 2). The average precision (AP) of pododermatitis and neoplasia for Dataset A was higher than both Dataset B and C whereas the AP of healthy was similar for all three datasets (Figure 3). Tiny YOLOv7 has the highest mAP (0.973) compared to the other two models. We deployed these models on an edge device, where Tiny YOLOv7 was the fastest with 40 FPS, while Tiny YOLOv4 and YOLOv5s had 20 FPS.

Figure 4 shows typical images for the three classes Healthy, Pododermatitis, and Neoplasia, and their detections using bounding boxes with class probabilities for each of the three datasets (A, B, and C). Detections of the three classes from Dataset A (Figure 4A – 4C) show larger bounding boxes and higher class probabilities compared to smaller bounding boxes and lower class probabilities for Datasets B and C. No detections resulted from Dataset C for the healthy and neoplasia classes (Figure 4G and 4I). The smaller bounding boxes from Datasets B and C were occasionally drawn outside of the relevant paw area in the images from Pododermatitis (Figure 4E and 4H).

We conducted a preliminary limited deployment of the Tiny YOLOv4 model in our dermatology exam room (Figure 5). Figure 5A shows the Jetson Xavier NX single-board computer connected to a Luxonis OAK-1 camera for the real-time detection of the three classes Healthy, Pododermatitis, and Neoplasia. Figure 5B shows the resulting bounding box, class label, and class prediction probability during real-time detection of a pododermatitis lesion on a canine paw. Figure 5C shows the schematic set-up of the edge device used for deployment in the clinical setting.

DISCUSSION

In this study, we apply CNNs to detect and classify objects in images in real-time. The Pawgnosis model predicts whether a canine paw shows characteristics of pododermatitis or neoplasia, or is clinically normal, with high performance after learning from a limited number of images. We demonstrate that the Pawgnosis tool can detect the three classes with supervision in a veterinary clinical setting.

Models constrained to these three classes are not without limitations, and the applications of such models could be expanded by subdividing and increasing the number of these classes. The model could be broadened to detect early flares of pododermatitis due to atopic dermatitis. Further studies can be performed to show agreement between the model detections and the validated atopic dermatitis scoring system CADESI-4 or other scores (Olivry et al., 2014).

The differences in model performance between the three Datasets A, B, and C emphasize the importance of sample size and consistent labeling techniques for achieving optimized prediction performance of CV models. The model returned higher mean mAP when trained on Dataset B (575 images with multiple labellers) compared to training on Dataset C (301 images with multiple labellers. When comparing a single labeller (Dataset A) to multiple labellers (Dataset B and C), the model returned higher mean mAP values when trained on Dataset A compared to training on Datasets B and C (Figure 1). This holds true across all other performance metrics including precision and recall (Figure 2). A single labeller achieves a better performance than multiple persons labeling. Dataset A outperformed Datasets B and C for mAP (0.95 compared to 0.54 and 0.53, respectively), recall (0.93 compared to 0.55 and 0.40, respectively), and precision (0.82 compared to 0.71 and 0.66, respectively).

The mAP of Dataset A is higher than the mAPs of Datasets B and C for both the pododermatitis and neoplasia classes, while the difference in AP between Dataset A and the other two datasets is reduced for the healthy class (Figure 3). This suggests that labeling disproportionally affected paws with dermatologic lesions compared to healthy paws. Healthy paws were generally labelled with a single box around the entire paw across all datasets. Therefore, more boxes with less general features would result in lower accuracy.

The proportion of the performance for Dataset A attributed to a single labeller or an increase in size and simplicity of labeled boxes is unknown. We assume that the benefits of a single labeller attributed to a reduction of inter-labeller variation, resulting in improved labeling consistency within the dataset. An increase in size of the labeled boxes to include the whole paw also decreased the inter-labeller variation. Consistent labeling of the whole paw in Dataset A was more accurate than labeling of individual lesions in Datasets B and C. Further studies comparing datasets with the same labeling criteria and different number of labellers to datasets with different labeling criteria and the same number of labellers can be performed to determine the true relative contribution for each factor. For the current analysis, an increased labeling consistency and increased sample size of datasets increased the performance of object detection CV models for dermatologic lesions. Labeling under the supervision of a board-certified dermatologists is strongly recommended.

Currently, the use of AI in veterinary medicine is a relatively new and emerging field, mostly applied to large animals, clinical and anatomic pathology, and radiology (Appleby and Basran, 2022; Aubreville et al., 2020; Basran and Appleby, 2022; Borges Oliveira et al., 2021; Celniak et al., 2023; Cernek et al., 2020; Gupta et al., 2022; Joslyn and Alexander, 2022; Nagamori et al., 2021, 2020; Oczak et al., 2022; Wu et al., 2023; Zhang et al., 2019). However, the use of AI for skin diseases in dogs is rare (Habal et al., 2021; Hwang et al., 2022; Upadhyay et al., 2023). Previous studies used image classification to label entire images, while the current approach uses object detection to localize in addition to label individual objects within an image. A recent study reported a YOLOv5 model for dry eye disease in dogs with a very high mAP of 0.995 (Kim et al., 2022). Another group built a model for dermatophytes, mange, and fleas (Upadhyay et al., 2023). Another model evaluated 12 dog skin diseases (yeast infection, folliculitis, impetigo, seborrhea, dermatophytes, alopecia, mange, fleas, color changes, acral lick/granuloma, skin tumors, hot spots, and anal sac disease) (Habal et al., 2021). In 2022 a researchers evaluated images from dogs with fungal skin infection, bacterial dermatitis, and allergy (Hwang et al., 2022). No further details are provided regarding how the diagnoses were made. From a dermatological perspective, it was also unknown if the dogs had a secondary skin infection and what was the most common complication or not. Limited information was provided by all three studies regarding the individuals that made the diagnosis, and their knowledge and training in veterinary dermatology. Overall, these previous studies as well as the current study provide the first few steps and serve as a launch pad for the implementation of AI in veterinary dermatology.

Only YOLO models were implemented to compare the performance of the three labeled datasets and embedded on an edge device to detect the three classes in a clinical setting. Further studies can explore the differences between the three state-of-the-art YOLO models and other object detection models. Two-stage object detection models such as Faster R-CNNs and Cascade R-CNNs can increase accuracy for improved prediction on a stand-alone device (Howard et al., 2017; Soviany and Ionescu, 2018). Other one-stage object detection models such as SSD and SSD Lite can be light-weight and increase speed for improved prediction on a mobile platform (Howard et al., 2017; Soviany and Ionescu, 2018). However, YOLO models provided higher speed and similar mAP compared to SSD, ResNet, and other models in human dermatology (Ding et al., 2022; Tan et al., 2021).

Our next goal is to deploy the YOLO model on an edge device in a clinical setting for external validation study. CV can be deployed on a laptop or desktop, an edge device, or a smartphone cloud-based application. Depending on the use purpose of each model a different option can be deployed. These models can be used for teaching both veterinary students and veterinarians, aiding in the decision and assisting in the diagnosis.

CONCLUSION

The developed "Pawgnosis" tool is the first object detection model using CV in veterinary dermatology to our knowledge. It has the potential to become an accurate and fast CV model for the management of canine pododermatitis. The model can be further improved for real-time detection of pododermatitis and monitoring of progression or treatment effects. It can also make recommendations for future diagnostic steps. Implementing the Pawgnosis on portable devices with dermatologists will further optimize the model. Pawgnosis may be used either as a clinical, research, or didactic tool. Pawgnosis may improve the field of veterinary dermatology, the welfare of dermatologic patients, and increase compliance of pet owners with therapeutic plans to manage canine pododermatitis. Further studies are needed to expand its abilities and validate its generalizability and applicability in everyday clinical practice.

REFERENCES

Adarsh, P., Rathi, P., Kumar, M., 2020. YOLO v3-Tiny: Object Detection and Recognition using one stage improved model, in: 2020 6th International Conference on Advanced Computing and Communication Systems (ICACCS). Presented at the 2020 6th International Conference on Advanced Computing and Communication Systems (ICACCS), pp. 687–694. https://doi.org/10.1109/ICACCS48705.2020.9074315

Alexey, 2023. Yolo v4, v3 and v2 for Windows and Linux.

Appleby, R.B., Basran, P.S., 2022. Artificial intelligence in veterinary medicine. Journal of the American Veterinary Medical Association 260, 819–824. https://doi.org/10.2460/javma.22.03.0093

Aubreville, M., Bertram, C.A., Marzahl, C., Gurtner, C., Dettwiler, M., Schmidt, A., Bartenschlager, F., Merz, S., Fragoso, M., Kershaw, O., Klopfleisch, R., Maier, A., 2020. Deep learning algorithms out-perform veterinary pathologists in detecting the mitotically most active tumor region. Sci Rep 10, 16447. https://doi.org/10.1038/s41598-020-73246-2

Bajwa, J., 2016. Canine pododermatitis. Can Vet J 57, 991–993.

Basran, P.S., Appleby, R.B., 2022. The unmet potential of artificial intelligence in veterinary medicine. Am J Vet Res 83, 385–392. https://doi.org/10.2460/ajvr.22.03.0038

Bochkovskiy, A., Wang, C.-Y., Liao, H.-Y.M., 2020. YOLOv4: Optimal Speed and Accuracy of Object Detection. https://doi.org/10.48550/arXiv.2004.10934

Borges Oliveira, D.A., Ribeiro Pereira, L.G., Bresolin, T., Pontes Ferreira, R.E., Reboucas Dorea, J.R., 2021. A review of deep learning algorithms for computer vision systems in livestock. Livestock Science 253, 104700. https://doi.org/10.1016/j.livsci.2021.104700

Celniak, W., Wodziński, M., Jurgas, A., Burti, S., Zotti, A., Atzori, M., Müller, H., Banzato, T., 2023. Improving the classification of veterinary thoracic radiographs through inter-species and inter-pathology self-supervised pre-training of deep learning models. Sci Rep 13, 19518. https://doi.org/10.1038/s41598-023-46345-z

Cernek, P., Bollig, N., Anklam, K., Döpfer, D., 2020. Hot topic: Detecting digital dermatitis with computer vision. Journal of Dairy Science 103, 9110–9115. https://doi.org/10.3168/jds.2019-17478

Chollet, F., 2021. Deep Learning with Python, Second Edition. Simon and Schuster.

Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., Fei-Fei, L., 2009. ImageNet: A large-scale hierarchical image database, in: 2009 IEEE Conference on Computer Vision and Pattern Recognition. Presented at the 2009 IEEE Conference on Computer Vision and Pattern Recognition, pp. 248–255. https://doi.org/10.1109/CVPR.2009.5206848

Ding, H., Zhang, E., Fang, F., Liu, X., Zheng, H., Yang, H., Ge, Y., Yang, Y., Lin, T., 2022. Automatic identification of benign pigmented skin lesions from clinical images using deep convolutional neural network. BMC Biotechnol 22, 28. https://doi.org/10.1186/s12896-022-00755-5

Google, n.d. Google Colaboratory [WWW Document]. URL https://colab.research.google.com/#

Gupta, H., Jindal, P., Verma, O.P., Arya, R.K., Ateya, A.A., Soliman, N.F., Mohan, V., 2022. Computer Vision-Based Approach for Automatic Detection of Dairy Cow Breed. Electronics 11, 3791. https://doi.org/10.3390/electronics11223791

Habal, B.G.M., Tiong, P.E.S., Pasatiempo, J.R., Balen, M.J., Amarga, M.R., Juco, L., 2021. Dog Skin Disease Recognition Using Image Segmentation and GPU Enhanced Convolutional Neural Network, in: 2021 IEEE 13th International Conference on Humanoid, Nanotechnology, Information Technology, Communication and Control, Environment, and Management (HNICEM). Presented at the 2021 IEEE 13th International Conference on Humanoid, Nanotechnology, Information Technology, Communication and Control, Environment, and Management (HNICEM), pp. 1–5. https://doi.org/10.1109/HNICEM54116.2021.9731885

Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H., 2017. MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications. https://doi.org/10.48550/arXiv.1704.04861

Hwang, S., Shin, H.K., Park, J.M., Kwon, B., Kang, M.-G., 2022. Classification of dog skin diseases using deep learning with images captured from multispectral imaging device. Mol. Cell. Toxicol. 18, 299–309. https://doi.org/10.1007/s13273-022-00249-7

Jiang, Z., Zhao, L., Li, S., Jia, Y., 2020. Real-time object detection method based on improved YOLOv4-tiny. https://doi.org/10.48550/arXiv.2011.04244

Jocher, G., 2020. YOLOv5 by Ultralytics. https://doi.org/10.5281/zenodo.3908559

Joslyn, S., Alexander, K., 2022. Evaluating artificial intelligence algorithms for use in veterinary radiology. Veterinary Radiology & Ultrasound 63, 871–879. https://doi.org/10.1111/vru.13159

Kim, J.Y., Han, M.G., Chun, J.H., Huh, E.A., Lee, S.J., 2022. Developing a diagnosis model for dry eye disease in dogs using object detection. Sci Rep 12, 21351. https://doi.org/10.1038/s41598-022-25867-y

Lin, T., 2018. LabelImg.

Luxonis, n.d. DepthAI's Documentation — DepthAI documentation | Luxonis [WWW Document]. URL https://docs.luxonis.com/en/latest/

Nagamori, Y., Hall Sedlak, R., DeRosa, A., Pullins, A., Cree, T., Loenser, M., Larson, B.S., Smith, R.B., Goldstein, R., 2020. Evaluation of the VETSCAN IMAGYST: an in-clinic canine and feline fecal parasite detection system integrated with a deep learning algorithm. Parasites Vectors 13, 346. https://doi.org/10.1186/s13071-020-04215-x

Nagamori, Y., Sedlak, R.H., DeRosa, A., Pullins, A., Cree, T., Loenser, M., Larson, B.S., Smith, R.B., Penn, C., Goldstein, R., 2021. Further evaluation and validation of the VETSCAN IMAGYST: in-clinic feline and canine fecal parasite detection system integrated with a deep learning algorithm. Parasites & Vectors 14, 89. https://doi.org/10.1186/s13071-021-04591-y

NVIDIA, n.d. NVIDIA Jetson Xavier Series [WWW Document]. NVIDIA. URL https://www.nvidia.com/en-us/autonomous-machines/embedded-systems/jetson-xavier-series/

Oczak, M., Bayer, F., Vetter, S.G., Maschat, K., Baumgartner, J., 2022. Where is the sow's nose: RetinaNet object detector as a basis for monitoring the use of rack with nest-building material. Frontiers in Animal Science 3.

Olivry, T., Saridomichelakis, M., Nuttall, T., Bensignor, E., Griffin, C.E., Hill, P.B., 2014. Validation of the Canine Atopic Dermatitis Extent and Severity Index (CADESI)-4, a simplified severity scale for assessing skin lesions of atopic dermatitis in dogs. Veterinary Dermatology 25, 77-e25. https://doi.org/10.1111/vde.12107

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S., 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library, in: Advances in Neural Information Processing Systems. Curran Associates, Inc.

Pham, T.-C., Luong, C.-M., Hoang, V.-D., Doucet, A., 2021. AI outperformed every dermatologist in dermoscopic melanoma diagnosis, using an optimized deep-CNN architecture with custom mini-batch logic and loss function. Sci Rep 11, 17485. https://doi.org/10.1038/s41598-021-96707-8

Redmon, J., Divvala, S., Girshick, R., Farhadi, A., 2016. You Only Look Once: Unified, Real-Time Object Detection. Presented at the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 779–788.

Redmon, J., Farhadi, A., 2018. YOLOv3: An Incremental Improvement. https://doi.org/10.48550/arXiv.1804.02767

Soviany, P., Ionescu, R.T., 2018. Optimizing the Trade-Off between Single-Stage and Two-Stage Deep Object Detectors using Image Difficulty Prediction, in: 2018 20th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC). Presented at the 2018 20th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC), pp. 209–214. https://doi.org/10.1109/SYNASC.2018.00041

Tan, L., Huangfu, T., Wu, L., Chen, W., 2021. Comparison of RetinaNet, SSD, and YOLO v3 for real-time pill identification. BMC Med Inform Decis Mak 21, 324. https://doi.org/10.1186/s12911-021-01691-8

Techzizou, 2021. YOLOv4 VS YOLOv4-tiny. Analytics Vidhya. URL https://medium.com/analytics-vidhya/yolov4-vs-yolov4-tiny-97932b6ec8ec (accessed 2.24.22).

Upadhyay, A., Singh, G., Mhatre, S., Nadar, P., 2023. Dog Skin Diseases Detection and Identification Using Convolutional Neural Networks. SN COMPUT. SCI. 4, 250. https://doi.org/10.1007/s42979-022-01645-5

Vaidya, B., Paunwala, C., 2019. Deep Learning Architectures for Object Detection and Classification, in: Mishra, M.K., Mishra, B.S.P., Patel, Y.S., Misra, R. (Eds.), Smart Techniques for a Smarter Planet: Towards Smarter Algorithms, Studies in Fuzziness and Soft Computing. Springer International Publishing, Cham, pp. 53–79. https://doi.org/10.1007/978-3-030-03131-2_4

Van Rossum, G., Drake, F.L., others, 1995. Python reference manual. Centrum voor Wiskunde en Informatica Amsterdam.

Wang, C.-Y., Bochkovskiy, A., Liao, H.-Y.M., 2023. YOLOv7: Trainable Bag-of-Freebies Sets New State-of-the-Art for Real-Time Object Detectors. Presented at the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 7464–7475.

Weiss, K., Khoshgoftaar, T.M., Wang, D., 2016. A survey of transfer learning. Journal of Big Data 3, 9. https://doi.org/10.1186/s40537-016-0043-6

Wu, D., Han, M., Song, H., Song, L., Duan, Y., 2023. Monitoring the respiratory behavior of multiple cows based on computer vision and deep learning. Journal of Dairy Science 106, 2963–2979. https://doi.org/10.3168/jds.2022-22501

Wu, X., Sahoo, D., Hoi, S.C.H., 2020. Recent advances in deep learning for object detection. Neurocomputing 396, 39–64. https://doi.org/10.1016/j.neucom.2020.01.085

Yu, H., Yang, L.T., Zhang, Q., Armstrong, D., Deen, M.J., 2021. Convolutional neural networks for medical image analysis: State-of-the-art, comparisons, improvement and perspectives. Neurocomputing 444, 92–110. https://doi.org/10.1016/j.neucom.2020.04.157

Zhang, L., Gray, H., Ye, X., Collins, L., Allinson, N., 2019. Automatic Individual Pig Detection and Tracking in Pig Farms. Sensors 19, 1188. https://doi.org/10.3390/s19051188

Zhao, Z.-Q., Zheng, P., Xu, S.-T., Wu, X., 2019. Object Detection With Deep Learning: A Review. IEEE Transactions on Neural Networks and Learning Systems 30, 3212–3232. https://doi.org/10.1109/TNNLS.2018.2876865

TABLES AND FIGURES

Figure 5.1 Mean average precision (mAP) over iteration number for training Tiny YOLOv4 custom models. The plots are grouped by training dataset where the thin, transparent lines correspond to three different training sessions per dataset and the thick, opaque line corresponds to the means of mAP from the three training sessions per dataset: Dataset A (red), Dataset B (green), and Dataset C (blue).

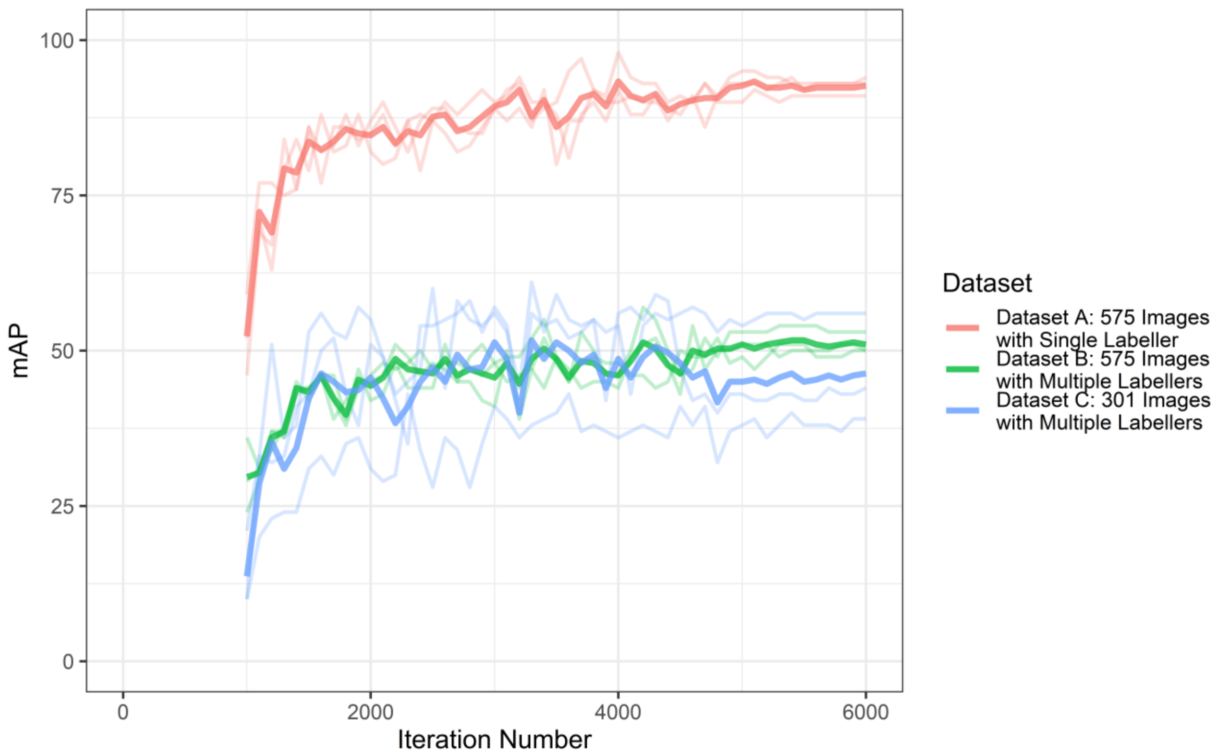

Figure 5.2 Mean and 95% confidence intervals for three performance metrics of each of the three datasets using Tiny YOLOv4. The plots are grouped by mean average precision (mAP), precision, and recall for each dataset. Small circles correspond to distinct training runs, large circles correspond to group means, and the vertical lines correspond to the 95% confidence intervals: Dataset A (red), Dataset B (green), and Dataset C (blue).

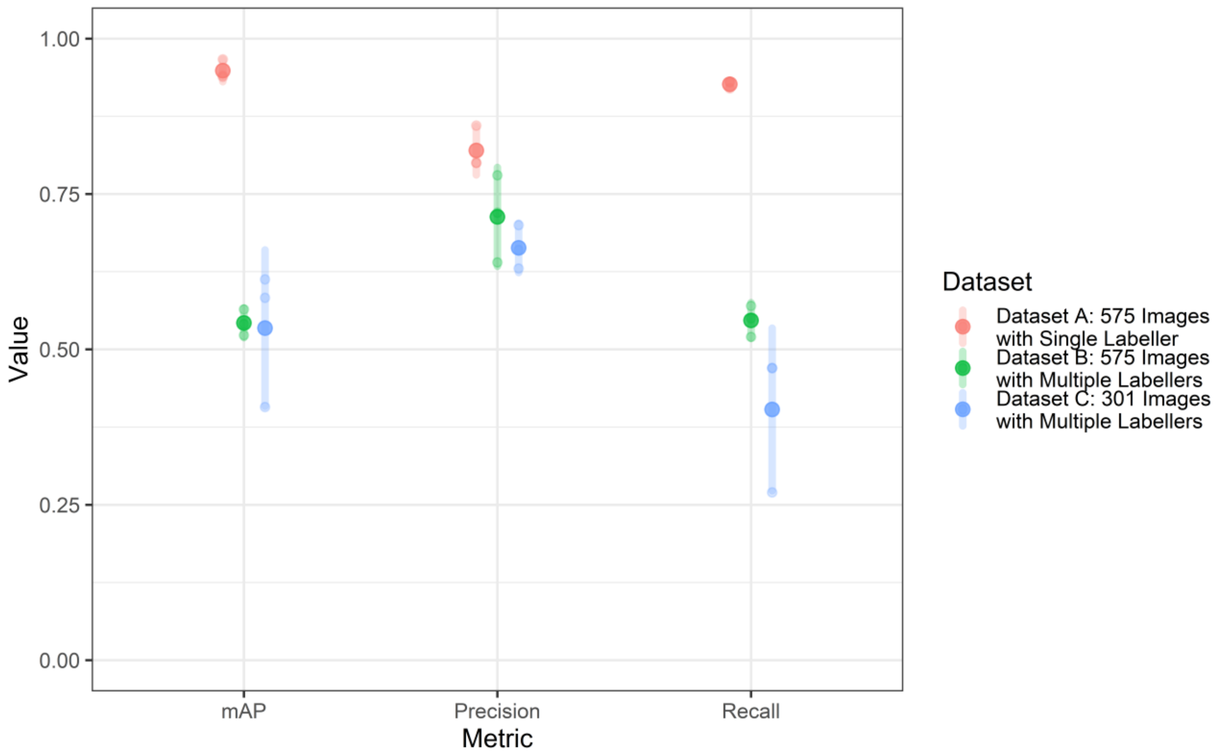

Figure 5.3 Mean and 95% confidence intervals for average precision (AP) of each of the three classes, Healthy, Pododermatitis, and Neoplasia, for each of the three datasets using Tiny YOLOv4. Small circles correspond to distinct runs, large circles correspond to group means, and the vertical lines correspond to the 95% confidence interval: Dataset A (red), Dataset B (green), and Dataset C (blue).

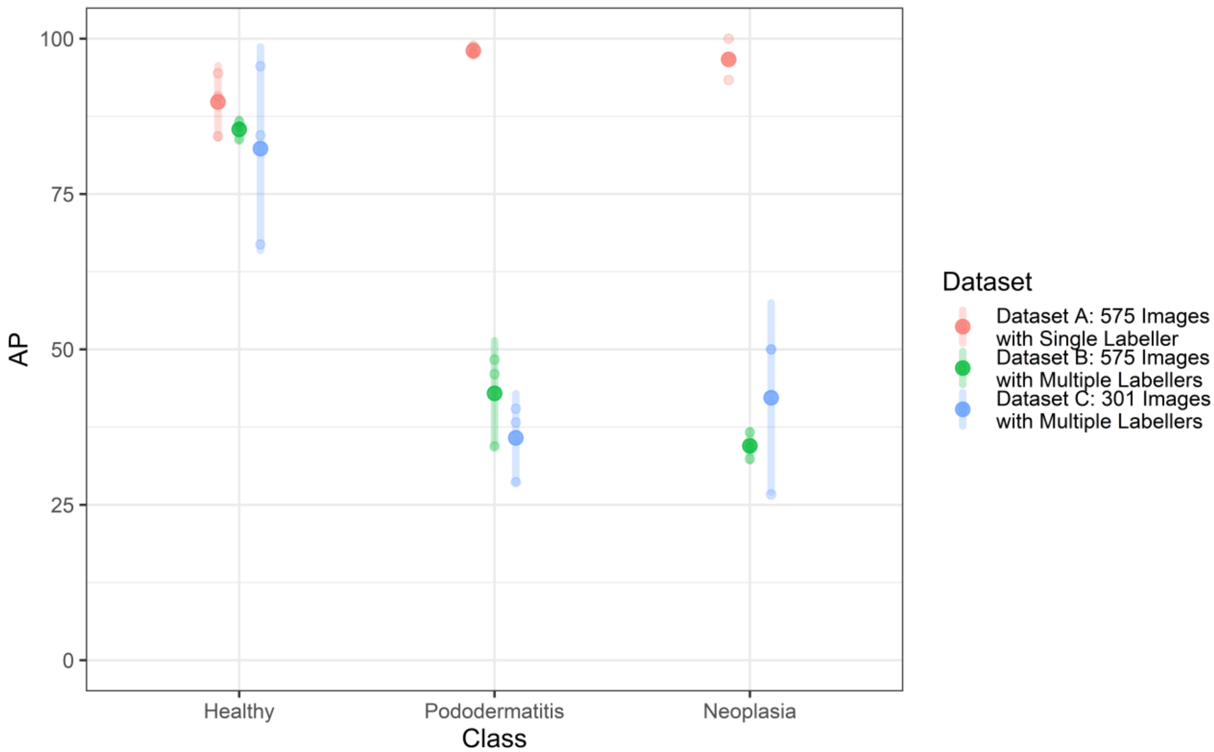

Figure 5.4 Model deployment of Tiny YOLOv4. Portable, battery-powered Jetson Xavier NX single-board computer connected to a Luxonis OAK-1 camera detecting three classes healthy, pododermatitis, and neoplasia in a clinical setting (Panel A). Canine pododermatitis bounding box with class label and class probability (Panel B). Schematic set-up for single-board computer, camera, LED monitor and power source (Panel C).

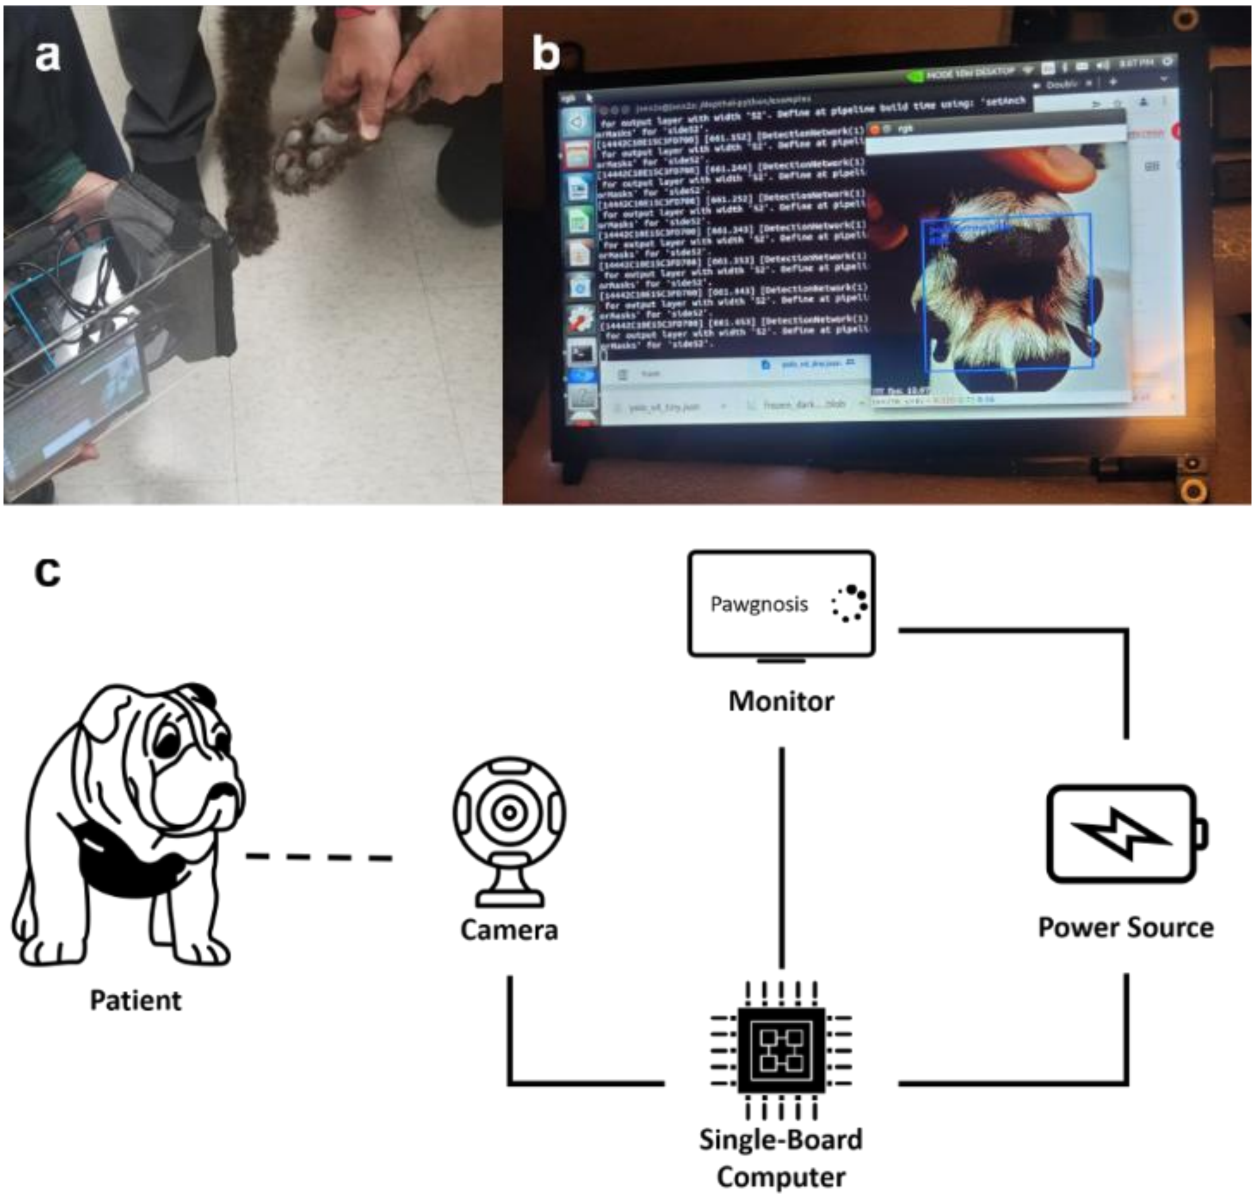

Figure 5.5 Image matrix of predictions. Image matrix of predictions from typical images of the healthy (left column), pododermatitis (middle column), and neoplasia (right column) classes using the best Tiny YOLOv4 model with the highest average mean average precision (mAP) value from the three training runs; Dataset A (top row), Dataset B (middle row), and Dataset C (bottom row). Matrix rows correspond to the training dataset and columns represent classes Healthy, Pododermatitis, and Neoplasia. The predictions are displayed using a bounding box and class label with corresponding prediction probability.

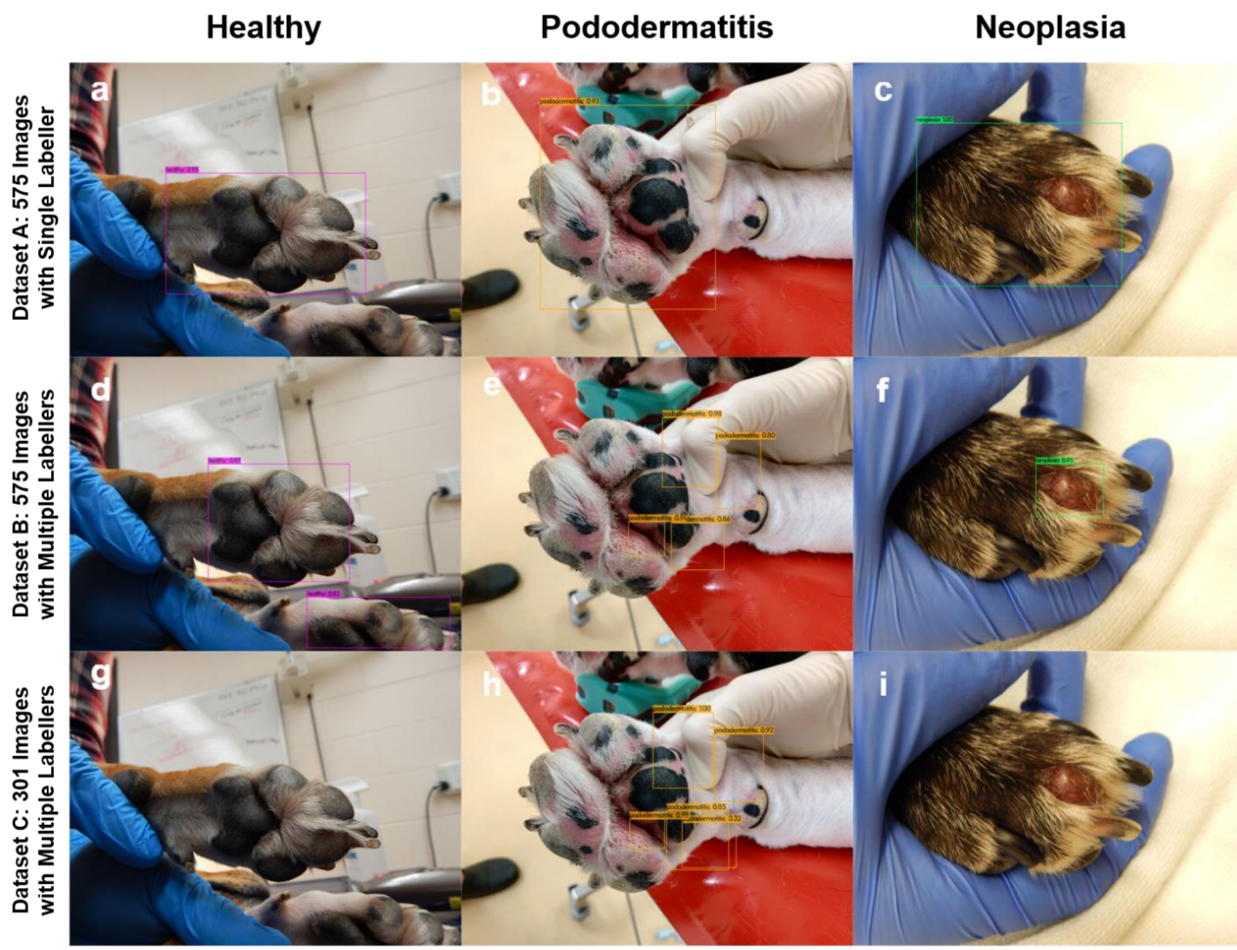

CHAPTER 6 DAIRYCOPILOT - AUTOMATED DATA COMPILATION AND ANALYSIS TOOLS FOR DAIRYCOMP DATA ASSETS

ABSTRACT

Modern dairy farm management requires meaningful data and careful analysis to maximize profitability, cow health, and welfare. Current data platforms, such as DairyComp, lack robust integrated data analysis tools. Producers and consultants need dedicated tools to turn collected data sets into assets for informed decision-making processes. The DairyCoPilot app allows users to rapidly extract health and production data from DairyComp, then compile and analyze the data using a menu-driven point-and-click approach. Prospects for training consultants in applied data analysis skills make DairyCoPilot a tool to identify farm management bottlenecks with less time spent for data analysis, improving cow health, and dairy production. The DairyCoPilot Dashboard R Shiny application is published using RStudio Connect: https://connect.doit.wisc.edu/dairy-copilot/.

INTRODUCTION

Managing dairy cows has become a data-intensive practice (Gargiulo et al., 2018; Morota et al., 2021). Increased number of cows per farm combined with enhanced data recording has expanded the availability of data for producers and consultants to make informed decisions ("USDA ERS - Milk Production Continues Shifting to Large-Scale Farms," n.d.). Deriving value from farm data and cow records requires software and trained users resulting in so-called 'data assets' aligned with evidence-based decision-making. Examples of dairy record management software are DairyComp, BoviSync, and animal by MILC. The current project will focus on data compilation and analysis from DairyComp (DC), a management package that is utilized by about 60% of dairy herds in the US ("5 Apps to Help Manage Your Dairy," 2022).

According to its promoters, DC has become one of the most popular dairy herd management software tools (VAS, n.d.). The software allows producers to store information about individual cows and cow-level events in a "cow card" to examine the cows' entire history, their health, reproduction, and production records. Cow management worksheets, monitoring reports, reproduction management, and other tools are integrated directly into DC. A set of built-in analysis tools for graphic displays of trends in health events, reproduction, and milk production assist with cow monitoring. However, DC does not currently possess internal tools for statistical inference controlled for confounders or provide a simple way to extract cow information for data analysis making it difficult for dairies to gain insights. Detailed health data analysis provides opportunities for adding value using knowledge about production and health challenges. Previously, detailed health analyses using DC data required significant manual data processing both in DC and in spreadsheet tool of user choice.

The proposed DairyCoPilot app is time efficient for single and repeated applications. We provide a singular, automated workflow for extracting health and event information from DC to DairyCoPilot followed by a graphical user interface application deployed on an R Shiny server. The initial setup requires minimal time in DC. The R Shiny application uploads the extracted file from DC to download a cleaned CSV file or perform data analysis.

MATERIALS AND METHODS

DairyComp Data Extraction

The data extraction process consists of two steps. First, a few so-called ‘items’ are created using the ALTER function in DC followed by a ‘protocol’ command to generate a CSV file. The protocol command can be stored in DC for repeated use. For users who are not familiar with the ALTER or SETUP options in DC, the full software reference guide is found in DC Reference Guide of https://dc-help.vas.com/ReferenceGuide/Home-DC305RefGuide.htm. Detailed instructions for extracting the raw data from DC are found in Appendix A or in the Documentation tab of https://connect.doit.wisc.edu/dairy-copilot/.

DairyCoPilot Data Cleaning

The CSV file from DC containing all variable names in the first row is read into the R Shiny application. The so-called ‘events’ from DC are automatically renamed e.g. 'SCOURS' to 'DIARRHEA' and missing events are filtered out. Missing columns required for further analysis in DairyCoPilot are created as placeholders. Remarks and protocols from DC are combined for completeness.

The data are grouped by cow identification (ID), lactation number (LACT), and event and the event number is appended to the event label. The dates and remarks are converted from a wide to a long format and the dates and remarks are appended to the event label respectively. The date and remarks are converted back from a long to a wide format with the variable names in the format of event, order, and if it is a date or remark. Missing events that are required for further analysis e.g. the first milk fever date for all ID and LACT are created as placeholders. The resulting variables are converted to the appropriate variable type, such as numeric or factor, for further analysis.

Cows who have not started their first lactation are filtered out and columns are created based on number of occurrences of an event for all ID and LACT. Columns are renamed as shown in Appendix B. The date and remark of a 'removed' event are created based on the sold and died events. Observations are excluded if the birth date, fresh date, or first calving date are invalid or missing. Predefined variables are created based on cutoff days post-calving or based on a value for a given event as described in Appendix B, for example the 'lactation group' is 3 if the 'lactation number' is greater than or equal to 3, otherwise, the lactation group is the lactation number. Additionally, event occurrence, the days in milk of the first occurrence in the lactation period, and predefined variable based on cutoff days post calving for a given event are created as described in Appendix B e.g. 'MLK_FVR<=7' is '1' if the days in milk of the first occurrence of milk fever in the lactation period is less than or equal to 7 DIM, otherwise 'MLK_FVR<=7' is '0'. Details regarding the cutoff value for the variables associated with transition cow health events are described in Appendix B and under the DairyCoPilot documentation tab. For binomial data, factor levels are reordered such that false or '0' is used as reference level and true or '1' as the alternative. For multinomial data, factor levels are reordered based on frequency from lowest to highest. For a given event, a rolling filter is applied such that if the next observation is recorded within a 3-day time lag of the previous observation, then only the previous observation is kept, else both observations are kept in the dataset. For example, if a cow has milk fever at 1 DIM, 2 DIM, and 4 DIM, then only the observations at 1 DIM and 4 DIM is kept. Detailed instructions for extracting the raw data from DC are found in Appendix C.

Software Requirements

All analyses are performed using R version 3.6.3 (R Core Team, 2020). The R packages *readr*, *dplyr*, and *tidyr* through *tidyverse* are used for data importation, data manipulation, and data tidying respectively (Wickham et al., 2022b, 2022a, 2019; Wickham and Girlich, 2022). The R package *stringr* is used for working with strings, *forcats* for handling categorical variables including reordering character vectors to improve display, and *lubridate* for working with dates and times (Grolemund and Wickham, 2011; Wickham, 2022a, 2022b). The R package *janitor* is used to generate two-way frequency tables for categorical analysis and *nnet* is used to fit multinomial log-linear models and to compute odds ratios (Firke, 2021; Venables and Ripley, 2002). The R package *DT* provides an R interface to the JavaScript library *DataTables* used to create an interactive data table and *skimr* is used to provide summary statistics about variables in a data table ("DataTables | Table plug-in for jQuery," n.d.; Waring et al., 2022; Xie et al., 2022). The R package *reactable* is used to create an interactive data table and *reactablefmtr* is used to streamline and enhance the styling and formatting of tables (Cuilla, 2022; Lin, 2022). The R packages *ggplot2*, *GGally*, and *plotly* are used to create data visualizations, pairs plots, and interactive graphing respectively (Schloerke et al., 2021; Sievert, 2020; Wickham, 2016). The R packages *shiny* and *shinydashboard* are used to build interactive web apps and dashboards (Chang and Ribeiro, 2021; Sievert, 2020). The DairyCoPilot Dashboard R Shiny application is published using RStudio Connect: https://connect.doit.wisc.edu/dairy-copilot/ (Aravamuthan et al., 2022). Connect assigns a distinct temporary directory to every process it initiates. Interactive applications such as Shiny are granted write access to the directory where the unprocessed code is stored. This directory serves as the working directory when launching an application, and any data written to it is accessible exclusively to processes linked with that particular application and is not visible to processes linked to other content. The data in the application directory remains accessible until the application is redeployed in Connect, which generates a new application directory exclusively containing the newly deployed content.

RESULTS

The following sections describe the four tabs of the DairyCoPilot Dashboard.

Input Tab

The Dashboard tab of the DairyCoPilot R Shiny app contains three panels. The Input panel is used to upload the downloaded CSV file from DC (Figure 1). Moreover, users are able to input the farm name, DairyComp extraction date, and earliest fresh date for analysis, most recent fresh date for analysis. Users can upload a 'raw' or previously unprocessed CSV, download a 'cleaned' CSV, make edits if there are errors in the DC records, and save these changes. Users can upload a previously cleaned CSV using the download button in the app or edited CSV in a spreadsheet app e.g. Microsoft Excel if necessary. The Contents panel is used to display the first 100 rows of the dataset providing filtering, pagination, and sorting (Figure 2). The Summary panel is used to display summary statistics the user can skim to understand the data (Figure 2). Results are printed horizontally with a section for each variable type and a row for each variable.

Pivot Table Tab

The Pivot Table tab is used for categorical analysis of multinomial data (Figure 3). The tab contains three columns for a total of six panels. The first column is the Input panel that is used to select the row variable and column variable for two-way tabulation where the user can also specify the reference level of both variables for further use. The Advanced Input panel allows users to specify a confidence level different than the default 95% level. The second column includes the Pivot Table panel of a two-way frequency table with the row and column sums, where the depth of the color mapping corresponds to counts, such that blue is equivalent to low counts and red is equivalent to high counts. The second column also includes the Odds Ratio panel of the odds ratio (OR) and corresponding confidence interval (CI) for each alternative level with respect to the reference levels previously specified (Dohoo et al., 2009). The OR and CI are computed by fitting multinomial a log-linear model via neural network (Venables and Ripley, 2002). The third column is used to visualize the data where the Comparative Bar Plot panel is used to display the count data in the pivot table and the Forest Plot panel is used to present the point estimate and interval data in the OR table. All plots are interactive with details on demand and can be downloaded as PNG files.

Regression Analysis Tab

The Regression Analysis tab is used to fit a regression model for an outcome variable of interest (Figure 4). Similar to the previous tab, the tab contains three columns for a total of five panels. The first column is the Input panel that is used to select the response and explanatory variables. The R Shiny app automatically recognizes if the response variable is a numeric or categorical variable as well as if the model is linear or logistic regression respectively (Chambers and Hastie, 1997; Dobson and Barnett, 2018; McCullagh and Nelder, 1998; Venables and Ripley, 2002). The second column includes the Summary panel that is the coefficients table for the regression model with estimate for slope, standard error, t-statistic, and p-value for each term in the regression model. The table can be used to determine the adjusted effect of each explanatory variable on the response variable and the statistical significance of the association. The second column also includes the ANOVA panel, that is the Analysis of Variance table of the regression model described above with degrees of freedom, sum of squared errors, mean squared errors, F-statistic, and p-value for each term in the regression model (Dohoo et al., 2009). The ANOVA table can be used to determine the variation of the response variable due to variation in the explanatory variable and the statistical significance of the association. The third column is used to visualize the data where the Generalized Pairs Plot panel automatically displays a comparative plot for each pair of variables in the regression model, depending on the type of data. The Coefficient Plot is used to visualize the estimate and 95% CI using the coefficient table. All plots are interactive with details on demand and can be downloaded as PNG files.

Documentation Tab

The Documentation tab provides instructions to generate the input CSV file in DC that is used for the data cleaning step by DairyCoPilot (Figure 5). A glossary of variables is provided for variables used to generate the output CSV file for the analysis in DairyCoPilot. Finally, an example dataset is provided for download for interested users to try out and test the app.

Software Usage and Example

A case study is performed to demonstrate the potential uses of the DairyCoPilot R Shiny app. A 1000-cow dairy in Wisconsin, USA provided access to DC records collected in 2022. Following data upload, summary statistics and cow records are displayed (Figure 2). Transformed data can be downloaded directly as a CSV file and additional analysis is performed using the Pivot Table and Regression Analysis tabs.

In the Pivot Table tab, the input selected for the Row variable is LCTGP with a reference level of 1 and the column variable is REMVD with a reference level of 0. The Pivot Table panel cross-tabulates categorical variables by level (Figure 3). In the LCTGP by REMVD example, the three lactation groups are split into 1st, 2nd, and 3rd+ lactations as rows in the panel and removed status (1) or not removed (0) as columns. The Odds Ratio panel provides a measure of association between two categorical risk factors. The odds of removal as a second lactation cow compared to a first lactation cow is 1.24×10^{0} (95% CI: 8.59×10^{-1} - 1.78×10^{0}) indicating no significant difference in removal between first and second lactation at a 95% confidence level. The 3+ cows have a significantly greater odds of removal than first lactation animals OR: 2.78×10^{0} (95% CI: 1.70×10^{0} - 3.05×10^{0}). The pivot table is visualized in a Comparative Barplot panel for total counts of removed and non-removed by lactation and OR is represented in the Forest Plot panel. Both of these graphs provide summary statistics by hovering over the image. Additionally, the graphics can be downloaded as an image file.

The Regression Analysis tab is used to quantify the associations between first-test butterfat percent (FSTBF), retained placenta (RP), lactation (LACT), and twins (TWIN). The summary panel provides estimates for regression coefficients and statistical significance (Figure 4). FSTBF is not significantly associated with lactation (estimate: 0.00; p-value: 0.96), negatively associated with twins (estimate: -0.21; p-value: 0.02), and FSTBF is negatively associated with the occurrence of retained placenta (estimate: -0.25; p-value: 0.02). Similar interpretations are generated using the ANOVA table. The Generalized Pairs Plot panel presents a representation of associations between coefficients. The Coefficient Plot panel provides a visualization of the regression parameter estimates. Both of these graphs provide summary statistics by hovering over the image. Additionally, the graphics can be downloaded as an image file.

Any choices of analysis for association of risk factors and outcome variables can be quantified using the DairyCoPilot tool, except for associations where limited numbers of observations per category result in non-convergence of the regression models. For such cases, the Pivot Table tab is preferred.

Example Analysis

An example analysis is completed to demonstrate a basic analysis of a dairy herd and can be replicated by downloading the example.csv found in the documentation tab. The CSV file was extracted from DairyComp on 10/13/22, and earliest fresh date for analysis was selected 10/12/20 to capture 2 years of date and last date was selected as 10/12/23. The data included 2106 rows. The file was then downloaded as a CSV file for storage and availability outside of DairyCoPilot. As an example, the relationship between twins, health, and milk production were examined using the Pivot Table tab. The association between twins (TWINS) and lactation group (LCTGP) were quantified and visualized. The reference level for LCTGP was selected as lactation 1 and for TWINS as 0. The OR for TWINS between LCTGP 1 and 2 was 1.73×10^{0} (95% CI: 8.17×10^{-1} - 3.673×10^{0}) and LCTGP 1 and 3 was 4.62×10^{0} (95% CI: 2.53×10^{0} - 8.44×10^{0}). The OR between removal from the herd before 60 DIM and twins was 3.415 (95% CI: 1.87×10^{0} - 6.24×10^{0}). The numerical variables and multiple variable model was inspected the Regression Analysis tab. First test butterfat was selected as an outcome variable for this example. The effect of Twins as an explanatory value for FSTBF was added as an explanatory variable. The estimated effect of twins on FSTBF was -0.26 with a p-value less than 0.01. In the categorical analysis the association of LCTGP and TWINS was noted. Therefore, an additional explanatory variable was added to the regression equation for FSTBF = TWINS + LCTGP. The resulting estimate for FSTBF changed to -0.28 with a p-value of 0.0. Finally, ketosis (KET) was added to the regression equation resulting in FSTBF = TWINS + LCTGP + KET. The p-value for KET was 0.48 and the user may elect to remove it from the regression model. Multiple other factors can be added to the regression model to improve user understanding and prediction performance. The resulting evaluation shows that twins are associated with lower FSTBF even when controlling for LACT at this dairy.

DISCUSSION

Dairy Data Collection

A significant proportion of dairy data is manually acquired and recorded. Events such as disease diagnosis and subsequent recording are subject to human error. Errors in data recording will propagate forward into subsequent analysis. The DairyCoPilot tool provides a method for reducing time to analysis but cannot improve the accuracy of input data. However, it is user responsibility to correctly validate data before data analysis.

Data Analysis Tools

Agricultural production is a data-rich environment that requires advanced analysis for informed decision-making processes (Antle et al., 2017; Moore et al., 2022). Advanced training can be accomplished through outreach and extension services with fit-for-purpose tools that reduce barriers of application by end users. Training and analysis tools such as DairyCoPilot presented in this study, fill a need in the dairy industry that is currently unmet. Decision-making processes using graphic analysis or trend observation alone are not powerful enough for modern production systems (Correll et al., 2012; Whitley and Ball, 2002). More advanced multiple variable statistical tools allow users to control multiple confounding variables such as lactation or milk production simultaneously (Kahlert et al., 2017). Data analysis tools used for local on-farm decision support should be customizable to user demands. The proposed DairyCoPilot application allows for customization of analysis choices and guarantees data ownership resulting in a locally secure tool.

Platforms providing data analysis including visualization, dashboarding, and economic evaluation of cow health and production include milking services companies e.g. Lely T4C, DeLaval DelPro, etc.; herd management platforms e.g. DC, animal, Bovisync, Dairy Data Wearhouse, etc.; and cow monitoring technologies e.g. SmaXtec, Connectera, CowManager, etc. (Lely, n.d.; DeLaval, n.d.; VAS, n.d.; Bovisync LLC, n.d.; “SmaXtec,” n.d.; Connectera, n.d.; CowManager B.V., n.d.; DDW, n.d.). Tools for exporting graphs and reports are limited, and statistical data analysis can be absent from these applications. Licenses must be paid, or large purchases must be made from these companies to access the data analysis tools. The DairyCoPilot application has advantages for data assets used for on-farm decision-making processes, because the app is free, easy-to-use, secure, and customizable.

Big Data Approaches

The alternative to modular data analysis tools is a big data analysis approach. These analysis systems take in data from multiple diverse sources, compile them, and generate inferences. For example, Dairy Data Warehouse provides integrated data storage, analytics, and forecasting for enrolled dairies (DDW, n.d.). These services are billed on a per-cow subscription fee. Another example of a big data approach is the UW Madison Dairy Brain and Data Hub (“Dairy Brain,” n.d.).

These tools enhance data analysis on farms. Questions regarding data ownership and security are important considerations regarding these tools (Runck et al., 2022; Wilgenbusch et al., 2022).

Food production stakeholders, including commercial and governmental organizations, will need to reach agreements about data sharing, ownership, and management. These challenges indicate that small data approaches, such as the proposed DairyCoPilot application, are still a viable option that allows improved and customized control by individual end users. Small data applications can be integrated into big data applications in the future.

Strengths

First and foremost, DC does not provide advanced data analysis required for dairies to gain insights and make data-driven decisions for complex problems. The user interface is extremely basic and customization is insufficient where the software is only developed for Windows desktops. The management application has a limited set of features lacking the advanced functionality to manipulate data, create graphics, and generate figures and tables for estimation and prediction.

Conversely, DairyCoPilot provides statistical tools for both categorical data analysis and regression analysis. The program provides an elegant and powerful web framework where the user interface is dynamic to conditionally generate input controls and uses reactive programming to automatically update outputs when inputs change. The web application can clean data for downstream data analysis, create interactive graphics for exploratory and expository visualization, and export high quality figures using interactive views for documents, reports, and presentations. Lastly, DairyCoPilot is mobile-friendly, desktop progressive web application for all device platforms.

Limitations

Currently, DairyCoPilot only handles a subset of events and items as detailed in Appendix B. The program is limited to analysis of cow data and is not yet available for calf data but can be added in the future. The implementation of software requires appropriate statistical training which must be obtained separately from the supplied software tool. Consultants must still use a holistic approach to evaluating dairies in addition to statistical tools where visual appraisal of dairies can provide information not contained in records and can even contradict records. The main drawback of the data analysis tool is slow performance for extremely large data sets containing multiple recordings of events compared to other languages. Accordingly, the lack of scalability compared to more popular frameworks such as Node.js or React can be an issue for dairies that need to handle a lot of traffic or requests. Currently, the regression tab does not provide the option to add random effects in the model. Additionally, the application does not provide the functionality to check the assumptions of statistical methods e.g. diagnostic plots for linear models and generalized linear models. This will be available in forthcoming updates.

CONCLUSION

Tools for automated data analysis are necessary for training the next generation of food animal consultants and life scientists. Tools, such as DairyCoPilot, help shift the burden of data formatting and editing in terms of time and effort to highly trained individuals. The food animal consultant does not need to be a statistical expert if automated statistical analysis tools are accessible and if such analysis is part of their daily discussion and decisions.

The outlook for local data analysis tools such as the DairyCoPilot is the expansion of variable formatting into customized categorical variables, interaction terms and random effects models in the near future. DairyCoPilot is the first step in developing advanced freely available tools for rapid analysis of records from dairy farms that are translatable to all other animal species and healthcare settings. Additionally, adoption of these tools will facilitate commercial providers, such as VAS, to include advanced graphical and statistical tools directly into their software updates and to improve workflows for data extraction to support other tools. The DairyCoPilot application represents an important step in turning farm records into data assets for more customized decision-making processes by informed consultants in the life sciences.

ACKNOWLEDGEMENTS

Dr. Gary Oetzel is thanked for his previous work analyzing health data in DairyComp. The Food Animal Production Medicine group at UW - Madison shared feedback and collaborated with us to utilize the initial versions of this tool in academic settings. We would like to thank the multiple dairy farms that provided access to data and described challenges that helped align the tool to their needs.

REFERENCES

5 Apps to Help Manage Your Dairy [WWW Document], 2022. . Dairy Herd. URL https://www.dairyherd.com/news/business/5-apps-help-manage-your-dairy (accessed 12.8.22).

Antle, J.M., Basso, B., Conant, R.T., Godfray, H.C.J., Jones, J.W., Herrero, M., Howitt, R.E., Keating, B.A., Munoz-Carpena, R., Rosenzweig, C., Tittonell, P., Wheeler, T.R., 2017. Towards a new generation of agricultural system data, models and knowledge products: Design and improvement. Agric Syst 155, 255–268. https://doi.org/10.1016/j.agsy.2016.10.002

Aravamuthan, S., Walleser, E., Dopfer, D., 2022. DairyCoPilot. DairyCoPilot.

Bovisync LLC, n.d. Bovisync [WWW Document]. Bovisync. URL https://bovisync.com/ (accessed 12.8.22).

Chang, W., Ribeiro, B.B., 2021. shinydashboard: Create Dashboards with "Shiny."

Connectera, n.d. ida [WWW Document]. URL https://www.connecterra.io/ (accessed 12.8.22).

Correll, M., Albers, D., Franconeri, S., Gleicher, M., 2012. Comparing averages in time series data, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM, Austin Texas USA, pp. 1095–1104. https://doi.org/10.1145/2207676.2208556

CowManager B.V., n.d. CowManager [WWW Document]. CowManager. URL https://www.cowmanager.com/ (accessed 12.8.22).

Cuilla, K., 2022. reactablefmtr: Streamlined Table Styling and Formatting for Reactable.

Dairy Brain [WWW Document], n.d. . Dairy Brain. URL https://dairybrain.wisc.edu/ (accessed 12.9.22).

DataTables | Table plug-in for jQuery [WWW Document], n.d. URL https://datatables.net/ (accessed 12.9.22).

DDW, n.d. Dairy Data Warehouse [WWW Document]. Dairy Data Warehouse. URL https://www.dairydatawarehouse.com/ (accessed 12.8.22).

DeLaval, n.d. DelPro [WWW Document]. URL https://www.delaval.com/en-us/explore-our-farm-solutions/delaval-delpro/ (accessed 12.8.22).

Firke, S., 2021. janitor: Simple Tools for Examining and Cleaning Dirty Data.

Gargiulo, J.I., Eastwood, C.R., Garcia, S.C., Lyons, N.A., 2018. Dairy farmers with larger herd sizes adopt more precision dairy technologies. Journal of Dairy Science 101, 5466–5473. https://doi.org/10.3168/jds.2017-13324

Grolemund, G., Wickham, H., 2011. Dates and Times Made Easy with lubridate. Journal of Statistical Software 40, 1–25.

Kahlert, J., Gribsholt, S.B., Gammelager, H., Dekkers, O.M., Luta, G., 2017. Control of confounding in the analysis phase – an overview for clinicians. Clin Epidemiol 9, 195–204. https://doi.org/10.2147/CLEP.S129886

Lely, n.d. T4C management System [WWW Document]. URL https://www.lely.com/farming-insights/t4c-management-system/ (accessed 12.8.22).

Lin, G., 2022. reactable: Interactive Data Tables Based on "React Table."

Moore, E.K., Kriesberg, A., Schroeder, S., Geil, K., Haugen, I., Barford, C., Johns, E.M., Arthur, D., Sheffield, M., Ritchie, S.M., Jackson, C., Parr, C., 2022. Agricultural data management and sharing: Best practices and case study. Agronomy Journal 114, 2624–2634. https://doi.org/10.1002/agj2.20639

Morota, G., Cheng, H., Cook, D., Tanaka, E., 2021. ASAS-NANP SYMPOSIUM: prospects for interactive and dynamic graphics in the era of data-rich animal science1. Journal of Animal Science 99, skaa402. https://doi.org/10.1093/jas/skaa402

R Core Team, 2020. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria.

Runck, B.C., Joglekar, A., Silverstein, K.A.T., Chan-Kang, C., Pardey, P.G., Wilgenbusch, J.C., 2022. Digital agriculture platforms: Driving data-enabled agricultural innovation in a world fraught with privacy and security concerns. Agronomy Journal 114, 2635–2643. https://doi.org/10.1002/agj2.20873

Schloerke, B., Cook, D., Larmarange, J., Briatte, F., Marbach, M., Thoen, E., Elberg, A., Crowley, J., 2021. GGally: Extension to "ggplot2."

Sievert, C., 2020. Interactive Web-Based Data Visualization with R, plotly, and shiny. Chapman and Hall/CRC.

SmaXtec [WWW Document], n.d. URL https://smaxtec.com/en/ (accessed 12.8.22).

USDA ERS - Milk Production Continues Shifting to Large-Scale Farms [WWW Document], n.d. URL https://web.archive.org/web/20150323151533/http:/www.ers.usda.gov/amber-waves/2014-december/milk-production-continues-shifting-to-large-scale-farms.aspx#.VRHBM2MR-JF (accessed 10.3.22).

VAS, n.d. DairyComp [WWW Document]. URL https://vas.com/dairycomp/

Venables, W.N., Ripley, B.D., 2002. Modern Applied Statistics with S, Fourth. ed. Springer, New York.

Waring, E., Quinn, M., McNamara, A., Rubia, E.A. de la, Zhu, H., Ellis, S., 2022. skimr: Compact and Flexible Summaries of Data.

Whitley, E., Ball, J., 2002. Statistics review 5: Comparison of means. Crit Care 6, 424–428.

Wickham, H., 2022a. stringr: Simple, Consistent Wrappers for Common String Operations.

Wickham, H., 2022b. forcats: Tools for Working with Categorical Variables (Factors).

Wickham, H., 2016. ggplot2: Elegant Graphics for Data Analysis. Springer-Verlag New York.

Wickham, H., Averick, M., Bryan, J., Chang, W., McGowan, L., François, R., Grolemund, G., Hayes, A., Henry, L., Hester, J., Kuhn, M., Pedersen, T., Miller, E., Bache, S., Müller, K., Ooms, J., Robinson, D., Seidel, D., Spinu, V., Takahashi, K., Vaughan, D., Wilke, C., Woo, K., Yutani, H., 2019. Welcome to the Tidyverse. JOSS 4, 1686. https://doi.org/10.21105/joss.01686

Wickham, H., François, R., Henry, L., Müller, K., 2022a. dplyr: A Grammar of Data Manipulation.

Wickham, H., Girlich, M., 2022. tidyr: Tidy Messy Data.

Wickham, H., Hester, J., Bryan, J., 2022b. readr: Read Rectangular Text Data.

Wilgenbusch, J.C., Pardey, P.G., Bergstrom, A., 2022. Big data promises and obstacles: Agricultural data ownership and privacy. Agronomy Journal 114, 2619–2623. https://doi.org/10.1002/agj2.21182

Xie, Y., Cheng, J., Tan, X., 2022. DT: A Wrapper of the JavaScript Library “DataTables.”

TABLES AND FIGURES

Figure 6.1 Dashboard Tab of the DairyCoPilot application before data upload.

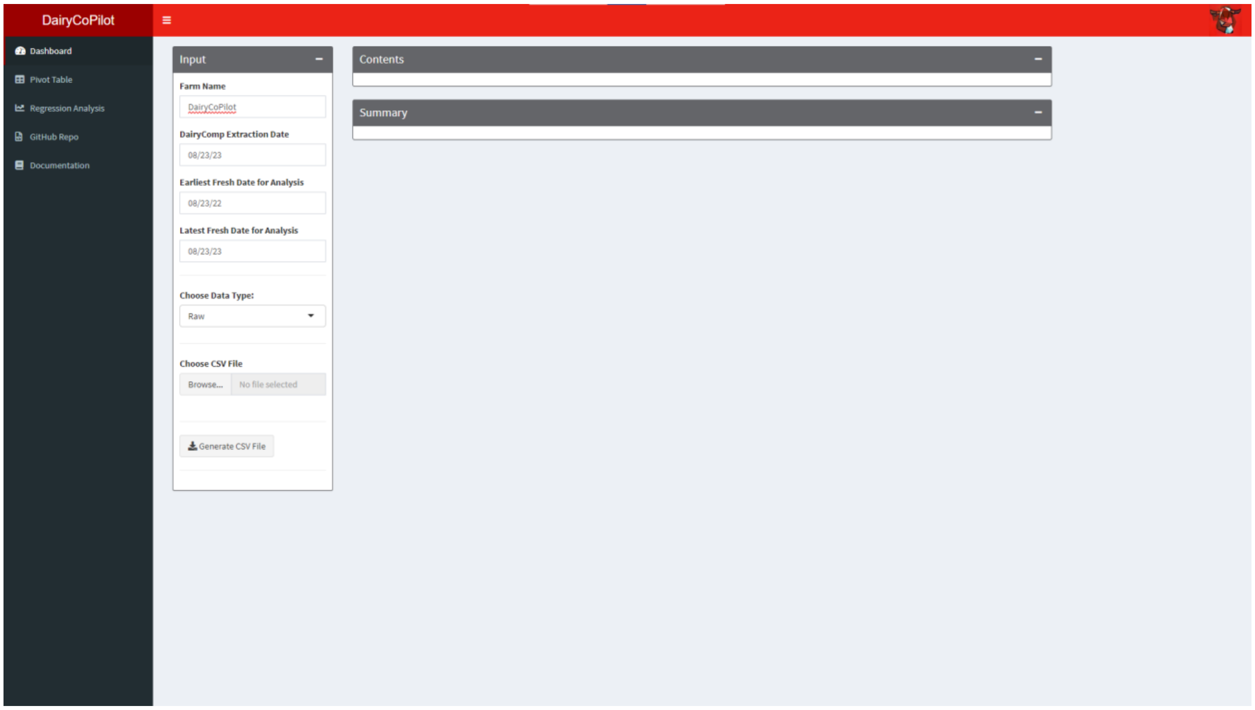

Figure 6.2 Dashboard Tab of the DairyCoPilot application after data upload. The user can input the farm name, the date of data extraction from DairyComp, and starting and ending fresh dates for inclusion in the data analysis. The Contents panel displays the transformed dataset and is searchable. The Summary panel includes information about all variables including descriptive statistics and missing variables.

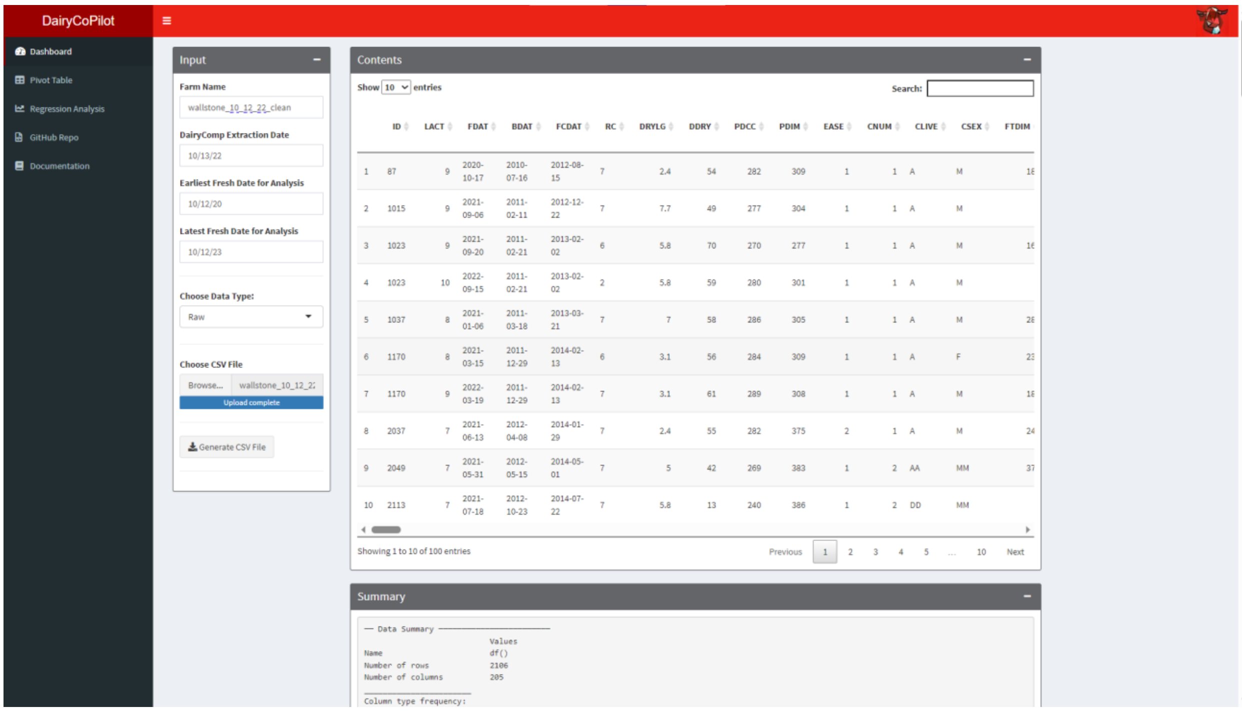

Figure 6.3 Pivot Table Tab. The Pivot Table of the DairyCoPilot application allows the user to perform statistical analysis for categorical variables and visualize two-way associations. Odd ratios and 95% confidence intervals allow the user to quantify pairwise associations and assess the statistical significance.

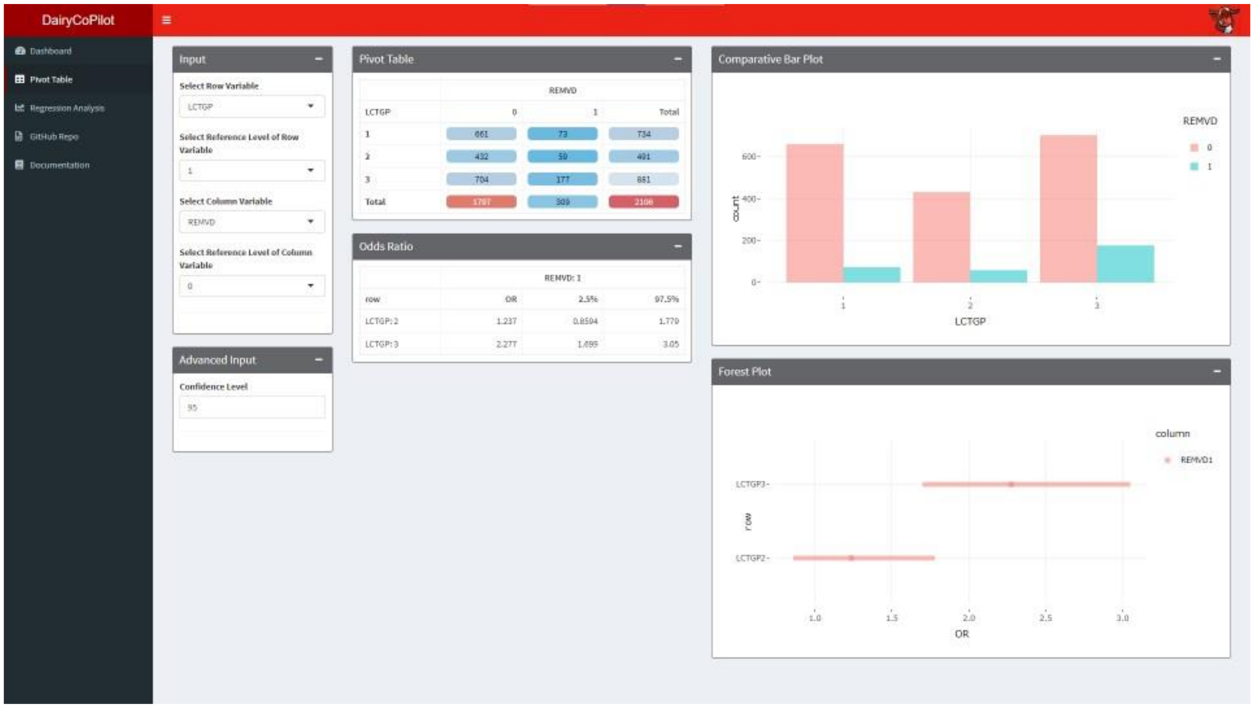

Figure 6.4 Regression Analysis Tab. Regression Analysis in the DairyCoPilot application allows the user to perform linear and logistic regression. The user can select a response variable and multiple explanatory variables followed by different variables to control for confounding across groups. Graphical analysis in the Generalized Pairs Plot and Coefficient plot panels visualize the relationships between variables in the Summary and ANOVA panels.

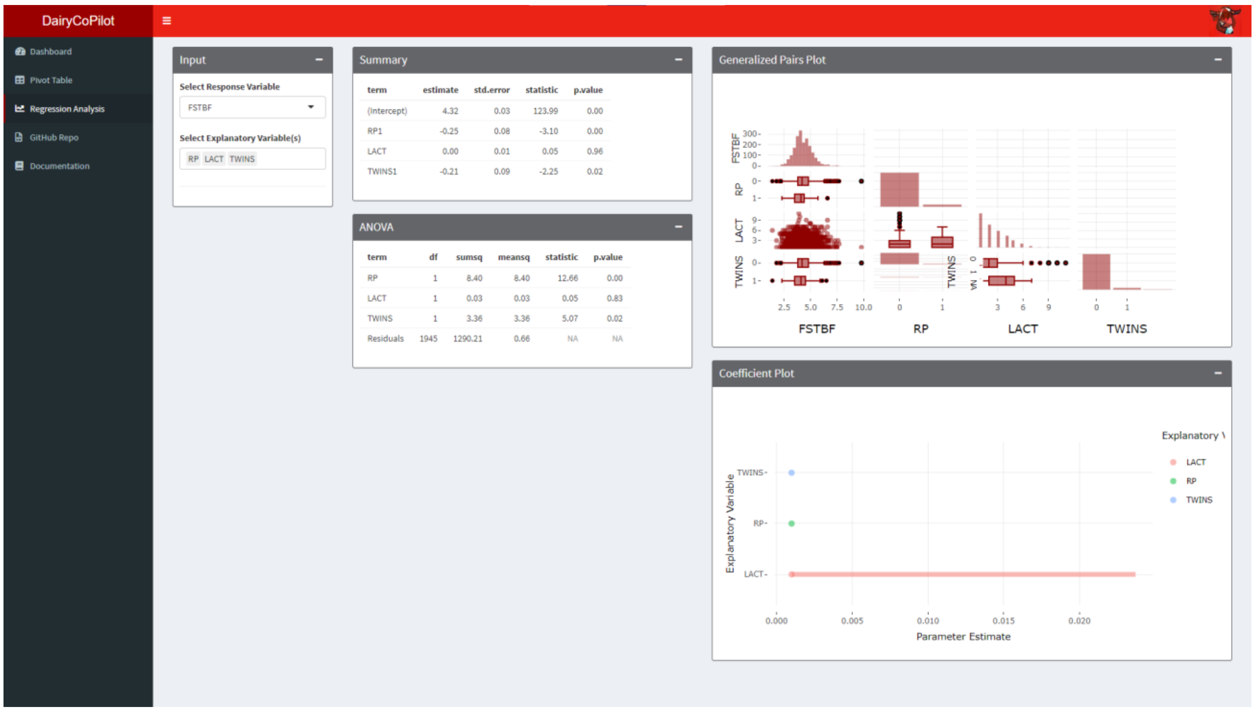

Figure 6.5 Documentation Tab. Documentation in the DairyCoPilot application provides instructions to generate a CSV file for analysis using DairyCoPilot.

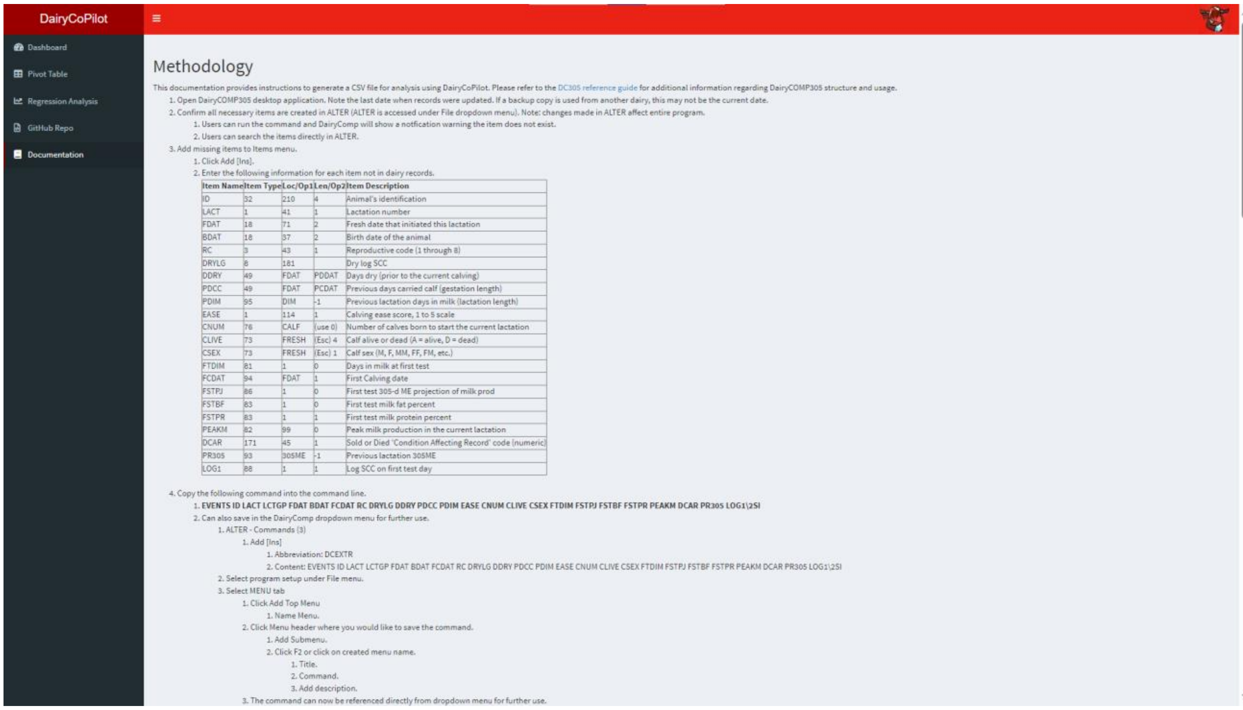

DairyCoPilot

Dashboard

Pivot Table

Regression Analysis

GitHub Repo

Documentation

Methodology

This documentation provides instructions to generate a CSV file for analysis using DairyCoPilot. Please refer to the DC305 reference guide for additional information regarding DairyCOMP305 structure and usage.

1. Open DairyCOMP305 desktop application. Note the last date when records were updated. If a backup copy is used from another dairy, this may not be the current date.
2. Confirm all necessary items are created in ALTER (ALTER is accessed under File dropdown menu). Note: changes made in ALTER affect entire program.
    1. Users can run the command and DairyComp will show a notfication warning the item does not exist.
    2. Users can search the items directly in ALTER.
3. Add missing items to Items menu.
    1. Click Add [Ins].
    2. Enter the following information for each item not in dairy records.

Item Name	Item Type	Loc/Op1	Len/Op2	Item Description
ID	32	210	4	Animal's identification
LACT	1	41	1	Lactation number
FDAT	18	71	2	Fresh date that initiated this lactation
BDAT	18	37	2	Birth date of the animal
RC	3	43	1	Reproductive code (1 through 8)
DRYLG	8	181		Dry log SCC
DDRY	49	FDAT	PDDAT	Days dry (prior to the current calving)
PDCC	49	FDAT	PCDAT	Previous days carried calf (gestation length)
PDIM	95	DIM	-1	Previous lactation days in milk (lactation length)
EASE	1	114	1	Calving ease score, 1 to 5 scale
CNUM	76	CALF	(use 0)	Number of calves born to start the current lactation
CLIVE	73	FRESH	(Esc) 4	Calf alive or dead (A = alive, D = dead)
CSEX	73	FRESH	(Esc) 1	Calf sex (M, F, MM, FF, FM, etc.)
FTDIM	81	1	0	Days in milk at first test
FCDAT	94	FDAT	1	First Calving date
FSTPJ	86	1	0	First test 305-d ME projection of milk prod
FSTBF	83	1	0	First test milk fat percent
FSTPR	83	1	1	First test milk protein percent
PEAKM	82	99	0	Peak milk production in the current lactation
DCAR	171	45	1	Sold or Died 'Condition Affecting Record' code (numeric)
PR305	93	305ME	-1	Previous lactation 305ME
LOG1	88	1	1	Log SCC on first test day

4. Copy the following command into the command line.
    1. **EVENTS ID LACT LCTGP FDAT BDAT FCDAT RC DRYLG DDRY PDCC PDIM EASE CNUM CLIVE CSEX FTDIM FSTPJ FSTBF FSTPR PEAKM DCAR PR305 LOG1\2SI**
    2. Can also save in the DairyComp dropdown menu for further use.
        1. ALTER - Commands (3)
            1. Add [Ins]
                1. Abbreviation: DCEXTR
                2. Content: EVENTS ID LACT LCTGP FDAT BDAT FCDAT RC DRYLG DDRY PDCC PDIM EASE CNUM CLIVE CSEX FTDIM FSTPJ FSTBF FSTPR PEAKM DCAR PR305 LOG1\2SI
        2. Select program setup under File menu.
        3. Select MENU tab
            1. Click Add Top Menu
                1. Name Menu.
            2. Click Menu header where you would like to save the command.
                1. Add Submenu.
                2. Click F2 or click on created menu name.
                    1. Title.
                    2. Command.
                    3. Add description.
            3. The command can now be referenced directly from dropdown menu for further use.

APPENDIX A

Instructions:

1. Open DairyComp desktop software. Make a note of the last date records were updated. If using a backup copy from another dairy this may not be the current date.
2. Ensure all necessary items are created in ALTER (ALTER is accessed under File dropdown menu):

   CAUTION: Changes made in ALTER affect the entire program, be cautious

   1. Users can run the following command and DairyComp will warn their item does not exist
   2. Items can be searched directly in ALTER
   3. Add missing items to Items menu
      1. Click Add [Ins]
      2. Enter information in the chart below for each item not found in dairy records

Item Name	Item Type	Loc/Op1	Len/Op2	Item Description
ID	32	210	4	Animal's identification
LACT	1	41	1	Lactation number
FDAT	18	71	2	Fresh date that initiated this lactation
BDAT	18	37	2	Birth date of the animal
RC	3	43	1	Reproductive code (1 through 8)
DRYLG	8	181		Dry log SCC
DDRY	49	FDAT	PDDAT	Days dry (prior to the current calving)
PDCC	49	FDAT	PCDAT	Previous days carried calf (gestation length)
PDIM	95	DIM	-1	Previous lactation days in milk (lactation length)
EASE	1	114	1	Calving ease score, 1 to 5 scale
CNUM	76	CALF	(use 0)	Number of calves born to start the current lactation
CLIVE	73	FRESH	(Esc) 4	Calf alive or dead (A = alive, D = dead)

Item Name	Item Type	Loc/Op1	Len/Op2	Item Description
CSEX	73	FRESH	(Esc) 1	Calf sex (M, F, MM, FF, FM, etc.)
FTDIM	81	1	0	Days in milk at first test
FCDAT	94	FDAT	1	First Calving date
FSTPJ	86	1	1	First test 305-d ME projection of milk prod
FSTBF	83	1	0	First test milk fat percent
FSTPR	83	1	1	First test milk protein percent
PEAKM	82	99	0	Peak milk production in the current lactation
DCAR	171	45	1	Sold or Died "Condition Affecting Record" code (numeric)
PR305	93	305ME	-1	Previous lactation 305ME
LOG1	88	1	1	Log SCC on first test day

Table 1:

Items to be created in the DairyComp file for the dairy to be analyzed.

4. Copy the following command into the command line:
    1. **EVENTS ID LACT LCTGP FDAT BDAT FCDAT RC DRYLG DDRY PDCC PDIM EASE CNUM CLIVE CSEX FTDIM FSTPJ FSTBF FSTPR PEAKM DCAR PR305 LOG1\2SI**
    2. This command can also be saved in the DairyComp dropdown menu for future use:
        1. ALTER – Commands (3)
            1. Add [Ins]
                1. Abbreviation: DCEXTR
                2. Content: EVENTS ID LACT LCTGP FDAT BDAT FCDAT RC DRYLG DDRY PDCC PDIM EASE CNUM CLIVE CSEX FTDIM FSTPJ FSTBF FSTPR PEAKM DCAR PR305 LOG1\2SI
        2. Under File select program setup
        3. Select MENU tab

            1. Click Add Top Menu
                1. Name menu
            2. Click Menu header you would like command to be saved under
                1. Add Submenu
                2. Click F2 or click in created menu name
                    1. Title
                    2. Command (DCEXTR)
                    3. Add description
            3. The Command can now be referenced directly from dropdown menu in the future
    5. Select a date 2-2.5 years prior to DC date
    6. Select Last date you wish to analyze data
    7. Select Events
        1. Select All will contain the most information and is preferred
        2. Can individually choose Events as well
            1. MUST include FRESH
        3. Select optional REM pattern:
            1. Select default "none" by clicking "OK"
        4. This may take a few minutes
    8. Write data to CSV file
        1. Find floppy disk icon above command line
        2. Click downward facing triangle to the right
        3. Save as CSV (Do Not save as XLS)
        4. Choose location and name for CSV
4. The downloaded CSV is now ready to be analyzed using DairyCoPilot app

APPENDIX B

Item Name	**Item Description**
ID	Animal's identification
LACT	Lactation number
LCTGP	Lactation groups (1: lactation number is 1; lactation number is 2; 3: lactation number is greater than or equal to 3)
FDAT	Fresh date that initiated this lactation
BDAT	Birth date of the animal
TOCU1	Date of last move to close-up pen
TOCU2	Date of previous move to close-up pen
RC	Reproductive code (1 through 8)
DRYLG	Dry log SCC
DDRY	Days dry (prior to the current calving)
PDCC	Previous days carried calf (gestation length)
PDIM	Previous lactation days in milk (lactation length)
EASE	Calving ease score, 1 to 5 scale
CNUM	Number of calves born to start the current lactation
CLIVE	Calf alive or dead (A: alive; D: dead)

CSEX	Calf sex (M, F, MM, FF, FM, etc.)
FTDIM	Days in milk at first test
FSTPJ	First test 305-d ME projection of milk prod
FSTBF	First test milk fat percent
FSTPR	First test milk protein percent
PEAKM	Peak milk production in the current lactation
DCAR	Sold or Died "Condition Affecting Record" code (numeric)
DATS	Date sold
REMS	Remark of the sold event
DATD	Date died
REMD	Remark of the died event
1DMF	Date of 1st milk fever in the current lactation
REMMF	Remark of 1st milk fever in the current lactation
XMF	Count of all milk fever events in the current lactation
1DRP	Date of 1st retained placenta in the current lactation
REMRP	Remark of 1st retained placenta in the current lactation

XRP	Count of all RP events in the current lactation
1DME	Date of 1st metritis event in the current lactation
REMME	Remark of 1st metritis event in the current lactation
XMETR	Count of all metritis events in the current lactation
1DKE	Date of 1st ketosis event in the current lactation
REMKE	Remark of 1st ketosis event in the current lactation
XKET	Count of all ketosis events in the current lactation
1DDA	Date of 1st DA event in the current lactation
REMDA	Remark of 1st DA event in the current lactation
XDA	Count of all DA events in the current lactation
1DPN	Date of 1st pneumonia event in the current lactation
REMPN	Remark of pneumonia event in the current lactation
XPNEU	Count of all pneumonia events in the current lactation

1DDIA	Date of 1st diarrhea event in the current lactation
REMDI	Remark of 1st diarrhea event in the current lactation
XDIAR	Count of all diarrhea events in the current lactation
1DLAM	Date of 1st lame event in the current lactation
REMLM	Remark of 1st lame event in the current lactation
XLAME	Count of all lame events in the current lactation
1DFT	Date of 1st foot trim event in the current lactation
REMFT	Remark of 1st foot trim event in the current lactation
XFTRM	Count of all foot trim events in the current lactation
1DMA	Date of 1st mastitis event in the current lactation
REMMA	Remark of 1st mastitis event in the current lactation
XMAST	Count of all mastitis events in the current lactation
1DIL	Date of 1st ILLMISC event in the current lactation

REMIL	Remark of 1st ILLMISC event in the current lactation
XILL	Count of all ILLMISC events in the current lactation
DATS	Date removed
REMS	Remark of the removed event
FRESH_MONTH	Fresh date that initiated this lactation (YYYY-MM)
DIM	Days in milk (difference between fresh date and data extraction date or date of removal if removed)
FRESH 372 TO 7	Days in milk between 7 and 372
FRESH 386 TO 21	Days in milk between 21 and 386
FRESH 425 TO 60	Days in milk between 60 and 425
BIRTH_DATE	Birth date of the animal
AGE_AT_CALVING_(MONTHS)	Age at calving (from birth date to fresh date; months)
AGE_AT_CALVING_1ST_LACT_(MONTHS)	Age at calving (from birth date to first calving date; months)
ABORT	Abort (1: gestation length less than 260; 0: otherwise)
CALVING_EASE_>=2	Calving ease score greater than or equal to 2
CALVING_EASE_>=3	Calving ease score greater than or equal to 3

TWINS	Twins (1: number of calves born to start the current lactation greater than 1; 0: number of calves born to start the current lactation equal to 1)
STILLBIRTH	Stillborn (1: dead; 0: otherwise)
CALF_SEX	Calf sex
MALE_CALF	Male calf
FPR	Fat protein ratio
FPR>1.4	Fat protein ratio greater than 1.4 (1: fat protein ratio greater than or equal to 1.4; 0: otherwise) (Duffield et al., 1997)
SOLD	Sold
SOLD_DIM	Days in milk at sale
SOLD<=60	Days in milk at sale less than or equal to 60 (1: sale less than or equal to 60; 0: otherwise)
DIED	Died
DIED_DIM	Days in milk at death
DIED<=60	Days in milk at death less than or equal to 60 (1: death less than or equal to 60; 0: otherwise)
REMVD	Removed
REMVD_DIM	Days in milk at removal
REMVD<=60	Days in milk at removal less than or equal to 60 (1: removal less than or equal to 60; 0: otherwise)

MLK_FVR	Milk Fever
MLK_FVR_DIM	Days in milk at 1st milk fever event
MLK_FVR<=7	Days in milk at 1st milk fever event less than or equal to 7 (1: days in milk less than or equal to 7; 0: otherwise)
RP	Retained placenta
RP_DIM	Days in milk at 1st retained placenta event
RP<=7	Days in milk at 1st retained placenta event less than or equal to 7 (1: days in milk less than or equal to 7; 0: otherwise)
METR	Metritis
METR_DIM	Days in milk at 1st metritis event
METR<=21	Days in milk at 1st metritis event less than or equal to 21 (1: days in milk less than or equal to 21; 0: otherwise)
KET	Ketosis
KET_DIM	Days in milk at 1st ketosis event
KET<=60	Days in milk at 1st ketosis event less than or equal to 60 (1: days in milk less than or equal to 60; 0: otherwise)
DA	DA
DA_DIM	Days in milk at 1st DA event

DA<=60	Days in milk at 1st DA event less than or equal to 60 (1: days in milk less than or equal to 60; 0: otherwise)
PNEU	Pneumonia
PNEU_DIM	Days in milk at 1st pneumonia event
PNEU<=60	Days in milk at 1st pneumonia event less than or equal to 60 (1: days in milk less than or equal to 60; 0: otherwise)
DIARR	Diarrhea
DIARR_DIM	Days in milk at 1st diarrhea event
DIARR<=60	Days in milk at 1st diarrhea event less than or equal to 60 (1: days in milk less than or equal to 60; 0: otherwise)
LAME	Lameness
LAME_DIM	Days in milk at 1st lameness event
LAME<=60	Days in milk at 1st lameness event less than or equal to 60 (1: days in milk less than or equal to 60; 0: otherwise)
FTRM	Foot trim
FTRM_DIM	Days in milk at 1st foot trim event
FTRM<=60	Days in milk at 1st foot trim event less than or equal to 60 (1: days in milk less than or equal to 60; 0: otherwise)
MAST	Mastitis
MAST_DIM	Days in milk at 1st mastitis event

MAST<=60	Days in milk at 1st mastitis event less than or equal to 60 (1: days in milk less than or equal to 60; 0: otherwise)
ILLMISC	Miscellaneous illness
ILLMISC_DIM	Days in milk at 1st miscellaneous illness event
ILLMISC<=60	Days in milk at 1st miscellaneous illness event less than or equal to 60 (1: days in milk less than or equal to 60; 0: otherwise)
Cutoff values were based high-risk time periods defined by previous research and opinion: immediately after calving (<=7), during the transition period (<=21), and during early portion of lactation (<=6) (Drackley, 1999; McArt et al., 2012; Barragan et al., 2018; McArt and Neves, 2020)	

APPENDIX C

Instructions:

1. Event dates are filtered between the earliest and latest fresh date for analysis
2. Observations are grouped by animal ID, lactation number, and event. Observations are arranged by event date in animal ID, lactation number, and event. Event dates are filtered where the next observation is removed if it is within a 3-day time lag of the previous observation. Observations are ungrouped.
3. Event names are modified by case when

Match Event Name	**Replacement Event Name**
FRESH	FRESH
ABORT	ABORT
SOLD	SOLD
DIED	DIED
MF	MF
RP	RP
METR	METR
KETOSIS	KETOSIS
DA	DA
LDA	DA
RDE	DA
PNEU	PNEU
SCOURS	DIARRHEA
LAME	LAME
FOOTRIM	FOOTRIM

ILLMISC	ILLMISC
RESP	PNEU
PNEUMON	PNEU
MFEVER	MF
MILKFEV	MF
MILKFVR	MF
MLKFEVR	MF
MLKFVR	MF
RETAINP	RP
MAST	MAST
MET	METR
METRITIS	METR
KET	KETOSIS
DIARRH	DIARRHEA
DIARHEA	DIARRHEA
HOOFROT	LAME
FOOTROT	LAME
FOOTRMK	FOOTRIM
TRIM	FOOTRIM
ILL	ILLMISC

Otherwise, event names default to as is. Missing event names are removed.

4. Modify FTDIM, FSTPJ, FSTBF, FSTPR, PEAKM, DCAR, and PR305 if value is '0', then replace with 'NA'. Otherwise, defaults to as is.
5. Modify LOG1 if FSTPJ is 'NA', then LOG1 is 'NA'. Otherwise, LOG1 defaults to as is.
6. Add ID, LACT, FDAT, BDAT, FCDAT, RC, DRYLG, DDRY, PDCC, PDIM, EASE, CNUM, CLIVE, CSEX, FTDIM, FSTPJ, FSTBF, FSTPR, PEAKM, DCAR, PR305, LOG1, Event, Date, Remark, Protocols, and Technician as placeholders if missing where all values are 'NA'.
7. Keep ID, LACT, FDAT, BDAT, FCDAT, RC, DRYLG, DDRY, PDCC, PDIM, EASE, CNUM, CLIVE, CSEX, FTDIM, FSTPJ, FSTBF, FSTPR, PEAKM, DCAR, PR305, LOG1, Event, Date, Remark, Protocols, and Technician. Drop all other variables.
8. Modify Remark where Remark, Protocols, and Technician are concatenated. Drop Protocols and Technician.
9. Event names are filtered to keep FRESH, ABORT, SOLD, DIED, MF, RP, METR, KETOSIS, DA, PNEU, DIARRHEA, LAME, FOOTRIM, MAST, and ILLMISC. Drop all other observations.
10. Observations are grouped by animal ID, lactation number, and event. Modify event names where event name and event number are concatenated. Observations are ungrouped.
11. Pivot data from wide to long format by Date and Remark. Modify event names where event name and if it is a Date or Remark are concatenated. Pivot data from long to wide format by event name.
12. Add MF_1_Date, RP_1_Date, METR_1_Date, KETOSIS_1_Date, PNEU_1_Date, DIARRHEA_1_Date, LAME_1_Date, FOOTRIM_1_Date, MAST_1_Date, and ILLMISC_1_Date as placeholders if missing where all values are 'NA'.
13. Lactation number is filtered to keep all LACT greater than 0. Drop all other observations. Fresh dates are filtered between the earliest and latest fresh date for analysis. Dates are transformed to Date objects. Remarks are transformed to character objects.
14. Observations are grouped by row. Create XMF, XRP, XMETR, XKET, XDA, XPNEU, XDIAR, XLAME, XFTRM, XMAST, and XILL to count the number of non-missing events for each animal ID and lactation number. Observations are ungrouped.
15. Create DATS as SOLD_1_Date, REMS as SOLD_1_Remark, DATD as DIED_1_Date, and REMD as DIED_1_Remark. Create DATR and REMR where if DATS is not missing, then DATR is DATS and REMR is REMS; else if DATD is not missing, then DATR is DATD and REMR is REMD. Otherwise, DATR and REMR are 'NA'.
16. Birth date and fresh are filtered to keep strings between 8 to 10 characters for valid dates. Modify FCDAT if string is not between 8 to 10 characters, then replace with 'NA'. Otherwise, FCDAT defaults to as is.
17. ID and RC are transformed to factor objects. BDAT, FDAT, FCDAT are transformed to Date objects. CLIVE and CSEX are transformed to character objects. LACT, FTDIM,

FSTPJ, FSTBF, FSTPR, PEAKM, DCAR, PR305, DRYLG, LOG1, LACT, DDRY, PDCC, PDIM, EASE, and CNUM are transformed to numeric objects.

18. Create LACTGP where if LACT is greater than or equal to three, then LACTGP is ‘3’. Otherwise, LACTGP is LACT. Create FRESH_MONTH by formatting FDAT in ISO8601 character format and represent with year and month format. Create DIM where if DATR is missing, then DIM is the difference between DC extraction date and fresh date. Otherwise, DIM is the difference between date removed and fresh date.

19. Create FRESH 372 TO 7, FRESH 386 TO 21, and FRESH 425 TO 60 where if DIM is between 7 and 372, then FRESH 372 TO 7 is ‘1’; if DIM is between 21 and 386, then FRESH 386 TO 21 is ‘1’; and if DIM is between 60 and 425, then FRESH 425 TO 60 is ‘1’ respectively. Otherwise, the FRESH 372 TO 7, FRESH 386 TO 21, and FRESH 425 TO 60 are ‘0’.

20. Create BIRTH_DATE as BDAT. Create AGE_AT_CALVING_(MONTHS) as the number of months between BIRTH_DATE and FDAT. Create AGE_AT_CALVING_1ST_LACT_(MONTHS) as the number of months between BIRTH_DATE and FCDAT.

21. Create DDRY where if LACT is not equal to one, then DDRY is ‘NA’. Otherwise, DDRY defaults to as is. Create ABORT where if PDCC less than 260, then ABORT is ‘1’. Otherwise, ABORT is ‘0’. Create PDIM where if LACT is not equal to one, then PDIM is ‘NA’. Otherwise, PDIM defaults to as is.

22. Create CALVING_EASE_>=2 and CALVING_EASE_>=3 where if EASE is greater than or equal to two, then CALVING_EASE_>=2 is ‘1’ and if EASE is greater than or equal to three, then CALVING_EASE_>=3 is ‘1’. Otherwise, CALVING_EASE_>=2 and CALVING_EASE_>=3 are ‘0’.

23. Modify CNUM if CNUM is equal to one and EASE is not equal to zero, then CNUM is ‘NA’. Otherwise, CNUM defaults to as is. Modify CSEX if CNUM is missing, then CSEX is ‘NA’. Otherwise, CSEX defaults to as is. Modify CLIVE if CNUM is missing, then CLIVE is ‘NA’. Otherwise, CLIVE defaults to as is.

24. Create TWINS where if CNUM is greater than one, then TWINS is ‘1’ and CNUM is equal to one, then TWINS is ‘0’. Otherwise, TWINS is ‘NA’. Create STILLBIRTH where if any element of the string CLIVE matches the pattern ‘D’, then STILLBIRTH is ‘1’. Otherwise, STILLBIRTH is ‘0’. Create CALF_SEX as CSEX. Create MALE_CALF where if any element of the string CSEX matches the pattern ‘M’, then MALE_CALF is ‘1’. Otherwise, MALE_CALF is ‘0’.

25. Create FPR where it is the ratio between FSTBF and FSTPR. Create FPR>1.4 where if FPR is greater than 1.4, then FPR>1.4 is ‘1’. Otherwise, FPR>1.4 is ‘0’.

26. Create sold, died, removed, milk fever, retained placenta, metritis, ketosis, DA, pneumonia, diarrhea, lameness, foot trim, mastitis, and miscellaneous illness and other associated variables according to the definitions and thresholds in **S1 Appendix B**.

27. Days in milk are transformed to numeric objects. CLIVE, CSEX, CALF_SEX, FRESH_MONTH, FRESH 372 TO 7, FRESH 386 TO 21, FRESH 425 TO 60, ABORT,

CALVING_EASE_>=2, CALVING_EASE_>=3, TWINS, STILLBIRTH, MALE_CALF, FPR>1.4, and all variable between SOLD and ILLMISC<=60 that does not end with ‘DIM’ are transformed to factor objects. The levels of RC, CLIVE, CSEX, CALF_SEX, FRESH_MONTH are sorted in alphabetical order. The levels of FRESH_MONTH, FRESH 372 TO 7, FRESH 386 TO 21, FRESH 425 TO 60, ABORT, CALVING_EASE_>=2, CALVING_EASE_>=3, TWINS, STILLBIRTH, MALE_CALF, FPR>1.4, and all variable between SOLD and ILLMISC<=60 that does not end with ‘DIM’ are defined such that ‘0’ is the reference level.

CHAPTER 7 THE DDCHECKPLUS APP – PREVENTION AND CONTROL OF DIGITAL DERMATITIS IN DAIRY HERDS USING EARLY DETECTION AND AUTOMATED ANALYSIS

ABSTRACT

Digital dermatitis (DD) is an infectious bovine claw disease, leading to lameness. The progression of the disease is characterized by five clinical stages, denoted as M-stages, representing distinct severities, clinical traits, and outcomes. The monitoring of the proportions of cows for each M-stage is critical for understanding and addressing DD in addition to identifying risk factors of the disease within a herd.

Changes in the distribution of cows across M-stages over time or between groups may be indicative of variations in management practices, environmental conditions, or treatment interventions. These changes can significantly impact the future claw health of the herd. However, detecting trends in claw health related to DD is not intuitive without statistical analysis of detailed records.

The aim of the study is to update and improve the DD Check app for individuals with limited statistical training or experience. The DDCheckPlus app standardizes M-stage data recording, automates comprehensive data analysis including trends over time, calculates predictions, and assigns Cow Types (I – III) based on the presence or absence of active DD lesions. The app is designed to include a DD detection module where cows can be scored for M-stages using a custom object detection model. Additionally, the app is developed to streamline data analysis for automated prediction of current DD trends and forecasting of future DD proportions. All plots

are interactive with details on demand and all tables are interactive with filtering, pagination, and sorting.

The predictions by the app are generated from stationary distributions in a class-structured multi-state Markov chain population model, commonly used for modeling endemic diseases. We hope that the flexibility in recording details at various levels will aid in the discovery of significant trends in M-stage prevalence and assist the user to make informed data-driven decisions for the prevention and control of DD on-farm.

The DDCheckPlus App was tested using two distinct datasets: a video file of the plantar aspect of standing feet of dairy cows for DD detection and a tabular format of cows scored for M-stages and signs of chronicity on a commercial farm for data analysis, demonstrating versatility and value of the app. The app facilitated easy recording of M-stages in diverse environments at different levels of detail. Results indicated that the tool effectively identifies trends in M-stage proportions, predicts potential DD outbreaks, and facilitates comparisons among Cow Types, signs of chronicity, scorers, or pens.

The DDCheckPlus app also provided a list of cows requiring treatment with individualized Cow Types to inform prognosis decisions and instruct treatment plans. This proactive approach allows producers to effectively control DD within their herds. The DDCheckPlus app serves as a powerful tool for the democratization the knowledge and insights of veterinary epidemiology, representing a significant stride in leveraging technology for monitoring, controlling, and preventing complex diseases in the cattle industry.

INTRODUCTION

Digital dermatitis (DD) primarily results from *Treponema* spp. and is a prevalent cause of lameness within the dairy industry (Evans et al., 2008; Gomez et al., 2012; Klitgaard et al., 2008; Yano et al., 2010; Zinicola et al., 2015). Beyond welfare concerns, it leads to reduced profits, caused by involuntary culling, diminished production, and fertility issues (Amstel and Shearer, 2008; Cha et al., 2010). DD is a multifaceted, polymicrobial disease exhibiting varying severities and clinical stages with distinct significance (Read and Walker, 1998). It impacts individual cows differently, and the disease manifestation can also vary across parities and lactation stages (Palmer and O'Connell, 2015; Schöpke et al., 2015).

Presently, common treatment and control methods involve topical applications administered during routine hoof trims or in response to severe lameness and regular hoof baths for the entire herd. The current approach to preventing and controlling DD assumes uniform reactions for all cows to interventions. Typically, prevention and control measures are reactive responses to outbreaks rather than proactive strategies based on continuous monitoring of herd dynamics. Effective proactive prevention and control strategies necessitate investigations into the progression of DD lesions over time.

Similar to monitoring udder health, tracking trends in DD prevalence at the individual cow level would improve our understanding and awareness of DD dynamics within herds that are endemically affected. Currently, numerous cases of DD go unnoticed, as the focus of detection is primarily on cows exhibiting lameness with large, active DD lesions (Rodriguez-Lainz et al., 1996). This detection strategy ignores many instances of DD where not all cows affected by DD display lameness (Tadich et al., 2010).

The M-stage classification system, has been employed to characterize the clinical course of DD, considering variations in disease severity (Berry et al., 2012; Döpfer et al., 1997). The system is recognized as a valuable tool in the fight against DD (Schöpke et al., 2015). Additionally, signs of chronicity such as hyperkeratosis and skin proliferations affected by DD are crucial, as such factors influence the resulting lameness, infectious potential, treatment requirements, and prognosis of DD lesions (Gomez et al., 2014). A proactive and systematic approach to recording DD lesions and indicators of chronicity should be an integral part of an inclusive prevention and control strategy against DD, including all age groups of cattle, including pre-calving heifers (Gomez et al., 2015).

Changes in M-stage prevalence over time at the herd level are linked to variations in management practices e.g. hygiene, environmental conditions e.g. temperature, and control and treatment methods (Cramer et al., 2009; Holzhauer et al., 2006; Rodriguez-Lainz et al., 1999; Somers et al., 2005; Wells et al., 1999). The prediction of these trends can be accomplished through different levels of record detail, ranging from M-stage records to simple presence or absence records.

The forecasting these trends requires the stationary distributions of transitions between DD states, such as M-stages or indicators of chronicity, using a class-structured population model (Otto and Day, 2007). Such models enable the proactive prediction of outbreaks and the anticipated effects of interventions before the occurrence and recurrence of DD in real time. The significance of these trends can be subjected to statistical testing using the uncertainty measures inherent in the outcomes of the model (Chernick, 2011).

At the individual cow level, the recurrence of active M2 DD lesions facilitates the categorization of cows into defined Cow Types I - III. This classification is determined by the history of M2 lesions (Döpfer et al., 2004; Gomez et al., 2014; Holzhauer et al., 2008, 2006). Cow Types exhibit distinct immunologic responses to *Treponema* spp., making it crucial for predicting treatment outcomes and guiding genetic selection against DD (Schöpke et al., 2015).

Advancements in technology have revolutionized the monitoring of dairy cattle, providing reliable and cost-effective solutions. In the past century, records were traditionally maintained manually, in notebooks or on cards above each cow. The labor-intensive nature of recording and calculating new information constrained the feasibility of advanced analytics. Today, dairy herd management software applications empower users to input data once and automatically generate future insights. These tools enable users to monitor thousands of animals, capturing specific information for each one.

Despite these advancements, the development of health analytics has faced challenges characterized by manual data entry and limited insights, perpetually constrained by the quality of data recording (Gonçalves et al., 2022; Wenz and Giebel, 2012). Current analysis options within herd management software lack statistical decision tools, relying predominantly on visual trend analysis for decision-making. Producers and consultants have been restricted in the questions of interest and corresponding answers due to limitations in data accessibility and available tools. Users typically invest significant time in completing a single analysis, and the average user lacks the data skills required for extracting this information. The development of tools for dairy professionals, regardless of their proficiency in data transformation, would unlock new insights and improve management practices.

In precision cattle farming, managerial decisions hinge on quantitative data (Mahmud et al., 2021). Various sensing technologies are employed for data collection where the dataset is subsequently analysis using advanced algorithms. Moreover, real-time quantitative data can be collected by utilizing sensors like accelerometers or gyroscopes to monitor cattle behavior or movement. Herd management software leverage data to facilitate in making informed decisions (Mahmud et al., 2021). This information aids in determining the individual needs of each animal, enabling personalized attention to enhance overall production outcomes (Banhazi and Black, 2009). For a comprehensive utilization of data for decision-making, the integration of various artificial intelligence and machine learning algorithms can be considered. This integration serves to automate and improve the accuracy and precision of the decision-making processes (Banhazi et al., 2012).

Computer vision (CV) is a subset within artificial intelligence to identify and understand objects in images and videos. Initially, the algorithm is trained using a dedicated dataset of labeled images, followed by validation employing a distinct dataset. The trained parameters are subsequently used for outcome prediction and decision-making processes. In precision cattle farming, CV is already applied to address diverse challenges (Mahmud et al., 2021). Such examples include the detection of flies on cattle bodies (Psota et al., 2021), identification of individual body parts (Jiang et al., 2019), breed recognition (Weber et al., 2020), lameness detection (Kang et al., 2020), and mastitis prediction (Xudong et al., 2020) using ground-based images. Additionally, CV is employed for predicting body weight (Gjergji et al., 2020) and counting (Xu et al., 2020) using of unmanned aerial vehicle images.

Object detection tasks have been employed for various purposes in studies involving livestock animals, where the primary objective is to identify one or more objects within an image or video (Borges Oliveira et al., 2021). Object detection algorithms for animal detection, have been used in previous studies focused on livestock (Cowton et al., 2019; Lee et al., 2019; Psota et al., 2019;

Seo et al., 2020). These algorithms have also been applied to detect lameness (Kang et al., 2020) and DD (Cernek et al., 2020) in dairy cattle.

Cernek et al. used YOLOv2 in RGB images, achieving an accuracy of 88% (Cernek et al., 2020). These findings underscore the significant potential of computer vision systems in identifying cows with DD, thereby reducing DD prevalence and improving animal welfare. Kang et al. developed a lameness scoring system for dairy cows using the Receptive Field Block Net Single Shot Detector deep learning network (Kang et al., 2020). This system effectively located cow hooves in videos with a mean average precision of 87.0%. Subsequently, the identified legs served as input for an algorithm to calculate the supporting phase, a critical metric derived from the difference between hoof lifting time and hoof load time.

There exists many herd management software designed to assist farmers and livestock managers in overseeing and optimizing various facets of animal husbandry (Afimilk, n.d., n.d.; AgriWebb, n.d.; CattleMax, n.d.; CattlePro, n.d.; CattleWorks, n.d.; CowManager, n.d.; Cownect, n.d.; DeLaval, n.d.; DTS, n.d.; Farmbrite, n.d.; FarmLogic, n.d.; HerdApp, n.d.; Herdlync, n.d., n.d.; Herdmaster, n.d.; Herdtrax, n.d.; Herdwatch, n.d.; Keogh, n.d.; Peltjes, n.d.; SMARTBOW, n.d.; StockManager, n.d.; UNIFORM-Agri, n.d.; VAS, n.d.). These software solutions are tailored for livestock operations, offering features that provide users to monitor health, track breeding and reproductive activities, manage feeding programs, and maintain accurate records. There are only a few applications that provides analysis tools and reporting features for real-time decision-making. Moreover, there is no commercial software that combines data collection using object detection or any other CV algorithm in addition to statistical analysis for predicting trends and forecasting health events.

While mathematical and statistical tools are capable of predicting DD states based on historical data and trends exist, it may not be readily accessible to dairy farmers and veterinarians. The objective of the study is to update and improve the DD Check app to generate standardized DD records, assign Cow Types based on DD history, and automate both descriptive and predictive data analyses (Tremblay et al., 2016). The DDCheckPlus app aims to assist the user with limited statistical training, scoring experience, or access to supporting programs. We hope that the app can be applied at various levels of record detail to identify statistically significant trends in the prevalence of M-stages, aiding in informed decision-making for on-farm DD prevention and control.

Given the intricate nature of DD and its widespread prevalence in affected herds and cows, automated tools for description and prediction are essential for long-term prevention and control strategies against DD. Similar apps designed to automate statistical analyses of standardized field datasets will enhance awareness and understanding among producers, managers, veterinarians, and owners, ultimately leading to improved prevention and control of endemic production diseases.

MATERIALS AND METHODS

The DDCheckPlus App is designed for iOS devices and can be easily accessed through the Apple App Store or iTunes, compatible with both iPad and iPhone platforms. The application is structured into two integral components: the mobile-based and web-based modules.

The mobile-based module is developed using Xcode, a powerful integrated development environment distributed by Apple (Apple Inc, n.d.). Xcode is a suite of tools developers use to build apps for Apple platforms, creating seamless deployment and user-friendly experience for iPad and iPhone users.

The web-based module is developed using R Shiny, an interactive and dynamic web application framework designed for the R programming language and software environment for statistical computing and graphics (R Core Team, n.d.). The framework facilitates data analysis and visualization, improving the app performance for detailed record-keeping and strategic decision-making.

The DD detection module is implemented in Swift to build a live video object detection iOS app using the YOLOv5 model (Hietala, 2022; hietalajulius, 2022). The module is created to setup video capture, output video stream, and visual layers for displaying the detections and inference time. Apple Vision framework is setup to perform various standard computer vision tasks on iOS devices using custom Core ML models (Apple Developer, 2023a). Core ML supports machine learning models on iOS devices to build, train, or convert completely custom models (Apple Developer, 2023b). Core ML optimizes on-device performance by leveraging the CPU, GPU, and Neural Engine while minimizing its memory footprint and power consumption (Apple Developer, 2023b).

The camera is configured for capture to optimal performance with Vision algorithms (Apple Developer, 2023c). The device and session resolutions are setup such that Vision can process results more efficiently. The video input is added to the session by adding the camera as a device and the video output is added to the session when specifying the pixel format. The device orientation is specified such that frames are processed relative to the orientation of the capture device (Apple Developer, 2023c). The class labels are designated in the Core ML model and the detected objects are parsed to draw a bounding box around the object and display the classification and confidence in a textual overlay (Apple Developer, 2023c).

The user interface is implemented in Objective C using a storyboard runtime for assisted DD scoring and standardized data entry. The module is created to provide a user-friendly approach for inputting data, segue to and from the DD detection module, summarize the inputted data, and generate a dataset for further data analysis. The module saves the dataset locally as a comma-separated value (CSV) file and makes a URL request to redirect the user to the mobile-based R Shiny app.

This dual-component design not only leverages the native capabilities of iOS devices, but also extends the app functionality to web browsers, providing the user with a versatile and accessible tool for managing and monitoring DD data. The DDCheckPlus App delivers a comprehensive solution to generate standardized DD records, assign Cow Types based on historical data, and automate both descriptive and predictive data analyses for effective on-farm prevention and control strategies.

Mobile-Based Module

The mobile component features a user-friendly interface designed for inputting the herd name, herd code, scorer details, date, pen number, cow ID, foot information, M-stage lesion details, and indicators of chronicity (Figure 1). The signs of chronicity are categorized as follows: None for smooth skin without thickening; Hyperkeratotic for thickened skin; and Proliferative for overgrown epidermal tissues.

The application provides visual aids including images and detailed descriptions to assist the user in accurately scoring DD lesions using the M-stage system (Figure 1). The M-stages of DD are categorized as follows: M0 represents normal digital skin; M1 represents an early, small circumscribed red to gray epithelial defect measuring less than 2 cm in diameter; M2 represents acute, active ulcerative or granulomatous digital skin alteration greater than 2 cm in diameter; M3 represents the healing stage within 1 – 2 days after topical treatment where the acute DD lesion is covered by a firm scab-like material; M4 represents the chronic stage of DD, which can be either hyperkeratotic or proliferative; and M4.1, represents a chronic stage with an early or intermediate stage (M1) (Berry et al., 2012; Evans et al., 2016).

The iPad version, benefitting from a larger screen, prominently displays DD lesions on the scoring buttons, enhancing the user ability to accurately assessments of DD lesions. The interface supports the scoring of cattle using all five M-stages (Berry et al., 2012), simplified M-stages where M3 and M4 are combined (Relun et al., 2011), or just two M-stages to indicate the absence or presence of DD (Cramer et al., 2018). Users can also employ any other combination based on their preferences, provided consistency is maintained throughout the dataset. In cases where a cow presents multiple lesions during scoring at the individual level, priority is given to M2 lesions over M3 lesions, M3 lesions over M4 lesions, and M4 lesions over M1 lesions. When multiple signs of chronicity are observed, proliferation takes precedence over hyperkeratosis.

Scoring for DD can be performed in various settings, including the milking parlor, using a restraint chute, or during pen walks where cattle are secured in head locks along the feed aisle, with the scorer walking behind to assess the hind feet (Jacobs et al., 2015; Relun et al., 2011). For free-ranging cattle in freestalls, scoring can occur as they walk past the investigator one by one in the alleys. Moreover, the app can be effectively employed to generate records in beef cattle operations, capturing data while cattle feed at the bunk or during what is referred to as DD alley checks. DD alley checks are feasible in any feedlot equipped with the necessary facilities, such as connecting alleys, and with handlers possessing the requisite experience and training.

The user can refer to the real-time DD detection module for assistance by pressing the "Detect" button. The module performs DD detection using the built-in, rear-facing camera. The lesion is circumscribed by a bounding box with class label for the respective M-stage classification and a confidence score for the prediction (Figure 2). The user can return to the previous module for scoring by pressing the "Close" button.

After scoring a cow, the user can save the data by pressing the "Save/New Cow" button, which updates a CSV file and resets the input section for the next cow. The recorded data is conveniently displayed in the lower section of the app and can be filtered by the current day, all days, all farms, or only M2 observations of a given day. For a comprehensive overview, the "Summary" button provides users with a list of individual records and a summary that includes the total number of records, percentages of feet scored, and percentages of lesions scored.

Web-Based Module

The web-based application facilitates the direct submission of data from the mobile app to a cloud-based server. Notably, the data is not retained on the server, ensuring confidentiality. This web-based tool provides relative frequencies of M-stages or signs of chronicity per scoring event and predictions of future frequencies using the DD Infection Model. The user interface of the web-based app is divided into three tabs: the Dashboard tab for data upload, view contents, and data summary; the Results tab for tables and plots; and the Documentation tab for dataset requirements and disclaimer of limited use.

The Dashboard tab of the DDCheckPlus R Shiny app contains three panels. The Input panel is used to upload the downloaded CSV file from DC (Figures 3 - 4). Moreover, users are able to input bootstrap sample size for predictions and confidence level used for further analysis. Specific weeks of interest from the dataset can be selected by the user using the slider in the interface. A drop-down menu featuring variables such as herd name, herd code, pen, signs of chronicity, Cow Type, and scorer is available for user to subset the data and view the results separately. This intuitive design ensures flexibility and customization in data analysis for enhanced user experience.

After the user uploads the CSV file of the dataset, inputs the values in the Input panel, the data is dynamically analyzed. The app generates an ID, a concatenation of cow ID and foot if the specific foot is provided. Otherwise, records are identified at the cow level. Duplicate observations of cow ID for a week or empty cow ID records are excluded. Dates are transformed to weeks, and each cow is assigned a Cow Type. Type I cows were never scored with an M2 lesion, Type II cows have one M2 lesion, and Type III cows have repeat M2 lesions. Cow Type is assigned based on all available records where a minimum of two records is required to define Cow Type. The assignment of Cow Type is at the cow level, even if records are detailed at the foot level.

There is no maximum time limit between consecutive recordings for assigning a Cow Type. However, maintaining consistent time intervals between frequent scoring events is ideal for the reliability, accuracy, and interpretability of the assigned Cow Type. A subset of the data is then created based on the weeks selected using the slider in the user interface. The cleaned dataset with the Cow Type can be downloaded using the "Generate CSV File" button of the Input panel (Figure 4). The Contents panel is used to display the first 100 rows of the dataset providing filtering, pagination, and sorting (Figure 4). The Summary panel is used to display summary statistics the user can skim to understand their data (Figure 4). Results are printed horizontally with a section for each variable type and a row for each variable.

The DD Infection Model generates predicted relative frequencies of DD disease states, such as M-stages or signs of chronicity, using class-structured Multi-state Markov Chain Models (Otto and Day, 2007). Sequential scoring events are used to identify transitions among disease states, and the total number of transitions between disease states per dataset is employed to construct a transition matrix. These transition matrices are then converted into proportions. The predictions are derived from the stationary distributions of the probability transition matrix, extracted from the first right eigenvector belonging to the dominant eigenvector (Otto and Day, 2007). Confidence intervals for the predictions are generated using bootstrapping methods (Caswell, 2000). For each bootstrap sample, cow ID are randomly sub-sampled with replacement to create a subset. Each subset is used to generate predictions via the DD Infection Model. The variance for all subset predictions is used to generate the 95% confidence intervals.

Given the high computational times associated with bootstrapping, the app initially sets the number of bootstraps to 10. However, 1000 bootstrap samples are recommended. Increasing the number of bootstraps will extend computational time contingent on the size of the dataset. While there are no strict rules for the number of bootstraps, increasing the number will improve the reliability of the 95% confidence intervals (Henderson, 2005). Maintaining consistent time intervals between frequent scoring events is advisable to improve the accuracy of predictions. To prevent predictions on data sets with low power e.g., small sample size or too few cows repeatedly scored, a limit of 15 transitions was set before results would appear. Error messages have been incorporated to assist users in addressing common mistakes or issues.

The Results tab is used for data analysis and reporting (Figure 5). The observed proportions of disease states for each week in addition to predicted proportions are visually presented using bar plots in the Plots panel. The proportion is represented as a bar and includes a 95% confidence interval represented as a line range. The confidence intervals serve as an indirect measure of significance for weeks or groups (du Prel et al., 2009). Non-overlapping confidence intervals indicate a statistically significant difference in proportions, whereas overlapping intervals suggest that the difference in proportion is not statistically significant. The predictions used to generate the plots are also accessible using tabular format in the Tables panel for a comprehensive view of the data. All plots are interactive with details on demand and can be downloaded as PNG files. Similarly, all tables are interactive with filtering, pagination, and sorting and can be downloaded as CSV files the "Generate CSV File" button.

There are four ways to explore the results of the data analysis, each accessible by selecting different tabs in the web-based interface of either the Plots or Tables panel. The first tab in the Plots panel or the second tab in the Tables panel is M-stage and provides the results categorized by M-stage and week. The second tab in the Plots panel or the third tab in the Tables panel is Chronicity and provides the results categorized by the sign of chronicity and week. The third tab in the Plots panel or the fourth tab in the Tables panel is Variable of Interest and provides the results categorized by the user-selected variable of interest, M-stage, and week. Additionally, the first tab in the Tables panel is M2 Lesions and presents a list of cows that require further treatment i.e., cows with M2 lesions during the most recent scoring event, including previous M-stage and Cow Type to aid in the prognosis of topical treatment. While M1, M2, and M4.1 lesions are all considered active, the treatment list focuses on M2 lesions since M1 and M4.1 lesions are not as painful as M2 lesions (Holzhauer et al., 2008).

The Documentation tab provides requirements to upload an input CSV file and to generate an output CSV file in DDCheckPlus (Figure 6). A disclaimer of limited use is provided for variables used to generate the output CSV file and for the analysis in DDCheckPlus. Finally, an example dataset is provided for download for interested users to try out and test the app.

RESULTS

Mobile-Based Module

A case study was performed to demonstrate the potential uses of the DDCheckPlus DD detection module. The sample data contained over 100 cows on a commercial dairy taken during July 2023. A GoPro Hero 5 Black cameras was used to take MP4 video recordings of the backside of the hind feet at claw level. The evaluation of network performance and efficiency of the implementations included two main components: accuracy and inference speed. Cohen's kappa was used to evaluate the agreement between the model predictions and the ground truth labels, providing valuable insights into the model's ability to make accurate predictions in real-world scenarios (Landis and Koch, 1977; Viera and Garrett, 2005). The Cohen’s kappa for DD detection module compared to the trained investigator was determined to be 0.625 and interpreted as substantial agreement between the two raters ($z = 10.7$; $p < 0.001$). The DD detection module was able to detect all five M-stages of DD on live streaming video. The prediction accuracy for DD detection module was 0.745 with moderate detection for M0, M2P, and M4P, strong detection for M4H, and almost perfect detection for M2 (Figure 7).

The inference speed was measured in frames per second (FPS), calculating the number of images that the model has processed and inferred within one second. A higher FPS value implied a faster and more efficient model, wanted in real-time applications or scenarios with limited computing resources. The average inference time was 20 ms or 50 FPS. The inference time far exceeded the minimum threshold for image processing by a human visual system at approximately 10 FPS. Moreover, the inference time always exceeded the minimum threshold for real-time detection at approximately 30 FPS.

Web-Based Module

A case study was performed to demonstrate the potential uses of the DDCheckPlus R Shiny app. The sample dataset contained 300 observations for 100 cows on a commercial farm. The sample dataset included herd code, herd name, scorer name, cow ID, pen, M-stage lesions of 0, 2, and 4, and signs of chronicity of 0, 1, and 2 at the cow level. Cows were scored by a trained investigator during three pen walks on 4/24/2014, 5/2/2014, and 5/10/2014 and individual Cow Types were defined based on week 1 to week 3.

The dataset was uploaded in the Dashboard tab where the app defaulted to a bootstrap sample of 10 and a confidence level of 95%. The weeks of 1 to 3 was selected using the slider and a variable of interest of chronicity was selected using the drop-down menu. The tables and plots were automatically generated in Tables and Plots panels of the Results tabs. There were 12 predicted M2 lesions where the two cows transitioned from M0 to M2 lesions, two cows transitioned from M4 to M2 lesions, and the outstanding 8 cows with M2 lesions remained constant. Note that the four cows that transitioned to M2 lesions were a cow type of II and the 8 cows with M2 lesions that remained constant were a cow type of III.

The predicted proportion of cows with M2 lesions based on weeks 1 to 3 was 0.124 (95% CI: 0.073 – 0.176) (Figure 8). Similarly, the predicted proportion of cows with M0 lesions based on weeks 1 to 3 was 0.617 (95% CI: 0.528 – 0.705) and the predicted proportion of cows with M4 lesions based on weeks 1 to 3 was 0.259 (95% CI: 0.175 – 0.343). The relative proportions of cows with M2 lesions were not significantly different in weeks 2 and 3 compared to week 1. The relative proportions of cows with M0 lesions were significantly higher in weeks 2 and 3 compared to week 1. The relative proportions of cows with M4 lesions were significantly lower in weeks 2 and 3 compared to week 1.

The predicted proportion of cows with proliferative lesions based on weeks 1 to 3 was 0.121 (95% CI: 0.070 – 0.171) (Figure 9). Likewise, the predicted proportion of cows with hyperkeratotic lesions based on weeks 1 to 3 was 0.667 (95% CI: 0.587 – 0.747) and the predicted proportion of cows with no chronic lesions based on weeks 1 to 3 was 0.123 (95% CI: 0.148 – 0.277). The relative proportions of cows with proliferative lesions were not significantly different in weeks 2 and 3 compared to week 1. The relative proportions of cows with hyperkeratotic lesions were significantly higher in weeks 2 and 3 compared to week 1. The relative proportions of cows with no chronic lesions were significantly lower in weeks 2 and 3 compared to week 1.

Lastly, the proportion of cows over weeks was faceted by the type of M-stage lesions and the variable of interest of signs of chronicity (Figure 10). The relative proportions of cows with M2 lesions and no chronicity were significantly lower in week 3 compared to weeks 1 and 2. All others relative proportions of cows between M-stage lesions and chronicity were not significantly different in weeks 2 and 3 compared to week 1. The resulting evaluation shows that M2 lesions are more chronic over time. Any analyses for association between the proportion of cows over weeks by the type of M-stage lesions controlling for a factor can be studied using the DDCheckPlus app if the number of observations for each level is sufficiently large.

DISCUSSION

The availability of applications on iOS devices has revolutionized access to information and tools, providing individuals with constant connectivity. For the dairy industry, there is an ongoing and progressive trend of integrating such technology in management practices and decision-making. These applications serve as valuable tools for dairy producers, enabling them to make informed improvements to the herds welfare and farm well-being. Moreover, the standardized datasets generated by these applications contribute to broader research on cattle husbandry and production.

A noteworthy example of this technological integration in dairy farming is evident in the current analysis of datasets collected using the DDCheckPlus app from various farms (Tremblay et al., 2016). This analysis demonstrates the automation of data analysis for decision-making processes. The insights derived from this automated analysis can be shared between dairy farmers and veterinarians, fostering collaboration and facilitating the interpretation of statistical differences. This collaborative approach allows for a deeper understanding of trends, interventions, and the impact of control measures in precision farming. Overall, the utilization of handheld devices and specialized applications marks a transformative shift in the way data is employed for the improvement of both individual herds and the broader agricultural industry.

Data Analysis Tools

Within the context of agricultural production, there exists an abundance of data that requires sophisticated analysis for well-informed decision-making processes (Antle et al., 2017; Moore et al., 2022). User engagement and acceptance can be facilitated through outreach and extension services, employing tools specifically designed for their intended purposes, thereby reducing barriers to application by end users. This study introduces the DDCheckPlus app to address an unmet need in the dairy industry. The current reliance on decision-making processes based on graphic analysis or trend observation is insufficient for contemporary production systems (Correll et al., 2012).

To meet the demands of modern agricultural practices, more advanced tools are imperative. The DDCheckPlus app offers a comprehensive solution by combining data collection, data cleaning, data visualization, and statistical analysis. Recognizing the diverse needs of users engaged in local on-farm decision support, it is critical that data analysis tools are customizable according to individual preferences. The DDCheckPlus app allows users to tailor their analytical choices. This emphasis on customization addresses the unique requirements of agricultural stakeholders, providing them with a powerful and adaptable resource for decision support within the complex landscape of modern agricultural production.

Various platforms offer data analysis features, encompassing visualization, dashboarding, and economic evaluation of cow health and production. These include milking services companies like Lely T4C and DeLaval DelPro, herd management platforms such as DC, Animal, Bovisync, Dairy Data Wearhouse, and cow monitoring technologies like SmaXtec, Connectera, and CowManager (BoviSync, n.d.; Connecterra, n.d.; CowManager, n.d.; DDW, n.d.; DeLaval, n.d.; Lely, 2016; smaXtec, n.d.; VAS, n.d.). Frequently, these applications have limitations in tools for exporting graphs and reports, and some may lack comprehensive statistical data analysis features. Users often need to pay for licenses or make significant purchases from these companies to access the full suite of data analysis tools. In contrast, the DDCheckPlus app presents distinct advantages for on-farm decision-making processes. The DDCheckPlus app is notable for customization options and user-friendly interface. This makes it a valuable and accessible tool for individuals in cow monitoring and herd management, offering a cost-effective and flexible solution to other commercial platforms.

Mobile-Based Module

The development of this application significantly enhances the efficiency of tracking DD, eliminating the traditional method of recording information with pen and paper, followed by manual data transfer to create electronic files. This app introduces a streamlined process that not only eliminates these time-consuming steps but also provides additional features that contribute to its effectiveness. One noteworthy feature is the inclusion of detailed descriptions and illustrative sample photographs of various lesions to promote consistency and introduce a standardized approach between users. Additionally, it serves as a valuable educational tool for training herd managers on the clinical course and different aspects of DD. While the app prioritizes consistency in identifying M-stages, its mobile interface remains sufficiently flexible to accommodate diverse environments within the dairy industry.

Practical applications of the app have been observed in a range of settings within the dairy sector, including scoring cattle in restraint chutes, milking parlors, during pen walks, and alley checks. Its versatility is evident in its use by various stakeholders, such as hoof trimmers, managers, veterinarians, and researchers, each with distinct objectives. Users leverage the DDCheckPlus App for purposes ranging from conducting treatment and control studies to producing treatment lists and enhancing monitoring of DD treatment and control effects on both commercial and research farms.

While the app is designed to ideally monitor DD at the M-stage level, its flexibility allows less experienced users to focus on simply detecting the presence or absence of the disease. The granularity of data recorded using M-stages can be tailored to individual preferences, presenting options such as including all five M-stages, only three M-stages, or including scores at the foot level. Furthermore, the app allows users to customize the interval between scoring events, enabling users the autonomy to adapt the tool to their demanding schedules and specific needs. Despite its broad applicability, the flexibility of the app may introduce a limitation in terms of comparing results across different farms. The variations in recording practices and user preferences may impact the direct comparability of data, emphasizing the need for careful interpretation and contextual understanding when analyzing results across different settings. Nevertheless, the overall adaptability and utility of the app make it a valuable tool for diverse users within the dairy industry.

Web-Based Module

Many producers face time constraints or lack the specialized training required for statistical analysis. The web-based component of the DDCheckPlus App serves as a valuable solution, streamlining the data analysis process and making it accessible to a wider audience. By seamlessly summarizing and distributing data collected through the mobile app, it addresses the common challenge of time constraints faced by producers.

This integration of the web-based platform not only simplifies the analytical process, but also serves as a powerful incentive for users to invest time in scoring cows and actively monitoring DD trends. This real-time feedback loop of data summarization after data collection, not only fosters engagement but also empowers users to make timely and informed decisions regarding herd management.

The web-based platform is designed to accommodate diverse user needs while assessing DD trends at various levels of detail. For beginners, a straightforward analysis of trends in the relative frequencies of M-stages provides a feasible starting point. This user-friendly approach allows individuals with limited experience in statistical analysis to derive valuable insights into the prevalence and progression of DD within the herds.

For more advanced users, the web-based platform offers advanced features for a deeper dive into the data. Such functionalities including examining predictions, confidence intervals, and exploring variations within subsets provide a more nuanced understanding of the dataset. This advanced level of analysis accommodates diverse user needs, providing them with the tools to explore predictive patterns, evaluate the uncertainty of results, and make informed decisions.

The transition between various lesions associated with DD is a dynamic process that unfolds over varying durations, ranging from a few days to many years (Berry et al., 2012; Döpfer et al., 1997; Krull et al., 2016). While it is highly recommended to conduct frequent and consistent scoring events, ideally on a weekly basis, to capture M-stage transitions effectively, the practicality of such short intervals is not always feasible on farms.

To enhance the accuracy of monitoring and to mitigate the risk of missing critical M-stage transitions between scoring dates, we recommend scoring intervals to be less than or equal to one month. Extended intervals increase the likelihood of missing two or more transitions that may have occurred between evaluations (Berry et al., 2012). It is important to note that the impact of interval duration is more pronounced on farms experiencing acute outbreaks of disease compared to those with endemic cases.

As scoring intervals increase or are variable, it is imperative to approach the interpretation of predictions with caution. In such instances, evaluating observed relative frequencies is a more reliable method. It's crucial to recognize that the predictions derived from the model are applicable at the average group level and may not precisely describe the prediction for an individual animal. However, decisions related to trends, outbreak prevention, and control are typically made at the population level. Therefore, utilizing the model predictions serves as a valuable tool for informed decisions regarding claw health management.

The model utility extends to individualized care, where specific animals can be identified for treatment. The M2 lesion treatment lists allow for the consideration of individual diagnosis and prognosis, contributing to a more targeted and effective approach in managing DD at both the population and individual levels. This dual approach assists in a comprehensive strategy for managing and preventing DD, providing a practical and adaptable tool for claw health management decisions on farms.

The precision in detecting significant differences between relative frequencies improves with larger sample sizes. While significant differences in observed prevalences can be inferred by assessing the overlap of confidence intervals, the efficacy of identifying statistically significant trends could be further improved through the integration of automated statistical tests, such as the chi-squared test or ANOVA.

In extensive herds, conducting regular scoring for every cow may pose logistical challenges. To address this challenge, farms can adopt a systematic approach by randomly selecting a proportion of the herd for scoring at regular intervals. This strategy provides an accurate representation of the impact of DD on the cows within a specific farm.

The determination of the sample size required for scoring is contingent on various factors, including the lowest M-stage prevalence, the desired level of precision, total population size, the accuracy of the scorer, and the scoring technique employed (Humphry et al., 2004; Naing et al., 2006; Thrusfield, 2018). It is necessary to minimize substantial changes in the population scored over time. Large variations in the population during different scoring events can artificially induce significant changes in proportions that may not be representative at the herd-level. Moreover, maintaining consistency in the population being scored is critical. A lack of uniformity in the population over time reduces the number of animals with consecutive scores contributing to the transition matrix and, consequently, the accuracy of predictions.

The application generates predictions for short-term trends of DD in farms with endemic infections. While the predictions offer potential benefits, it is important to note that there is inherent uncertainty in the statistical predictions. As such, further research is necessary to validate the accuracy and reliability of these results.

The choice of a class-structured multi-state Markov chain population model for the app was deliberate, aligning with its targeted use in herds characterized by endemic DD infections. The selection is based on the assumption of constant transition rates between M-stages over time, indicative of an endemic state within the infection dynamics, away from periods of outbreaks. Notably, outbreaks typically exhibit non-linear features with heightened transition rates due to the rapid transmission of disease. In such cases, dynamic models like compartmental SIR models are more suitable for accurately capturing the dynamics of infectious disease outbreaks (Anderson and May, 1991).

Infection transmission models are particularly applicable to the outbreak scenarios of DD (Döpfer et al., 2012). Future models, driven by clinical hypotheses about meaningful transitions, can combine M-stages and signs of chronicity within a single model. This approach would involve the selection of transitions that align exclusively with clinical observations, providing a more realistic representation of the disease progression. For instance, it does not clinically make sense for a healthy foot to transition directly into M2 without passing through other stages.

In addition to modeling transitions, integrating sojourn times or the duration spent in different clinical stages of the disease and relevant covariates such as lactation number or production levels can further enhance the outcomes and interpretation of the transition analysis (Caswell, 2000). This comprehensive approach ensures that the models not only reflect the clinical reality of DD progression but also provide valuable insights into the factors influencing the disease dynamics within a herd.

CONCLUSION

The DDCheckPlus App effectively standardizes the recording of M-stages of DD, automating both descriptive and predictive analyses of longitudinal data. The flexibility allows users to employ various levels of record detail in diverse situations, facilitating the identification of statistically significant trends in observed or predicted prevalence M-stage. The app aids in making informed decisions for the prevention and control of DD on-farm.

The app performs multiple functions, including assigning Cow Types, generating treatment lists, and producing data sets for distribution among herd managers. The app supports the comparison of different subsets of cows. The user-friendly interface allows individuals without a statistical background to predict near-future trends of DD in farms with endemic infections. Cattle scoring can be seamlessly integrated into routine farm activities such as during pen walks, at the milking parlor, or using a cattle chute.

Similar apps geared towards standardizing datasets have the potential to increase awareness and understanding within the dairy industry, ultimately leading to improved prevention and control of endemic diseases. Future applications would include economic models to estimate the cost and return associated with treating specific DD lesions or implementing prevention and control methods. Integrating prevalence and cow characteristics from farm management programs into the model offers the opportunity to visualize the association of cow and farm-level risk factors with DD. This multifaceted approach presents wide-ranging implications for both the efficiency and effectiveness of DD management strategies in precision farming.

ACKNOWLEDGEMENTS

Funding support for the research was provided by the US Department of Agriculture through the National Institute for Food and Agriculture - Animal Health Grant (WIS03082). We would also like to thank Jamie Sullivan (Carman, MB, Canada), Chris Bauer (Mondovi, WI), and all other dairy producers for contributing to the images of hoof lesions. The authors have no financial and personal relationships with other people or organizations that could inappropriately influence (bias) their work.

REFERENCES

Afimilk, n.d. Top Dairy Technology for Better Farm Management [WWW Document]. Afimilk. URL https://www.afimilk.com/

Afimilk, n.d. AfiFarm - The Ultimate Farm Management System. Afimilk. URL https://www.afimilk.com/afifarm/

AgriWebb, n.d. AgriWebb [WWW Document]. AgriWebb. URL https://www.agriwebb.com/

Amstel, S. van, Shearer, J., 2008. Manual for Treatment and Control of Lameness in Cattle. John Wiley & Sons.

Anderson, R.M., May, R.M., 1991. Infectious Diseases of Humans: Dynamics and Control. OUP Oxford.

Antle, J.M., Basso, B., Conant, R.T., Godfray, H.C.J., Jones, J.W., Herrero, M., Howitt, R.E., Keating, B.A., Munoz-Carpena, R., Rosenzweig, C., Tittonell, P., Wheeler, T.R., 2017. Towards a new generation of agricultural system data, models and knowledge products: Design and improvement. Agricultural Systems 155, 255–268. https://doi.org/10.1016/j.agsy.2016.10.002

Apple Developer, 2023a. Vision [WWW Document]. Apple Developer Documentation. URL https://developer.apple.com/documentation/vision/

Apple Developer, 2023b. Core ML [WWW Document]. Apple Developer Documentation. URL https://developer.apple.com/documentation/coreml

Apple Developer, 2023c. Recognizing Objects in Live Capture [WWW Document]. Apple Developer Documentation. URL https://developer.apple.com/documentation/vision/recognizing_objects_in_live_capture

Apple Inc, n.d. Xcode 15 [WWW Document]. Apple Developer. URL https://developer.apple.com/xcode/ (accessed 11.15.23).

Banhazi, T.M., Black, J.L., 2009. Precision Livestock Farming: A Suite of Electronic Systems to Ensure the Application of Best Practice Management on Livestock Farms. Australian Journal of Multi-Disciplinary Engineering 7, 1–14. https://doi.org/10.1080/14488388.2009.11464794

Banhazi, T.M., Lehr, H., Black, J.L., Crabtree, H., Schofield, P., Tscharke, M., Berckmans, D., 2012. Precision Livestock Farming: An international review of scientific and commercial aspects. International Journal of Agricultural and Biological Engineering 5, 1–9. https://doi.org/10.25165/ijabe.v5i3.599

Berry, S.L., Read, D.H., Famula, T.R., Mongini, A., Döpfer, D., 2012. Long-term observations on the dynamics of bovine digital dermatitis lesions on a California dairy after topical treatment with lincomycin HCl. The Veterinary Journal, Special Issue:Bovine Lameness 193, 654–658. https://doi.org/10.1016/j.tvjl.2012.06.048

Borges Oliveira, D.A., Ribeiro Pereira, L.G., Bresolin, T., Pontes Ferreira, R.E., Reboucas Dorea, J.R., 2021. A review of deep learning algorithms for computer vision systems in livestock. Livestock Science 253, 104700. https://doi.org/10.1016/j.livsci.2021.104700

BoviSync, n.d. Dairy Herd Management Software that streamlines dairies [WWW Document]. BoviSync. URL https://bovisync.com/ (accessed 11.20.23).

Caswell, H., 2000. Matrix population models. Sinauer Sunderland, MA.

CattleMax, n.d. Cattle Management Software by CattleMax [WWW Document]. CattleMax. URL https://www.cattlemax.com/

CattlePro, n.d. CattlePro.com | Beef Cattle Records Management and Performance Analysis Software [WWW Document]. URL https://www.cattlepro.com/

CattleWorks, n.d. CattleWorks® [WWW Document]. CattleWorks®. URL https://cattleworks.com/

Cernek, P., Bollig, N., Anklam, K., Döpfer, D., 2020. Hot topic: Detecting digital dermatitis with computer vision. Journal of Dairy Science 103, 9110–9115. https://doi.org/10.3168/jds.2019-17478

Cha, E., Hertl, J.A., Bar, D., Gröhn, Y.T., 2010. The cost of different types of lameness in dairy cows calculated by dynamic programming. Preventive Veterinary Medicine 97, 1–8. https://doi.org/10.1016/j.prevetmed.2010.07.011

Chernick, M.R., 2011. Bootstrap Methods: A Guide for Practitioners and Researchers. John Wiley & Sons.

Connecterra, n.d. Connecterra - intelligent platform for the dairy industry [WWW Document]. Connecterra. URL https://connecterra.ai/

Correll, M., Albers, D., Franconeri, S., Gleicher, M., 2012. Comparing averages in time series data, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI ’12. Association for Computing Machinery, New York, NY, USA, pp. 1095–1104. https://doi.org/10.1145/2207676.2208556

CowManager, n.d. CowManager | Ahead of the Herd Together [WWW Document]. CowManager. URL https://www.cowmanager.com/

Cownect, n.d. Cownect. URL https://www.itk.fr/en/research-and-development-projects/cownect/

Cowton, J., Kyriazakis, I., Bacardit, J., 2019. Automated Individual Pig Localisation, Tracking and Behaviour Metric Extraction Using Deep Learning. IEEE Access 7, 108049–108060. https://doi.org/10.1109/ACCESS.2019.2933060

Cramer, G., Lissemore, K.D., Guard, C.L., Leslie, K.E., Kelton, D.F., 2009. Herd-level risk factors for seven different foot lesions in Ontario Holstein cattle housed in tie stalls or free stalls. Journal of Dairy Science 92, 1404–1411. https://doi.org/10.3168/jds.2008-1134

Cramer, G., Winders, T., Solano, L., Kleinschmit, D.H., 2018. Evaluation of agreement among digital dermatitis scoring methods in the milking parlor, pen, and hoof trimming chute. Journal of Dairy Science 101, 2406–2414. https://doi.org/10.3168/jds.2017-13755

DDW, n.d. DDW [WWW Document]. URL https://www.dairydatawarehouse.com/

DeLaval, n.d. DeLaval DelPro™ - DeLaval [WWW Document]. URL https://www.delaval.com/en-us/explore-our-farm-solutions/farm-management/delaval-delpro/

Döpfer, D., Holzhauer, M., Boven, M. van, 2012. The dynamics of digital dermatitis in populations of dairy cattle: Model-based estimates of transition rates and implications for control. The Veterinary Journal, Special Issue:Bovine Lameness 193, 648–653. https://doi.org/10.1016/j.tvjl.2012.06.047

Döpfer, D., Huurne, A.A.H.M. ter, Cornelisse, J.L., Asten, A.J.A.M. van, Koopmans, A., Meijer, F.A., Schukken, Y.H., Szakáll, I., Klee, W., Bosma, R.B., 1997. Histological and bacteriological evaluation of digital dermatitis in cattle, with special reference to spirochaetes and Campylobacter faecalis. Veterinary Record 140, 620–623. https://doi.org/10.1136/vr.140.24.620

Döpfer, D., Van Boven, M., Jong, D., 2004. A mathematical model for the dynamics of digital dermatitis in dairy cattle. Proceedings of the 13th ICPD.

DTS, n.d. Dairy Herd Management Software [WWW Document]. Herd-Man. URL https://www.dts.co.nz/products-and-services/dairy-automation/dairy-herd-management-software/

du Prel, J.-B., Hommel, G., Röhrig, B., Blettner, M., 2009. Confidence Interval or P-Value? Dtsch Arztebl Int 106, 335–339. https://doi.org/10.3238/arztebl.2009.0335

Evans, N.J., Brown, J.M., Demirkan, I., Murray, R.D., Vink, W.D., Blowey, R.W., Hart, C.A., Carter, S.D., 2008. Three unique groups of spirochetes isolated from digital dermatitis lesions in UK cattle. Veterinary Microbiology 130, 141–150. https://doi.org/10.1016/j.vetmic.2007.12.019

Evans, N.J., Murray, R.D., Carter, S.D., 2016. Bovine digital dermatitis: Current concepts from laboratory to farm. The Veterinary Journal 211, 3–13. https://doi.org/10.1016/j.tvjl.2015.10.028

Farmbrite, n.d. Farm Management Software for your whole farm [WWW Document]. Farmbrite. URL https://www.farmbrite.com

FarmLogic, n.d. FarmLogic. FarmLogic. URL https://farmlogic.com/

Gjergji, M., de Moraes Weber, V., Otávio Campos Silva, L., da Costa Gomes, R., Luís Alves Campos de Araújo, T., Pistori, H., Alvarez, M., 2020. Deep Learning Techniques for Beef Cattle Body Weight Prediction, in: 2020 International Joint Conference on Neural Networks (IJCNN). Presented at the 2020 International Joint Conference on Neural Networks (IJCNN), pp. 1–8. https://doi.org/10.1109/IJCNN48605.2020.9207624

Gomez, A., Anklam, K.S., Cook, N.B., Rieman, J., Dunbar, K.A., Cooley, K.E., Socha, M.T., Döpfer, D., 2014. Immune response against Treponema spp. and ELISA detection of digital dermatitis. Journal of Dairy Science 97, 4864–4875. https://doi.org/10.3168/jds.2013-7616

Gomez, A., Cook, N.B., Bernardoni, N.D., Rieman, J., Dusick, A.F., Hartshorn, R., Socha, M.T., Read, D.H., Döpfer, D., 2012. An experimental infection model to induce digital dermatitis infection in cattle. Journal of Dairy Science 95, 1821–1830. https://doi.org/10.3168/jds.2011-4754

Gomez, A., Cook, N.B., Socha, M.T., Döpfer, D., 2015. First-lactation performance in cows affected by digital dermatitis during the rearing period. Journal of Dairy Science 98, 4487–4498. https://doi.org/10.3168/jds.2014-9041

Gonçalves, J.L., de Campos, J.L., Steinberger, A.J., Safdar, N., Kates, A., Sethi, A., Shutske, J., Suen, G., Goldberg, T., Cue, R.I., Ruegg, P.L., 2022. Incidence and Treatments of Bovine Mastitis and Other Diseases on 37 Dairy Farms in Wisconsin. Pathogens 11, 1282. https://doi.org/10.3390/pathogens11111282

Henderson, A.R., 2005. The bootstrap: A technique for data-driven statistics. Using computer-intensive analyses to explore experimental data. Clinica Chimica Acta 359, 1–26. https://doi.org/10.1016/j.cccn.2005.04.002

HerdApp, n.d. HerdApp [WWW Document]. URL https://app.agrinet.ie/HerdApp/

Herdlync, n.d. Farmwizard [WWW Document]. URL https://www.farmwizard.com/Account/Login

Herdlync, n.d. Herdlync [WWW Document]. Herdlync. URL https://herdlync.com/

Herdmaster, n.d. Herdmaster [WWW Document]. URL https://app.agrinet.ie/HerdMaster/#/F_Start

Herdtrax, n.d. Herdtrax | TELUS Agriculture & Consumer Goods [WWW Document]. TELUS Agriculture & Consumer Goods. URL https://www.telus.com/agcg/blog-resources/herdtrax

Herdwatch, n.d. Herdwatch Farming App for Calf Registrations and Farm Management Software [WWW Document]. Herdwatch.com. URL https://herdwatch.com/

Hietala, J., 2022. How To Build a YOLOv5 Object Detection App on iOS [WWW Document]. Medium. URL https://betterprogramming.pub/how-to-build-a-yolov5-object-detection-app-on-ios-39c8c77dfe58 (accessed 7.31.23).

hietalajulius, 2022. ios-demo-app/ObjectDetection-CoreML [WWW Document]. GitHub. URL https://github.com/hietalajulius/ios-demo-app/tree/main/ObjectDetection-CoreML

Holzhauer, M., Bartels, C.J.M., Döpfer, D., van Schaik, G., 2008. Clinical course of digital dermatitis lesions in an endemically infected herd without preventive herd strategies. The Veterinary Journal 177, 222–230. https://doi.org/10.1016/j.tvjl.2007.05.004

Holzhauer, M., Hardenberg, C., Bartels, C.J.M., Frankena, K., 2006. Herd- and Cow-Level Prevalence of Digital Dermatitis in The Netherlands and Associated Risk Factors. Journal of Dairy Science 89, 580–588. https://doi.org/10.3168/jds.S0022-0302(06)72121-X

Humphry, R.W., Cameron, A., Gunn, G.J., 2004. A practical approach to calculate sample size for herd prevalence surveys. Preventive Veterinary Medicine 65, 173–188. https://doi.org/10.1016/j.prevetmed.2004.07.003

Jacobs, C., Orsel, K., Morabito, E., Barkema, H., 2015. Pen walks as a tool to determine the prevalence of digital dermatitis in young stock on Alberta dairy farms, in: Proc. Of 18th International Symposium/10th International Conference on Lameness in Ruminants. p. 129.

Jiang, B., Wu, Q., Yin, X., Wu, D., Song, H., He, D., 2019. FLYOLOv3 deep learning for key parts of dairy cow body detection. Computers and Electronics in Agriculture 166, 104982. https://doi.org/10.1016/j.compag.2019.104982

Kang, X., Zhang, X.D., Liu, G., 2020. Accurate detection of lameness in dairy cattle with computer vision: A new and individualized detection strategy based on the analysis of the supporting phase. Journal of Dairy Science 103, 10628–10638. https://doi.org/10.3168/jds.2020-18288

Keogh, J., n.d. Livestock Monitoring Systems [WWW Document]. Moocall. URL https://www.moocall.com/

Klitgaard, K., Boye, M., Capion, N., Jensen, T.K., 2008. Evidence of Multiple Treponema Phylotypes Involved in Bovine Digital Dermatitis as Shown by 16S rRNA Gene Analysis and Fluorescence In Situ Hybridization. Journal of Clinical Microbiology 46, 3012–3020. https://doi.org/10.1128/jcm.00670-08

Krull, A.C., Shearer, J.K., Gorden, P.J., Scott, H.M., Plummer, P.J., 2016. Digital dermatitis: Natural lesion progression and regression in Holstein dairy cattle over 3 years. Journal of Dairy Science 99, 3718–3731. https://doi.org/10.3168/jds.2015-10535

Landis, J.R., Koch, G.G., 1977. The Measurement of Observer Agreement for Categorical Data. Biometrics 33, 159–174. https://doi.org/10.2307/2529310

Lee, S., Ahn, H., Seo, J., Chung, Y., Park, D., Pan, S., 2019. Practical Monitoring of Undergrown Pigs for IoT-Based Large-Scale Smart Farm. IEEE Access 7, 173796–173810. https://doi.org/10.1109/ACCESS.2019.2955761

Lely, 2016. T4C management system - Lely [WWW Document]. URL https://www.lely.com/farming-insights/t4c-management-system/

Mahmud, M.S., Zahid, A., Das, A.K., Muzammil, M., Khan, M.U., 2021. A systematic literature review on deep learning applications for precision cattle farming. Computers and Electronics in Agriculture 187, 106313. https://doi.org/10.1016/j.compag.2021.106313

Moore, E.K., Kriesberg, A., Schroeder, S., Geil, K., Haugen, I., Barford, C., Johns, E.M., Arthur, D., Sheffield, M., Ritchie, S.M., Jackson, C., Parr, C., 2022. Agricultural data management and sharing: Best practices and case study. Agronomy Journal 114, 2624–2634. https://doi.org/10.1002/agj2.20639

Naing, L., Winn, T., Rusli, B., 2006. Practical issues in calculating the sample size for prevalence studies. Archives of orofacial Sciences 1, 9–14.

Otto, S.P., Day, T., 2007. A Biologist's Guide to Mathematical Modeling in Ecology and Evolution.

Palmer, M.A., O'Connell, N.E., 2015. Digital Dermatitis in Dairy Cows: A Review of Risk Factors and Potential Sources of Between-Animal Variation in Susceptibility. Animals 5, 512–535. https://doi.org/10.3390/ani5030369

Peltjes, B., n.d. Dairylive. UNIFORM-Agri. URL https://www.uniform-agri.com/dairylive/

Psota, E.T., Luc, E.K., Pighetti, G.M., Schneider, L.G., Trout Fryxell, R.T., Keele, J.W., Kuehn, L.A., 2021. Development and validation of a neural network for the automated detection of horn flies on cattle. Computers and Electronics in Agriculture 180, 105927. https://doi.org/10.1016/j.compag.2020.105927

Psota, E.T., Mittek, M., Pérez, L.C., Schmidt, T., Mote, B., 2019. Multi-Pig Part Detection and Association with a Fully-Convolutional Network. Sensors 19, 852. https://doi.org/10.3390/s19040852

R Core Team, n.d. R: The R Project for Statistical Computing [WWW Document]. URL https://www.r-project.org/ (accessed 11.16.23).

Read, D.H., Walker, R.L., 1998. Papillomatous Digital Dermatitis (Footwarts) in California Dairy Cattle: Clinical and Gross Pathologic Findings. J VET Diagn Invest 10, 67–76. https://doi.org/10.1177/104063879801000112

Relun, A., Guatteo, R., Roussel, P., Bareille, N., 2011. A simple method to score digital dermatitis in dairy cows in the milking parlor. Journal of Dairy Science 94, 5424–5434. https://doi.org/10.3168/jds.2010-4054

Rodriguez-Lainz, A., Melendez-Retamal, P., Hird, D., 1996. Prevalence of papillomatous digital dermatitis in dairy herds in Chile, in: Proceedings of 77th Annual Meeting of Conference of Research Workers in Animal Diseases (Abstract), Chicago.

Rodriguez-Lainz, A., Melendez-Retamal, P., Hird, D.W., Read, D.H., Walker, R.L., 1999. Farm- and host-level risk factors for papillomatous digital dermatitis in Chilean dairy cattle. Preventive Veterinary Medicine 42, 87–97. https://doi.org/10.1016/S0167-5877(99)00067-7

Schöpke, K., Gomez, A., Dunbar, K.A., Swalve, H.H., Döpfer, D., 2015. Investigating the genetic background of bovine digital dermatitis using improved definitions of clinical status. Journal of Dairy Science 98, 8164–8174. https://doi.org/10.3168/jds.2015-9485

Seo, J., Ahn, H., Kim, D., Lee, S., Chung, Y., Park, D., 2020. EmbeddedPigDet—Fast and Accurate Pig Detection for Embedded Board Implementations. Applied Sciences 10, 2878. https://doi.org/10.3390/app10082878

SMARTBOW, n.d. SMARTBOW: Advanced Dairy Cow Monitoring [WWW Document]. URL https://www.smartbow.com/

smaXtec, n.d. smaXtec [WWW Document]. smaXtec. URL https://smaxtec.com/en/

Somers, J.G.C.J., Frankena, K., Noordhuizen-Stassen, E.N., Metz, J.H.M., 2005. Risk factors for digital dermatitis in dairy cows kept in cubicle houses in The Netherlands. Preventive Veterinary Medicine 71, 11–21. https://doi.org/10.1016/j.prevetmed.2005.05.002

StockManager, n.d. StockManager [WWW Document]. StockManager. URL https://www.stockmanager.co.nz

Tadich, N., Flor, E., Green, L., 2010. Associations between hoof lesions and locomotion score in 1098 unsound dairy cows. The Veterinary Journal 184, 60–65. https://doi.org/10.1016/j.tvjl.2009.01.005

Thrusfield, M., 2018. Veterinary Epidemiology. John Wiley & Sons.

Tremblay, M., Bennett, T., Döpfer, D., 2016. The DD Check App for prevention and control of digital dermatitis in dairy herds. Preventive Veterinary Medicine 132, 1–13. https://doi.org/10.1016/j.prevetmed.2016.07.016

UNIFORM-Agri, n.d. DairyLive Anywhere [WWW Document]. URL https://www.dairylive.com/dl/DairyLiveAnywhere.asp

VAS, n.d. Dairy Herd Management Program - DairyComp. VAS. URL https://vas.com/dairycomp/

Viera, A.J., Garrett, J.M., 2005. Understanding interobserver agreement: the kappa statistic. Fam Med 37, 360–363.

Weber, F. de L., Weber, V.A. de M., Menezes, G.V., Oliveira Junior, A. da S., Alves, D.A., de Oliveira, M.V.M., Matsubara, E.T., Pistori, H., Abreu, U.G.P. de, 2020. Recognition of Pantaneira cattle breed using computer vision and convolutional neural networks. Computers and Electronics in Agriculture 175, 105548. https://doi.org/10.1016/j.compag.2020.105548

Wells, S.J., Garber, L.P., Wagner, B.A., 1999. Papillomatous digital dermatitis and associated risk factors in US dairy herds. Preventive Veterinary Medicine 38, 11–24. https://doi.org/10.1016/S0167-5877(98)00132-9

Wenz, J.R., Giebel, S.K., 2012. Retrospective evaluation of health event data recording on 50 dairies using Dairy Comp 305. J Dairy Sci 95, 4699–4706. https://doi.org/10.3168/jds.2011-5312

Xu, B., Wang, W., Falzon, G., Kwan, P., Guo, L., Chen, G., Tait, A., Schneider, D., 2020. Automated cattle counting using Mask R-CNN in quadcopter vision system. Computers and Electronics in Agriculture 171, 105300. https://doi.org/10.1016/j.compag.2020.105300

Xudong, Z., Xi, K., Ningning, F., Gang, L., 2020. Automatic recognition of dairy cow mastitis from thermal images by a deep learning detector. Computers and Electronics in Agriculture 178, 105754. https://doi.org/10.1016/j.compag.2020.105754

Yano, T., Moe, K.K., Yamazaki, K., Ooka, T., Hayashi, T., Misawa, N., 2010. Identification of candidate pathogens of papillomatous digital dermatitis in dairy cattle from quantitative 16S rRNA clonal analysis. Veterinary Microbiology 143, 352–362. https://doi.org/10.1016/j.vetmic.2009.12.009

Zinicola, M., Lima, F., Lima, S., Machado, V., Gomez, M., Döpfer, D., Guard, C., Bicalho, R., 2015. Altered Microbiomes in Bovine Digital Dermatitis Lesions, and the Gut as a Pathogen Reservoir. PLOS ONE 10, e0120504. https://doi.org/10.1371/journal.pone.0120504

TABLES AND FIGURES

Figure 7.1 The DDCheckPlus app user interface. The user can enter data including the herd name, herd code, scorer details, date, pen number, cow ID, foot information, M-stage lesion details, and indicators of chronicity. The application provides visual aids including images and detailed descriptions to assist the user in accurately scoring DD lesions using the M-stage system.

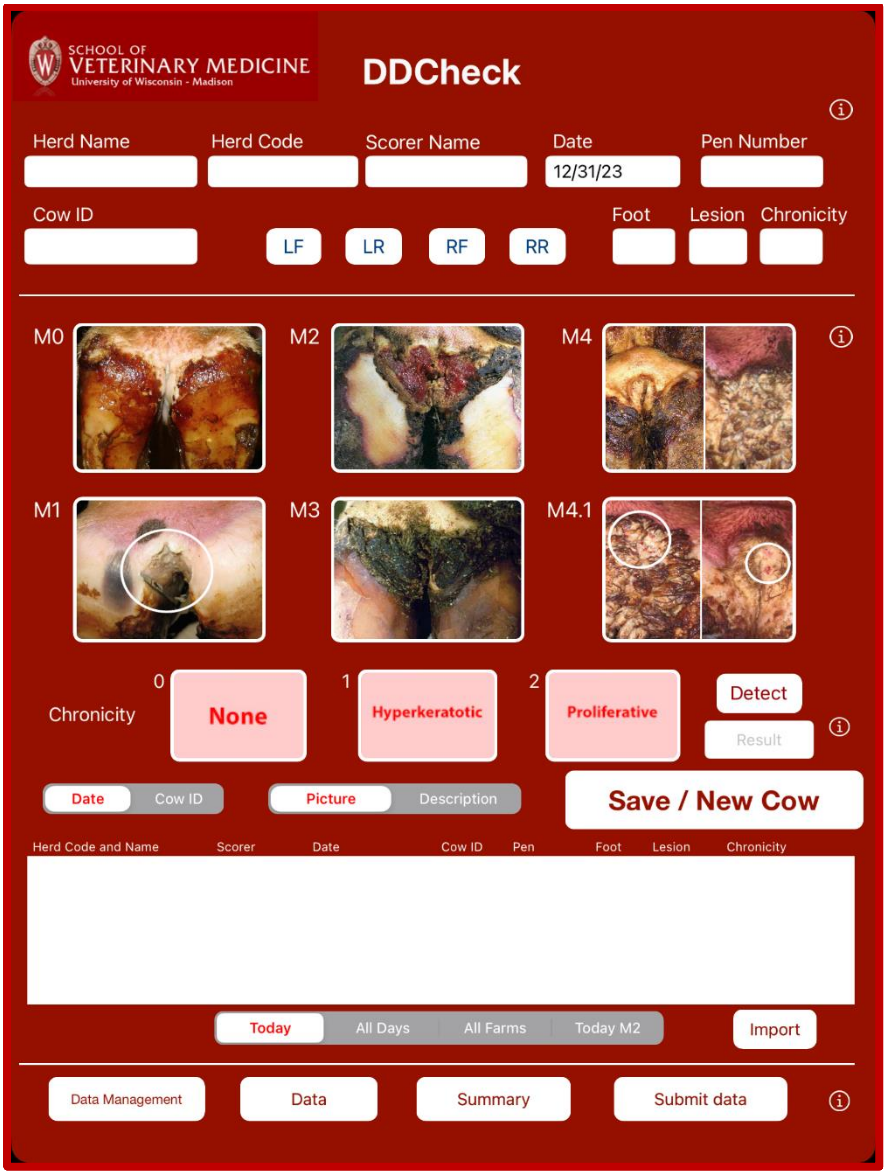

Figure 7.2 The DDCheckPlus app DD detection module. The lesion is circumscribed by a bounding box with class label for the respective M-stage classification and a confidence score for the prediction.

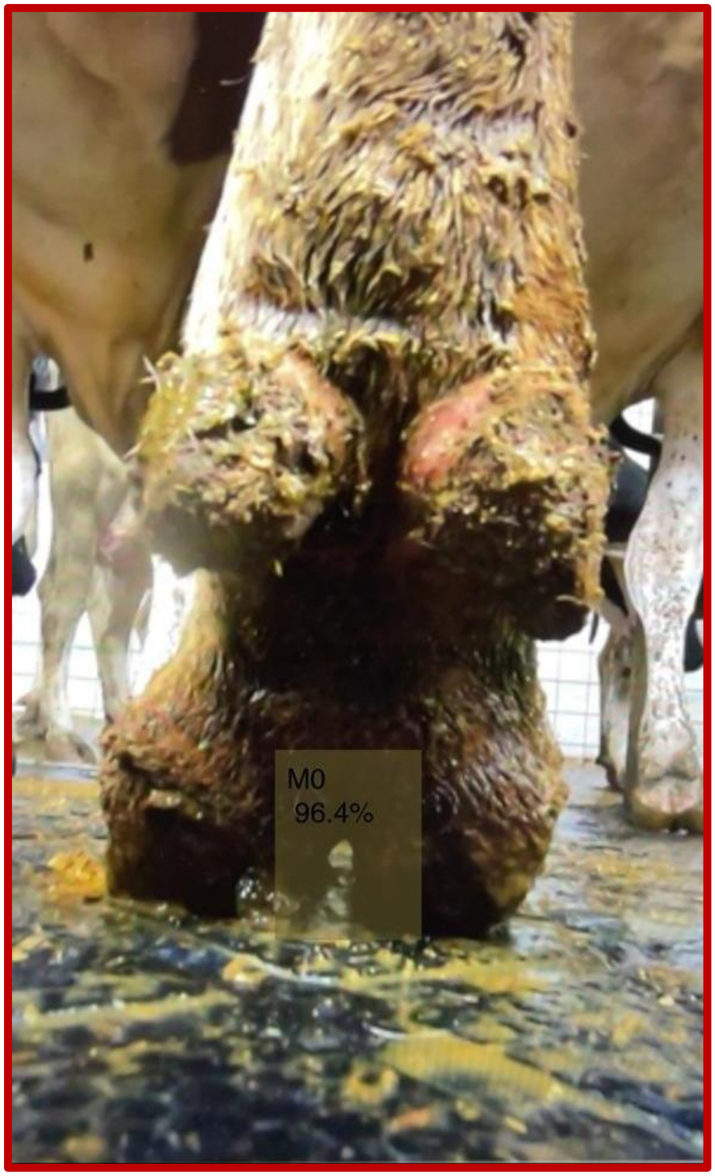

Figure 7.3 The DDCheckPlus app Dashboard tab before data upload. The user can upload the dataset, then input the number of bootstrap samples and the confidence level.

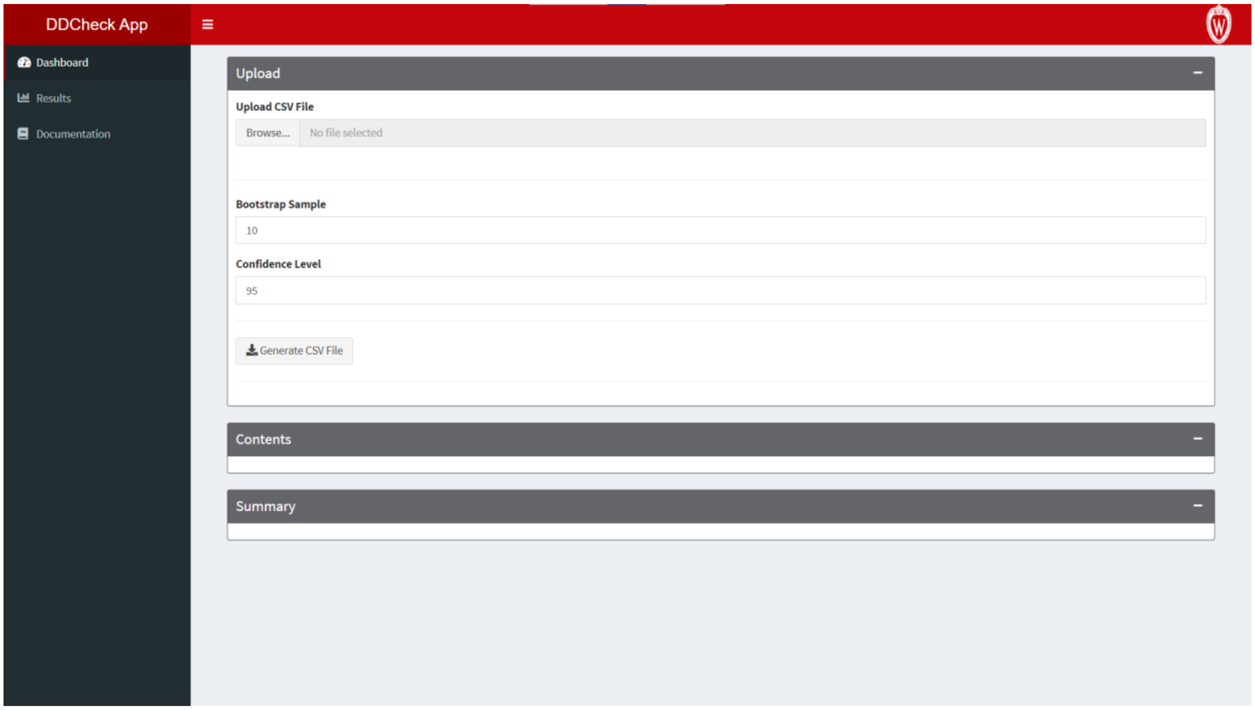

Figure 7.4 The DDCheckPlus app Dashboard tab after data upload. In the Upload panel, the user can select the week range using the slider, the variable of interest using the drop-down menu, and generate a CSV file with Cow Type using the action button (top-center). The Contents panel displays the first 100 rows of the dataset where the user can filter, search, and sort the data in the table (bottom-left). The Summary panel displays the data structure and data summary including the summary statistics, the number of missing values, the rate of complete values, and a low-level, unicode rendering of a histogram (bottom-right).

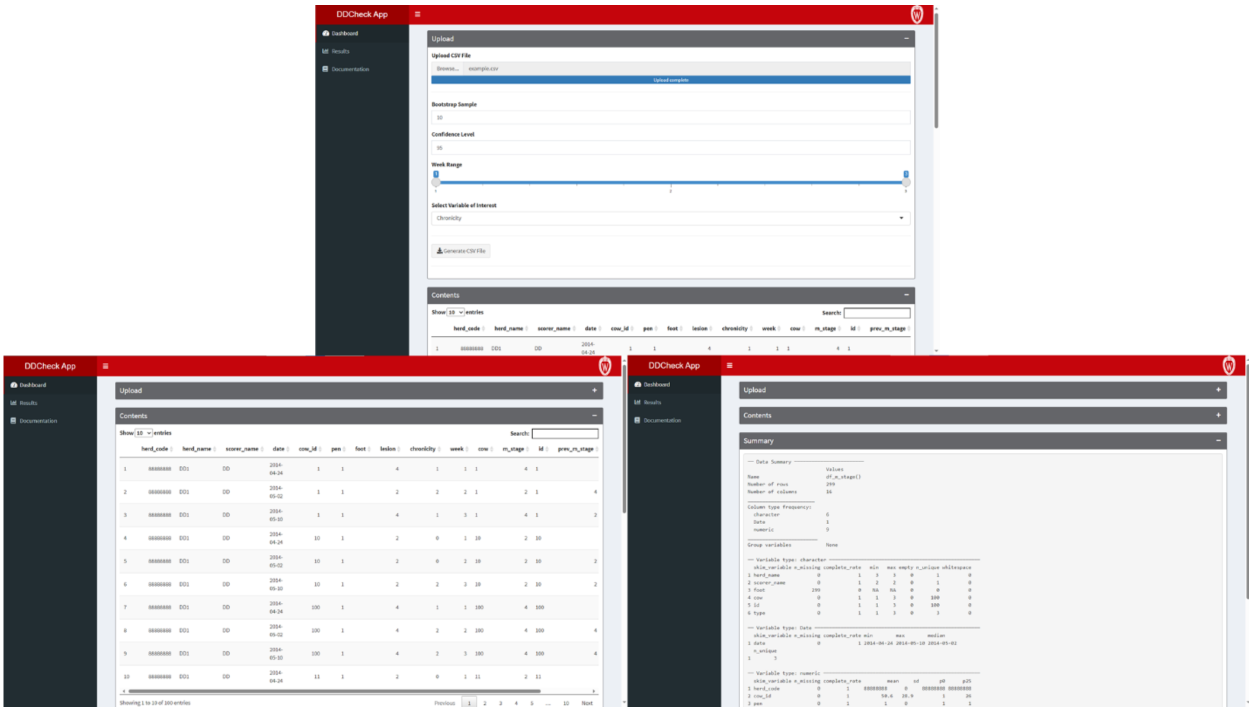

Figure 7.5 The DDCheckPlus app Results tab. The user can select the interactive box plot and corresponding interactive table to visualize and understand DD trends. The user can hover over the box plot for details on demand and filter, search, and sort the data in the table.

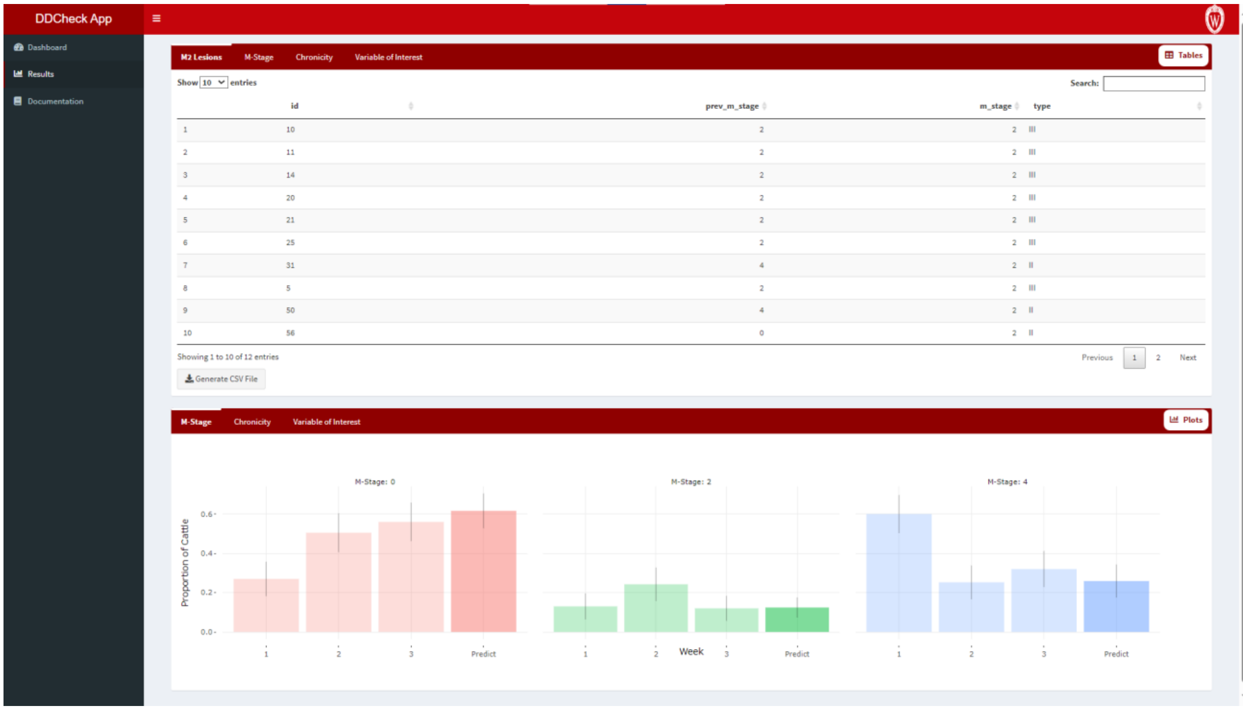

Figure 7.6 The DDCheckPlus app Documentation tab. The DDCheckPlus app provides instructions to generate a CSV file for analysis.

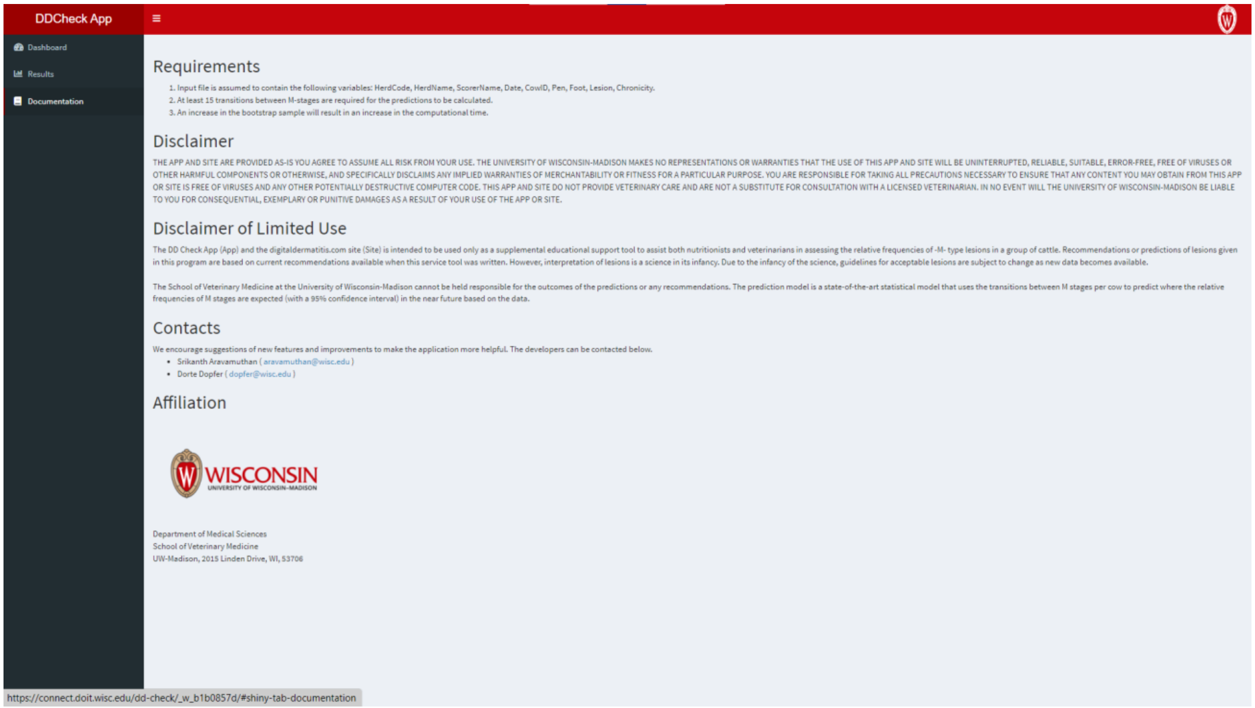

Figure 7.7 Accuracy of the five M-stages by the DD detection module on the DDCheckPlus app. Accuracy is measured as the number of correct predictions and calculated for each of the five M-stages: M0, M2, M2P, M4H, and M4P.

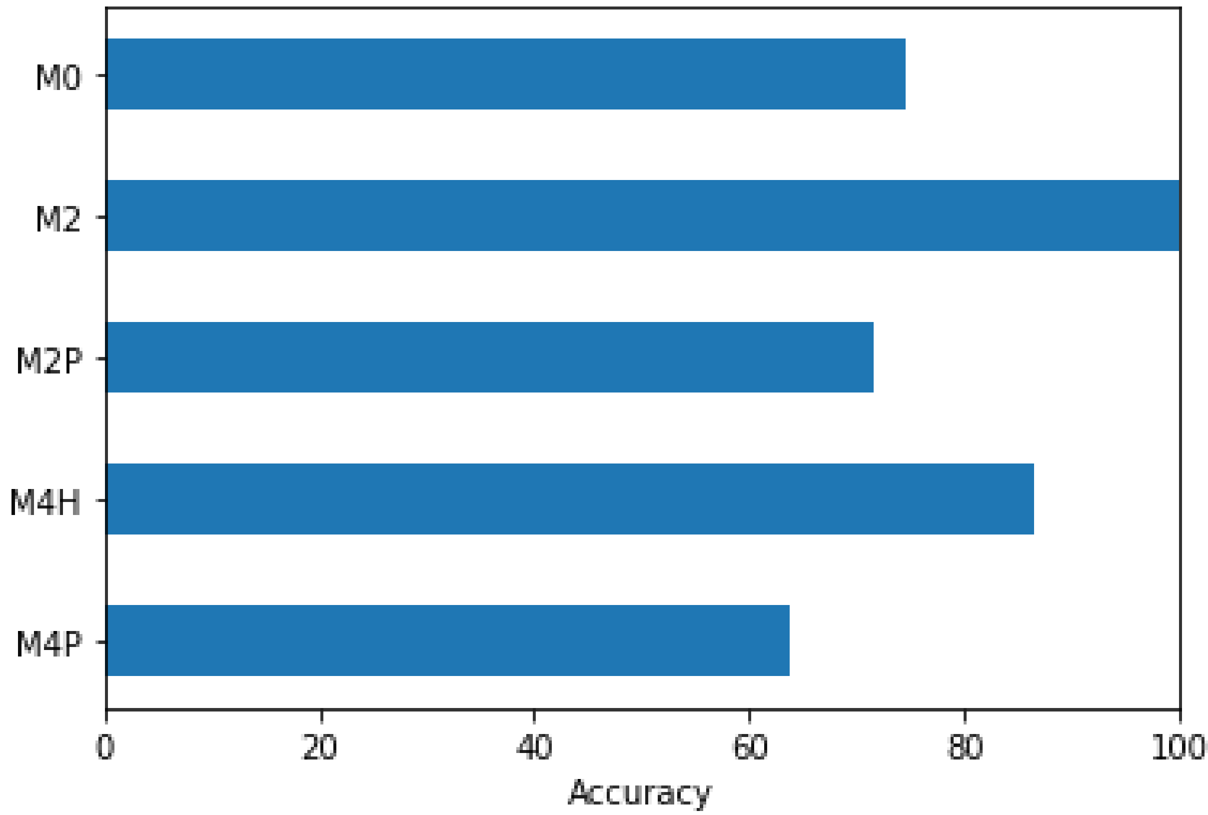

Figure 7.8 Faceted box plot of proportion of cattle over weeks by M-stage classification. The box represents the predicted prevalence and the line range represents the 95% confidence interval. The panels are faceted by M-stage classification where the first panel is M0 lesions, the second panel is M2 lesions, and the third panel is M4 lesions.

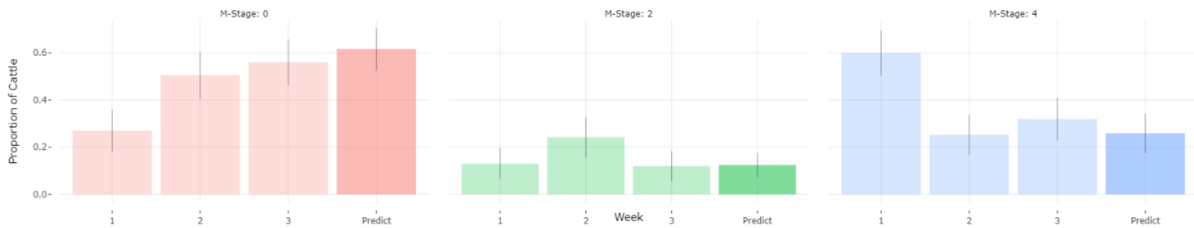

Figure 7.9 Faceted box plot of proportion of cattle over weeks by signs of chronicity. The box represents the predicted prevalence and the line range represents the 95% confidence interval. The panels are faceted by signs of chronicity where the first panel is no chronicity, the second panel is hyperkeratotic lesions, and the third panel is proliferative lesions.

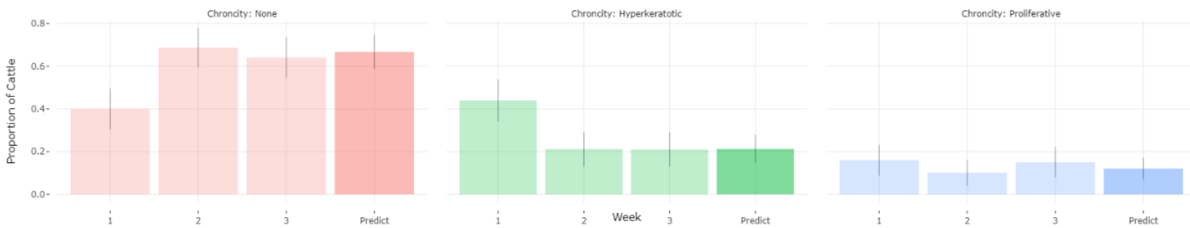

Figure 7.10 Faceted box plot of proportion of cattle over weeks by M-stage classification and signs of chronicity. The box represents the predicted prevalence and the line range represents the 95% confidence interval. The panels are faceted by M-stage classification and signs of chronicity where the first column is M0 lesions, the second column is M2 lesions, and the third column is M4 lesions while the first row is no chronicity, the second row is hyperkeratotic lesions, and the third row is proliferative lesions.

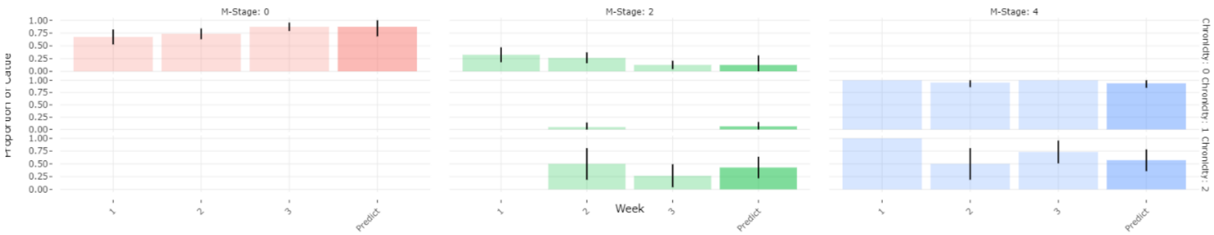

CHAPTER 8 REAL-TIME DETECTION OF DIGITAL DERMATITIS IN DAIRY COWS ON ANDROID AND IOS APPS USING COMPUTER VISION TECHNIQUES

ABSTRACT

The aim of the study is to deploy Android or iOS mobile applications models for real-time detection of digital dermatitis (DD) lesions in dairy cows using computer vision (CV). Early detection of DD lesions in dairy cows is critical for prompt treatment. Android and iOS apps can facilitate early detection of DD at both dairy and beef farms. Dairy farmers can implement preventive and control methods, including foot baths, topical treatment, hoof trimming, or quarantining cows affected by DD to prevent the spread of the disease. Early detection and prompt treatment decreases the severity of the condition, but also increases the overall productivity of cows. We applied transfer learning to DD image data for 5 classes, M0, M4H, M2, M2P, and M4P, on pretrained YOLOv5 model architecture using COCO-128 pretrained weights. The combination of localization loss, classification loss, and objectness loss was used for the optimization of prediction performance. This custom DD detection model was trained on 363 images of size 416x416 pixels and tested on 46 images. During model training, data were augmented to increase model robustness in different environments. The model was converted into TFLite format for Android devices and CoreML format for iOS devices. These models were deployed as Android and iOS applications for real-time DD detection using Android Studio and XCode software. Techniques such as quantization were implemented to improve inference speed in real-world settings. The DD models achieved an average mean average precision (mAP) of 0.95 on the test dataset. When tested in real-time, iOS devices resulted in Cohen's kappa value of 0.57 averaged across 5 classes denoting the moderate agreement of the model detection with human investigators. The Android device resulted in a Cohen's kappa value of 0.38 denoting fair

agreement between model and investigator. Combining M2 and M2P classes and M4H and M4P classes resulted in a Cohen's kappa value of 0.65 and 0.46, for Android and iOS devices respectively. For the two-class model i.e. presence an absence of DD lesion, a Cohen's kappa value of 0.74 and 0.65 was achieved for iOS and Android devices. The iOS app achieved an inference time of 20ms, compared to 57 ms on the Android app. Additionally, we deployed models on Ultralytics iOS and Android apps and our custom apps surpassed the Ultralytics apps in terms of Cohen's kappa and confidence score.

CHAPTER 9 WORKFLOW FOR IMPLEMENTING AN OBJECTION MODEL WITH CUSTOM DATA

Summary

Object detection is a fundamental task in the field of computer vision (CV). It is a valuable tool with industrial applications across diverse domains including medical imaging, agriculture, and other automated systems. State-of-the-art object detection algorithms harness the power of convolutional neural networks (CNNs) to achieve high accuracy and speed (Girshick, 2015; Girshick et al., 2014; He et al., 2017; LeCun et al., 2015; Liu et al., 2020, 2016; Redmon et al., 2016; Ren et al., 2015). Achieving high performance requires extensive training on a significant collection of labeled images (Goodfellow et al., 2015; Liu et al., 2020; Ouyang et al., 2017; Zhu et al., 2020). This task is energy-intensive and time-consuming where images are manually collected, annotated, and processed prior to model training and can be specific to model selection (Liu et al., 2020). Given the wide array of tasks, the ability to quickly curate domain-specific datasets and custom-trained models is a critical constraint (Liu et al., 2020). Hence, there exists a demand for instantaneous and automated processing approaches.

The workflow for object detection involves several steps to train and deploy a model that can accurately detect and localize objects of interest within images or videos. The steps can be further organized into three phases: pre-training, intra-training, and post-training (Figure 1). Initially, pre-training includes data collection, data preprocessing, and image annotation. Intra-training includes model selection, model training, and model evaluation. Lastly, post-training includes fine-tuning and optimization; inference and deployment; and monitoring and maintenance.

Pre-Training Phase

The pre-training phase is a pivotal process in the development of accurate and effective object detection models. The first step of data collection involves sourcing a wide array of images that represent or simulate real-world scenarios for model learning. This library of images containing various environmental conditions, lighting situations, angles, and scales ensures model versatility. An exhaustive dataset not only facilitates the robustness of the model, but also enhances its ability to detect objects accurately across different contexts (Goodfellow et al., 2015; Liu et al., 2020; Zhao et al., 2019).

However, data collection alone is insufficient. The data labeling is also needed for model learning. Image annotation involves labeling the images with relevant information that helps the model identify and differentiate the objects of interest. For object detection, this includes drawing bounding boxes around the objects and tagging it with the appropriate class label.

Image annotation can be a manual or semi-automated process, depending on the complexity of the task and the available tools (Liu et al., 2020). Manual annotation involves human annotators meticulously drawing bounding boxes or creating masks around each object. While this can be time-consuming, it ensures high accuracy and granularity (Liu et al., 2020). On the other hand, semi-automated approaches use AI-assisted tools to speed up the process by suggesting annotations that human annotators can then refine.

The quality of data annotation significantly influences the performance of the final model (Liu et al., 2020). Inaccurate or imprecise annotations can lead to false positives or negatives, affecting the reliability of the model to recognize objects (Liu et al., 2020). Rigorous quality control measures, such as regular reviews and inter-annotator agreement assessments, are critical to maintaining annotation accuracy.

Data collection and annotation are iterative processes. As the model evolves and encounters new scenarios, the dataset must also evolve to reflect these changes. Continuous data collection, augmentation, and annotation refinement ensure that the model remains proficient at detecting objects in evolving real-world contexts. The success of any object detection model relies on the rigorous of data collection and annotation, which form the cornerstone of accurate and reliable object recognition.

Data preprocessing is a critical phase in the development of object detection models that significantly impacts the performance and efficiency. It involves a series of steps to clean, enhance, and prepare raw data before model training. Effective data preprocessing not only improves the model accuracy, but also accelerates the training process and increases the generalizability in real-world situations.

Resizing is a fundamental step in data preprocessing for object detection. Images in the dataset may vary in size, and images are resized to a consistent dimension such that the model can uniformly process images. This step makes the data more manageable, but also makes the model less biased towards specific image sizes during training (Huang et al., 2017; Liu et al., 2016). Normalization is another step where pixel values are scaled to a common range. The model is less sensitive to variations in pixel intensities because of lighting or shading conditions (Huang et al., 2017; Liu et al., 2016). This standardization aids in faster convergence during training and assists in generalizability to unseen data (Huang et al., 2017; Liu et al., 2016).

Augmentation techniques are employed to artificially increase the diversity of the dataset when the available data is limited to prevent overfitting to the available data (Chatfield et al., 2014; Girshick, 2015; Girshick et al., 2014). Augmentation involves applying transformations to the images, such as rotations, reflections, brightness, blur, translations, cropping, and mosaics (Liu et al., 2020). This generates variations of the same instance, effectively expanding the dataset and making the model more robust to the variations in inference. Augmentation helps the model learn to detect objects under different conditions, such as varying lighting, viewpoints, and occlusions.

Annotation handling is an integral part of data preprocessing. The annotations should be adjusted and transformed to match any changes to the corresponding images during resizing or augmentation. Otherwise, the bounding boxes or masks would be inaccurate after the preprocessing steps. Additionally, noisy, blurry, obscured, obstructed or otherwise trivial data samples are filtered out. Lastly, class distribution can be balanced to prevent the model from being biased towards dominant or majority classes.

Intra-Training Phase

Selecting an appropriate model for object detection is a critical decision that significantly impacts the success of the computer vision project. With the rapid advancement of deep learning, there is a wide range of models available, each with its strengths, architectures, and performance characteristics. The choice of model depends on various factors including the specific requirements of the application, the complexity of the detection task, the available computational resources, and the expected accuracy.

One popular category of object detection models is the family of Faster Region-based Convolutional Neural Network (R-CNN) and its variants (Ren et al., 2015). Faster R-CNN introduced the concept of using a Region Proposal Network (RPN) to generate potential object proposals, followed by classifying and refining these proposals. Single Shot MultiBox Detector (SSD) is another popular model that emphasizes speed while maintaining accuracy (Liu et al., 2016). The SSD algorithm performs object detection at multiple scales using a set of predefined anchor boxes. You Only Look Once (YOLO) is another prominent model that has gained popularity because of its impressive speed and real-time performance (Redmon et al., 2016). The YOLO algorithm divides the image into a grid and directly predicts bounding boxes and class probabilities for each grid cell. It is the ideal choice for real-time applications requiring rapid detection and relatively high accuracy, such as video analysis and surveillance. The size of the model and the computational requirements should also align with the available resources (Diwan et al., 2023; Zaidi et al., 2022; Zhao et al., 2019). Larger models provide high accuracy, but demand more computational power and memory during training and inference. If deployed on resource-constrained devices, lightweight architectures such as SSD Lite, Tiny YOLOv4, and YOLOv5s may be appropriate (Bochkovskiy et al., 2020; Jocher, 2020; Liu et al., 2016; Redmon and Farhadi, 2018, 2017; Rodriguez-Conde et al., 2021).

Many pre-trained models are available that have been trained on large-scale datasets. Pre-trained models on large datasets such as ImageNet or MS COCO provide an initial starting point since general features have been learned and can be fine-tuned for domain-specific object detection tasks (Lin et al., 2014; Russakovsky et al., 2015). Fine-tuning models with a smaller dataset specific to the task can significantly speed up training and improve performance. Fine-tuning models using transfer learning can significantly expedite the training process and improve accuracy for specific tasks as well (Girshick et al., 2016; Sharif Razavian et al., 2014; Shin et al., 2016; Zhou et al., 2014). During transfer learning, the initial layers are usually frozen, preventing them from being updated during training. Only the final layers are fine-tuned using a smaller, domain-specific dataset. Moreover, architectures designed for object detection e.g., bounding boxes may not be appropriate other tasks e.g., masks. Therefore, model selection should also take into account the specific task requirements. Model selection for object detection is a balance between accuracy, speed, and resource constraints. Careful consideration of the task requirements, real-time constraints, and available datasets is necessary to make an informed decision. Ultimately, conducting comparative analyses and evaluating model performance on representative data can help determine the best model for the object detection application.

Model training is involves teaching the selected model to accurately recognize and locate objects within images. Proper training requires well-preprocessed data, an appropriate loss function, and a careful balance of hyperparameters to achieve optimal results. Model training aims to minimize a defined loss function. This function quantifies the difference between the predicted bounding boxes and class probabilities and the ground-truth annotations. Commonly used loss functions include the mean squared error (MSE) for bounding box regression and the cross-entropy loss for class prediction (Krizhevsky et al., 2012; Ren et al., 2015; Zhao et al., 2019). However, object detection models involve more complex loss formulations that combine multiple components, such as localization loss, objectness loss, and confidence loss (Alexe et al., 2012; Rahtu et al., 2011; Redmon et al., 2016; Ren et al., 2015; Wu et al., 2020; Zaidi et al., 2022). Other commonly used loss functions for object detection include the intersection over union (IoU) loss, focal loss, and smooth L1 loss (Girshick, 2015; Lin et al., 2017; Wu et al., 2020; Yu et al., 2016; Zaidi et al., 2022).

Model training is computationally intensive, often requiring powerful GPUs or cloud resources. Training time varies depending on factors like dataset size, model complexity, and hardware. Hyperparameters, such as learning rate, batch size, and regularization strength, can significantly impact the training process and the final model performance (Huang et al., 2017; Liu et al., 2020; Pandiya et al., 2020; Peng et al., 2018; Wu et al., 2020; Zhao et al., 2019). Grid search or random search can be used to explore various hyperparameter combinations and identify the ones that yield the best results (Huang et al., 2017; Pandiya et al., 2020). Iterative training is common, where the model is trained over multiple epochs while updating the weights over differing learning rates (He et al., 2016; Pandiya et al., 2020; Shmelkov et al., 2017; Wu et al., 2020; Zhao et al., 2019). During training, the model performance is evaluated on a validation dataset. Early stopping is employed if the validation performance plateaus or starts to decrease, preventing the model from overfitting to the training data (Arulprakash and Aruldoss, 2022; Goodfellow et al., 2016; Prechelt, 1998). Techniques including learning rate scheduling and early stopping can improve training stability and efficiency. Metrics, such as mean average precision (mAP), precision, recall, and average precision, provide insights into the model accuracy and robustness (Liu et al., 2020; Zaidi et al., 2022). The training process can be fine-tuned based on validation results to achieve the required detection accuracy.

Model evaluation is used to assess the performance and effectiveness of object detection models. It involves measuring the model accuracy, precision, recall, and other metrics to determine the ability to detect and localize objects within images. A comprehensive evaluation provides insights into the strengths, weaknesses, and areas for improvement.

As previously mentioned, one commonly used metric for object detection is mAP. It combines precision and recall across different levels of confidence thresholds to provide an overall measure of model accuracy (Liu et al., 2020). The mAP considers not only if an object is detected but also the magnitude of alignment between the detected bounding box and the ground truth (Liu et al., 2020). Higher mAP implies better detection performance. Precision and recall are fundamental metrics for object detection evaluation. Precision measures the ratio of correctly detected objects to all detected objects, while recall measures the ratio of correctly detected objects to all actual objects (Borji et al., 2019; Liu et al., 2020). Achieving a balance between precision and recall is crucial. A high-precision model minimizes false positives, while a high-recall model minimizes false negatives (Borji et al., 2019; Liu et al., 2020). F1 score is another metric that combines precision and recall, providing a single value that reflects the overall performance (Borji et al., 2019). It is particularly useful when there is an imbalance between positive and negative samples in the dataset. Additionally, the receiver operating characteristic (ROC) curve and area under the curve (AUC) provide insights into the model's trade-off between true positive rate and false positive rate as well (Borji et al., 2019). Visualizing the model predictions is essential for understanding its behavior. Visualization tools can overlay predicted bounding boxes on images and color-code them based on the class labels or confidence scores. This helps identify cases where the model struggles and highlights potential areas for refinement.

Evaluation should be performed on a separate validation test dataset that the model has not previously seen during training or validation. Cross-validation techniques can also be employed to assess the model performance across multiple subsets of the data. Comparing the model performance against baseline models or existing state-of-the-art models is valuable for understanding its relative strengths and weaknesses. Regularly monitoring and updating the evaluation as the model is refined or deployed in different environments is crucial to maintain its performance over time. In summary, model evaluation in object detection involves assessing metrics such as mAP, precision, recall, F1 score, and visualizations to gain insights into the model performance. It is a crucial step in determining how well the model can detect and localize objects and driving improvements in its design and development.

Post-Training Phase

Model optimization is used to refine the model performance and efficiency. Optimization aims to enhance the model speed, accuracy, and resource utilization while minimizing computational complexity and memory requirements. This process is essential such that object detection models can be deployed effectively in real-world applications, especially on resource-constrained devices, including edge devices and cloud servers.

Accelerating inference speed is frequently the primary optimization goal. Optimization techniques such as quantization and pruning contribute to a lighter architecture and faster model (Diwan et al., 2023; Zaidi et al., 2022). Quantization is used to reduce the precision of the model's weights and activations from floating-point values to lower-bit representations, such as 8-bit integers (Courbariaux et al., 2015; Han et al., 2016). This significantly reduces the computational complexity and memory requirements without significantly compromising accuracy. Such models run faster and consume less power, making it ideal for model deployment on resource-constrained devices. Architectural optimization involves modifying the current model design where efficient model architectures are designed with respect to optimization. This includes customization or pruning to reduce the number of parameters, resulting in a lighter architecture and faster model (Diwan et al., 2023; Zaidi et al., 2022). Pruning is a technique to remove weights or neurons from the model (Chen et al., 2021; Denker, 1989; Diwan et al., 2023; Hassibi et al., 1993; Zaidi et al., 2022). This can be achieved through weight pruning to remove weights with low magnitudes. Structured pruning involves removing entire channels or layers, leading to more efficient models. Pruning combined with quantization provide an effective solution for achieving lightweight, but accurate object detection models (Diwan et al., 2023; Zaidi et al., 2022).

Hardware-specific optimizations are critical, specifically for model deployment on edge devices or embedded systems (Feng et al., 2019; Reuther et al., 2019; Rodriguez-Conde et al., 2021). Model quantization and weight sharing leverage specialized hardware accelerators such as GPUs, TPUs, VPUs, NPUs or dedicated neural network inference chips (Feng et al., 2019; Reuther et al., 2019; Rodriguez-Conde et al., 2021). Additionally, techniques like model parallelism and multi-threading can take advantage of parallel processing capabilities to further speed up inference (Cortés Gallardo Medina et al., 2021; Garland et al., 2008; Hosny and Salah, 2023; Murthy et al., 2020; Stone et al., 2010; Thabet et al., 2014). Efficient post-processing techniques also contribute to model optimization. Non-Maximum Suppression (NMS) algorithms suppress overlapping bounding boxes to retain the highest confidence bounding box, reducing redundancy and improving performance. This can be optimized to reduce redundant computations and improve speed (Bodla et al., 2017; Hosang et al., 2017; Liu et al., 2020; Zou et al., 2023). Finally, optimizing data pipelines and input preprocessing can also contribute to faster inference. As previously stated, techniques such as batching and data prefetching help reduce processing overhead (He et al., 2015; Ioffe and Szegedy, 2015; Liu et al., 2020; Peng et al., 2018; Wu et al., 2020; Zaidi et al., 2022; Zhao et al., 2019). Additionally, resizing images to the optimal input size and applying proper normalization techniques can improve efficiency (Chin et al., 2019; Dai et al., 2016; Felzenszwalb et al., 2010; Girshick et al., 2014; He et al., 2016; Hu and Ramanan, 2017; Liu et al., 2020, 2017; Singh and Davis, 2018; Wu et al., 2020; Yang et al., 2016; Zhu et al., 2020; Zou et al., 2023). Tailoring the model size and optimization techniques to the target platform is important to balancing accuracy and speed.

Model optimization for object detection involves a combination of architectural design, quantization, pruning, hardware-specific optimizations, and efficient post-processing techniques. Efficient inference is a critical consideration, especially when deploying models on resource-constrained devices such as edge devices, smartphones, or embedded systems. Optimization techniques such as model quantization, pruning, and hardware-specific accelerators are employed to improve speed and reduce memory and computation requirements. These strategies collectively aim to produce lightweight, fast, and accurate object detection models such that it can be deployed efficiently across various platforms and applications.

Inference is the process of applying a trained model to new, unseen data to make predictions about the presence, location, and class of objects within images or videos. This stage is necessary for real-time application and real-world applicability, where the model needs to perform accurately and efficiently (Goodfellow et al., 2016). During inference, the trained object detection model takes an input image or video frame and the network computes predictions for bounding box coordinates, class labels, and confidence scores for each object detected (Goodfellow et al., 2016). The predictions are based on the features learned during the model training phase (Goodfellow et al., 2016). The bounding boxes outline the spatial extent of the objects, class labels indicate the type of object, and confidence scores indicate the confidence in the predictions (Goodfellow et al., 2016).

The output of the inference process is then visualized by overlaying the predicted bounding boxes, associated class labels, and corresponding confidence score on the input image. This visualization provides valuable insights into the model performance and helps users overall understanding of the model ability to identify and classify objects within the images. Note that model performance during inference directly depends on the quality of the training data, the effectiveness of the data preprocessing steps, and the model architecture to generalize to new, unseen data (Rodriguez-Conde et al., 2021; Zaidi et al., 2022; Zou et al., 2023). Inference can be performed on various platforms, from high-performance servers to edge devices like smartphones, embedded systems, and even drones (Rodriguez-Conde et al., 2021; Zaidi et al., 2022; Zou et al., 2023).

Deployment is the pivotal phase in the lifecycle of an object detection model, transitioning from development and testing to real-world applications to perform the task of identifying and localizing objects within images or video streams. Effective deployment involves a series of steps for the model to work seamlessly, efficiently, and reliably in the intended environment on new, unseen data. A well-executed deployment process is critical to achieving the desired outcomes and impact in various domains, such as surveillance, robotics, agriculture, and healthcare.

One of the key considerations in deployment is selecting the appropriate hardware platform. The choice of hardware depends on factors such as the desired speed of inference, resource constraints, and the target application. Depending on the specific use case, object detection models can be deployed on a range of devices, from cloud servers and edge devices to mobile phones and specialized hardware accelerators (Feng et al., 2019; Reuther et al., 2019; Rodriguez-Conde et al., 2021). Some applications may require high-performance GPUs for real-time processing, while others might focus on energy-efficient solutions like edge devices or specialized hardware accelerators (Feng et al., 2019; Reuther et al., 2019; Rodriguez-Conde et al., 2021; Yang et al., 2019; Zhu et al., 2020).

Software integration is another critical aspect of deployment. The object detection model needs to seamlessly integrate into the existing software infrastructure of the application. This may involve developing Application Programming Interfaces (APIs), libraries, or other interfaces that allow developers to easily integrate the model's predictions into their applications or systems (Garland et al., 2008; Harper et al., 2023; Kamath and Renuka, 2023; Murthy et al., 2020; NVIDIA, 2023a, 2023b, 2018, 2017, 2015, 2014; Stone et al., 2010). Compatibility with programming languages, frameworks, and platforms commonly used in the application domain is essential to ensure a smooth integration process (NVIDIA, 2023c, 2016).

Model updates and maintenance are continuous aspects of long-term deployment. Models are periodically retrained on new data to adapt to changes in object appearance or environmental conditions for improved accuracy (Ahmad and Rahimi, 2022; Cao et al., 2017; Girshick et al., 2014; Murthy et al., 2020; Rodriguez-Conde et al., 2021; Yao et al., 2020; Zhao et al., 2019; Zhu et al., 2020). It is important to establish a system for updating the deployed models without disrupting the application. Continuous monitoring of model performance and adapting to changing conditions help identify potential issues and maintain the effectiveness of the object detection system over time (Ahmad and Rahimi, 2022; Cao et al., 2017; Girshick et al., 2014; Rodriguez-Conde et al., 2021; Zhao et al., 2019). Real-world testing and validation are critical before a model is fully deployed (Liu et al., 2020; Rodriguez-Conde et al., 2021). Testing the model in diverse and challenging scenarios helps identify any potential issues and fine-tune the system for optimal performance (Liu et al., 2020; Rodriguez-Conde et al., 2021). It is important to simulate a wide range of scenarios that the model may encounter in an operational context for robustness and reliability (Kamath and Renuka, 2023; Murthy et al., 2020).

Ultimately, successful deployment of object detection models requires a holistic approach that include hardware selection, software integration, optimization, ongoing maintenance, and validation. By carefully addressing these aspects, organizations can leverage the power of object detection to enhance a wide range of applications, from surveillance and healthcare to manufacturing and autonomous systems.

REFERENCES

Ahmad, H.M., Rahimi, A., 2022. Deep learning methods for object detection in smart manufacturing: A survey. Journal of Manufacturing Systems 64, 181–196. https://doi.org/10.1016/j.jmsy.2022.06.011

Alexe, B., Deselaers, T., Ferrari, V., 2012. Measuring the Objectness of Image Windows. IEEE Transactions on Pattern Analysis and Machine Intelligence 34, 2189–2202. https://doi.org/10.1109/TPAMI.2012.28

Arulprakash, E., Aruldoss, M., 2022. A study on generic object detection with emphasis on future research directions. Journal of King Saud University - Computer and Information Sciences 34, 7347–7365. https://doi.org/10.1016/j.jksuci.2021.08.001

Bochkovskiy, A., Wang, C.-Y., Liao, H.-Y.M., 2020. YOLOv4: Optimal Speed and Accuracy of Object Detection. arXiv:2004.10934 [cs, eess].

Bodla, N., Singh, B., Chellappa, R., Davis, L.S., 2017. Soft-NMS -- Improving Object Detection With One Line of Code. Presented at the Proceedings of the IEEE International Conference on Computer Vision, pp. 5561–5569.

Borji, A., Cheng, M.-M., Hou, Q., Jiang, H., Li, J., 2019. Salient object detection: A survey. Comp. Visual Media 5, 117–150. https://doi.org/10.1007/s41095-019-0149-9

Cao, Z., Simon, T., Wei, S.-E., Sheikh, Y., 2017. Realtime Multi-Person 2D Pose Estimation Using Part Affinity Fields. Presented at the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 7291–7299.

Chatfield, K., Simonyan, K., Vedaldi, A., Zisserman, A., 2014. Return of the Devil in the Details: Delving Deep into Convolutional Nets. https://doi.org/10.48550/arXiv.1405.3531

Chen, K., Franko, K., Sang, R., 2021. Structured Model Pruning of Convolutional Networks on Tensor Processing Units. https://doi.org/10.48550/arXiv.2107.04191

Chin, T.-W., Ding, R., Marculescu, D., 2019. AdaScale: Towards Real-time Video Object Detection using Adaptive Scaling. Proceedings of Machine Learning and Systems 1, 431–441.

Cortés Gallardo Medina, E., Velazquez Espitia, V.M., Chípuli Silva, D., Fernández Ruiz de las Cuevas, S., Palacios Hirata, M., Zhu Chen, A., González González, J.Á., Bustamante-Bello, R., Moreno-García, C.F., 2021. Object Detection, Distributed Cloud Computing and Parallelization Techniques for Autonomous Driving Systems. Applied Sciences 11, 2925. https://doi.org/10.3390/app11072925

Courbariaux, M., Bengio, Y., David, J.-P., 2015. BinaryConnect: Training Deep Neural Networks with binary weights during propagations, in: Advances in Neural Information Processing Systems. Curran Associates, Inc.

Dai, J., Li, Y., He, K., Sun, J., 2016. R-FCN: Object Detection via Region-based Fully Convolutional Networks, in: Advances in Neural Information Processing Systems. Curran Associates, Inc.

Denker, J.S., 1989. Optimal brain damage. Advances in neural information processing systems 2, 598–605.

Diwan, T., Anirudh, G., Tembhurne, J.V., 2023. Object detection using YOLO: challenges, architectural successors, datasets and applications. Multimed Tools Appl 82, 9243–9275. https://doi.org/10.1007/s11042-022-13644-y

Felzenszwalb, P.F., Girshick, R.B., McAllester, D., Ramanan, D., 2010. Object Detection with Discriminatively Trained Part-Based Models. IEEE Transactions on Pattern Analysis and Machine Intelligence 32, 1627–1645. https://doi.org/10.1109/TPAMI.2009.167

Feng, X., Jiang, Y., Yang, X., Du, M., Li, X., 2019. Computer vision algorithms and hardware implementations: A survey. Integration 69, 309–320. https://doi.org/10.1016/j.vlsi.2019.07.005

Garland, M., Le Grand, S., Nickolls, J., Anderson, J., Hardwick, J., Morton, S., Phillips, E., Zhang, Y., Volkov, V., 2008. Parallel Computing Experiences with CUDA. IEEE Micro 28, 13–27. https://doi.org/10.1109/MM.2008.57

Girshick, R., 2015. Fast R-CNN. Presented at the Proceedings of the IEEE International Conference on Computer Vision, pp. 1440–1448.

Girshick, R., Donahue, J., Darrell, T., Malik, J., 2016. Region-Based Convolutional Networks for Accurate Object Detection and Segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence 38, 142–158. https://doi.org/10.1109/TPAMI.2015.2437384

Girshick, R., Donahue, J., Darrell, T., Malik, J., 2014. Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation. Presented at the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 580–587.

Goodfellow, I., Bengio, Y., Courville, A., 2016. Deep Learning, Illustrated edition. ed. The MIT Press, Cambridge, Massachusetts.

Goodfellow, I.J., Shlens, J., Szegedy, C., 2015. Explaining and Harnessing Adversarial Examples. https://doi.org/10.48550/arXiv.1412.6572

Han, S., Mao, H., Dally, W.J., 2016. Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding. https://doi.org/10.48550/arXiv.1510.00149

Harper, E., Majumdar, S., Kuchaiev, O., Jason, L., Zhang, Y., Bakhturina, E., Noroozi, V., Subramanian, S., Nithin, K., Jocelyn, H., Jia, F., Balam, J., Yang, X., Livne, M., Dong, Y., Naren, S., Ginsburg, B., 2023. NeMo: a toolkit for Conversational AI and Large Language Models.

Hassibi, B., Stork, D.G., Wolff, G.J., 1993. Optimal Brain Surgeon and general network pruning, in: IEEE International Conference on Neural Networks. Presented at the IEEE International Conference on Neural Networks, pp. 293–299 vol.1. https://doi.org/10.1109/ICNN.1993.298572

He, K., Gkioxari, G., Dollar, P., Girshick, R., 2017. Mask R-CNN. Presented at the Proceedings of the IEEE International Conference on Computer Vision, pp. 2961–2969.

He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep Residual Learning for Image Recognition. Presented at the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778.

He, K., Zhang, X., Ren, S., Sun, J., 2015. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification. Presented at the Proceedings of the IEEE International Conference on Computer Vision, pp. 1026–1034.

Hosang, J., Benenson, R., Schiele, B., 2017. Learning Non-Maximum Suppression. Presented at the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4507–4515.

Hosny, K.M., Salah, A., 2023. Recent Advances in Computer Vision Applications Using Parallel Processing. Springer Nature.

Hu, P., Ramanan, D., 2017. Finding Tiny Faces. Presented at the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 951–959.

Huang, J., Rathod, V., Sun, C., Zhu, M., Korattikara, A., Fathi, A., Fischer, I., Wojna, Z., Song, Y., Guadarrama, S., Murphy, K., 2017. Speed/Accuracy Trade-Offs for Modern Convolutional Object Detectors. Presented at the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 7310–7311.

Ioffe, S., Szegedy, C., 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift, in: Proceedings of the 32nd International Conference on Machine Learning. Presented at the International Conference on Machine Learning, PMLR, pp. 448–456.

Jocher, G., 2020. YOLOv5 by Ultralytics. https://doi.org/10.5281/zenodo.3908559

Kamath, V., Renuka, A., 2023. Deep learning based object detection for resource constrained devices: Systematic review, future trends and challenges ahead. Neurocomputing 531, 34–60. https://doi.org/10.1016/j.neucom.2023.02.006

Krizhevsky, A., Sutskever, I., Hinton, G.E., 2012. ImageNet Classification with Deep Convolutional Neural Networks, in: Advances in Neural Information Processing Systems. Curran Associates, Inc.

LeCun, Y., Bengio, Y., Hinton, G., 2015. Deep learning. Nature 521, 436–444. https://doi.org/10.1038/nature14539

Lin, T.-Y., Goyal, P., Girshick, R., He, K., Dollar, P., 2017. Focal Loss for Dense Object Detection. Presented at the Proceedings of the IEEE International Conference on Computer Vision, pp. 2980–2988.

Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L., 2014. Microsoft COCO: Common Objects in Context, in: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (Eds.), Computer Vision – ECCV 2014, Lecture Notes in Computer Science. Springer International Publishing, Cham, pp. 740–755. https://doi.org/10.1007/978-3-319-10602-1_48

Liu, L., Ouyang, W., Wang, X., Fieguth, P., Chen, J., Liu, X., Pietikäinen, M., 2020. Deep Learning for Generic Object Detection: A Survey. Int J Comput Vis 128, 261–318. https://doi.org/10.1007/s11263-019-01247-4

Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.-Y., Berg, A.C., 2016. SSD: Single Shot MultiBox Detector, in: Leibe, B., Matas, J., Sebe, N., Welling, M. (Eds.), Computer Vision – ECCV 2016, Lecture Notes in Computer Science. Springer International Publishing, Cham, pp. 21–37. https://doi.org/10.1007/978-3-319-46448-0_2

Liu, Y., Li, H., Yan, J., Wei, F., Wang, X., Tang, X., 2017. Recurrent Scale Approximation for Object Detection in CNN. Presented at the Proceedings of the IEEE International Conference on Computer Vision, pp. 571–579.

Murthy, C.B., Hashmi, M.F., Bokde, N.D., Geem, Z.W., 2020. Investigations of Object Detection in Images/Videos Using Various Deep Learning Techniques and Embedded Platforms—A Comprehensive Review. Applied Sciences 10, 3280. https://doi.org/10.3390/app10093280

NVIDIA, 2023a. DALI. NVIDIA Corporation.

NVIDIA, 2023b. NVIDIA Clara for Medical Devices [WWW Document]. NVIDIA. URL https://www.nvidia.com/en-us/clara/medical-devices/ (accessed 9.3.23).

NVIDIA, 2023c. Data Science, Machine Learning, AI, HPC Containers | NVIDIA NGC [WWW Document]. NVIDIA NGC Catalog. URL https://catalog.ngc.nvidia.com/containers (accessed 9.3.23).

NVIDIA, 2018. NVIDIA TAO Toolkit [WWW Document]. NVIDIA Developer. URL https://developer.nvidia.com/tao-toolkit (accessed 9.3.23).

NVIDIA, 2017. NVIDIA Collective Communications Library (NCCL) [WWW Document]. NVIDIA Developer. URL https://developer.nvidia.com/nccl (accessed 9.3.23).

NVIDIA, 2016. DL Frameworks [WWW Document]. NVIDIA Developer. URL https://developer.nvidia.com/deep-learning-frameworks (accessed 9.3.23).

NVIDIA, 2015. NVIDIA DIGITS [WWW Document]. NVIDIA Developer. URL https://developer.nvidia.com/digits (accessed 9.3.23).

NVIDIA, 2014. CUDA Deep Neural Network [WWW Document]. NVIDIA Developer. URL https://developer.nvidia.com/cudnn (accessed 9.3.23).

Ouyang, W., Zeng, X., Wang, X., Qiu, S., Luo, P., Tian, Y., Li, Hongsheng, Yang, S., Wang, Z., Li, Hongyang, Wang, K., Yan, J., Loy, C.-C., Tang, X., 2017. DeepID-Net: Object Detection with Deformable Part Based Convolutional Neural Networks. IEEE Transactions on Pattern Analysis and Machine Intelligence 39, 1320–1334. https://doi.org/10.1109/TPAMI.2016.2587642

Pandiya, M., Dassani, S., Mangalraj, P., 2020. Analysis of Deep Learning Architectures for Object Detection - A Critical Review, in: 2020 IEEE-HYDCON. Presented at the 2020 IEEE-HYDCON, pp. 1–6. https://doi.org/10.1109/HYDCON48903.2020.9242776

Peng, C., Xiao, T., Li, Z., Jiang, Y., Zhang, X., Jia, K., Yu, G., Sun, J., 2018. MegDet: A Large Mini-Batch Object Detector. Presented at the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 6181–6189.

Prechelt, L., 1998. Early Stopping - But When?, in: Orr, G.B., Müller, K.-R. (Eds.), Neural Networks: Tricks of the Trade, Lecture Notes in Computer Science. Springer, Berlin, Heidelberg, pp. 55–69. https://doi.org/10.1007/3-540-49430-8_3

Rahtu, E., Kannala, J., Blaschko, M., 2011. Learning a category independent object detection cascade, in: 2011 International Conference on Computer Vision. Presented at the 2011 International Conference on Computer Vision, pp. 1052–1059. https://doi.org/10.1109/ICCV.2011.6126351

Redmon, J., Divvala, S., Girshick, R., Farhadi, A., 2016. You Only Look Once: Unified, Real-Time Object Detection. Presented at the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 779–788.

Redmon, J., Farhadi, A., 2018. YOLOv3: An Incremental Improvement. arXiv:1804.02767 [cs].

Redmon, J., Farhadi, A., 2017. Yolo9000: better, faster, stronger arXiv preprint. arXiv preprint arXiv 1612.

Ren, S., He, K., Girshick, R., Sun, J., 2015. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks, in: Advances in Neural Information Processing Systems. Curran Associates, Inc.

Reuther, A., Michaleas, P., Jones, M., Gadepally, V., Samsi, S., Kepner, J., 2019. Survey and Benchmarking of Machine Learning Accelerators, in: 2019 IEEE High Performance Extreme Computing Conference (HPEC). Presented at the 2019 IEEE High Performance Extreme Computing Conference (HPEC), pp. 1–9. https://doi.org/10.1109/HPEC.2019.8916327

Rodriguez-Conde, I., Campos, C., Fdez-Riverola, F., 2021. On-Device Object Detection for More Efficient and Privacy-Compliant Visual Perception in Context-Aware Systems. Applied Sciences 11, 9173. https://doi.org/10.3390/app11199173

Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L., 2015. ImageNet Large Scale Visual Recognition Challenge. Int J Comput Vis 115, 211–252. https://doi.org/10.1007/s11263-015-0816-y

Sharif Razavian, A., Azizpour, H., Sullivan, J., Carlsson, S., 2014. CNN Features Off-the-Shelf: An Astounding Baseline for Recognition. Presented at the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, pp. 806–813.

Shin, H.-C., Roth, H.R., Gao, M., Lu, L., Xu, Z., Nogues, I., Yao, J., Mollura, D., Summers, R.M., 2016. Deep Convolutional Neural Networks for Computer-Aided Detection: CNN Architectures, Dataset Characteristics and Transfer Learning. IEEE Transactions on Medical Imaging 35, 1285–1298. https://doi.org/10.1109/TMI.2016.2528162

Shmelkov, K., Schmid, C., Alahari, K., 2017. Incremental Learning of Object Detectors Without Catastrophic Forgetting. Presented at the Proceedings of the IEEE International Conference on Computer Vision, pp. 3400–3409.

Singh, B., Davis, L.S., 2018. An Analysis of Scale Invariance in Object Detection - SNIP. Presented at the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 3578–3587.

Stone, J.E., Gohara, D., Shi, G., 2010. OpenCL: A Parallel Programming Standard for Heterogeneous Computing Systems. Comput Sci Eng 12, 66–72. https://doi.org/10.1109/MCSE.2010.69

Thabet, R., Mahmoudi, R., Bedoui, M.H., 2014. Image processing on mobile devices: An overview, in: International Image Processing, Applications and Systems Conference. Presented at the International Image Processing, Applications and Systems Conference, pp. 1–8. https://doi.org/10.1109/IPAS.2014.7043267

Wu, X., Sahoo, D., Hoi, S.C.H., 2020. Recent advances in deep learning for object detection. Neurocomputing 396, 39–64. https://doi.org/10.1016/j.neucom.2020.01.085

Yang, F., Choi, W., Lin, Y., 2016. Exploit All the Layers: Fast and Accurate CNN Object Detector With Scale Dependent Pooling and Cascaded Rejection Classifiers. Presented at the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2129–2137.

Yang, W., Liu, B., Li, W., Yu, N., 2019. Tracking Assisted Faster Video Object Detection, in: 2019 IEEE International Conference on Multimedia and Expo (ICME). Presented at the 2019 IEEE International Conference on Multimedia and Expo (ICME), pp. 1750–1755. https://doi.org/10.1109/ICME.2019.00301

Yao, Y., Wang, Y., Guo, Y., Lin, J., Qin, H., Yan, J., 2020. Cross-dataset Training for Class Increasing Object Detection. https://doi.org/10.48550/arXiv.2001.04621

Yu, J., Jiang, Y., Wang, Z., Cao, Z., Huang, T., 2016. UnitBox: An Advanced Object Detection Network, in: Proceedings of the 24th ACM International Conference on Multimedia, MM ’16. Association for Computing Machinery, New York, NY, USA, pp. 516–520. https://doi.org/10.1145/2964284.2967274

Zaidi, S.S.A., Ansari, M.S., Aslam, A., Kanwal, N., Asghar, M., Lee, B., 2022. A survey of modern deep learning based object detection models. Digital Signal Processing 126, 103514. https://doi.org/10.1016/j.dsp.2022.103514

Zhao, Z.-Q., Zheng, P., Xu, S.-T., Wu, X., 2019. Object Detection With Deep Learning: A Review. IEEE Transactions on Neural Networks and Learning Systems 30, 3212–3232. https://doi.org/10.1109/TNNLS.2018.2876865

Zhou, B., Lapedriza, A., Xiao, J., Torralba, A., Oliva, A., 2014. Learning Deep Features for Scene Recognition using Places Database, in: Advances in Neural Information Processing Systems. Curran Associates, Inc.

Zhu, H., Wei, H., Li, B., Yuan, X., Kehtarnavaz, N., 2020. A Review of Video Object Detection: Datasets, Metrics and Methods. Applied Sciences 10, 7834. https://doi.org/10.3390/app10217834

Zou, Z., Chen, K., Shi, Z., Guo, Y., Ye, J., 2023. Object Detection in 20 Years: A Survey. Proceedings of the IEEE 111, 257–276. https://doi.org/10.1109/JPROC.2023.3238524

TABLES AND FIGURES

Figure 9.1 Information graphic of workflow for object detection. The workflow is subdivided into three phases: pre-training, intra-training, and post-training. Each phase is further subdivided into three steps.

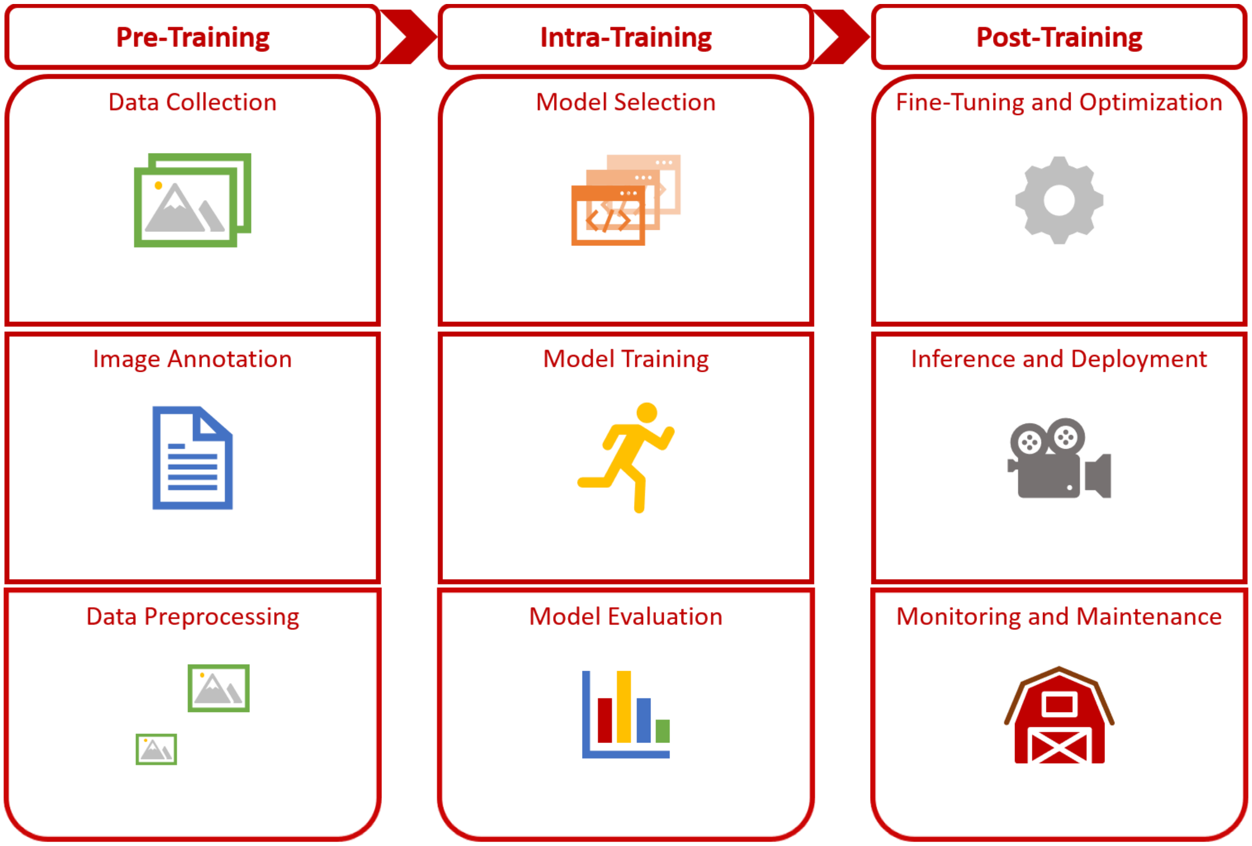

CHAPTER 10 DISCUSSION

Previous Studies for Early Detection of Digital Dermatitis

The early detection of digital dermatitis (DD) plays a critical role in facilitating prompt treatment and minimizing the overall severity of the disease (Alsaaod et al., 2014; Stokes et al., 2012a). By identifying such conditions at the initial stages, it is possible to intervene early, preventing chronically infected cattle and the progression of the disease (Leach et al., 2012) by reducing the infectious reservoir and outbreak courses within the herd (Döpfer, 2009; Döpfer et al., 2012; Stokes et al., 2012a).

The gold standard for diagnosing DD involves visually inspecting the foot in a hoof trimming chute. It has traditionally been conducted by restraining the cow in a claw-trimming chute and lifting the foot during routine claw trimming (Holzhauer et al., 2006; Manske et al., 2002; Thomsen et al., 2008b). This method, while effective, it is not efficient since it can be time-consuming and labor-intensive and cause disruption in a cow's time budget, limiting the number of cows that can be examined within a short period (Relun et al., 2011; Stokes et al., 2012b). As a result, researchers have explored alternative scoring areas and methods to identify cows with DD in more practical settings, such as the milking parlor, headlocks, or during pen walks (Jacobs et al., 2017; Relun et al., 2011; Solano et al., 2017; Stokes et al., 2012b; Thomsen et al., 2008a).

Visual inspection of the feet in the parlor can significantly improve the efficiency of DD diagnosis, as it eliminates the need for cows to be individually restrained in the hoof-trimming chute (Relun et al., 2011; Stokes et al., 2012b). It enables quick examination, reducing the time and labor required for each cow. However, early-stage or small DD lesions that can be easily overlooked during the visual inspection (Solano et al., 2017). Additionally, the presence of gross contamination on the distal limb can obscure DD lesions, thereby making it difficult to detect. To mitigate this issue, feet may need to be washed before inspection to ensure clearer visibility, but this could potentially compromise udder hygiene (Oliveira et al., 2017).

Therefore, while inspecting hind feet in the milking parlor can improve the efficiency of DD diagnosis, it is essential to be aware of its limitations, especially early-stage lesions and potential hygiene concerns. Efforts are made to develop non-invasive and practical diagnostic methods that can be seamlessly integrated into existing milking routines without compromising udder health. Such advancements would not only improve DD diagnosis but also contribute to overall productivity and welfare in dairy farming.

Various methods have explored alternative approaches to score cows in the milking parlor, including the use of mirrors (Relun et al., 2011; Solano et al., 2017) or borescopes (Laven, 1999; Stokes et al., 2012b), without any tools (Oliveira et al., 2017; Rodriguez-Lainz et al., 1998; Thomsen et al., 2008a), or without prior washing of feet (Oliveira et al., 2017). Several studies compared the DD scores through these methods to the gold standard of hoof trimming chute inspection (Cramer et al., 2018). However, some methods were considered impractical due to the associated cost or their impact on milking duration (Laven, 1999; Stokes et al., 2012b).

The accuracy of the detection method was influenced by factors such as milking parlor design and the difficulty in accessing hind feet in certain configurations (Stokes et al., 2012b; Thomsen et al., 2008a). In general, studies with more detailed descriptions of lesions attempted to assess DD (e.g., by color, depth, or stage) had lower agreement and test characteristics compared to the gold standard (Cramer et al., 2018). On the other hand, the highest agreement was generally observed when DD was simplified into "present" and "absent" categories (Relun et al., 2011; Solano et al., 2017; Stokes et al., 2012b).

Most alternative scoring method studies have focused on scoring DD in the milking parlor. However, it could be beneficial to explore alternative scoring areas, considering that different parlor designs may affect the accuracy of DD detection and scoring in the milking parlor can be disruptive to the milking process (Thomsen et al., 2008a). Additionally, with the increasing prevalence of automatic milking systems globally, there is a growing need to identify alternative scoring areas for efficient DD detection.

Assessing and quantifying pain in animals can be challenging, but it is important to understand the welfare impact of DD in dairy cattle. Inference of pain can be made by examining production parameters, physiological responses, and behavior (Prunier et al., 2013). Accurately quantifying the pain associated with DD is essential for developing pain prevention protocols, determining the need for pain mitigation, and promoting animal welfare.

Previous studies have employed various methods to quantify DD-associated pain in dairy cattle, including observing behaviors such as limb withdrawal, kicking, and falling (Stilwell et al., 2019), evaluating changes in locomotion (Laven and Logue, 2006), grouping behaviors into subjective pain scores (Britt et al., 1999; Shearer and Hernandez, 2000), measuring nociceptive threshold (Dyer et al., 2007; Whay et al., 2005), and using infrared thermography (IRT) as a proxy for inflammation (D.S. LokeshBabu et al., 2018; Stokes et al., 2012a). However, these alternative methods often exhibit a lower diagnostic capacity when compared to the gold standard (Orsel et al., 2018).

Changes in locomotion or gait characteristics are often among the first noticeable signs of lameness. Dairy cattle affected with DD frequently modify their gait and posture in response to pain including reduced mobility, lifting or shaking of the affected leg, or walking with a toe-down posture to avoid contact with the floor (Bassett et al., 1990; Blowey and Sharp, 1988; Read and Walker, 1998; Rodriguez-Lainz et al., 1998; Shearer et al., 2005). Dairy cattle affected by DD can experience changes in the heel area, which may contribute to the persistence and occurrence of heel horn erosion (Gomez et al., 2015). Locomotion scoring is a commonly used behavioral indicator to assess pain associated with foot disorders in cattle (Gigliuto et al., 2014). Visual gait scoring methods have been developed and commonly used to assess lameness in dairy cattle. Herd mobility scoring is a widely adopted screening tool used to identify lame animals, followed by clinical investigation to determine the underlying cause and administer treatment. It has been demonstrated to have a correlation with the severity of foot lesions (O'Callaghan et al., 2003; Whay et al., 1997; Winckler and Willen, 2001). These scoring systems enable producers to visually observe and assess the gait of their cattle, providing valuable insights into the presence and severity of lameness-associated diseases (Flower and Weary, 2006; Manson and Leaver, 1988; Sprecher et al., 1997).

This approach has its limitations as it is time-consuming and may not be sensitive to detect DD lesions (Tadich et al., 2010). This is because a significant proportion of cows with DD do not exhibit lameness (Cramer et al., 2018). These scores are based on subjective scales, and their proper usage and interpretation require training and consistent intra- and inter-observer agreement over time (Archer et al., 2010; Channon et al., 2009; Engel et al., 2003; Flower and Weary, 2009; Whay, 2002). There is a need for more specific and sensitive methods to rapidly detect and score DD in cattle, allowing for timely intervention and improved management decisions. The presence of DD does not constantly manifest as visible signs of lameness, such that lameness may only be evident once DD lesions have reached a severe stage, potentially leading to underreporting of the issue (Krull et al., 2016; Laven and Proven, 2000; Stokes et al., 2009). Consequently, early detection of DD is critical to prevent further progression and promote timely and effective treatment (Döpfer et al., 2012; Shearer and Hernandez, 2000).

As farm sizes continue to grow, traditional subjective scoring systems for measuring locomotion and gait become labor-intensive, highlighting the need for automated and objective measurement technologies. Accelerometers, force platforms, and other technologies have been developed and employed in dairy production systems to evaluate gait and locomotion in adult cows (Van Nuffel et al., 2015). These technologies have primarily been developed and used within dairy production systems for adult cows. For feedlot cattle, the application of similar technologies can be cost-prohibitive and impractical, especially for evaluating beef cattle in feedlot settings.

Given the significance of early detection, there is not only a growing demand, but a pressing need for non-invasive and practical methods that facilitates fast and frequent screening for the presence of DD at both the foot and cow levels in real-time. By developing such methods, farmers and veterinarians can take proactive measures to manage and treat DD, ultimately improving the welfare and health outcomes of dairy cattle.

Advancements in computer vision and deep learning technologies have shown promising potential in the field of precision agriculture, including the detection of DD. Implementing automated computer vision systems that can efficiently analyze and interpret visual data can revolutionize the early detection of DD. These technologies, when combined with other data sources, such as sensor data and health records, may provide a comprehensive and holistic approach to monitoring and managing DD in dairy cattle.

Early detection of DD is paramount for effective disease management, leading to improved animal welfare and reduced economic losses. Developing and integrating advanced computer vision techniques, along with traditional gait scoring methods, can significantly enhance the precision and efficiency of DD detection, allowing for timely intervention and improved overall herd health.

Detection Digital Dermatitis with Computer Vision using YOLOv2

Cernek et al. studied computer vision (CV) models as a potential solution for early detection of DD on commercial dairy farms (Cernek et al., 2020). The primary objective was to develop and deploy a novel CV tool capable of identifying DD lesions on a commercial dairy farm setting. A large database containing more than 3,500 images of DD lesions was used to train a CV model using the YOLOv2 architecture including 1,177 M0 images (936 screenshot JPG images and 241 JPG images), 1,414 M4H images (914 screenshot JPG images and 500 JPG images), and 1,050 M2 images (660 screenshot JPG images and 390 JPG images). The YOLOv2 model was trained to detect between 2 DD lesions: M0/M4H and M2 lesions using 2,591 M0/M4H images and 1,050 M2 images.

For internal validation, the YOLOv2 CV model achieved a detection accuracy of 71% for DD lesions. Additionally, the agreement between the model's predictions and a human evaluator was quantified as "moderate" using Cohen's kappa. For the external validation, the YOLOv2 CV model demonstrated even better performance, achieving a detection accuracy of 88% for DD lesions. However, the agreement between the model's predictions and a human evaluator was quantified as "fair" by Cohen's kappa because of the data imbalance between the two M-stage classifications.

These results indicate that the YOLOv2 model shows potential for identifying DD lesions on commercial dairy farms. Early detection of DD can have significant benefits for the welfare of the animals and the overall management of the dairy herd. The application of CV technology in this context has the potential to revolutionize the monitoring and detection of hoof health issues in dairy cattle, leading to more proactive and effective management practices.

Based on the accuracy and agreement measured by Cohen's kappa in the previous study, the YOLOv2 approach proved to be viable option for detecting DD in a milking parlor setting. However, there are opportunities to further improve the model performance and explore other CV model architectures for different application domains.

Cernek et al. proposed adapting the current model from YOLOv2 to YOLOv3 using TensorFlow version 2.0 and updated CUDA drivers (Cernek et al., 2020). YOLOv3 builds on YOLOv2 with several improvements, leveraging the latest version with updated drivers would potentially lead to improved detection performance, increased accuracy, and better use of available computing resources for real-time applications (Redmon and Farhadi, 2018).

In addition, Cernek et al. proposed future studies to conduct comparative analyses of other available CV model architectures to identify the most effective approaches in various application domains (Cernek et al., 2020). For example, segmentation-based models such as Mask R-CNN provide clearer detection boundaries, making them more appropriate for scenarios where hooves may overlap, such as in heifer raising facilities or feedlots (He et al., 2018). For applications where hooves are distinctly separated, such as in a milking parlor, YOLO remains an excellent choice because of its ability to efficiently and accurately classify lesions. By identifying the most appropriate CV model for different scenarios, researchers and practitioners can tailor their approach to specific use cases, ensuring optimal performance and precision in each setting.

The combined accuracy from the external validation set exceeded the combined accuracy from the internal validation for the previous study, but the prediction accuracies within specific image categories were similar between the two validations. Notably, the model misclassified M2 lesions more frequently than M0/M4H lesions, primarily due to the lower prevalence of M2 lesions (3.5%) compared to M0/M4H lesions (96.5%) on the farm. The combined accuracy in the study is misleading with respect to overall performance, while the Cohen's kappa presented a more appropriate measure of accuracy and provided lower values during the external validation, where the detection accuracy of M2 lesions is expected to improve with a higher prevalence of M2 lesions in the training the training dataset. The observed performance degradation in the study can be attributed to the differences in image settings for each validation set. The internal validation images were randomly selected from the training dataset, while the external validation images were randomly selected from webcam footage at a single farm location.

Cernek et al. recommended to generate a more representative training dataset that is similar to the expected images during external validation (Cernek et al., 2020). This can be achieved by including more GoPro-derived images on the farm in the training dataset and implementing the model on additional commercial dairy farms. This approach would provide the model with a more diverse range of real-world scenarios, improving its ability to accurately detect DD lesions on various farms.

Additionally, the issue of distance discrepancy between the camera and the object is another factor that may have contributed to detection problems with the YOLOv2 model (Redmon and Farhadi, 2018). The perspective of the images in the training set was closer to the feet than what a camera would typically capture in a real-world application. This disparity in distance between training data and real-world scenarios may contribute to the lower-than-expected accuracy of the model. By incorporating images captured at different camera distances, the model can better adapt to real-world situations and improve its ability to accurately detect DD lesions under varying conditions. This strategy can help bridge the gap between the training data and the practical implementation of the model, leading to more reliable and robust performance in real-world settings.

Cernek et al. recommended an ideal CV model for detecting DD would greatly benefit from two essential components: an expanded training database with a broader range of images and increased on-farm implementation across multiple farms (Cernek et al., 2020). Hoof diseases, including DD, can affect various types of production animals, and enhancing the training database by including additional M-stages of DD, such as M4P, as well as other hoof diseases like reverse corkscrew claws, white line disease, and foot rot, would significantly enhance the model's versatility and practicality.

The YOLOv2 DD detection model developed in this study demonstrated a fair level of accuracy in identifying and classifying DD lesions on a commercial dairy farm. This marks a significant step in applying CV to the fields of veterinary medicine and agriculture. However, it is important to view this achievement as a foundational milestone for further advancements. Further optimization of DD detection methods has the potential of improving treatment strategies and enhancing overall animal welfare.

Comparative Analysis of Computer Vision Algorithms

The study aims to mitigate the negative impact of DD and lameness in all cattle by enabling early detection and prompt treatment. Eight state-of-the-art CV models including both one-stage and two-sage detectors were trained to identify DD lesions and score M-stages. The CV models were compared using various performance metrics and inference time to choose the most effective and efficient model for future application of DD detection.

Out of the eight CV models tested, six demonstrated the ability to detect DD lesions in dairy cattle from images. For detection accuracy, both YOLOv4 and Tiny YOLOv4 outperformed all other models with high precision, perfect recall, and the highest mAP. The next closest models were Cascade R-CNN, Faster R-CNN, YOLOv3, and Tiny YOLOv3, but still trailed behind YOLOv4 and Tiny YOLOv4 with respect to all three performance metrics. For detection speed, Tiny YOLOv4 outperformed all other models by a wide margin, achieving the highest inference time at approximately 333 FPS.

For Dataset 1 or the simpler dataset, YOLOv4 and Tiny YOLOv4 achieved the same precision and recall, but Tiny YOLOv4 yielded higher mAP. However, for Dataset 2 or the more complex dataset, both models achieved similar mAP, but Tiny YOLOv4 demonstrated higher precision and recall in addition to faster speed. Given its impressive results, Tiny YOLOv4 emerged as the top-performing CV model for the real-time detection of DD in streaming video. Notably, both YOLOv3 and YOLOv4 models surpassed the performance of the YOLOv2 model when processing the same input data of the simpler dataset addressing the limitation of the Cernek et al. study. As a proof-of-concept, the Tiny YOLOv4 model successfully detected various M-stages of DD on image and video files, demonstrating its practicality and appropriateness for real-world applications.

YOLO models can be converted into TensorFlow Lite, TensorRT, or MyriadX BLOB files, leading to a reduction in model size and power consumption (david8862, 2023; Hùng, 2023; Jung, 2023; Kurniawan, 2023). Consequently, these lightweight models can be deployed on edge devices such as mobiles or microcontrollers (tensorflow, 2021). Moreover, the PyTorch framework has shown remarkable speed and efficiency compared to some of its alternatives (Paszke et al., 2019; pytorch, 2021). Notably, YOLOv3, which is a popular object detection model, has already been implemented in PyTorch by Ultralytics (ultralytics, 2021). Additionally, YOLOv5, the PyTorch implementation of YOLOv4, delivers comparable speed and accuracy with the added benefit of reduced training time (ultralytics, 2022a).

Increasing accessibility to CV tools is essential to the widespread adoption of these techniques in veterinary medicine. By making CV models compatible with portable platforms, such as Android or iOS applications for mobile devices (akhilkailas2001, n.d.; Apple Developer, 2023a; hietalajulius, 2022a; TensorFlow, 2022a), and deploying them as cloud-based applications through Docker containers (Docker, 2021), the adoption threshold for these models can be significantly reduced. Furthermore, collaborating with cattle professionals and integrating the CV models into handheld devices will facilitate the generation of a diverse image library, resulting in the optimization and validation of CV models, leading to more accurate and efficient applications. The proposed CV tool, once implemented and widely adopted, has the potential to greatly enhance animal welfare and increase production in large-scale cattle facilities. As a successful example in veterinary medicine and agriculture, this approach may pave the way for similar advancements in other domains and applications.

Real-time DD detection application has the potential to revolutionize the disease monitoring and management on dairy farms. By leveraging state-of-the-art CV models, the application will provide fast and accurate detection of DD lesions, leading to prompt veterinary care and treatment. Ultimately, this technology aims to enhance the welfare and productivity of dairy cattle while minimizing the physical pain and economic losses associated with DD and lameness.

Benchmarking Analysis of Edge Device Implementation

Lightweight CV models were trained for the real-time detection of DD in dairy cows. The proposed CV models were designed to identify DD lesions and score M-stages on cattle feet. The CV models were compared using various performance metrics and inference time to choose the most effective and efficient model for DD detection. The best CV model was optimized for constrained environments, deployed on portable edge devices, and automated for real-time detection using live streaming video. The CV model was evaluated for agreement and inference time to determine feasibility of the implementation for real-time detection on a portable device.

The DD detection model implemented on an edge device demonstrated real-time detection and accurately identifying DD lesions and classifying M-stages. The Tiny YOLOv4 model exceeded the threshold for real-time detection, achieving an inference speed of approximately 40 FPS. Notably, it exhibited a high mAP on Google Colab and a high Cohen's kappa on the edge device, implying reliable and accurate detection performance. The Tiny YOLOv4 model on an edge device in this study outperformed the YOLOv2 model on a local machine in Cernek et al. study. While YOLOv2 model only used two classes to mitigate performance on a simpler dataset, Tiny YOLOv4 used five distinct M-stages of DD on a more complex dataset, representing practical scenarios. As a proof-of-concept, the Tiny YOLOv4 model on an edge device effectively detected various M-stages of DD on images, videos, and live streaming video.

The implementation uses Tiny YOLOv4 and the TensorFlow 1.0 framework, demonstrating impressive speed and performance. However, the model can transition to TensorFlow 2.0 framework, providing enhanced usability and performance (NVIDIA Corporation, 2022; tensorflow, 2017). Additionally, TensorFlow models can be converted to TensorFlow Lite models, reducing the model size and power consumption, facilitating fast deployment and detection on edge devices (tensorflow, 2021). Alternatively, the model architecture can be further improved by migrating to the PyTorch framework and adopting YOLOv5 or YOLOv7, yielding better speed and performance (Paszke et al., 2019; pytorch, 2021). These updates would enable the model to take advantage of the latest advancements in deep learning and CV for higher accuracy and speed in DD detection. By exploring these alternatives, the model can be optimized for real-time detection of DD lesions, creating a valuable tool for on-farm applications.

Benchmarking Analysis of Cloud Computing Implementations

The study focuses on training lightweight CV models intended for cloud deployment, with the goal of evaluating their effectiveness in real-time detection of DD in dairy cows. The CV models are specifically trained to detect and score DD lesions, optimized for deployment, and subsequently compared using various performance metrics and inference time. Furthermore, the process is automated to facilitate real-time detection using images, videos, and video streams on cloud devices.

First and foremost, cloud deployment provides the ability for scalability, ensuring that the system can effectively handle varying workloads as the demand increases. Additionally, the cloud infrastructure allows for greater flexibility, enabling updates and improvements to the models without disruptions. Furthermore, cloud deployment enhances overall system efficiency, optimizing computational resources and delivering faster response times. Lastly, cloud services can also offer enhanced privacy and security measures, ensuring that sensitive data remains protected throughout the process.

The DD live detection applications developed showed impressive ability to accurately identify DD lesions and classify M-stages in real-time. The Docker container using an IP camera via RTSP demonstrated the fastest inference speed at approximately 25 FPS, closely followed by the Docker container using an IP camera via HTTP. All deployments exceeded the minimum threshold required for image processing by a human visual system at approximately 10 FPS. All implementations exhibited lower inference speed compared to a detector running locally with a webcam at approximately 40 FPS. This disparity is to be expected, as the model deployment involves additional steps related to network connectivity and associated latency, leading to a slight reduction in performance.

For detection accuracy, the Tensorflow.js application was highly effective in classifying the M-stages of DD with a Cohen's kappa value of 0.763 and an overall accuracy of 0.842. Comparatively, the YOLOv5s model deployed in the web browser using TensorFlow.js, outperformed the YOLOv2 model on a local machine in the Cernek et al study, while detecting additional M-stages of DD on a more complex dataset, representing practical scenarios. As a proof-of-concept, the YOLOv5s model successfully detected all M-stages of DD across all sources, including images, videos, and streaming video, demonstrating its versatility and potential for real-world applications.

The latest YOLOv5 model was developed using the PyTorch framework, resulting in notable improvements in speed and overall performance. An additional advantage of using PyTorch is the ability to convert the models to TensorFlow Lite models in addition to TensorFlow.js Layer models as well as ONXX, OpenVino, TensorRT, and CoreML models for AI development and ML applications. This conversion leads to reduced model size and latency, making it easier to deploy the models on cloud platforms and mobile devices, thereby enhancing accessibility and versatility (tensorflow, 2021).

The model architecture can be upgraded to YOLOv7 or YOLOv8 (Paszke et al., 2019; pytorch, 2021). YOLOv7 or YOLOv8 would significantly increase the speed and performance for real-time detection. Additionally, the model architecture and the format of the exported model open up opportunities for further optimization, allowing for the expansion of possible deployment options. This progress would contribute to a more efficient and effective detection system, tailoring a broader range of applications to a wide array of use cases.

Improving accessibility to these models can be achieved by developing mobile applications, such as TensorFlow Lite, Ultralytics Hub Android App, or Docker containers, which offer lightweight and standalone packages (Docker, 2022a; *Docker*, 2022b; PyTorch, 2023; techzizou, 2021; Ultralytics, 2023a, 2023b). Increased engagement with these models will lead to the generation of a complex and comprehensive library of images that can be used for the optimization and validation of computer vision models. This availability and rapid deployment of CV tools will have significant benefits for agriculture applications and will provide enhanced access to valuable information and accelerate advancements in the field.

Mobile Deployment of YOLO Model on Android and iOS Devices

Mobile deployment of YOLO models involves optimizing and converting the model to run efficiently on mobile devices, such as smartphones and tablets, to perform real-time object detection. YOLO models are optimized to reduce its size and computational complexity while maintaining its accuracy. This can be achieved through various techniques, such as model quantization to reduce precision, model pruning to remove unnecessary parameters, and model compression to create a lightweight version for mobile devices.

YOLO models are typically trained using deep learning frameworks like TensorFlow or PyTorch (Alexey, 2023; Jocher, 2020). For mobile deployment, the models are converted to a compatible format with mobile frameworks such as TensorFlow Lite for Android or Core ML for iOS. AI accelerators are used to achieve real-time performance on mobile devices. Many modern mobile devices have dedicated neural processing units (NPUs) or graphics processing units (GPUs). YOLO models can be optimized to leverage these AI accelerators efficiently to speed up inference tasks.

TensorFlow Lite is a lightweight and optimized version of the TensorFlow machine learning framework designed specifically for mobile and embedded devices (TensorFlow, 2023a, 2023b, 2023c, 2023d). It allows developers to run TensorFlow models efficiently on edge devices, such as smartphones, tablets, IoT devices, and other resource-constrained platforms, without the need for a constant internet connection or cloud-based processing.

TensorFlow Lite achieves this efficiency by using techniques like model quantization, which reduces the precision of model parameters to decrease memory and computation requirements while maintaining acceptable accuracy (TensorFlow, 2021a, 2022b, 2023e, 2023f). It also supports hardware acceleration through libraries such as Android Neural Networks API and Apple Core ML, leveraging specialized hardware on devices to further boost inference speed (TensorFlow, 2023g, 2022c, 2022d, 2021b). Similarly, it supports hardware acceleration through GPU and Neural Processing Units (NPUs), to accelerate inference on supported devices, necessary for real-time detection.

TensorFlow Lite supports a variety of model formats, including TensorFlow models and models from other popular deep learning frameworks including PyTorch, Keras, and Caffe (TensorFlow, 2022c, 2022e, 2022f). Additionally, TensorFlow Lite supports various platforms including Android, iOS, Linux, and Windows (TensorFlow, 2023a, 2022c). This versatility allows developers to easily convert models and deploy it across different platforms.

The main goal of TensorFlow Lite is to provide a seamless experience for developers to deploy machine learning models on mobile devices, enabling on-device inference and real-time processing (TensorFlow, 2023a, 2023b, 2023c, 2023d). This capability opens up a wide range of applications, including image and speech recognition, natural language processing, object detection, and more, all performed directly on the device without relying on cloud-based servers. By enabling machine learning on the edge, TensorFlow Lite empowers developers to create efficient and privacy-conscious applications, while also making AI accessible to a broader range of users in various domains.

There exists tutorials to build an Android apps using TensorFlow Lite to continuously detect objects in frames captured by a device camera (TensorFlow, 2022a). Similarly, there exists code examples to setup the object detection app and run it using Android Studio (akhilkailas2001, n.d.). YOLOv5 can be directly exported from a PyTorch weights to a TensorFlow Lite file, while the gradle file can be update to reflect the changes to the underlying model. Running the app using the COCO dataset presents an inference time of 10 FPS on CPU and 20 FPS on GPU using live streaming video, greater than the threshold for human visual systems and sufficient for DD detection.

Core ML is a machine learning framework developed by Apple for developers to integrate models into iOS, iPadOS, macOS, watchOS, and tvOS apps (Apple Developer, 2023b; Apple

Inc, 2023). It provides a seamless way to run pre-trained machine learning models on Apple devices for developers to add smart features to applications without internet connectivity. Core ML supports a variety of machine learning model formats, including models from popular frameworks like TensorFlow, Keras, scikit-learn, and PyTorch (apple, 2023a, 2023b). These models can be converted to the Core ML format using specialized tools provided by Apple. Core ML enables on-device inference where machine learning models run directly on the device without the need for internet access or cloud servers. This ensures privacy and data security for real-time processing without relying on external servers.

Core ML takes advantage of the underlying hardware on Apple devices, such as the GPU and Neural Engine, to accelerate the inference process and improve performance (Apple Developer, 2023b; Apple Inc, 2023). The tools include optimizations to reduce model size and improve efficiency for mobile deployment, ensuring that models run smoothly on resource-constrained devices. It supports various machine learning tasks, including image classification, object detection, and instance segmentation. Core ML seamlessly integrates with the Apple Vision framework to simplify CV tasks like object detection, text recognition, and barcode scanning (Apple Developer, 2023c). By leveraging Apple Vision, developers can create sophisticated applications that can understand and process visual information. The framework's ease of use and integration with other Apple technologies make it a powerful tool for implementing CV applications on Apple devices.

There exist demos of iOS object detection apps with YOLOv5 and Core ML using the COCO dataset (Apple Developer, 2023a; Hietala, 2022). Additionally, there exists a Python script to export a PyTorch weights as a Core ML file for a YOLOv5 custom model (hietalajulius, 2022b, 2022c, 2022a). The Core ML file includes non-maximum suppression (NMS) at the end of the model to support Apple Vision. Running the finished app using the COCO dataset presents an inference time of approximate 20 ms or 50 FPS on live streaming video, greater than the threshold for real-time detection and better than all other implementations described in previous studies for DD detection.

Therefore, the mobile deployment of YOLOv5 model on Android and iOS device is not only possible, but very much achievable. Additionally, mobile devices with the upgraded camera and chip specification can provide greater margin of error for distance from object and increased inference time. Mobile applications can be used to monitor the health of individual animals or the entire herd where early detection of health issues can prevent disease outbreaks and improve animal welfare. Real-time data analysis allows for timely decision-making and immediate action to address any emerging issues. Mobile applications allow farmers to remotely monitor their livestock and farm operations from anywhere at any time. This is especially valuable for farmers managing large-scale operations in rural, remote areas. Consequently, object detection apps have the potential to transform disease monitoring and management on dairy farms. The application will enable quick and precise detection of DD lesions, leading to prompt veterinary care and treatment. The ultimate goal is to enhance the welfare and productivity of dairy cattle while reducing the physical pain and economic losses associated with DD and lameness. By deploying object detection models and other CV technology on mobile devices, these applications can impact the welfare of dairy cattle and the production of dairy farming.

Sensors

The most commonly used image sensor devices in various applications are standard digital cameras and surveillance cameras that capture electromagnetic waves within the visible light spectra to produce digital images in color or grayscale (Helmers and Schellenberg, 2003). These cameras are widely accessible and offer a broad range of applications in photography, surveillance, and many other fields.

In addition to standard digital cameras, there exist other specialized imaging technologies tailored for specific applications. Infrared imaging uses infrared radiation to capture images and detect heat signatures and operate in low-light conditions (Lavers et al., 2005; McManus et al., 2016). This technology finds application in fields like night vision, thermal imaging, and detecting heat patterns in various industries. Ultrasound-based imaging devices uses high-frequency sound waves to generate images in medical applications, such as prenatal imaging or non-invasive testing (Fernandes et al., 2020). Furthermore, ionizing radiation-based devices are commonly employed in medical imaging techniques like X-ray and CT scans, enabling visualization of internal structures, commonly used in medical diagnostics (Fernandes et al., 2020).

Advanced imaging technologies can generate more complex arrays of images. Three-dimensional (3D) imaging produces detailed spatial representations, assisting in fields like CV (Giancola et al., 2018; Zanuttigh et al., 2016). Hyperspectral imaging captures images in multiple narrow spectral bands and provides rich data about the chemical composition of materials (Fernandes et al., 2020). This technology has applications in agriculture, environmental monitoring, and remote sensing. However, specialized imaging technologies are more expensive and less common than standard digital cameras. Consequently, their adoption is limited to specific industries or research applications where the benefits outweigh the costs.

Overall, the broad spectrum of imaging technologies available provides a powerful toolbox for addressing various challenges and fulfilling specific requirements in different domains.

Understanding the strengths and limitations of each technology allows professionals to choose the most appropriate imaging approach for their particular applications.

Internet of Things

Internet of Things (IoT) has emerged as a transformative technology in precision livestock farming, revolutionizing how livestock operations are managed and monitored (Nóbrega et al., 2019; Shi et al., 2019; Wurtz et al., 2019). IoT encompasses a network of interconnected devices, sensors, and data analytics that collect and exchange data over the internet, enabling real-time monitoring and decision-making for livestock farmers.

An important application of IoT in precision livestock farming is the monitoring of animal health and behavior (Benhai et al., 2015; Ilapakurti and Vuppalapati, 2015; Jingqiu et al., 2017; Kim and Choi, 2014; Shinde and Prasad, 2017; Stutz et al., 2019). Wearable sensors, such as smart collars, ear tags, and leg bands, are equipped with various biometric and environmental sensors to continuously track vital signs, activity levels, and environmental conditions (Benaissa et al., 2020; Chaudhry et al., 2020; Rahman et al., 2018; Subeesh and Mehta, 2021; Williams et al., 2020). This data provides valuable insights into individual animal health, detecting early signs of diseases, estrus cycles, and stress indicators. By identifying health issues promptly, farmers can administer timely treatments and prevent disease outbreaks, leading to improved animal welfare and reduced veterinary costs.

Internet of Things also plays a critical role in optimizing feeding strategies and nutrition management (Memon et al., 2016). Smart feeding systems equipped with Radio-Frequency Identification (RFID) technology and automated feeders can precisely deliver personalized feed rations to individual animals based on their nutritional requirements and production goals (Ariff and Ismail, 2018; Brown-Brandl et al., 2019). This targeted approach enhances feed efficiency, reduces waste, and ensures optimal growth and production in livestock.

Internet of Things facilitates remote monitoring of environmental conditions within livestock facilities (Nóbrega et al., 2019; Shi et al., 2019). Sensors installed in barns and pens can track temperature, humidity, ventilation, and air quality parameters (Chaudhry et al., 2020). Automated systems can then adjust climate control and ventilation settings accordingly, creating a comfortable and stress-free environment for animals. This not only promotes better animal health but also enhances productivity and reduces energy consumption.

Another significant application of IoT in precision livestock farming is the tracking of animal movements and behavior (Ariff and Ismail, 2018; Bailey et al., 2018; Brown-Brandl et al., 2019; Wolfger et al., 2017). Global Positioning System (GPS) collars and tags allow farmers to monitor grazing patterns, herd movements, and pasture utilization (Bailey et al., 2018). This information aids in making informed decisions regarding rotational grazing, optimizing forage usage, and ensuring sustainable land management practices.

The integration of IoT with data analytics and cloud computing further enhances the potential of precision livestock farming. The vast amount of data generated by IoT sensors can be analyzed in real-time, providing farmers with actionable insights and predictive analytics. By leveraging machine learning algorithms, farmers can gain a deeper understanding of livestock behavior, identify trends, and make data-driven decisions to improve farm efficiency.

Internet of Things enables seamless connectivity between different components of the farm, such as milking parlors, feeding systems, and environmental controls. Centralized data management platforms allow farmers to access critical information from multiple devices and systems, streamlining farm management and fostering greater operational efficiency.

Some challenges exist in the widespread adoption of IoT in precision livestock farming. Data security and privacy concerns are critical considerations, as the interconnected nature of IoT devices increases the risk of data breaches. Moreover, ensuring compatibility and standardization among different IoT devices from various manufacturers is essential to achieve seamless integration and interoperability.

Internet of Things is transforming precision livestock farming by providing real-time monitoring, data-driven decision-making, and enhanced operational efficiency. The integration of IoT technologies with advanced analytics offers tremendous potential for optimizing livestock management, promoting animal welfare, and contributing to sustainable and profitable farming practices in the modern agricultural landscape.

Unmanned Aerial Vehicle

Unmanned Aerial Vehicles (UAVs), commonly known as drones, have also found valuable applications in precision livestock farming (Aquilani et al., 2022; Yousefi et al., 2022). These aerial devices equipped with various sensors and cameras offer significant benefits in monitoring and managing livestock, contributing to improved animal welfare, optimized herd management, and enhanced farm productivity (Barbedo and Koenigkan, 2018; Jung and Ariyur, 2017; Li et al., 2020; Yinka-Banjo and Ajayi, 2019).

A current application of UAVs in precision livestock farming is the aerial monitoring of extensive grazing areas (Andrew et al., 2017; de Lima Weber et al., 2019; García et al., 2020; Nielsen, 2013). UAVs equipped with thermal cameras can detect heat stress in livestock, enabling early intervention and preventive measures to maintain animal health during extreme weather conditions (Herlin et al., 2021; Tzanidakis et al., 2023). Additionally, aerial surveys help identify potential hazards, such as predators or damaged fences, allowing farmers to respond promptly and mitigate risks to their livestock (di Virgilio et al., 2018; Li et al., 2020; Reddy Maddikunta et al., 2021).

Unmanned Aerial Vehicles can be employed for individual animal tracking and identification, particularly in large-scale farms where monitoring each animal manually can be labor-intensive and time-consuming (Goolsby et al., 2016; Li and Xing, 2019; Vayssade et al., 2019; Wamuyu, 2017). By using specialized tracking technology, farmers can keep track of livestock movements, feeding patterns, and behavior, allowing for better insights into individual animal health and performance.

Unmanned Aerial Vehicles offer the potential for disease detection in livestock (Moysiadis et al., 2021). Thermal cameras can be used to identify animals with elevated body temperatures, which may indicate the presence of infectious diseases (Herlin et al., 2021; Tzanidakis et al., 2023). Early detection enables rapid isolation and treatment, reducing the risk of disease spread and supporting overall herd health.

Integration with data analytics and decision-support systems is necessary. The data collected by UAVs can be processed using CV algorithms and AI techniques, generating valuable insights and recommendations for farm management (Anderson et al., 2013). These data-driven approaches assist farmers in making more informed decisions related to animal health, nutrition, and overall farm efficiency.

Some challenges exist in the adoption of UAVs in precision livestock farming. Compliance with aviation regulations, data privacy concerns, and the need for skilled operators are important aspects to consider. Additionally, robust communication networks are essential for transmitting data between UAVs and farm systems in real-time.

Unmanned Aerial Vehicles have become valuable tools in precision livestock farming for real-time monitoring, individual animal tracking, and enhanced decision-making capabilities for farmers. The combination of UAV technology with CV tools holds great potential to further advance livestock management practices, leading to improved animal welfare, increased farm productivity, and sustainable agriculture practices in the future.

Unmanned Aerial Vehicles can be dropped in the current workflow by cloud deployment. Real-Time Messaging Protocol (RTMP) is a protocol that drones use to transmit video. YOLOv5 can be employed to run inference using a drone camera via RTMP stream similar to HTTP and RTSP. However, in remote regions, the protocol may lack sufficient cellular network connectivity for support. Nevertheless, the implementation of a single, small edge device using a fast, lightweight model is ideal for low resource settings over large areas.

YOLOv7

YOLO remains a prominent object detection network because of its accessibility, accuracy, and computationally cost. These characteristics have made YOLO widely accepted and applied, beyond the data science community, because of its practicality. YOLOv7 was developed by the creators of YOLOv4, Scaled-YOLOv4, and YOLOR and represents the latest version of this popular algorithm with significant improvements over its previous iterations in the YOLO family (Wang et al., 2022a; WongKinYiu, 2022).

YOLOv7 employs Extended Efficient Layer Aggregation Networks (E-ELAN) for model re-parametrization during the inference stage to accelerate inference time (Wang et al., 2022a; WongKinYiu, 2022). E-ELAN leverages expand, shuffle, and merge cardinality techniques to enhance network adaptability and learning without impacting the original gradient path (Wang et al., 2022b). This approach applies the same group parameter and channel multiplier to each computational block in the layer, followed by shuffling and combining the feature maps into a set number of groups. The number of channels in each group of feature maps is the same as the original architecture. The groups are added to merge cardinality such that the transition layer is left unaffected and the gradient path is fixed.

Scaling is a common practice in object detection models to release a series of models for different use cases. YOLOv7 achieves scaling by simultaneously adjusting network depth and width while concatenating layers (Wang et al., 2022a; WongKinYiu, 2022). This approach maintains an optimal model architecture for different sizes, mitigating potential performance drawbacks when scaling certain aspects, such as the input and output channel ratio in transition layers.

YOLOv7 employs gradient flow propagation paths for the combination of re-parameterized convolution with different networks to improve inference (Wang et al., 2022a; WongKinYiu, 2022). Model training process is split into multiple modules and the outputs are ensembled to obtain the final model. Specifically, the 3×3 convolution layer of the E-ELAN computational block is replaced with the RepConv layer. YOLOv7 also performs re-parameterization on Convolution Batch Normalization (Conv-BN), Online Convolutional Re-parameterization (OREPA), and YOLO-R to get the best results (Hu et al., 2022; Wang et al., 2021).

Deep supervision is a method that introduces an additional auxiliary head within the intermediate layers of the network. This auxiliary head leverages shallow network weights with an assistant loss to guide the learning process. This technique is particularly valuable in scenarios where model weights tend to converge. The lead head responsible for the final output, while the auxiliary head used to assist in training. YOLOv7 uses the lead head predictions as guidance to generate coarse-to-fine hierarchical labels (Wang et al., 2022a; WongKinYiu, 2022). These labels are then used for learning in both the auxiliary head and lead head, respectively. This approach enhances network learning and helps achieve improved performance even in challenging convergence scenarios.

These various improvements result in significant advancements in performance and reductions in cost compared to its previous versions. YOLOv7 provides enhanced accuracy and efficiency, making it a possible choice for object detection tasks in various applications moving forward.

YOLOv7 exceeds all known object detectors in both speed and accuracy in the range from 5 FPS to 160 FPS and has the highest accuracy 56.8% AP for all known real-time object detectors with 30 FPS or higher on GPU V100 (Wang et al., 2022a; WongKinYiu, 2022). YOLOv7-E6 object detector (56 FPS V100, 55.9% AP) outperforms both state-of-the-art detector SWIN-L Cascade-Mask R-CNN (9.2 FPS A100, 53.9% AP) by 509% in speed and 2% in accuracy, and ConvNeXt-XL Cascade-Mask R-CNN (8.6 FPS A100, 55.2% AP) by 551% in speed and 0.7% AP in accuracy (Wang et al., 2022a; WongKinYiu, 2022). Additionally, YOLOv7 outperforms YOLOR, YOLOX, Scaled-YOLOv4, YOLOv5, DETR, Deformable DETR, DINO-5scale-R50, ViT-Adapter-B and many other object detectors in speed and accuracy (Wang et al., 2022a; WongKinYiu, 2022).

YOLOv8

YOLOv8 is the newest state-of-the-art YOLO model designed for various tasks such as object detection, image classification, and instance segmentation. YOLOv8 was developed by Ultralytics, who also created YOLOv5, and introduces a wide range of architectural and developer experience improvements compared to YOLOv5 (Jocher et al., 2023; Ultralytics, 2023c).

YOLOv8 employs an anchor-free approach by directly predicting the centroid of the object instead of calculating the offset from predetermined anchor boxes (Jocher et al., 2023; Ultralytics, 2023c). In previous YOLO models, anchor boxes posed challenges as they may reflect the distribution of reference boxes in a benchmark dataset, not necessarily the distribution of a specific custom dataset. This anchor-free strategy reduces the quantity of box predictions, resulting in faster processing of NMS to improve efficiency.

The Bottleneck architecture of YOLOv8 is similar to that of YOLOv5, except the first convolutional layer's kernel size has been changed from 1x1 to 3x3, reverting back to the ResNet block design (Jocher et al., 2023; Ultralytics, 2023c). The Bottleneck module directly concatenates features without enforcing the same channel dimensions. This results in a reduction in the number of parameters and the overall size of the tensors, resulting in improved efficiency and computational performance.

While deep learning research often prioritizes model architecture, the training routine plays a critical role in the success of YOLOv5 and YOLOv8. In the case of YOLOv8, images are augmented during the training process (Jocher et al., 2023; Ultralytics, 2023c). At each epoch, the model encounters slightly different variations of the images it has been given for training.

This dynamic augmentation strategy improves model performance and generalizability to unseen data.

Mosaic augmentation is a notable technique used in the YOLOv5 and YOLOv8 models (Jocher et al., 2023; Ultralytics, 2023c). This approach combines four images into a mosaic to expose and challenge the model to new object locations, partial occlusions, and varying surrounding pixels. However, it has been observed that continuously applying mosaic augmentation throughout the entire training process can lead to a decline in performance. Therefore, YOLOv8 disables mosaic augmentation for the last ten training epochs (Jocher et al., 2023; Ultralytics, 2023c).

YOLOv8 outperforms YOLOv7, YOLOv6-2.0, and YOLOv5-7.0 with respect to mAP, size, and latency during training (Jocher et al., 2023; Ultralytics, 2023c). The mAP increases as the size, speed, and FLOPs increase. The largest YOLOv5 model, YOLOv5x, achieved a maximum mAP value of 50.7 (Jocher et al., 2023; Ultralytics, 2023c). YOLOv8 achieved an impressive improvement of 2.2 units in mAP, signifying a significant enhancement in capabilities (Jocher et al., 2023; Ultralytics, 2023c). This is conserved for all model sizes where YOLOv8 models consistently outperforming YOLOv5. Overall, YOLOv8 signifies a significant step from YOLOv5 and other competing frameworks.

Instance Segmentation

Instance segmentation is a computer vision task that combines object detection and semantic segmentation to provide a more detailed understanding of an image. While semantic segmentation, which assigns a single class label to each pixel in an image, instance segmentation identifies and differentiates individual objects within the same class. This means that objects of the same class, such as multiple cows or lesions, are treated as separate instances with unique segmentation masks.

The goal of instance segmentation is to accurately outline and delineate the boundaries of each object instance in an image for precise object localization and separation of overlapping instances. This fine-grained level of detail is valuable in various applications including medical imaging.

Instance segmentation algorithms typically use deep learning models, such as Mask R-CNN, DeepLab, and YOLOv5 with ProtoNet (Chen et al., 2017a, 2017b; He et al., 2017; ultralytics, 2022a, 2022b, p. 10258). These models leverage CNNs to extract rich feature representations from the input image and sophisticated post-processing techniques to refine and improve the segmentation results. These models combine object detection with semantic segmentation to simultaneously predict bounding boxes and pixel-level masks for each instance. This task is particularly useful in scenarios where objects of the same class may have different shapes, sizes, and orientations, making it challenging for traditional object detection and semantic segmentation methods.

The YOLOv5 instance segmentation architecture consists of two main components: the YOLOv5 object detection head and a small fully connected neural network called ProtoNet (Roboflow, 2023; Ultralytics, 2022; ultralytics, 2022a, 2022b, 2022c). The ProtoNet architecture generating prototype masks for the segmentation model (Snell et al., 2017). In the post-processing phase, after the detection boxes and segmentation masks are obtained, the segmentation masks are clipped to fit neatly within each detected bounding box to ensure proper alignment. This prevents the segmentation masks from flowing out of the boundaries of the corresponding bounding boxes.

Image segmentation is the least explored of the three main CV tasks for animal science and precision livestock applications (Oliveira et al., 2017). Instance segmentation allows for precise tracking of animal movement and behavior to provide insights into growth rates, social interactions, and overall health (Nye et al., 2020). It can aid in monitoring water consumption and other feeding patterns. It can assist in automating livestock management tasks, such as counting and sorting animals, for record-keeping, compliance, and business operations (Qiao et al., 2019; Wu et al., 2020). Instance segmentation can aid in early detection of small changes in posture, gait, or behavior and assist in scoring of lameness, pain, or discomfort (Brünger et al., 2020). It can identify specific lesions or injuries, helping farmers and veterinarians take timely action, leading to early intervention, and resulting in better animal welfare outcomes. Ultimately, it can be employed to detect DD lesions and used by hoof trimmers for therapeutic trimming of sole ulcers and other bovine claw diseases to manage and monitor herd outbreaks.

While instance segmentation holds great potential, there are challenges to consider in its implementation. Instance segmentation models is computationally intensive and may require robust hardware infrastructure for implementation. Dairy parlors and beef feedlots can be diverse, with varying lighting conditions, animal poses, and occlusions. Instance segmentation models need to handle these challenges for reliable results. Developing accurate instance segmentation models requires large and diverse annotated datasets, which can be labor-intensive to create. For real-time applications, instance segmentation models need to deliver fast and efficient performance. Successful deployment of instance segmentation in precision livestock requires integration with other systems, such as data management platforms and automation tools. Nevertheless, the development and adoption of instance segmentation in precision livestock have the potential to revolutionize animal monitoring, leading to improved animal welfare, increased productivity, and more sustainable farming practices.

Pose Estimation

While YOLO is primarily designed for object detection tasks, it can be extended to estimate the body or keypoint locations of subjects, effectively turning it into a pose estimation model (Jocher et al., 2023; Munawar, 2023; Ultralytics, 2023d; WongKinYiu, 2022). To train a pose estimation model with YOLO, a dataset of images or videos containing annotated keypoints or body parts is required (Lynn, 2023). This dataset should include images or frames with labeled coordinates of keypoints, such as joints of a human body or specific body parts of animals (Lynn, 2023).

The YOLO architecture is originally designed for object detection, where it predicts bounding boxes and class probabilities for detected objects. For pose estimation, the model needs to be modified to predict the keypoints or body part coordinates instead of bounding boxes. The training dataset must have annotations of keypoints on the subjects of interest. These annotations specify the coordinates of the keypoints within the image or frame.

The YOLO model is then trained on the annotated dataset to accurately predict the keypoints. Model training involves adjusting the model parameters (weights and biases) to minimize the error between the predicted keypoints and the ground-truth keypoints. YOLO model can be used for pose estimation on new, unseen images or video frames. The model takes an input image or frame and predicts the coordinates of keypoints for each subject present in the image. The output of the model will be a set of predicted keypoints for each subject (Lynn, 2023). Post-processing techniques can be applied to refine the keypoints, smooth the results, and classify the poses based on the prediction (Lynn, 2023).

Pose estimation using YOLO has various applications, including human pose estimation for action recognition, gesture recognition, sports analysis, and fitness tracking. In veterinary sciences, it can be used for animal pose estimation to analyze behavior, gait, and posture, which can be valuable for monitoring animal health and well-being. Note that while YOLO can be adapted for pose estimation, there are other dedicated pose estimation models and algorithms that may provide more accurate and specialized results. Popular pose estimation models include OpenPose, Hourglass, EfficientPose, and MediaPipe (Bukschat and Vetter, 2020; Cao et al., 2019; MediaPipe, 2023; Newell et al., 2016). The choice of the model depends on the specific requirements and the complexity of the pose estimation task at hand.

Action Recognition

Action recognition is a computer vision task that involves identifying and classifying activities or actions from images or videos. The goal of action recognition is to develop algorithms and models that can automatically detect and categorize various actions performed by individuals or objects in a scene. In precision livestock farming, action recognition can play a critical role in monitoring and managing the overall welfare of the animals and economic well-being of the farm. By automatically detecting and classifying specific actions performed by animals, farmers and researchers can gain valuable insights into their behavior, health, and overall welfare.

Action recognition can be used to identify abnormal behaviors or actions that may indicate potential health issues in animals (Fuentes et al., 2023, 2020; Liang et al., 2018; Ma et al., 2022; McDonagh et al., 2021; Nguyen et al., 2021). For example, recognizing unusual movements or postures in cattle may help detect lameness, while identifying changes in feeding behavior could indicate signs of injury or illness (Jiang et al., 2020). Action recognition can aid in tracking reproductive behaviors in livestock, such as estrus detection in dairy cows. Recognizing specific behaviors associated with estrus can help optimize breeding programs and improve reproductive efficiency.

Monitoring actions related to activity levels and stress can provide insights into the overall welfare of livestock (Fuentes et al., 2023; Liang et al., 2018; Nguyen et al., 2021). An increase in stress-related behaviors can indicate environmental or management issues to be addressed. Additionally, monitoring actions associated with water consumption and other feeding patterns, can provide insights into the overall welfare of livestock as well (Cangar et al., 2008). This information can be used to adjust feeding strategies and ensure proper nutrition for each animal. By automating of action recognition, farmers and researchers can collect large amounts of data efficiently and consistently. This data can be used for further analysis and decision-making in precision livestock farming.

Accurate pose estimation is often a prerequisite for action recognition (Song et al., 2021). Pose estimation introduces additional challenges, such as occlusion and varying body sizes and shapes. Actions can be observed from multiple viewpoints and angles, significantly affecting the appearance of the action. The model needs to be robust to variations in views to accurately recognize actions. Proper camera placement is important for clear and unobstructed views of animal actions. This requires careful planning and adjustment to capture relevant behaviors effectively.

Livestock animals exhibit a wide range of behaviors, and there exists variability in actions by different individuals and species. Actions can be performed in different ways by different individuals and in various contexts. Variability in pose, speed, and appearance of actions makes it challenging for the recognition model to generalize across different instances of the same action. Actions occur in occluded scenes and dynamic environments with complex backgrounds.

Distinguishing the action of interest from the background noise can be difficult and may lead to misclassifications.

Building a labeled dataset for training the action recognition model can be time-consuming, labor-intensive process, requiring additional domain expertise to accurately annotate actions (Tian et al., 2020). Similarly, building a large and diverse dataset for training action recognition models can be time-consuming and resource-intensive process. Limited data may result in overfitting or reduced performance.

Data augmentation techniques are commonly used to artificially increase the size of the training dataset. However, finding effective data augmentation strategies for action recognition can be challenging as well. For real-time monitoring and decision-making, the action recognition system needs to process video streams efficiently and in real-time. This can be computationally demanding, particularly for large-scale datasets or high-resolution videos. Optimizing models for efficient inference is essential, especially in resource-constrained environments.

Despite these challenges, action recognition in precision livestock farming has the potential to significantly improve animal welfare, optimize management practices, and enhance the overall efficiency and productivity of livestock operations. As computer vision techniques continue to advance, the application of action recognition in precision livestock farming is expected to be prevalent and prominent.

Computer Vision in Precision Farming

The concept of developing CV for automatic monitoring and measuring of animal traits of interest has been an area of interest for researchers and practitioners. Early advancements in digital image analysis and computer vision have demonstrated the potential of using images to evaluate various animal characteristics, including behavior, gait, body weight, and other traits, particularly in controlled experimental settings. More recent studies have extended these applications to on-farm scenarios, showing the feasibility of using computer vision in real-world agricultural practices.

Research has revealed that certain imaging technologies are better suited for specific applications in animal health and food sciences. For instance, infrared thermography (IRT) has proven effective in identifying mastitis and digital dermatitis in dairy cattle by detecting variations in surface temperature associated with these conditions. Spectral and hyperspectral imaging have also shown promise in food sciences, enabling detailed analysis of food quality and safety parameters.

Despite the success of specialized imaging technologies, there is a growing interest in developing CV based on more accessible and widely available technologies, such as standard digital cameras and 3D cameras. These devices offer practical and cost-effective solutions for precision livestock monitoring, making them more feasible for broader adoption in diverse settings.

Standard digital cameras, which capture images within the visible light spectrum, have been widely used in various applications. Their simplicity and affordability make them suitable for routine monitoring of animal behavior, assessing body condition, and tracking animal movements. These cameras have shown great potential in enhancing farm management practices and improving animal welfare.

3D cameras, on the other hand, provide a more sophisticated approach to data acquisition, allowing for the creation of three-dimensional images of animals and their surroundings. This advanced level of spatial information opens up new opportunities for precise measurements and analysis of animal traits, such as volume, surface, and gait. The adoption of 3D cameras is steadily increasing as technology advances and costs decrease, making them a promising tool for precision livestock monitoring.

The development of CV based on standard digital cameras and 3D cameras offers several benefits for agriculture and animal husbandry. Firstly, the widespread availability of these devices makes them easily accessible to farmers and researchers alike, reducing the barrier to entry for utilizing computer vision technologies in the field. Secondly, their practicality and ease of use make them suitable for real-time monitoring and decision-making, enabling prompt interventions when necessary.

The versatility of these imaging technologies allows for a wide range of applications, from simple behavior analysis to more complex morphological measurements. For instance, CV using standard digital cameras can be employed in animal behavior studies and identifying abnormal behaviors that may indicate health issues. On the other hand, 3D cameras can be used for precise measurements of body dimensions, assisting in animal selection and breeding decisions.

As research and technology in computer vision continue to advance, the capabilities and functionalities of CV based on standard digital cameras and 3D cameras will likely expand. Integrating these technologies into precision livestock management systems will provide valuable insights into animal health, performance, and overall well-being. By leveraging the power of computer vision in combination with accessible imaging devices, the agriculture industry can make significant strides towards more efficient and sustainable practices while ensuring the welfare and productivity of livestock.

Applications of CV in animal and veterinary sciences is an emerging field of research. While some commercial products already exist for monitoring groups of live animals, there are still significant challenges for the successful development and future implementation of practical solutions.

A current challenge is the development of reliable CV for autonomously acquiring data on single or multiple traits in farm conditions. Some studies have explored these systems, but there remains a need for further validation using diverse datasets that include different animal breeds or species within the same farm or across multiple farms. Accuracy and robustness of these systems in real-world farming environments is essential for their widespread adoption and practical application.

Another area of interest is individual animal identification and tracking. Existing methods for animal identification are prone to errors, leading to potential issues with data accuracy and reliability. The precision and efficiency of individual animal tracking within groups is critical for obtaining accurate and meaningful results for farm management and research purposes.

Furthermore, as the number of devices used for different applications increases, the demand for effective methods to integrate and connect these devices increases as well. This integration would lead to the implementation of more sophisticated algorithms that can consider multiple inputs and outputs. By leveraging data from diverse sources, CV can improve the accuracy of joint predictions of multiple traits, providing valuable insights for animal welfare and production systems.

Another area of focus is the development of user-friendly applications to deliver the data generated by CV to connected systems. The clear and concise dissemination of information is essential for presenting actionable insights to farmers and managers, ultimately improving decision-making and optimizing farming practices. Technologies such as big data and IoT will be indispensable in creating a comprehensive ecosystem where CV-generated data can be analyzed, combined with other relevant data, and used to optimize animal breeding, health, and overall production.

Successful implementation of CV can lead to enhanced animal welfare, reduced labor-intensive monitoring, early disease detection, improved breeding programs, and increased production efficiency in agriculture and animal husbandry. The development and deployment of reliable CV in animal and veterinary sciences require continued research and collaboration among researchers, practitioners, and technology developers. Addressing the challenges of real-world data validation, individual animal identification, device integration, and data delivery will pave the way for the broader adoption of CV in animal breeding programs and production systems, ultimately improving the well-being and productivity of livestock and contributing to sustainable and efficient agricultural practices.

REFERENCES

akhilkailas2001, n.d. TFLite-Object-Detection-Android-App-YOLOv5 [WWW Document]. This repository teaches you how to create android app for your yolov5 model using tflite. To know more download the rar file in this repository and extract it. URL https://github.com/akhilkailas2001/TFLite-Object-Detection-Android-App-YOLOv5 (accessed 7.31.23).

Alexey, 2023. Yolo v4, v3 and v2 for Windows and Linux.

Alsaaod, M., Syring, C., Dietrich, J., Doherr, M.G., Gujan, T., Steiner, A., 2014. A field trial of infrared thermography as a non-invasive diagnostic tool for early detection of digital dermatitis in dairy cows. The Veterinary Journal 199, 281–285. https://doi.org/10.1016/j.tvjl.2013.11.028

Anderson, D.M., Estell, R.E., Cibils, A.F., 2013. Spatiotemporal Cattle Data—A Plea for Protocol Standardization. Positioning 4, 115–136. https://doi.org/10.4236/pos.2013.41012

Andrew, W., Greatwood, C., Burghardt, T., 2017. Visual Localisation and Individual Identification of Holstein Friesian Cattle via Deep Learning. Presented at the Proceedings of the IEEE International Conference on Computer Vision Workshops, pp. 2850–2859.

apple, 2023a. Core ML Tools. Apple.

apple, 2023b. Core ML Model [WWW Document]. Core ML Format Reference documentation. URL https://apple.github.io/coremltools/mlmodel/Format/Model.html (accessed 8.4.23).

Apple Developer, 2023a. Recognizing Objects in Live Capture [WWW Document]. Apple Developer Documentation. URL https://developer.apple.com/documentation/vision/recognizing_objects_in_live_capture (accessed 7.31.23).

Apple Developer, 2023b. Core ML [WWW Document]. Apple Developer Documentation. URL https://developer.apple.com/documentation/coreml (accessed 8.4.23).

Apple Developer, 2023c. Vision [WWW Document]. Apple Developer Documentation. URL https://developer.apple.com/documentation/vision (accessed 8.4.23).

Apple Inc, 2023. Core ML - Machine Learning [WWW Document]. Apple Developer. URL https://developer.apple.com/machine-learning/core-ml/ (accessed 8.4.23).

Aquilani, C., Confessore, A., Bozzi, R., Sirtori, F., Pugliese, C., 2022. Review: Precision Livestock Farming technologies in pasture-based livestock systems. Animal 16, 100429. https://doi.org/10.1016/j.animal.2021.100429

Archer, S., Bell, N., Huxley, J., 2010. Lameness in UK dairy cows: a review of the current status. In Practice 32, 492–504.

Ariff, M., Ismail, I., 2018. Rfid application development for a livestock monitoring system, in: Bioresources Technology in Sustainable Agriculture. Apple Academic Press, pp. 81–94.

Bailey, D.W., Trotter, M.G., Knight, C.W., Thomas, M.G., 2018. Use of GPS tracking collars and accelerometers for rangeland livestock production research. Translational Animal Science 2, 81–88.

Barbedo, J.G.A., Koenigkan, L.V., 2018. Perspectives on the use of unmanned aerial systems to monitor cattle. Outlook on agriculture 47, 214–222.

Bassett, H.F., Monaghan, M.L., Lenhan, P., Doherty, M.L., Carter, M.E., 1990. Bovine digital dermatitis. Veterinary Record 126, 164–165.

Benaissa, S., Tuyttens, F.A.M., Plets, D., Trogh, J., Martens, L., Vandaele, L., Joseph, W., Sonck, B., 2020. Calving and estrus detection in dairy cattle using a combination of indoor localization and accelerometer sensors. Computers and Electronics in Agriculture 168, 105153. https://doi.org/10.1016/j.compag.2019.105153

Benhai, X., Liang, Y., Xiaohua, P., Zhengang, Y., 2015. Review on application of Internet of Things technology in animal husbandry in China. Transactions of the Chinese Society of Agricultural Engineering 31.

Blowey, R.W., Sharp, M.W., 1988. Digital dermatitis in dairy cattle. Vet Rec 122, 505–508. https://doi.org/10.1136/vr.122.21.505

Britt, J.S., Berry, S.L., Shearer, J., Hemling, T., Steevens, B., Dreher, M., 1999. A uniform protocol for evaluating response to treatment of papillomatous digital dermatitis lesions. The Bovine Practitioner 149–154.

Brown-Brandl, T.M., Adrion, F., Maselyne, J., Kapun, A., Hessel, E., Saeys, W., Van Nuffel, A., Gallmann, E., 2019. A review of passive radio frequency identification systems for animal monitoring in livestock facilities. Applied Engineering in Agriculture 35, 579–591.

Brünger, J., Gentz, M., Traulsen, I., Koch, R., 2020. Panoptic Segmentation of Individual Pigs for Posture Recognition. Sensors 20, 3710. https://doi.org/10.3390/s20133710

Bukschat, Y., Vetter, M., 2020. EfficientPose: An efficient, accurate and scalable end-to-end 6D multi object pose estimation approach. https://doi.org/10.48550/arXiv.2011.04307

Cangar, Ö., Leroy, T., Guarino, M., Vranken, E., Fallon, R., Lenehan, J., Mee, J., Berckmans, D., 2008. Automatic real-time monitoring of locomotion and posture behaviour of pregnant cows prior to calving using online image analysis. Computers and Electronics in Agriculture, Smart Sensors in precision livestock farming 64, 53–60. https://doi.org/10.1016/j.compag.2008.05.014

Cao, Z., Hidalgo, G., Simon, T., Wei, S.-E., Sheikh, Y., 2019. OpenPose: Realtime Multi-Person 2D Pose Estimation using Part Affinity Fields. https://doi.org/10.48550/arXiv.1812.08008

Cernek, P., Bollig, N., Anklam, K., Döpfer, D., 2020. Hot topic: Detecting digital dermatitis with computer vision. Journal of Dairy Science 103, 9110–9115. https://doi.org/10.3168/jds.2019-17478

Channon, A.J., Walker, A.M., Pfau, T., Sheldon, I.M., Wilson, A.M., 2009. Variability of Manson and Leaver locomotion scores assigned to dairy cows by different observers. Veterinary Record 164, 388–392. https://doi.org/10.1136/vr.164.13.388

Chaudhry, A.A., Mumtaz, R., Hassan Zaidi, S.M., Tahir, M.A., Muzammil School, S.H., 2020. Internet of Things (IoT) and Machine Learning (ML) enabled Livestock Monitoring, in: 2020 IEEE 17th International Conference on Smart Communities: Improving Quality of Life Using ICT, IoT and AI (HONET). Presented at the 2020 IEEE 17th International Conference on Smart Communities: Improving Quality of Life Using ICT, IoT and AI (HONET), pp. 151–155. https://doi.org/10.1109/HONET50430.2020.9322666

Chen, L.-C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L., 2017a. DeepLab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs.

Chen, L.-C., Papandreou, G., Schroff, F., Adam, H., 2017b. Rethinking Atrous Convolution for Semantic Image Segmentation. https://doi.org/10.48550/arXiv.1706.05587

Cramer, G., Winders, T., Solano, L., Kleinschmit, D.H., 2018. Evaluation of agreement among digital dermatitis scoring methods in the milking parlor, pen, and hoof trimming chute. Journal of Dairy Science 101, 2406–2414. https://doi.org/10.3168/jds.2017-13755

david8862, 2023. TF Keras YOLOv4/v3/v2 Modelset.

de Lima Weber, F., Istela Cagnin, M., Barroso Paiva, D.M., de Moraes Weber, V.A., de Medeiros, S.R., da Costa Gomes, R., Pistori, H., 2019. Use of computational vision and UAVs in livestock: a Literature review, in: 2019 14th Iberian Conference on Information Systems and Technologies (CISTI). Presented at the 2019 14th Iberian Conference on Information Systems and Technologies (CISTI), pp. 1–7. https://doi.org/10.23919/CISTI.2019.8760955

di Virgilio, A., Morales, J.M., Lambertucci, S.A., Shepard, E.L., Wilson, R.P., 2018. Multi-dimensional Precision Livestock Farming: A potential toolbox for sustainable rangeland management. PeerJ 6, e4867.

Docker, 2022a. Docker: Accelerated, Containerized Application Development [WWW Document]. URL https://www.docker.com/ (accessed 7.23.23).

Docker, 2022b.

Docker, 2021. Docker Documentation [WWW Document]. Docker Documentation. URL https://docs.docker.com/

Döpfer, D., 2009. The dynamics of digital dermatitis in dairy cattle and the manageable state of disease, in: Proceedings of the CanWest Veterinary Conference. pp. 17–20.

Döpfer, D., Holzhauer, M., Boven, M. van, 2012. The dynamics of digital dermatitis in populations of dairy cattle: Model-based estimates of transition rates and implications for control. The Veterinary Journal, Special Issue:Bovine Lameness 193, 648–653. https://doi.org/10.1016/j.tvjl.2012.06.047

D.S. LokeshBabu, Jeyakumar, S., Vasant, P.J., Sathiyabarathi, M., Manimaran, A., Kumaresan, A., Pushpadass, H.A., Sivaram, M., Ramesha, K.P., Kataktalware, M.A., Siddaramanna, 2018. Monitoring foot surface temperature using infrared thermal imaging for assessment of hoof health status in cattle: A review. Journal of Thermal Biology 78, 10–21. https://doi.org/10.1016/j.jtherbio.2018.08.021

Dyer, R.M., Neerchal, N.K., Tasch, U., Wu, Y., Dyer, P., Rajkondawar, P.G., 2007. Objective Determination of Claw Pain and Its Relationship to Limb Locomotion Score in Dairy Cattle. Journal of Dairy Science 90, 4592–4602. https://doi.org/10.3168/jds.2007-0006

Engel, B., Bruin, G., Andre, G., Buist, W., 2003. Assessment of observer performance in a subjective scoring system: visual classification of the gait of cows. The Journal of Agricultural Science 140, 317–333. https://doi.org/10.1017/S0021859603002983

Fernandes, A.F.A., Dórea, J.R.R., Rosa, G.J. de M., 2020. Image Analysis and Computer Vision Applications in Animal Sciences: An Overview. Frontiers in Veterinary Science 7.

Flower, F.C., Weary, D.M., 2009. Gait assessment in dairy cattle. Animal 3, 87–95. https://doi.org/10.1017/S1751731108003194

Flower, F.C., Weary, D.M., 2006. Effect of Hoof Pathologies on Subjective Assessments of Dairy Cow Gait. Journal of Dairy Science 89, 139–146. https://doi.org/10.3168/jds.S0022-0302(06)72077-X

Fuentes, A., Han, S., Nasir, M.F., Park, J., Yoon, S., Park, D.S., 2023. Multiview Monitoring of Individual Cattle Behavior Based on Action Recognition in Closed Barns Using Deep Learning. Animals 13, 2020. https://doi.org/10.3390/ani13122020

Fuentes, A., Yoon, S., Park, J., Park, D.S., 2020. Deep learning-based hierarchical cattle behavior recognition with spatio-temporal information. Computers and Electronics in Agriculture 177, 105627. https://doi.org/10.1016/j.compag.2020.105627

García, R., Aguilar, J., Toro, M., Pinto, A., Rodríguez, P., 2020. A systematic literature review on the use of machine learning in precision livestock farming. Computers and Electronics in Agriculture 179, 105826. https://doi.org/10.1016/j.compag.2020.105826

Giancola, S., Valenti, M., Sala, R., 2018. A survey on 3D cameras: Metrological comparison of time-of-flight, structured-light and active stereoscopy technologies. Springer.

Gigliuto, C., De Gregori, M., Malafoglia, V., Raffaeli, W., Compagnone, C., Visai, L., Petrini, P., Avanzini, M.A., Muscoli, C., Viganò, J., Calabrese, F., Dominioni, T., Allegri, M., Cobianchi, L., 2014. Pain assessment in animal models: do we need further studies? Journal of Pain Research 7, 227–236. https://doi.org/10.2147/JPR.S59161

Gomez, A., Cook, N.B., Rieman, J., Dunbar, K.A., Cooley, K.E., Socha, M.T., Döpfer, D., 2015. The effect of digital dermatitis on hoof conformation. Journal of Dairy Science 98, 927–936. https://doi.org/10.3168/jds.2014-8483

Goolsby, J., Jung, J., Landivar, J., McCutcheon, W., Lacewell, R., Duhaime, R., Schwartz, A., 2016. Evaluation of unmanned aerial vehicles (UAVs) for detection of cattle in the cattle fever tick permanent quarantine zone. Subtropical Agriculture and Environments 67, 24–27.

He, K., Gkioxari, G., Dollár, P., Girshick, R., 2018. Mask R-CNN. https://doi.org/10.48550/arXiv.1703.06870

He, K., Gkioxari, G., Dollár, P., Girshick, R., 2017. Mask r-cnn, in: Proceedings of the IEEE International Conference on Computer Vision. pp. 2961–2969.

Helmers, H., Schellenberg, M., 2003. CMOS vs. CCD sensors in speckle interferometry. Optics & Laser Technology 35, 587–595. https://doi.org/10.1016/S0030-3992(03)00078-1

Herlin, A., Brunberg, E., Hultgren, J., Högberg, N., Rydberg, A., Skarin, A., 2021. Animal Welfare Implications of Digital Tools for Monitoring and Management of Cattle and Sheep on Pasture. Animals 11, 829. https://doi.org/10.3390/ani11030829

Hietala, J., 2022. How To Build a YOLOv5 Object Detection App on iOS [WWW Document]. Medium. URL https://betterprogramming.pub/how-to-build-a-yolov5-object-detection-app-on-ios-39c8c77dfe58 (accessed 7.31.23).

hietalajulius, 2022a. ios-demo-app [WWW Document]. PyTorch iOS examples. URL https://github.com/hietalajulius/ios-demo-app (accessed 7.31.23).

hietalajulius, 2022b. yolov5 [WWW Document]. URL https://github.com/hietalajulius/yolov5 (accessed 8.6.23).

hietalajulius, 2022c. ios-demo-app/ObjectDetection-CoreML [WWW Document]. URL https://github.com/hietalajulius/ios-demo-app/tree/main/ObjectDetection-CoreML (accessed 8.6.23).

Holzhauer, M., Hardenberg, C., Bartels, C.J.M., Frankena, K., 2006. Herd- and Cow-Level Prevalence of Digital Dermatitis in The Netherlands and Associated Risk Factors. Journal of Dairy Science 89, 580–588. https://doi.org/10.3168/jds.S0022-0302(06)72121-X

Hu, M., Feng, J., Hua, J., Lai, B., Huang, J., Gong, X., Hua, X., 2022. Online Convolutional Re-parameterization.

Hùng, V., 2023. tensorflow-yolov4-tflite.

Ilapakurti, A., Vuppalapati, C., 2015. Building an IoT Framework for Connected Dairy, in: 2015 IEEE First International Conference on Big Data Computing Service and Applications. Presented at the 2015 IEEE First International Conference on Big Data Computing Service and Applications, pp. 275–285. https://doi.org/10.1109/BigDataService.2015.39

Jacobs, C., Orsel, K., Barkema, H.W., 2017. Prevalence of digital dermatitis in young stock in Alberta, Canada, using pen walks. Journal of Dairy Science 100, 9234–9244. https://doi.org/10.3168/jds.2017-13044

Jiang, B., Yin, X., Song, H., 2020. Single-stream long-term optical flow convolution network for action recognition of lameness dairy cow. Computers and Electronics in Agriculture 175, 105536. https://doi.org/10.1016/j.compag.2020.105536

Jingqiu, G., Zhihai, W., Ronghua, G., Huarui, W., 2017. Cow behavior recognition based on image analysis and activities. International Journal of Agricultural and Biological Engineering 10, 165–174. https://doi.org/10.25165/ijabe.v10i3.3080

Jocher, G., 2020. YOLOv5 by Ultralytics. https://doi.org/10.5281/zenodo.3908559

Jocher, G., Chaurasia, A., Qiu, J., 2023. YOLO by Ultralytics.

Jung, J.K., 2023. tensorrt_demos.

Jung, S., Ariyur, K., 2017. Strategic Cattle Roundup using Multiple Quadrotor UAVs. International Journal of Aeronautical and Space Sciences 18, 315–326. https://doi.org/10.5139/IJASS.2017.18.2.315

Kim, Y.-B., Choi, D.-O., 2014. Design of business management system for livestock pens based of IoT. J. Korean Entertain. Ind. Assoc 8, 207–216.

Krull, A.C., Shearer, J.K., Gorden, P.J., Scott, H.M., Plummer, P.J., 2016. Digital dermatitis: Natural lesion progression and regression in Holstein dairy cattle over 3 years. Journal of Dairy Science 99, 3718–3731. https://doi.org/10.3168/jds.2015-10535

Kurniawan, I., 2023. YOLOv4 with TensorRT engine.

Laven, R., 1999. The environment and digital dermatitis. Cattle Practice 7, 349–354.

Laven, R.A., Logue, D.N., 2006. Treatment strategies for digital dermatitis for the UK. The Veterinary Journal 171, 79–88. https://doi.org/10.1016/j.tvjl.2004.08.009

Laven, R.A., Proven, M.J., 2000. Use of an antibiotic footbath in the treatment of bovine digital dermatitis. Veterinary Record 147, 503–506. https://doi.org/10.1136/vr.147.18.503

Lavers, C., Franks, K., Floyd, M., Plowman, A., 2005. Application of remote thermal imaging and night vision technology to improve endangered wildlife resource management with minimal animal distress and hazard to humans, in: Journal of Physics: Conference Series. IOP Publishing, p. 207.

Leach, K.A., Tisdall, D.A., Bell, N.J., Main, D.C.J., Green, L.E., 2012. The effects of early treatment for hindlimb lameness in dairy cows on four commercial UK farms. The Veterinary Journal, Special Issue:Bovine Lameness 193, 626–632. https://doi.org/10.1016/j.tvjl.2012.06.043

Li, D., Wang, C., Yan, T., Wang, Q., Wang, J., Bing, W., 2020. Cloud Grazing Management and Decision System Based on WebGIS, in: Zhang, X., Liu, G., Qiu, M., Xiang, W., Huang, T. (Eds.), Cloud Computing, Smart Grid and Innovative Frontiers in Telecommunications, Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering. Springer International Publishing, Cham, pp. 424–436. https://doi.org/10.1007/978-3-030-48513-9_34

Li, X., Xing, L., 2019. Reactive Deployment of Autonomous Drones for Livestock Monitoring Based on Density-based Clustering, in: 2019 IEEE International Conference on Robotics and Biomimetics (ROBIO). Presented at the 2019 IEEE International Conference on Robotics and Biomimetics (ROBIO), pp. 2421–2426. https://doi.org/10.1109/ROBIO49542.2019.8961763

Liang, Y., Xue, F., Chen, Xiaoming, Wu, Z., Chen, Xiangji, 2018. A Benchmark for Action Recognition of Large Animals, in: 2018 7th International Conference on Digital Home (ICDH). Presented at the 2018 7th International Conference on Digital Home (ICDH), pp. 64–71. https://doi.org/10.1109/ICDH.2018.00020

Lynn, T., 2023. Pose Estimation Algorithms: History and Evolution [WWW Document]. Roboflow Blog. URL https://blog.roboflow.com/pose-estimation-algorithms-history/ (accessed 8.6.23).

Ma, S., Zhang, Q., Li, T., Song, H., 2022. Basic motion behavior recognition of single dairy cow based on improved Rexnet 3D network. Computers and Electronics in Agriculture 194, 106772. https://doi.org/10.1016/j.compag.2022.106772

Manske, T., Hultgren, J., Bergsten, C., 2002. Topical treatment of digital dermatitis associated with severe heel-horn erosion in a Swedish dairy herd. Preventive Veterinary Medicine 53, 215–231. https://doi.org/10.1016/S0167-5877(01)00268-9

Manson, F.J., Leaver, J.D., 1988. The influence of concentrate amount on locomotion and clinical lameness in dairy cattle. Animal Science 47, 185–190. https://doi.org/10.1017/S0003356100003251

McDonagh, J., Tzimiropoulos, G., Slinger, K.R., Huggett, Z.J., Down, P.M., Bell, M.J., 2021. Detecting Dairy Cow Behavior Using Vision Technology. Agriculture 11, 675. https://doi.org/10.3390/agriculture11070675

McManus, C., Tanure, C.B., Peripolli, V., Seixas, L., Fischer, V., Gabbi, A.M., Menegassi, S.R.O., Stumpf, M.T., Kolling, G.J., Dias, E., Costa, J.B.G., 2016. Infrared thermography in animal production: An overview. Computers and Electronics in Agriculture 123, 10–16. https://doi.org/10.1016/j.compag.2016.01.027

MediaPipe, 2023. Pose landmark detection guide [WWW Document]. Google for Developers. URL https://developers.google.com/mediapipe/solutions/vision/pose_landmarker (accessed 8.6.23).

Memon, M.H., Kumar, W., Memon, A., Chowdhry, B.S., Aamir, M., Kumar, P., 2016. Internet of Things (IoT) enabled smart animal farm, in: 2016 3rd International Conference on Computing for Sustainable Global Development (INDIACom). Presented at the 2016 3rd International Conference on Computing for Sustainable Global Development (INDIACom), pp. 2067–2072.

Moysiadis, V., Sarigiannidis, P., Vitsas, V., Khelifi, A., 2021. Smart Farming in Europe. Computer Science Review 39, 100345. https://doi.org/10.1016/j.cosrev.2020.100345

Munawar, M.R., 2023. yolov7-pose-estimation.

Newell, A., Yang, K., Deng, J., 2016. Stacked Hourglass Networks for Human Pose Estimation. https://doi.org/10.48550/arXiv.1603.06937

Nguyen, C., Wang, D., Von Richter, K., Valencia, P., Alvarenga, F.A.P., Bishop–Hurley, G., 2021. Video-based cattle identification and action recognition, in: 2021 Digital Image Computing: Techniques and Applications (DICTA). Presented at the 2021 Digital Image Computing: Techniques and Applications (DICTA), pp. 01–05. https://doi.org/10.1109/DICTA52665.2021.9647417

Nielsen, P.P., 2013. Automatic registration of grazing behaviour in dairy cows using 3D activity loggers. Applied Animal Behaviour Science 148, 179–184. https://doi.org/10.1016/j.applanim.2013.09.001

Nóbrega, L., Gonçalves, P., Pedreiras, P., Pereira, J., 2019. An IoT-Based Solution for Intelligent Farming. Sensors 19, 603. https://doi.org/10.3390/s19030603

NVIDIA Corporation, 2022. TensorRT Open Source Software. NVIDIA Corporation.

Nye, J., Zingaretti, L.M., Pérez-Enciso, M., 2020. Estimating Conformational Traits in Dairy Cattle With DeepAPS: A Two-Step Deep Learning Automated Phenotyping and Segmentation Approach. Front Genet 11, 513. https://doi.org/10.3389/fgene.2020.00513

O'Callaghan, K.A., Cripps, P.J., Downham, D.Y., Murray, R.D., 2003. Subjective and Objective Assessment of Pain and Discomfort Due to Lameness in Dairy Cattle. Animal Welfare 12, 605–610. https://doi.org/10.1017/S0962728600026257

Oliveira, V.H.S., Sørensen, J.T., Thomsen, P.T., 2017. Can digital dermatitis be detected in the milking parlor without washing cows' feet? Research in Veterinary Science 115, 325–326. https://doi.org/10.1016/j.rvsc.2017.06.017

Orsel, K., Plummer, P., Shearer, J., De Buck, J., Carter, S.D., Guatteo, R., Barkema, H.W., 2018. Missing pieces of the puzzle to effectively control digital dermatitis. Transboundary and Emerging Diseases 65, 186–198. https://doi.org/10.1111/tbed.12729

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S., 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library, in: Advances in Neural Information Processing Systems. Curran Associates, Inc.

Prunier, A., Mounier, L., Neindre, P.L., Leterrier, C., Mormède, P., Paulmier, V., Prunet, P., Terlouw, C., Guatteo, R., 2013. Identifying and monitoring pain in farm animals: a review. animal 7, 998–1010. https://doi.org/10.1017/S1751731112002406

PyTorch, 2023. android-demo-app/ObjectDetection at master · pytorch/android-demo-app [WWW Document]. GitHub. URL https://github.com/pytorch/android-demo-app/tree/master/ObjectDetection (accessed 7.24.23).

pytorch, 2021. pytorch/pytorch. pytorch.

Qiao, Y., Truman, M., Sukkarieh, S., 2019. Cattle segmentation and contour extraction based on Mask R-CNN for precision livestock farming. Computers and Electronics in Agriculture 165, 104958. https://doi.org/10.1016/j.compag.2019.104958

Rahman, A., Smith, D.V., Little, B., Ingham, A.B., Greenwood, P.L., Bishop-Hurley, G.J., 2018. Cattle behaviour classification from collar, halter, and ear tag sensors. Information Processing in Agriculture 5, 124–133. https://doi.org/10.1016/j.inpa.2017.10.001

Read, D.H., Walker, R.L., 1998. Papillomatous Digital Dermatitis (Footwarts) in California Dairy Cattle: Clinical and Gross Pathologic Findings. J VET Diagn Invest 10, 67–76. https://doi.org/10.1177/104063879801000112

Reddy Maddikunta, P.K., Hakak, S., Alazab, M., Bhattacharya, S., Gadekallu, T.R., Khan, W.Z., Pham, Q.-V., 2021. Unmanned Aerial Vehicles in Smart Agriculture: Applications, Requirements, and Challenges. IEEE Sensors Journal 21, 17608–17619. https://doi.org/10.1109/JSEN.2021.3049471

Redmon, J., Farhadi, A., 2018. YOLOv3: An Incremental Improvement. arXiv:1804.02767 [cs].

Relun, A., Guatteo, R., Roussel, P., Bareille, N., 2011. A simple method to score digital dermatitis in dairy cows in the milking parlor. Journal of Dairy Science 94, 5424–5434. https://doi.org/10.3168/jds.2010-4054

Roboflow, 2023. YOLOv5 Instance Segmentation Instance Segmentation Model [WWW Document]. URL https://roboflow.com/model/yolov5-instance-segmentation (accessed 8.6.23).

Rodriguez-Lainz, A., Melendez-Retamal, P., Hird, D.W., Read, D.H., 1998. Papillomatous digital dermatitis in Chilean dairies and evaluation of a screening method. Preventive Veterinary Medicine 37, 197–207. https://doi.org/10.1016/S0167-5877(98)00091-9

Shearer, J.K., Hernandez, J., 2000. Efficacy of Two Modified Nonantibiotic Formulations (Victory) for Treatment of Papillomatous Digital Dermatitis in Dairy Cows. Journal of Dairy Science 83, 741–745. https://doi.org/10.3168/jds.S0022-0302(00)74936-8

Shearer, J.K., Van Amstel, S.R., Gonzalez, A., 2005. Manual of foot care in cattle. Hoard's Dairyman Books.

Shi, X., An, X., Zhao, Q., Liu, H., Xia, L., Sun, X., Guo, Y., 2019. State-of-the-Art Internet of Things in Protected Agriculture. Sensors 19, 1833. https://doi.org/10.3390/s19081833

Shinde, T.A., Prasad, D.J.R., 2017. IoT based Animal Health Monitoring with Naive Bayes Classification.

Snell, J., Swersky, K., Zemel, R., 2017. Prototypical Networks for Few-shot Learning, in: Advances in Neural Information Processing Systems. Curran Associates, Inc.

Solano, L., Barkema, H.W., Jacobs, C., Orsel, K., 2017. Validation of the M-stage scoring system for digital dermatitis on dairy cows in the milking parlor. Journal of Dairy Science 100, 1592–1603. https://doi.org/10.3168/jds.2016-11365

Song, L., Yu, G., Yuan, J., Liu, Z., 2021. Human pose estimation and its application to action recognition: A survey. Journal of Visual Communication and Image Representation 76, 103055. https://doi.org/10.1016/j.jvcir.2021.103055

Sprecher, D.J., Hostetler, D.E., Kaneene, J.B., 1997. A lameness scoring system that uses posture and gait to predict dairy cattle reproductive performance. Theriogenology 47, 1179–1187. https://doi.org/10.1016/S0093-691X(97)00098-8

Stilwell, G.T., Ferrador, A.M., Santos, M.S., Domingues, J.M., Carolino, N., 2019. Use of topical local anesthetics to control pain during treatment of hoof lesions in dairy cows. Journal of Dairy Science 102, 6383–6390. https://doi.org/10.3168/jds.2018-15820

Stokes, J., Leach, K., Main, D., Whay, H., 2009. Can locomotion scoring be used to identify cows with digital dermatitis? Cattle Lameness Conference, Sutton Bonnington, Nottingham, 25 March.

Stokes, J.E., Leach, K.A., Main, D.C.J., Whay, H.R., 2012a. An investigation into the use of infrared thermography (IRT) as a rapid diagnostic tool for foot lesions in dairy cattle. The Veterinary Journal, Special Issue:Bovine Lameness 193, 674–678. https://doi.org/10.1016/j.tvjl.2012.06.052

Stokes, J.E., Leach, K.A., Main, D.C.J., Whay, H.R., 2012b. The reliability of detecting digital dermatitis in the milking parlour. The Veterinary Journal, Special Issue:Bovine Lameness 193, 679–684. https://doi.org/10.1016/j.tvjl.2012.06.053

Stutz, P.V., Golani, L.K., Witkin, J.M., 2019. Animal models of fatigue in major depressive disorder. Physiology & Behavior 199, 300–305. https://doi.org/10.1016/j.physbeh.2018.11.042

Subeesh, A., Mehta, C.R., 2021. Automation and digitization of agriculture using artificial intelligence and internet of things. Artificial Intelligence in Agriculture 5, 278–291. https://doi.org/10.1016/j.aiia.2021.11.004

Tadich, N., Flor, E., Green, L., 2010. Associations between hoof lesions and locomotion score in 1098 unsound dairy cows. The Veterinary Journal 184, 60–65. https://doi.org/10.1016/j.tvjl.2009.01.005

techzizou, 2021. Build Android app for custom object detection (TF 2.x). TECHZIZOU. URL https://techzizou.com/build-android-app-for-custom-object-detection-using-tf2/ (accessed 7.24.23).

TensorFlow, 2023a. TensorFlow Lite | ML for Mobile and Edge Devices [WWW Document]. TensorFlow. URL https://www.tensorflow.org/lite (accessed 8.5.23).

TensorFlow, 2023b. TensorFlow Lite [WWW Document]. TensorFlow. URL https://www.tensorflow.org/lite/guide (accessed 8.5.23).

TensorFlow, 2023c. TensorFlow Lite for Android [WWW Document]. TensorFlow. URL https://www.tensorflow.org/lite/android (accessed 8.5.23).

TensorFlow, 2023d. TensorFlow Lite for Microcontrollers [WWW Document]. TensorFlow. URL https://www.tensorflow.org/lite/microcontrollers (accessed 8.5.23).

TensorFlow, 2023e. Post-training quantization | TensorFlow Lite [WWW Document]. TensorFlow. URL https://www.tensorflow.org/lite/performance/post_training_quantization (accessed 8.5.23).

TensorFlow, 2023f. Post-training dynamic range quantization | TensorFlow Lite [WWW Document]. TensorFlow. URL https://www.tensorflow.org/lite/performance/post_training_quant (accessed 8.5.23).

TensorFlow, 2023g. TensorFlow Lite Task Library [WWW Document]. TensorFlow. URL https://www.tensorflow.org/lite/inference_with_metadata/task_library/overview (accessed 8.5.23).

TensorFlow, 2022a. Object detection with Android | TensorFlow Lite [WWW Document]. URL https://www.tensorflow.org/lite/android/tutorials/object_detection (accessed 7.31.23).

TensorFlow, 2022b. TensorFlow model optimization | TensorFlow Model Optimization [WWW Document]. TensorFlow. URL https://www.tensorflow.org/model_optimization/guide (accessed 8.5.23).

TensorFlow, 2022c. TensorFlow Lite [WWW Document]. TensorFlow. URL https://www.tensorflow.org/lite/guide (accessed 8.5.23).

TensorFlow, 2022d. TensorFlow Lite NNAPI delegate [WWW Document]. TensorFlow. URL https://www.tensorflow.org/lite/android/delegates/nnapi (accessed 8.5.23).

TensorFlow, 2022e. Model conversion overview | TensorFlow Lite [WWW Document]. URL https://www.tensorflow.org/lite/models/convert (accessed 8.5.23).

TensorFlow, 2022f. Convert TensorFlow models | TensorFlow Lite [WWW Document]. URL https://www.tensorflow.org/lite/models/convert/convert_models (accessed 8.5.23).

tensorflow, 2021. TensorFlow Lite for Microcontrollers. tensorflow.

TensorFlow, 2021a. Model optimization | TensorFlow Lite [WWW Document]. TensorFlow. URL https://www.tensorflow.org/lite/performance/model_optimization (accessed 8.5.23).

TensorFlow, 2021b. TensorFlow Lite Delegates [WWW Document]. TensorFlow. URL https://www.tensorflow.org/lite/performance/delegates (accessed 8.5.23).

tensorflow, 2017. Documentation for TensorRT in TensorFlow (TF-TRT). tensorflow.

Thomsen, P.T., Klaas, I.C., Bach, K., 2008a. Short Communication: Scoring of Digital Dermatitis During Milking as an Alternative to Scoring in a Hoof Trimming Chute. Journal of Dairy Science 91, 4679–4682. https://doi.org/10.3168/jds.2008-1342

Thomsen, P.T., Sørensen, J.T., Ersbøll, A.K., 2008b. Evaluation of Three Commercial Hoof-Care Products Used in Footbaths in Danish Dairy Herds. Journal of Dairy Science 91, 1361–1365. https://doi.org/10.3168/jds.2007-0820

Tian, H., Wang, T., Liu, Y., Qiao, X., Li, Y., 2020. Computer vision technology in agricultural automation —A review. Information Processing in Agriculture 7, 1–19. https://doi.org/10.1016/j.inpa.2019.09.006

Tzanidakis, C., Tzamaloukas, O., Simitzis, P., Panagakis, P., 2023. Precision Livestock Farming Applications (PLF) for Grazing Animals. Agriculture 13, 288. https://doi.org/10.3390/agriculture13020288

Ultralytics, 2023a. Android [WWW Document]. URL https://docs.ultralytics.com/hub/app/android (accessed 7.24.23).

Ultralytics, 2023b. HUB App [WWW Document]. Ultralytics. URL https://ultralytics.com/app (accessed 7.24.23).

Ultralytics, 2023c. Ultralytics YOLOv8 Docs: Home [WWW Document]. URL https://docs.ultralytics.com/ (accessed 8.2.23).

Ultralytics, 2023d. Ultralytics YOLOv8 Docs: Pose [WWW Document]. URL https://docs.ultralytics.com/tasks/pose (accessed 8.6.23).

ultralytics, 2022a. ultralytics/yolov5. Ultralytics.

ultralytics, 2022b. v7.0 - YOLOv5 SOTA Realtime Instance Segmentation · ultralytics/yolov5 · Discussion #10258 [WWW Document]. GitHub. URL https://github.com/ultralytics/yolov5/discussions/10258 (accessed 8.6.23).

Ultralytics, 2022. Introducing Instance Segmentation in YOLOv5 v7.0 [WWW Document]. Ultralytics. URL https://ultralytics.com/article/Introducing-Instance-Segmentation-In-YOLOv5-v7.0 (accessed 8.6.23).

ultralytics, 2022c. yolov5/segment/tutorial.ipynb at master · ultralytics/yolov5 [WWW Document]. GitHub. URL https://github.com/ultralytics/yolov5/blob/master/segment/tutorial.ipynb (accessed 8.6.23).

ultralytics, 2021. ultralytics/yolov3. Ultralytics.

Van Nuffel, A., Zwertvaegher, I., Van Weyenberg, S., Pastell, M., Thorup, V.M., Bahr, C., Sonck, B., Saeys, W., 2015. Lameness Detection in Dairy Cows: Part 2. Use of Sensors to Automatically Register Changes in Locomotion or Behavior. Animals 5, 861–885. https://doi.org/10.3390/ani5030388

Vayssade, J.-A., Arquet, R., Bonneau, M., 2019. Automatic activity tracking of goats using drone camera. Computers and Electronics in Agriculture 162, 767–772. https://doi.org/10.1016/j.compag.2019.05.021

Wamuyu, P.K., 2017. A Conceptual Framework for Implementing a WSN Based Cattle Recovery System in Case of Cattle Rustling in Kenya. Technologies 5, 54. https://doi.org/10.3390/technologies5030054

Wang, C.-Y., Bochkovskiy, A., Liao, H.-Y.M., 2022a. YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. https://doi.org/10.48550/arXiv.2207.02696

Wang, C.-Y., Liao, H.-Y.M., Yeh, I.-H., 2022b. Designing Network Design Strategies Through Gradient Path Analysis.

Wang, C.-Y., Yeh, I.-H., Liao, H.-Y.M., 2021. You Only Learn One Representation: Unified Network for Multiple Tasks.

Whay, H., 2002. Locomotion scoring and lameness detection in dairy cattle. In Practice 24, 444–449. https://doi.org/10.1136/inpract.24.8.444

Whay, H.R., Waterman, A.E., Webster, A.J.F., 1997. Associations between locomotion, claw lesions andnociceptive threshold in dairy heifers during the peri-partum period. The Veterinary Journal 154, 155–161. https://doi.org/10.1016/S1090-0233(97)80053-6

Whay, H.R., Webster, A.J.F., Waterman-Pearson, A.E., 2005. Role of ketoprofen in the modulation of hyperalgesia associated with lameness in dairy cattle. Veterinary Record 157, 729–733. https://doi.org/10.1136/vr.157.23.729

Williams, L.R., Moore, S.T., Bishop-Hurley, G.J., Swain, D.L., 2020. A sensor-based solution to monitor grazing cattle drinking behaviour and water intake. Computers and Electronics in Agriculture 168, 105141. https://doi.org/10.1016/j.compag.2019.105141

Winckler, C., Willen, S., 2001. The Reliability and Repeatability of a Lameness Scoring System for Use as an Indicator of Welfare in Dairy Cattle. Acta Agriculturae Scandinavica, Section A — Animal Science 51, 103–107. https://doi.org/10.1080/090647001316923162

Wolfger, B., Jones, B.W., Orsel, K., Bewley, J.M., 2017. Technical note: Evaluation of an ear-attached real-time location monitoring system. Journal of Dairy Science 100, 2219–2224. https://doi.org/10.3168/jds.2016-11527

WongKinYiu, 2022. yolov7 [WWW Document]. Implementation of paper - YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. URL https://github.com/WongKinYiu/yolov7 (accessed 8.2.23).

Wu, D., Yin, X., Jiang, B., Jiang, M., Li, Z., Song, H., 2020. Detection of the respiratory rate of standing cows by combining the Deeplab V3+ semantic segmentation model with the phase-based video magnification algorithm. Biosystems Engineering 192, 72–89. https://doi.org/10.1016/j.biosystemseng.2020.01.012

Wurtz, K., Camerlink, I., D'Eath, R.B., Fernández, A.P., Norton, T., Steibel, J., Siegford, J., 2019. Recording behaviour of indoor-housed farm animals automatically using machine vision technology: A systematic review. PLOS ONE 14, e0226669. https://doi.org/10.1371/journal.pone.0226669

Yinka-Banjo, C., Ajayi, O., 2019. Sky-farmers: Applications of unmanned aerial vehicles (UAV) in agriculture. Autonomous vehicles 107–128.

Yousefi, D.B.M., Rafie, A.S.M., Al-Haddad, S.A.R., Azrad, S., 2022. A Systematic Literature Review on the Use of Deep Learning in Precision Livestock Detection and Localization Using Unmanned Aerial Vehicles. IEEE Access 10, 80071–80091. https://doi.org/10.1109/ACCESS.2022.3194507

Zanuttigh, P., Marin, G., Dal Mutto, C., Dominio, F., Minto, L., Cortelazzo, G.M., others, 2016. Time-of-flight and structured light depth cameras. Technology and Applications 978–3.