

The Role of Language Experience on Sentence Processing: The Effects and the Origins

By

Ya-ling Hsiao

A dissertation submitted in partial fulfillment of

The requirements for the degree of

Doctor of Philosophy

(Psychology)

at the

UNIVERSITY OF WISCONSIN-MADISON

2015

Date of final oral examination: 8/19/2015

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

Maryellen C. MacDonald, Professor, Psychology

Mark S. Seidenberg, Professor, Psychology

Martha W. Alibali, Professor, Psychology

Timothy T. Rogers, Associate Professor, Psychology

Benjamin Snyder, Assistant Professor, Computer Science

© Copyright by Ya-ling Hsiao 2015

All Rights Reserved

Table of Contents

List of Tables	ii
List of Figures	iii
Acknowledgement	v
Abstract	vii
Introduction.....	1
Chapter 1: Experience and Generalization in a Connectionist Model of Mandarin Chinese Relative Clause Processing	21
Chapter 2: Working Memory and Experience: A Connectionist Model of English Relative Clause Processing	78
Chapter 3: Production Predicts Comprehension: Animacy Effects in Mandarin Relative Clause Processing	106
Conclusion	166

List of Tables

Table 1, Chapter 1: The competing and neighbor structures of SRCs & ORCs and the at the two matrix positions.....	33
Table 2, Chapter 1: Token frequencies of overt subject simple sentences and pro-drop sentences found in Chinese Treebank 7.0.....	36
Table 3, Chapter 1: Token frequencies of subject- and object-modifying SRCs (transitive & intransitive) and ORCs at found in Chinese Treebank 7.0.....	37
Table 4, Chapter 1: Finite state grammar with corpus-based bigram transitional probabilities....	41
Table 5, Chapter 1: Comparison of experimental materials and major findings of representative human reading time studies	55
Table 1, Chapter 2: Token frequencies of simple sentences extracted from Wall Street Journal Treebank 2	89
Table 2, Chapter 2: Token frequencies of subject- and object-modifying subject relatives (transitive & intransitive) and object relatives (full & reduced) Wall Street Journal Treebank 2	89
Table 3, Chapter 2: Categories of words represented by the units and the number of units	91
Table 4, Chapter 2: Finite state grammar with corpus-based bigram transitional probabilities....	95
Table 1, Chapter 3: Frequency counts of object relatives and passive relatives in the Chinese Penn Treebank 7.0	129
Table 2, Chapter 3: Example items for the sentence completion task.....	134
Table 3, Chapter 3: Example items for the comprehension task	140

List of Figures

Figure 1, Chapter 1: Comparison of SRN performance in MacDonald & Christiansen (2002) and human reading times in Wells et al. (2009)	24
Figure 2, Chapter 1: Architecture of the Simple Recurrent Network used in Study 2	40
Figure 3, Chapter 1: Word-by-word GPEs of subject-modifying SRCs and ORCs in the critical RC region, the head, and next word after the head	46
Figure 4, Chapter 1: Word-by-word GPEs of subject-modifying SRCs and ORCs in the critical RC region, the head, and next word after the head	50
Figure 5, Chapter 1: GPEs of SRCs and ORCs at the head noun	54
Figure 1, Chapter 2: Architecture of the Simple Recurrent Network.	91
Figure 2, Chapter 2: GPEs of the two relative clause types at the subject-modifying position.....	97
Figure 3, Chapter 2: GPEs of the two relative clause types at the object-modifying position.	98
Figure 1, Chapter 3: An example experimental item for the production task in Experiment 1...121	
Figure 2, Chapter 3: An example picture that described the action “kiss” in the pre-training task.	122
Figure 3, Chapter 3: The percentage of structural choice as either in the active or passive relative clause forms when describing animate or inanimate entities.....	124
Figure 4, Chapter 3: Percentage of agent-drop among passive relative clause responses as a function of animacy of the head noun.....	124

Figure 5, Chapter 3: Percentage of completions with 1) relative clauses with animate heads, 2) relative clauses with inanimate heads, and 3) non-relative clause responses, given either the object relative fragment or the passive relative clause fragment at Early Gate and Late Gate ...	135
Figure 6, Chapter 3: The mean log residual reading times of active and passive relative clauses at each word position	143
Figure 7, Chapter 3: Regression results of log residual reading times at word positions being predicted by rate of completion at either the Early Gate (on the rate of sentence frame completed as a relative clause) or Late Gate (rate of relative clause completions with the matching head noun) for each item. (A):head noun, Early Gate, (B): head noun, Late Gate, (C): adverb, Late Gate, (D): main verb, Late Gate.....	148

Acknowledgement

This dissertation would have not been possible without the support of many people who participated in my life during my graduate school years. I would like to dedicate it to my parents, who have been giving me unconditional love and support throughout my life. It has not been easy for them to have to be far away from their daughter, living in another continent that they could only see once in a while. Although my mom was not able to see the completion of this dissertation, I know she would have been very pleased. My gratitude also goes to other family members: Ernesto Carrella, Leonardo Carrella, and other members of the Carrella and Hsiao family.

I would also express my tremendous gratitude to my mentor Maryellen MacDonald. She guided me throughout the journey with patience, kindness, and of course expertise in the field of psycholinguistics. She is never stingy on spending time with me discussing the big theoretical concepts and she can spend also hours working out details of experimental design or ways to structure a paper. She has opened up so many doors for me to different research topics and introduced me to other researchers and collaborators. She has helped me become a better scholar myself.

My thanks also go to fellow students in the Psychology Department, especially those in the Laboratory of Cognitive Neuroscience and Language. These people include former graduate students of the lab: Jessica Montage, Jon Willits, Lang Chen, and Michael Amoto, and current graduate students: Mark Koranda, Tianlin Wang, Christopher Cox, and April Murphy. Stimulating feedback and discussion has been very important in developing ideas and improving my works. Special thanks to the hard working research assistants over the years: Yannan Gao,

Tiantian Wang, Sarah Isaak, Qihong Lu, Joon Kim, Ping Yeh, Gang Yang, Danki Wang and people who have helped me in any ways. Without your effort in running participants and coding data, the works could never have been done. I would also like to mention names outside of the lab who have been great friends: Amanda Eggen and family, Bret Vlach, Xiaoming Ma, Yue Jiang, Xin Zhou, Hui-chun Chen, Chia-rung Li, Chien-ming Huang, and Chie-wei Chang.

I would also like to thank the following faculty members in the department: Mark Seidenberg, Tim Rogers, Martha Alibali, and Keith Kleunder (now in Purdue). I benefited a lot either from discussion with you about my own work or potential collaborative work with you, or a course that was offered by you that equipped me with the knowledge and skills for conducting my research.

Staff members in the department have also given me assistance in numerous occasions. These people include Kevin Belt, Gary Malchow, Benjamin Ball, John Govin and Weijin Shang. They guided me through paper work and helped me schedule and complete important milestones for my graduate studies.

Finally, I would like to thank my committee members: Maryellen MacDonald, Mark Seidenberg, Tim Rogers, Martha Alibali and Benjamin Snyder, for taking their time out to provide me feedback on this dissertation work.

Abstract

The dissertation examines the role of statistical distributions in shaping the comprehension behavior and further traces distributions to production choices that speakers make under the pressure of minimizing effort. The three studies presented here all involved the relative clause structure, a complex sentence type that has been extensively studied for the computational limitation of the parser. Among the three studies, two of them investigated the complex interactions of factors affecting processing difficulty besides relative clause structure type (subject relatives vs. object relatives). I implemented connectionist Simple Recurrent Networks trained on important statistics based on large corpora in Mandarin Chinese (Chapter 1) and English (Chapter 2). The model showed performance that reflected frequencies of not only the target relative clause structures but also those of neighboring structures similar in various properties. This generalization in learning implicates a different language processing account from other experience-based accounts. Moreover, the patterns of model performance replicated findings in previous empirical human studies, when manipulated factors were equated. The match between humans and a computer model without built in grammatical knowledge and a separate working memory system suggests that memory limitation is not separable from language experience. Moreover, previous studies that produced conflicting results, especially in Mandarin, could be subject to the limit of human experiments that only a limited number of factors can be examined. Computational modelling provides a venue to evaluate the complex interactions of factors.

After establishing that comprehenders tracks closely to fine grained aspects of language statistics, Chapter 3 examined the origins of the statistics through choices that speakers make. A series of studies, including on-line picture description task, corpus analyses, sentence completion

task, and self-paced reading task, demonstrated that producers make structure choices to alleviate difficulty and avoid interference given what is afforded by the language, and that comprehenders extract regularities from these choices that later become part of the language statistics. Using regression analyses, I showed that rate of production of a specific sentence type significantly predicted the reading time of sentences of the same type. The studies together establish a link among language production, distributions, and language comprehension.

Introduction

The language comprehension process is strongly modulated by ambiguities that exist in the sentence structure. When reading a sentence-initial fragment such as “The horse raced past the barn”, most people assume that this sentence is complete with a subject-verb-object (SVO) argument structure. However, the original interpretation needs to be revised when another verb appears after it, like “The horse raced past the barn fell”. The sentence is a reduced relative clause structure that could also be rephrased in a non-ambiguous way: “The horse that was raced past the barn fell”. Reanalysis from the simple sentence interpretation to the relative clause structure is cognitively taxing because the commitment to the simple SVO interpretation is so strong that it is difficult to recommit to a rarely encountered structure such as a relative clause. Such phenomenon is termed “the garden-path effect” (Bever, 1970). However, structural frequency is not the only factor involved in the garden-path effect. Semantic relationship in the lexical items that compose the argument structure is crucial in the formation of garden path interpretations. For example, “The horse ridden past the barn fell” is much less likely to be garden-pathed not only because “ridden” is in the past participle form and cannot be construed as the main clause verb but also because it is semantically implausible for a horse to ride on other things. The above examples of garden-path sentences illustrate the constraint-satisfaction characteristics of human language comprehension, where multiple aspects of linguistic experiences are drawn to make inferences about what the sentence means.

While reduced relative clauses are often studied for the role of ambiguity resolution in language comprehension, another set of relative clauses are also extensively investigated for another reason. Subject relative clauses and object relative clauses, exemplified in (1a) and (1b) respectively, are said to be unambiguous.

(1a) Subject Relative Clause: The candidate [who attacked the opponent] won this election.

(1b) Object Relative Clause: The candidate [who the opponent attacked] won this election.

Prior studies in English which used materials that matched in vocabulary and in event plausibility of the comparison between (1a) and (1b) found that subject relatives are easier to process than object relatives (e.g. Miller & Chomsky, 1963; Miller & Isard, 1964). Several language processing theories have emerged to account for this finding. Structural complexity has been attributed as the cause for difficulty with object relatives, which has the center-embedded structure that falls between the noun (“the candidate”) and its dependent verb phrases (“won this election”). Increasingly nested object relatives would increase comprehension difficulty, as exemplified by (2b). However, multiple embeddings of subject relatives, as in (2a), do not disrupt comprehension as much because the dependency relationships are easily integrated. Miller & Chomsky (1963) argued that our limited memory capacity fails us when trying to comprehend this sentence, even when we know it is grammatical according to our knowledge of grammatical rules.

(2a) The candidate [who attacked the opponent [who Mary met]] won this election.

(2b) The candidate [who the opponent [who Mary met] attacked] won this election.

This distinction of competence and performance has influenced the way of how linguists and psychologists theorize language acquisition and language processing for the past decades. The concept of competence grammar is deeply rooted in the tradition of generative grammar. Competence allows speakers to produce and understand an infinite number of sentences in their language. Performance, which can be affected by grammatically irrelevant conditions, such as memory limitations, attention, interest, and so on, has no relation to the speaker’s or hearer’s

competence (Chomsky, 1965). This idealized grammatical capacity is regarded as universal, innate and unlearnable, and linguists have adopted such notion to place syntax at the highest level from which semantics and phonology are developed. Syntactic complexity is thus considered central as the cause of processing difference between subject relatives and object relatives. However, more and more evidence has shown that learning does take place from tracking statistics at various levels of the language. Syntactic rules are the result of abstraction and internalization through exposure to linguistic input and all aspects in the past experience regarding syntax, semantics, prosody, discourse context and so on, can mutually constrain one another at the same time during language processing (e.g. MacDonald, Pearlmutter, & Seidenberg, 1994; Tanenhaus & Trueswell, 1995). Since grammatical knowledge emerges with language experiences, and such knowledge affects on-line performance, distinction between competence and performance is deemed unnecessary. Psycholinguistic accounts on language comprehension can be roughly defined to have developed upon these two distinctive perspectives: one that attributes comprehension difficulty to the inherent computation capacity, and the other to the multiple probabilistic constraints in the language experience. Next, I review these two accounts and supportive evidence from prior research.

Capacity-based Theories

The capacity-based theories consider syntactic complexity the major cause of processing cost. Comprehending complex sentence structures requires more working memory resources for storage and manipulation of previously encountered information (e.g., Baddeley, 1986; Engle, Cantor, & Carullo, 1992; Just & Carpenter, 1992), especially for structures that involve non-adjacent dependency relationship such as relative clauses. Gibson (1998) proposed a theory called “the Dependency Locality Theory”, which measures working memory demand in

separate ability called verbal working memory. Another major general criticism is related to the negligence of the role of language experience, which I discuss below.

Experience-based Theories

Contrary to the capacity-based perspective, developmental psychologists have provided general evidence for the role of experience in observed improvement, independent of age-related cognitive development. For example, children showed solid improvement after being exposed systematically to linguistic structures beyond their developmental grasp (Roth, 1984), suggesting that children's inability to understand and use relative clauses is the result of the input children receive. In the CHILDES corpus (Child Language Data Exchange System), which documents child-directed speech from care takers and children's production, Diessel & Tomasello, (2000) found that the majority of the earliest relative clauses in both input to children and children's speech are in the presentational form embedded in copular constructions (e.g. This is the sugar that goes in there.) Rather, the traditional forms used to test children's comprehension on relative clauses are attached to either the main clause subject (e.g. 1a) or the main clause object (e.g. 1b), both relatively rare in children's input and output. A similar finding comes from a written corpus analysis by Montag & MacDonald (2015), who analyzed the Corpus of Contemporary American English (COCA) for the texts from children's books. They found that children's production of passive relatives increased with age and such pattern matches higher occurrence of passive relatives in children's literature targeting older children.

Input frequency as a crucial factor is also suggested by results from corpus analyses that directly compare subject relatives and object relatives. Roland, Dick and Elman (2007) conducted a comprehensive analysis of several English spoken corpora and found a larger proportion of subject relatives than object relatives, again suggesting that human processing

difficulty of the latter may reflect the lack of practice with the structure. If direct experience with the target structures affects language processing, performance should significantly benefit from increased experience, especially for rare structures. Wells, Christiansen, Race, Acheson and MacDonald (2009) conducted a training study over several sessions by providing the participants in the experimental group large amount of exposure to relative clauses. They found that these participants' reading times on object relatives significantly decreased compared to the control group. Performance on subject relatives, however, did not differ across the two groups.

Wells et al. (2009)'s training results were replicated by MacDonald & Christiansen (2002) using a neural network. Additionally, they argued that the processing of relative clauses was not only affected by direct exposure but also by the experience with neighboring structures. Similar to the ambiguity resolution mechanism demonstrated by the garden-path sentence "the horse race past the barn fell", which shows competing effect from the SVO simple sentence structure to the reduced relative clause structure, the processing asymmetry of subject vs. object relatives may also be largely explained by the dominant word order in the language: subject relatives resemble the regular word order of subject-verb-object (SVO). Object relatives, on the other hand, have the unique SOV word order, which rarely occurs in the English language (see again 1a-b). Here, the processing difficulty of object relatives could be attributed to the lack of practice with the rare word order, on top of the rare structural frequency.

One important fact to note is MacDonald and Christiansen's SRN's training corpus — object and subject relatives were presented equally often (each type occurred 2.5% of time) but were not equally difficult, because the model generalized from simple transitive sentences to the similar subject relatives. Several other researchers have pursued this point, including Fitz et al. (2011), who used a dual-path SRN that models both sentence production and sentence semantics.

Although these models are not meant to be accounts of the acquisition process, similar generalization effects have been found in child language acquisition. For example, Diessel and colleagues found that the relative clauses that emerged earlier in English (subject relatives) do so because of their word order resemblance to the dominant word order in simple sentences (Diessel & Tomasello, 2005; Diessel, 2004, 2007).

In addition to structural distributions, lexical properties are also among the statistics that comprehenders track in linguistic input. Animacy of the head noun has been reported to have a significant modulating effect in the processing of the two relative clause types (Traxler, Morris, & Seely, 2002; Traxler, Williams, Blozis, & Morris, 2005). The difficulty associated with object relatives is mitigated by having an inanimate head noun (e.g. the movie that the director watched) and shown to be no harder to comprehend than subject relatives. Corpus analysis also displays a matching pattern, animate heads more often attached to subject relatives and inanimate heads to object relatives (Roland et al., 2007). This animacy effect appears to be universally true across languages, reflecting the world knowledge of animate entities mostly play the agent roles and inanimate entities often being the affected. Empirical evidence from human reading time data reflect the corpus distributional patterns (Dutch: Mak, Vonk, & Schriefers, 2002, 2006; Mandarin Chinese: Pu, 2007; Wu, Kaiser, & Andersen, 2012; Hsiao & MacDonald, 2013). The above findings imply that animacy serves an important cue for thematic role assignment, which in turn affects ambiguity resolution (Gennari & MacDonald, 2008).

In the next section, the capacity-based and experience-based processing accounts will be compared and evaluated.

Comparison of Capacity-based and Experience-based Theories

The capacity-based theories place emphasis on functional factors such as working memory constraints and integration or storage costs. The DLT theory is specifically based on the principle that the incomplete head dependencies creates demands on memory resources, making integration especially hard for object relatives due to the intervening discourse referents. The experience-based approach, on the other hand, explains the processing asymmetry between subject and object relatives through statistical regularities of the language. Syntactic processing is viewed to be constrained by a wide variety of probabilistic factors at various levels, such as syntax, semantics, lexical, contextual and so on. Comparing the two approaches, one may find the capacity-based theories appropriately account for syntactic complexity at the phrase structure level and the individual difference in cognitive resources, but one may also notice that the capacity-based theories are not able to address the significant lexical and semantic effects in empirical studies, such as animacy and pronoun use in relative clauses, and the tight relationship between computation capacity and reading experience. More and more recent psycholinguistic studies have shown supportive evidence for the experience-based approach, or a stance that incorporates both capacity and experience (e.g. Levy, Fedorenko, Breen, & Gibson, 2012; Levy & Gibson, 2013).

From Distribution to Production

As statistical learning from language distribution is supported by behavioral, computational and corpus studies on language comprehension, it is unclear how these distributional patterns come to be. To help answer this question, one may turn to another question: why do speakers make the production choices that they do? Production is a more demanding task than comprehension, requiring message to be translated into words, formulated

in grammatical sequences and finally articulated in a fluent manner. Understanding the connection between production and comprehension is crucial because these production patterns later constitute the distributional statistics that comprehenders encounter and implicitly learn.

Production research has identified several factors that affect structure choices. One of the most robust effects is accessibility. The more conceptually salient or more frequent words or phrases are often more easily retrieved from memory and thus produced early in the sentence or at more syntactically prominent positions (Bock, 1982; MacDonald, 2013). For example, animate nouns tend to be uttered early in the subject position, even when they are not thematically the agents (e.g. “The lightning struck the boy,” is often passivized into “The boy was struck by the lightning”) (Bock & Irwin, 1980; Bock, 1987; Bock & Warren, 1985; McDonald, Bock, & Kelly, 1993).

Another production strategy is avoidance of interference. Similarity-based interference is well-documented in language comprehension literature (Gordon, Hendrick, Johnson, & Lee, 2006). When two similar noun phrases need to be held in the working memory, like the case of object relatives, comprehension is likely interrupted. Similar effect exists in production, which requires longer planning time when two conceptually similar words or phrases are to be planned together. Smith & Wheeldon (2004) For example, people take more time to initiate the production of the phrase “the saw and the axe move down” than “the saw and the cat move down”. In real-life situations, speakers can usually choose among alternative structures to reduce such interference, such as a structure that separates similar nouns apart or omit one of the nouns completely. Gennari, Mirković, & MacDonald (2012) demonstrated that people avoided object relatives (e.g. the girl that the woman kissed) and produced the alternative passive voice (e.g. the girl that was kissed (by the woman)).

If the production system complies with the accessibility principle mentioned above at all time, the production patterns should be consistent whenever the circumstances are the same. Such production patterns persist across speakers and through time and gradually form abstract sequence patterns. Producers then re-use these patterns to further facilitate fluency. This phenomenon has been documented in many studies on syntactic priming, an effect of speakers repeating the type of sentences they use during language production (Pickering & Branigan, 1999). For example, when asked to describe a picture after hearing a semantically unrelated prime sentence in the passive voice, participants tend to respond with a passive sentence (Bock, 1986). Sometimes, the priming effect can be observed across structures that have similar surface forms but different semantic usage. For instance, the locative by-phrase in “The foreigner was loitering by the broken traffic light” can prime passive sentence involving an agentive by-phrase (Bock & Loebell, 1990). While these priming effects may seem temporary and fragile, research has shown that long-term implicit learning is in reality taking place whenever language users hear others’ or their own utterances that repeatedly adopt these syntactic patterns (Chang, Dell, Bock, & Griffin, 2000).

Connecting Production and Comprehension

MacDonald (2013) proposed a “Production-Distribution-Comprehension” framework, which accounts for the comprehension processes through production choices that make up the distributional statistics in the language. She termed the above three production biases as “Easy First”, “Reduce Interference” and “Plan Reuse”. These three factors jointly affect production choices when structure or word order flexibility is allowed. As mentioned previously, Gennari, Mirkovic and MacDonald (2012) investigated the production of object relatives (actives) and the alternative passive voice, manipulating the animacy of the modified noun. A majority of the

responses were passives when people described animate nouns. Moreover, the rate of omitting the by-phrase in the passives was higher for animates than for inanimates. The results may have reflected the bias of Reduce Interference, through setting the two semantically similar nouns apart by using passives, and further reducing the competition by dropping the by-phrase altogether. The choice of passives may have also reflected Easy First, which drove the participants to assign the animate head noun at the prominent subject position in the relative clause (i.e. it is both the head noun and the subject of the relative clause).

This study demonstrates the effects of accessibility and competition in on-line production, where producers compromise to an optimal structure in the language that best accommodates the constraints. In earlier studies, the team connected production choices to the actual statistical distributions of the language, extracted from corpus data (Gennari & MacDonald, 2009). They found that distributional patterns significantly predicted reading times in the comprehension task (Gennari and MacDoanld, 2008).

The above series of studies established not only the role of language experience on comprehension but also the interactive relationships among production, distribution and comprehension. However, the relationships have only been attested in English, which is very extensively studied as a language. The importance of cross-linguistic data lies in the fact that theories can be further examined for validity in other typologically different languages.

Next, I discuss the linguistic features and related past finding of relative clauses comprehension and production in Mandarin Chinese, a language I pursued for cross-linguistic data in this dissertation.

Mandarin Chinese Relative Clauses

Unlike English relative clauses, which are “head-first” (i.e. the relative clause appears after the head noun, as in “*the candidate [that attacked the opponent]*”), Mandarin relative clauses are “head-final”, such that the relative clause precedes the head noun, as in “[*attacked the opponent*] *the candidate*”. The word order difference in the two types of languages results in different predictions in language processing theories, both on comprehension and production. Since the head occurs on the opposite end of the clause, the working memory-based theories predict the head-final languages to have a contrastive comprehension pattern – namely, object relative clause advantage, regardless of the properties of the lexical items. Distributional frequencies are still predicted by the experience-based theories to be crucial in the statistical learning process of language comprehension.

To date, empirical studies have yielded conflicting results, with some studies finding the typical cross-linguistic pattern of subject relatives being easier than object relatives (e.g. Lin & Bever, 2006), and others finding the opposite result (e.g., Hsiao & Gibson, 2003). This reversal of the dominant cross-linguistic pattern (i.e. object relatives easier) finds a clear interpretation in the Dependency Locality account: subject relatives were more difficult than object relatives because of the longer distance between dependencies. However, multiple later studies failed to replicate this pattern (Jager et al., 2015; Lin & Bever, 2011; Vasishth et al., 2013) and corpus statistics show higher frequency of subject relatives over object relatives (Pu, 2007; Wu et al., 2012). While all these studies attempted to draw conclusions on whether subject relatives are harder than object relatives or vice versa, the real question lies in how these inconsistencies in the literature came to be. I proposed that, based on the insight we gained from English, comprehenders are sensitive to multiple statistical constraints in the language, which include

noun animacy, dominant word order and other properties, and this sensitivity may have resulted in the mixed results in prior studies, which used differing experimental materials. In fact, most of the previous studies examined only a narrow subset of relative clauses within a given type. As a result, the Mandarin relative clause research, which could in principle be extremely informative about cross-linguistic regularities and differences in complex sentence interpretation, is instead marginalized by a lack of consensus and by overly-broad claims based on a narrow range of stimulus materials.

Three Studies in this Dissertation

The conflicting findings in the Mandarin relative clause literature permit an opportunity to synthesize all the potential factors that have been not examined altogether previously. The fine-grained statistical knowledge of the language users hold likely allows complex and dynamic interactions among these factors in the comprehension process. Therefore, a comprehensive examination of the frequency distribution of these factors and how they jointly affect readers' expectation is necessary. Eventually these distributional patterns would be traced to the constraints and biases of the production system. The dissertation contained two major themes: one addressing the comprehension process with distributional information of important features (Chapter 1 and 2), the other tracing the distributional patterns to the production process (Chapter 3).

Specifically, Chapter 1 addresses the controversies in the comprehension literature on Mandarin relative clauses through corpus analyses and neural network modeling (Simple Recurrent Network, SRN). Model performance was then compared to human reading time data from previous studies. Because the corpus analyses and, therefore, the training set, were so detailed, we could compare experiment findings with model performance specifically for the

types of sentences used in several empirical studies. The model was able to resolve conflicting findings in the literature and identify broader themes in Mandarin relative clause processing.

Using the same methods, Chapter 2 is an SRN on English relative clause processing, also based on statistics from a large corpus. Different from MacDonald & Christiansen (2002)'s SRN, I examined more factors that potentially affect relative clause processing and implemented realistic statistics of structure frequencies gathered from the corpus. Another purpose of the model was to replicate the well-known finding of subject relative clause advantage in human studies and establish the validity of the learning mechanism of SRN. Moreover, I compare the English model to the Mandarin model in Chapter 1 for insight on cross-linguistic differences and commonalities in sentence comprehension more generally

After providing the evidence of language comprehension relying heavily on the statistical learning mechanism, I investigate the origin of the distributional patterns through the production system in Chapter 3, which includes a series of production and comprehension experiments that established the connection: spoken picture description task written sentence completion task, and self-paced reading task. The structural choice made in spoken production and sentence completion were compared and related to the patterns in reading time data. I argue that production constitutes an important driving force for the formation of linguistic distributional statistics, which later influence comprehenders' incremental prediction behavior during reading and listening.

Despite the theoretical linkage of the three chapters, each chapter can be read as a stand-alone paper. Chapter 1 has been published as Hsiao & MacDonald (2013), Chapter 2 is yet submitted for publication, and Chapter 3 is currently under review.

References

- Baddeley, A. D. (1986). *Working memory*. Oxford: Clarendon Press.
- Bever, T. G. (1970). The cognitive basis for linguistic structures. In : *J* (pp. 279–362). Editor, Cognition and the development of language, Wiley, New York R. Hayes.
- Bock, J. K. (1987). Meaning, sound and syntax: Lexical priming in sentence production. *Journal of Experimental Psychology*, *12*, 575–586.
- Bock, J. K., & Irwin, D. E. (1980). Syntactic effects of information availability in sentence production. *Journal of Verbal Learning and Verbal Behavior*, *19*, 467–484.
- Bock, J. K., & Loebell, H. (1990). Framing sentences. *Cognition*, *35*, 1–39.
- Bock, J. K. (1986). Syntactic persistence in language production. *Cognitive Psychology*, *18*, 355-387.
- Bock, K., & Warren, R. K. (1985). Conceptual accessibility and syntactic structure in sentence formulation. *Cognition*, *21*(1), 47–67.
- Chang, F., Dell, G. S., Bock, K., & Griffin, Z. M. (2000). Structural priming as implicit learning: A comparison of models of sentence production. *Journal of Psycholinguistic Research*, *29*(2), 217–230.
- Chomsky, N. (1965). *Aspects of the Theory of Syntax*. Cambridge, Mass.: MIT Press.
- Diessel, H., (2004). *The Acquisition of Complex Sentences*. Cambridge: Cambridge University Press, UK.
- Diessel, H. (2007). A construction-based analysis of the acquisition of East Asian relative clauses. *Studies in Second Language Acquisition*, *29*, 311-320.
- Diessel, H., & Tomasello, M. (2000). The development of relative clauses in spontaneous child speech. *Cognitive Linguistics*, *1*, 131–151.

- Diessel, H., & Tomasello, M. (2005). A new look at the acquisition of relative clauses. *Language*, 81, 1-25.
- Engle, R. W., Cantor, J., & Carullo, J. J. (1992). Individual differences in working memory and comprehension: A test of four hypotheses. *Journal of Experimental Psychology: Learning Memory and Cognition*, 18, 972–992.
- Fitz, H., Chang, F., & Christiansen, M. H. (2011). A connectionist account of the acquisition and processing of relative clauses. In E. Kidd (Ed.), *The Acquisition of Relative Clauses: Processing, Typology and Function*. (Vol 8) (pp. 39–60). Amsterdam: John Benjamins.
- Gennari, S. P., & MacDonald, M. C. (2008). Semantic indeterminacy in object relative clauses. *Journal of memory and language*, 58(2), 161–187.
- Gennari, S. P., & MacDonald, M. C. (2009). Linking production and comprehension processes: The case of relative clauses. *Cognition*, 111(1), 1–23.
- Gennari, S. P., Mirković, J., & MacDonald, M. C. (2012). Animacy and competition in relative clause production: A cross-linguistic investigation. *Cognitive Psychology*, 65(2), 141–176. doi:10.1016/j.cogpsych.2012.03.002
- Gibson, E. (1998). Linguistic complexity: Locality of syntactic dependencies. *Cognition*, 68(1), 1–76.
- Gibson, E., Desmet, T., Grodner, D., Watson, D., & Ko, K. (2005). Reading relative clauses in English. *Cognitive Linguistics*, 16(2), 313.
- Gordon, P. C., Hendrick, R., Johnson, M., & Lee, Y. (2006). Similarity-based interference during language comprehension: Evidence from eye tracking during reading. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 32, 1304–1321.

- Hsiao, F., & Gibson, E. (2003). Processing relative clauses in Chinese. *Cognition*, *90*(1), 3–27.
- Hsiao, Y., & MacDonald, M. C. (2013). Experience and generalization in a connectionist model of Mandarin Chinese relative clause processing. *Frontiers in Psychology*, *4*.
doi:10.3389/fpsyg.2013.00767
- Jager, L., Chen, Z., Li, Q., Lin, C.-J. C., & Vasishth, S. (2015). The subject-relative advantage in Chinese: Evidence for expectation-based processing. *Journal of Memory and Language*, *79-80*:97-120.
- Just, M. A., & Carpenter, P. A. (1992). A capacity theory of comprehension: Individual differences in working memory. *Psychological Review*, *99*, 122–149.
- King, J., & Just, M. A. (1991). Individual differences in syntactic processing: The role of working memory. *Journal of Memory and Language*, *30*(5), 580–602.
- Levy, R., Fedorenko, E., Breen, M., & Gibson, T. (2012). The processing of extraposed structures in English. *Cognition*, *12*, 12–36.
- Levy, R., & Gibson, E. (2013). Surprisal, the PDC, and the primary locus of processing difficulty in relative clauses. *Frontiers in Language Sciences*, *4*(229), 1–3.
- Lin, C.-J. C., & Bever, T. G. (2006). Subject preference in the processing of relative clauses in Chinese. In *Proceedings of the 25th west coast conference on formal linguistics* (pp. 254–260). Retrieved from <http://www.iub.edu/~lacl/att/paper1456.pdf>
- MacDonald, M. C. (2013). How language production shapes language form and comprehension. *Frontiers in Psychology*, *4*. doi:10.3389/fpsyg.2013.00226

- MacDonald, M. C., & Christiansen, M. H. (2002). Reassessing working memory: Comment on Just and Carpenter (1992) and Waters and Caplan (1996). *PSYCHOLOGICAL REVIEW-NEW YORK-*, *109*(1), 35–54.
- MacDonald, M. C., Pearlmutter, N. J., & Seidenberg, M. S. (1994). Lexical nature of syntactic ambiguity resolution. *Psychological review*, *101*(4), 676–703.
- McDonald, J. L., Bock, K., & Kelly, M. H. (1993). Word and world order: Semantic, Phonological and metrical determinants of serial order. *Cognitive Psychology*, *25*, 188–230.
- Mak, W. M., Vonk, W., & Schriefers, H. (2002). The influence of animacy on relative clause processing. *Journal of Memory and Language*, *47*, 50–68.
- Mak, W. M., Vonk, W., & Schriefers, H. (2006). Animacy in processing relative clauses: The hikers that rocks crush. *Journal of Memory and Language*, *54*(4), 466–490.
- Miller, G. A., & Chomsky, N. (1963). Finitary models of language users. In R. D. Luce, R. Bush, & E. Galanter (Eds.), *Handbook of Mathematical Psychology* (pp. 419–492). New York: John Wiley.
- Miller, G.A. & Isard, S. (1964). Free recall of self-embedded English sentences. *Information and Control*, *7*, 292--303.
- Montag, J. L., & MacDonald, M. C. (2015). Text exposure predicts spoken production of complex sentences in eight and twelve year old children and adults. *Journal of Experimental Psychology: General*, *144*, 447-468.

- Pickering, M. J., & Branigan, H. P. (1999). Syntactic priming in language production. *Trends in Cognitive Sciences*, 3(4), 136–141.
- Pu, M. M. (2007). The distribution of relative clauses in Chinese discourse. *Discourse Processes*, 43(1), 25–53.
- Roland, D., Dick, F., & Elman, J. L. (2007). Frequency of basic English grammatical structures: A corpus analysis. *Journal of Memory and Language*, 57(3), 348–379.
doi:10.1016/j.jml.2007.03.002
- Roth, F. P. (1984). Accelerating language learning in young children. *Journal of Child Language*, 11, 89–107.
- Smith, M. C., & Wheeldon, R. L. (2004). Horizontal Information Flow in Spoken Sentence Production. *Journal of Experimental Psychology: Learning Memory and Cognition*, 30, 675–686.
- Tanenhaus, M. K., & Trueswell, J. C. (1995). Sentence comprehension (pp. 217–262). San Diego: Academic Press.
- Traxler, M. J., Williams, R. S., Blozis, S. A., & Morris, R. K. (2005). Working memory, animacy, and verb class in the processing of relative clauses. *Journal of Memory and Language*, 53(2), 204–224.
- Traxler, Matthew J., Morris, R. K., & Seely, R. E. (2002). Processing subject and object relative clauses: Evidence from eye movements. *Journal of Memory and Language*, 47(1), 69–90.

Vasishth, S., Chen, Z., Li, Q., & Guo, G. (2013). Processing Chinese Relative Clauses: Evidence for the Subject-Relative Advantage. *PLoS ONE*, 8(10), e77006.

doi:10.1371/journal.pone.0077006

Wells, J. B., Christiansen, M. H., Race, D. S., Acheson, D. J., & MacDonald, M. C. (2009).

Experience and sentence processing: Statistical learning and relative clause comprehension. *Cognitive psychology*, 58(2), 250–271.

Wu, F., Kaiser, E., & Andersen, E. (2012). Animacy effects in Chinese relative clause processing.

Language and Cognitive Processes, 27(10), 1489–1524.

Chapter 1: Experience and Generalization in a Connectionist Model of Mandarin Chinese Relative Clause Processing

Sentence comprehension is generally considered to be a complex constraint satisfaction process integrating probabilistic information from syntactic, semantic, prosodic, and discourse sources (e.g. MacDonald et al., 1994; Tanenhaus & Trueswell, 1995). This emphasis on multiple probabilistic constraints in sentence comprehension demands precise accounts of how constraints of different types and different strengths are weighed, so as to yield clearly testable models of comprehension. Unfortunately, compared to a large number of empirical studies in sentence comprehension, there are relatively few implemented computational models of sentence processing phenomena, which could illuminate the interaction of complex probabilistic constraints in sentence comprehension.

One issue that has been addressed in computational models of sentence comprehension is the role of computational capacity in accounts of human comprehension behavior. In particular, a key question in comprehension research is the separability between linguistic knowledge and the capacity to use that knowledge in comprehension and other language behavior. These issues are familiar in the competence-performance distinction that has traditionally distinguished much of linguistic and psycholinguistic research (e.g. Miller & Chomsky, 1963), but it also arises within psycholinguistic accounts of complex sentence comprehension—how much is human comprehension difficulty attributable to limitations on human working memory capacity, independent of people’s experience with language? For example, several different accounts have been offered for the difference in comprehension difficulty for subject relative clauses and object relative clauses in English and many other languages. An example can be seen in sentences (1a-

b), where the object relatives (1b) are generally found to be more difficult than subject relatives (1a).

(1a) Subject Relative Clause: The candidate [**who₁ attacked₁ the opponent**] won this election.

(1b) Object Relative Clause: The candidate [**who₁ the opponent attacked₁**] won this election.

A common argument for the difference in comprehension difficulty between these two sentence types points to the role of working memory, that the object relatives (1b) place higher working memory demands on the comprehender than do the subject relatives (e.g. King & Just, 1991). In one variant of this view, the Dependency Locality Theory (Gibson, 1998), the working memory demands are tied to greater distance between related elements (shown with subscripts in 1a-b) in object relatives (1b) compared to subject relatives (1a). This additional distance in (1b) requires longer memory maintenance of the partially processed information (“the candidate”) until it can be integrated with the action (“attacked”). Failure to maintain or retrieve the information disrupts comprehension. Thus on this view, comprehension difficulty is directly tied to the capacity to maintain discontinuous elements during sentence comprehension.

These questions have also been addressed in computational models of sentence processing. Simple Recurrent Network (SRN) models of sentence comprehension have a computational capacity that is inherently tied to the model’s experience with linguistic input (MacDonald & Christiansen, 2002; this is true of connectionist models more generally, e.g. McClelland & Elman, 1986). As originally implemented by Elman (1990), an SRN is a partially recurrent network equipped with an additional context layer that stores the output of the hidden layer, which is paired with the next input to the network. The model’s task is to predict the next input word, and the presence of the context layer permits the prior linguistic context to influence

the prediction of later words. SRNs have been applied to several different types of complex sentences (e.g., Elman, 1991; 1993; Christiansen, 1994; Christiansen & Chater, 1999), including the contrast in difficulty between subject and object relatives, as in (1a-b). Direct comparisons of human and model performance reveal important similarities. Figure 1 shows MacDonald and Christiansen's (2002) SRN and Wells et al.'s (2009) human subjects' reading times for subject and object relatives in studies in which experience with the two structures was explicitly manipulated in the model and in the human readers. MacDonald and Christiansen's model had extensive experience with several kinds of simple sentences, while only 5% of the sentences in the training set contained relative clauses. Their model was tested at three different points in training to investigate a claim about how these models generalize from the common simple sentences, such as *The candidate attacked the opponent*, to relative clauses such as (1a-b). They hypothesized that interpretation of subject relatives would be aided by these sentences' similarities to simple sentences, and as a result, the model would rapidly learn to make accurate predictions for subject relatives, and it would show little effect of additional training, as it had already benefited from the overlap with the common simple sentence "neighbors". By contrast, object relatives have an idiosyncratic word order that is not aided by extensive experience with simple sentences, and MacDonald and Christiansen predicted that as a result, the model would be extremely sensitive to the degree of direct experience with object relatives. Figure 1 shows that these predictions were supported, and it also shows that human reading times in the Wells et al. study were similarly influenced by a training manipulation, in which additional exposure to subject and object relatives over the course of a month affected participants reading patterns for object relatives (right panel) but not subject relatives (left panel).

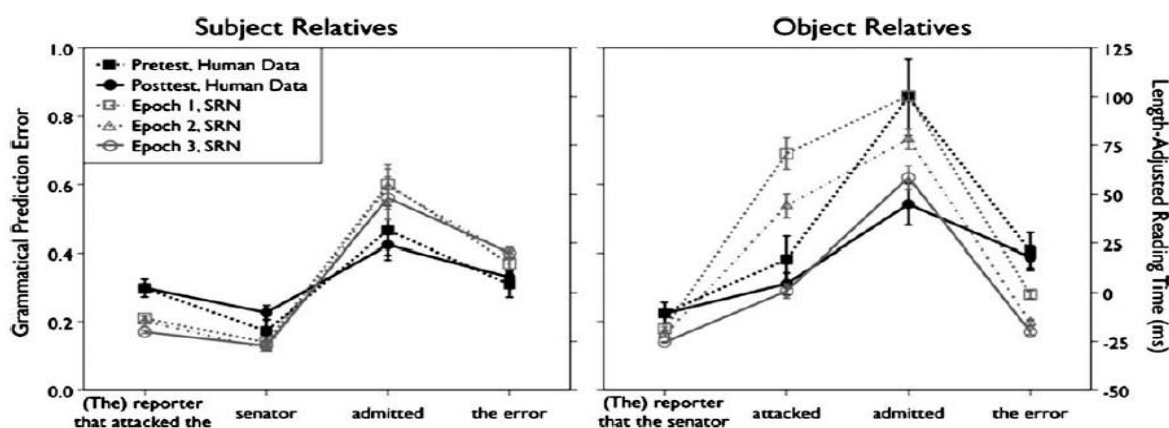


Figure 1. Comparison of SRN performance in MacDonald & Christiansen (2002) and human reading times in Wells et al. (2009)

An important feature of SRNs is that they generalize over their training experiences with individual sentences to find regularities across the training items (Elman, 1990; St. John & McClelland, 1990). The results in Figure 1 reflect generalization in MacDonald and Christiansen’s SRN—in that object and subject relatives were presented equally often but were not equally difficult, because the model generalized from simple transitive sentences to the similar subject relatives. Several other researchers have pursued this point, including Fitz et al. (2011), who used a dual-path SRN that models both sentence production and sentence semantics. Although these models are not meant to be accounts of the acquisition process, similar generalization effects have been found in child language acquisition. For example, Yip and Mathews (2007) found that the relative clauses that emerged earlier in Cantonese (object relatives) do so because of their word order resemblance to the dominant word order in simple sentences (for similar effects in other languages, see Diessel, 2004, 2007; Diessel & Tomasello, 2005; Ozeki & Shirai, 2007)

These results hold promise for SRN accounts of sentence comprehension, but to date, their application has been quite limited, and in particular the models have not typically

incorporated realistic frequencies of various sentence types, which would better allow researchers to understand how complex probabilistic constraints are weighed in comprehension. In this paper, we take steps toward meeting this challenge with an SRN model of relative clause comprehension that accurately represents critical elements of the distributional knowledge that humans bring to bear in interpreting these structures. Rather than elaborating accounts of English relative clause comprehension, about which an enormous body of research exists, we address relative clause comprehension in Mandarin Chinese. As we detail below, Mandarin relative clause comprehension is particularly interesting because a) the relative clause structure is quite different from English, b) some key findings about comprehension difficulty show potentially opposite patterns than in English and in many other languages, c) there is a fair amount of controversy concerning comprehenders' performance in various structures and discourse contexts and d) the number of important factors affecting comprehension appears to be too large to be manipulated together in empirical studies. Thus a computational account of Mandarin relative clause processing has the opportunity to have a substantial impact informing the nature of constraint satisfaction processes in Mandarin and more generally.

Mandarin Relative Clauses

Relative clauses in English and many other languages are “head-first”, meaning that the relative clause appears after the head noun, as in *the candidate [that attacked the opponent]*. In many of such languages, comprehenders show a clear pattern of finding subject relatives easier than object relatives (English: e.g. King & Just, 1991; Gibson et al., 2005; Dutch: e.g. Frazier, 1987; German: e.g. Schriefers et al., 1995; French: e.g. Frauenfelder et al., 1980). Other languages, including Korean, Japanese, and Mandarin, have a “head-final” relative clause structure, such that the relative clause precedes the head noun, as in the Mandarin examples in

(2). In these examples, in which the relative clause modifies the subject of the main clause, the relative clause begins the sentence, ending with the relativizer DE (equivalent to the English *that* in this context), followed by the noun phrase (head) being modified, in this case *candidate*. This head-final relative clause word order is a critical factor in accounts of cross-linguistic differences in relative clause processing, because the different relative clause structures create different degrees of distance between dependent elements, as shown by subscripts in (2).

(2a) Subject-modifying subject relative: [攻擊 競爭者 的] 候選人 贏了 這場 選舉

[e_1 attack opponent DE] candidate₁ won this election

The candidate who attacked the opponent won this election.

(2b) Subject-modifying object relative: [競爭者 攻擊 的] 候選人 贏了 這場 選舉

[opponent attack e_1 DE] candidate₁ won this election

The candidate who the opponent attacked won this election.

Most studies of relative clause processing in Korean and Japanese have suggested that subject relatives are easier than object relatives (Japanese: e.g. Miyamoto & Nakamura, 2003; Korean: Kwon et al., 2010), similar to English and other head-first languages. Some researchers suggest that this pattern points to a universal subject preference for relative clause processing, as first proposed by Keenan and Comrie under the term “Accessibility Hierarchy” (1977), which argues that noun phrases at the subject position are the easiest to be relativized due to their higher syntactic position, compared to noun phrases lower in the hierarchy, such as genitive (though cf. Ozeki & Shirai, 2007; Yip & Matthews, 2007). However, the pattern of subject preference does not appear to apply to Mandarin, which differs in several ways from the other languages considered here. Mandarin has the head-final relative clause structure, like Japanese

and Korean, but whereas Japanese and Korean have a rich system of case marking on nouns that presumably aids relative clause processing, case-marking is nonexistent in Mandarin. Mandarin is also different in that it has the dominant word order of subject-verb-object (SVO) in main clauses, like English and many European languages, but with the head-final relative clause structure that is absent in these languages. This combination of SVO basic word order, head final relative clause structure and absence of case marking is attested in the world's languages only in four Sino-Tibetan languages such as Bai, and other Chinese languages like Mandarin, Cantonese, and Hakka (Keenan, 1985; Haspelmath et al. 2005).

These features make Mandarin an interesting test case to disentangle the influences of various potential factors in relative clause processing. To date, empirical studies have yielded conflicting results, with some studies finding the typical cross-linguistic pattern of subject relatives being easier than object relatives (e.g. Lin & Bever, 2006), and others finding the opposite result (e.g., Hsiao & Gibson, 2003). This reversal of the dominant cross-linguistic pattern finds a clear interpretation in the Dependency Locality account. Hsiao and Gibson argued that subject relatives (e.g. 2a) were more difficult than object relatives (e.g. 2b) because they have higher storage and integration costs, owing to the longer distance between dependencies in subject relatives: there are more intervening words between the filler and the gap and thus more new discourse referents and incomplete dependencies in a subject relative than in an object relative. Support for this memory-based view comes from a study in which the added difficulty of subject relatives was higher in participants with lower working memory span (Chen et al., 2008) and patients with aphasia (Su et al., 2007). Several studies manipulating other factors, such as relative clause topicalization (Lin & Garnsey, 2011) and context (Gibson & Wu, 2013), also found a similar object relative advantage.

However, a reading time advantage for subject relatives has also been reported in several studies. Vasishth et al. (in press) modified Hsiao and Gibson (2003)'s materials and failed to replicate their effect, instead finding that object relatives were harder than subject relatives. Lin & Bever (2011) found no difference between the two types of relative clauses modifying the main clause subject (as in 2a-b) but found shorter reading times for subject relatives than object relatives when the relative clause was modifying the main clause object, such that subject relatives like (3a) were easier than object relatives like (3b). They also manipulated whether the participants were informed that they were reading relative clauses and about which noun positions they were modifying. They found that participants who were not informed had the most difficulty reading object-modifying ORCs, whose word order in combination with the matrix clause (i.e. the first three words in 3b) could lead the readers into a garden path of a simple sentence.

(3a) Object-modifying subject relative: 選民 支持 [攻擊 競爭者 的] 候選人

voter support [e_1 attack opponent DE] candidate₁

Voters support the candidate who attacked the opponent.

(3b) Object-modifying object relative: 選民 支持 [競爭者 攻擊 的] 候選人

voter support [opponent attack e_1 DE] candidate₁

Voters support the candidate who the opponent attacked.

Still other studies investigated the effect of lexical semantics and found it to have a modulating influence on the difficulty of the two relative clause types. Wu et al. (2012) manipulated the animacy of both the head noun and the relative clause noun and found that comprehension difficulty in both sentence types was strongly dependent on whether the sentence

contained a canonical animacy configuration, in which animate entities acted on inanimate ones. These results resonate well with the Mak et al. (2002; 2006) and Traxler et al. (2002, 2005) studies in other languages, in that animacy serves an important cue for thematic role assignment, which in turn affects ambiguity resolution (Gennari & MacDonald, 2008).

More generally, these results, together with absence of case marking and head-final relative clause structure in Mandarin, suggest that difficulty in relative clause comprehension may be strongly modulated by temporary ambiguities in sentences, where comprehenders initially interpret nouns and verbs in the input as being part of a main clause only to realize later that they had encountered a relative clause.

Summarizing over these various studies of Mandarin relative clause processing, it is impossible to draw conclusions about overall comprehension difficulty of the two sentence types. The inconsistency across studies and the sensitivity of the results to multiple factors such as animacy and modifying position may suggest that comprehenders are able to use very detailed information about patterns of relative clause use in comprehending these structures. Thus while many prior studies argued for one structure being absolutely easier than the other, a closer look at the materials in these studies suggests researchers' conclusions are limited by their methodological choices and stimulus items. For example, while claims about relative clause processing difficulty are typically phrased to cover all relative clauses of a certain type (such as subject relatives), in fact, most of the previous studies examined only a narrow subset of relative clauses within a given type. For example, most studies have examined only relative clauses that modified main clause subjects, and containing only animate head nouns and relative clause nouns. Such relative clauses are in reality rare in the linguistic environment (Pu, 2007; Wu et al., 2012; Vasishth et al., in press). As a result, the Mandarin relative clause research, which could

in principle be extremely informative about cross-linguistic regularities and differences in complex sentence interpretation, is instead marginalized by a lack of consensus and by overly-broad claims based on a narrow range of stimulus materials.

While it is effectively impossible to combine all the potentially important factors in a single empirical study that would allow us to observe complex interactions among them, it is possible to examine many of these effects in an SRN. We conducted our study in two phases, reflecting the fact that an attractive property of SRNs is their strong sensitivity to the distributional patterns in the language, a property that also appears in human comprehension. As described in Study 1, we examined the distributional statistics of relative clauses in a large corpus, extensively hand-coding the corpus data for features we believe to be critical in relative clause comprehension. In Study 2 we used this distributional information to develop a finite state grammar from which we developed training materials that faithfully represented key properties of relative clauses identified in the corpus. We used these training materials to train an SRN that was exposed to several different types of simple sentences and relative clauses. Model performance was then compared to human reading time data from previous studies. Because the corpus analysis and, therefore, the training set, were so detailed, we can compare experiment findings with model performance specifically for the types of sentences used in several empirical studies. In this way, we aim to use the model to help resolve the conflicting findings in the literature and be able to identify broader themes in Mandarin relative clause processing and cross-linguistic differences and commonalities in sentence comprehension more generally.

Study 1 – Corpus Analysis

Humans' comprehension of relative clauses is influenced not only by their prior experience with relative clauses but also by their experience with other sentences in the language. For example, MacDonald & Christiansen (2002) suggested that English speakers' experiences with simple transitive sentences aided comprehenders' interpretation of subject relative clauses, which are similar in that nouns in analogous sentence positions receive the same thematic role assignments. In this sense simple transitive sentences are helpful "neighbors" of subject relatives in English, as experience with these highly frequent simple sentences allows generalization to the rare subject relative structure (see also Fitz et al., 2011). MacDonald & Christiansen found a similar neighborhood effect in their SRN, attributable to the overlap in word order between the two sentence types, as the SRN does not assign thematic roles or interpret sentences. Conversely, humans' prior experiences with other kinds of sentences can increase the difficulty of relative clause comprehension. Gennari & MacDonald (2008) showed that object relative clauses in English contain temporary ambiguities for which the object relative clause is often an infrequent and disfavored interpretation. In this case, experience with other, more frequent sentence types affects the process of ambiguity resolution, leading to difficulty in interpreting object relatives. We will call these alternative interpretations "competitors," recognizing that both neighbor and competitor effects reflect comprehenders' generalizations over their prior linguistic experiences.

Mandarin relative clauses may similarly be affected by competitors and neighbors. First, Mandarin head-final relative clauses exhibit temporary ambiguities such that the unfolding linguistic input has several alternative interpretations. Second, certain Mandarin relative clauses have highly similar word orders to some more frequent simple sentences; generalization over these common neighbor sentences should help relative clause processing. Note that the same

sentence type may serve both competitor and neighbor functions at different points, in that it might provide an alternative interpretation that affects ambiguity resolution early in processing but that following a point of disambiguation, certain overlap with a relative clause may help in eventual relative clause interpretation (see Fitz et al., 2011, for further discussion). Appendix A describes some of the main neighbors and competitors for relative clauses in Mandarin.

To investigate the range of competitors and neighbors of Mandarin relative clauses, which should be informative for both human and computational work, and to develop a realistic training corpus for our model, we conducted a corpus analysis that enabled us to calculate the statistics of various types of relative clauses, crucial competitor structures that could increase comprehension difficulty, and neighboring structures that could reduce comprehension difficulty.

Here it is necessary to introduce some clarifying terminology, because it can be confusing to discuss both subject and object relative clauses (in which the modified head noun is the subject or object of the relative clause verb, respectively) crossed with the main clause subject vs. object modifying positions (in which the relative clause-modified noun is either the subject or object of the main clause; see examples 2a-b and 3a-b above). For the remainder of this paper, we will refer to relative clauses as RCs, and the subject and object relatives as SRCs and ORCs respectively. We will continue to spell out subject- vs. object-modifying positions in the main clause, so that full word descriptions are associated with main clause factors and acronyms are used for the embedded clauses.

Table 1 shows the competing and neighbor structures that we investigate in the present study and the ambiguities or facilitation they create at different points of an SRC and an ORC. The cells with gray shading indicate sentence types in the training set for which we do not expect

large effects on RC processing in the model. As the table shows, some sentences may serve both competitor and neighbor functions at different points in the sentence.

Table 1. The competing and neighbor structures (listed on the top) we examine in the current study and the ambiguities and facilitation they create at different points for SRCs and ORCs at the two matrix positions (listed on the left)

		Overt Subject simple Sentences: N V {N}	Pro-Drop Simple Sentences: V {N}	Subject-modifying intransitive SRCs: [V DE] N V...	Object-modifying intransitive SRCs: N V [V DE] N.
Subject-modifying	SRCs: [V N DE] N V...		Competitor: before DE Neighbor: after DE (for V N word order)	Neighbor: early (promotes RC interpretation of first V)	
	ORCs: [N V DE] N V...	Competitor: before DE (interpret the initial N V order as start of simple sentence) Neighbor: after DE (similar N V N order)			
Object-modifying	SRCs: N V [V N DE] N.				Neighbor: early (promotes RC interpretation of first V)
	ORCs: N V [N V DE] N.	Competitor: from beginning (interpret RC N as object N of simple sentence)			

Note: N=noun, V=verb, DE=relativizer, SRC=subject relative clause, ORC=object relative clause, [Square brackets] indicate the relative clause, {Curly brackets} indicate optional direct object N, which is present following transitive verbs and absent following intransitive verbs.

*Because the model did not include those verbs that can be both transitive and intransitive (e.g. "I ate" vs. "I ate an apple"), intransitive SRCs were not listed as a competitor to transitive SRCs.

Methods

We used Tgrep2 1.15 (Rohde, 2005) to extract sentences from a parsed corpus of spoken and written language, the Penn Chinese Treebank 7.0 (Xue, et al., 2010). The corpus consists of more than one million words in more than fifty thousand sentences, with sources from Chinese newswire, broadcast news, magazine news, broadcast conversation programs, web newsgroups, and others. The Tgrep2 search patterns used to extract the sentences are contained in Appendix B. Our aim was to retrieve every subject- and object-modifying ORC, every subject- and object-modifying SRC (with both transitive and intransitive verbs), every single-clause pro-drop construction, and every overt subject single-clause simple sentence in the corpus. A total of 6255 sentences were extracted.

We limited extraction of simple sentences to those with only a single clause, because our model was not designed to process other types of multi-clause sentences. However, we extracted all RCs that modified a sentential subject or object regardless of whether the rest of the sentence was a simple main clause or a more complex sentence. We cast this wider net for RCs because they are fairly rare in Mandarin, and extracting literally every RC in the corpus would give us the best estimate of the relative frequency of different RC types and animacy configurations. These methodological choices yielded a higher ratio of RC sentences to non-RC sentences than in the whole corpus (because many other non-RC multi-clause sentences were not extracted). Nonetheless, simple sentences outnumbered RCs at about 5:1 in the extracted set of 6255 sentences.

For each sentence, noun animacy was coded by hand, as NP animacy is known to affect Mandarin RC processing (Wu et al., 2012), and we expected that the animacy configuration of

simple sentences could also serve as important experience for RC processing. Coding followed the criteria that animate nouns refer to living entities that possess agency and volition to perform an action, while inanimate nouns refer entities without these properties (Hundt, 2004). For example, plants are living things but do not have volition, and thus they were coded as inanimate. Nouns that represented a group of people, such as organizations, countries, etc. (e.g. *the school that taught me math*), were considered animate. However, when these nouns were used purely as locations (e.g. *the school that I went to*), they were coded as inanimate. Coding was performed by two native speakers of Mandarin who were instructed in these coding criteria.

Results

Table 2 reports the frequencies of all simple sentences that were coded, and Table 3 reports all relative clauses. Inspection of these tables reveals the following patterns:

(1) Simple sentences with an overt subject (a potentially helpful neighbor of ORCs) were the most frequent (N=3309) were more frequent than pro-drop simple sentences (a competitor for some RC interpretations, N=1885), and sentences with RCs (N=1061) were less frequent than either of these simple sentence types. Thus both some competitor interpretations and helpful neighbor interpretations are more frequent than RCs themselves.

(2) The majority of (overt) main clause subjects were animate and the majority of main clause objects were inanimate. As main clauses are more common than RCs, these patterns of main clause animacy may influence expectations for RC animacy configurations.

(3) There were relatively more subject-modifying RCs (N=636, 60% of all RCs) than object-modifying RCs (N=415). SRCs (transitive and intransitive combined) occurred more often than

ORCs in both modifying positions. Transitive SRCs were the most frequent among the three types of RCs.

(4) ORCs, regardless of modifying position, mostly had inanimate head nouns and animate RC embedded nouns, consistent with patterns observed in English (Roland et al., 2007). Transitive SRCs had a higher proportion of animate heads and inanimate RC embedded nouns in the subject-modifying position, whereas there was not a big difference in head animacy but a preference of inanimate RC embedded nouns in the object-modifying position. Intransitive SRCs at both modifying positions had more inanimate heads than animate heads. These relatively high rates of inanimate head nouns, which generally followed the patterns of animacy usage in main clauses, could be attributed to the high percentage of sentences like “The growth rate rose,” in the newswire genre, which comprises a large subset of the Chinese Treebank 7.0 corpus.

Table 2. Token frequencies of overt subject simple sentences and pro-drop sentences found in Chinese Treebank 7.0

Simple Sentences with overt grammatical subject						
	Animate Subject Nouns			Inanimate Subject Nouns		
	Object Noun Type			Object Noun Type		
Verb Type	Animate	Inanimate	none	Animate	Inanimate	none
Intransitive			162			702
Transitive	355	1477		121	492	
Pro-drop Sentences (with grammatical subject omitted)						
	Object Noun Type					
Type	Animate		Inanimate		none	
Intransitive					295	
Transitive	539		1051			

Note: gray cells mark nonexistent combinations of verb transitivity and object type; for example, intransitive verbs by definition have no direct objects, and so cells representing the animacy coding of objects are not relevant for intransitive verbs.

Table 3. Token frequencies of subject- and object-modifying SRCs (transitive & intransitive) and ORCs at found in Chinese Treebank 7.0

	Subject-Modifying RCs					
	Animate Head Nouns			Inanimate Head Nouns		
	Relative Clause (RC) Noun Type			Relative Clause (RC) Noun Type		
RC-type	Animate	Inanimate	none	Animate	Inanimate	none
SRC, Intransitive			14			56
SRC, Transitive	61	209		68	31	
ORC	26	0		163	8	
	Object-Modifying RCs					
	Animate Head Nouns			Inanimate Head Nouns		
	Relative Clause (RC) Noun Type			Relative Clause (RC) Noun Type		
RC-type	Animate	Inanimate	none	Animate	Inanimate	none
SRC, Intransitive			34			74
SRC, Transitive	30	62		27	76	
ORC	15	0		79	18	

Note: Relative Clause Nouns are the Relative clause object in SRC Transitive sentences and the Relative clause subject in ORC sentences. Gray cells reflect nonexistent sentence types (because transitive verbs must have a relative clause noun and intransitive verbs do not have a direct object noun).

Discussion

This corpus analysis yielded several important patterns. First, the configuration of animacy in the corpus, in both main clauses and relative clauses, is strikingly similar to findings in other languages and also replicates and extends previous Mandarin corpus studies (Pu, 2007; Wu et al., 2012). Key results here include the tendency for main clause subjects to be animate and objects to be inanimate, the tendency of SRC heads to be animate and ORC heads to be inanimate, and the rarity of RCs with two nouns of the same animacy (both animate or both

inanimate, Wu et al., 2012). The general similarity of these patterns to main and relative clause usage in other languages (e.g., Bock & Warren, 1985; Gennari & MacDonald, 2009; Roland et al., 2007) suggests that the difference in patterns of relative clause interpretation in Mandarin vs. other languages does not lie in different animacy configurations and must instead reflect other critical cross-linguistic differences. Candidates for other important cross-linguistics differences include the very high degree of temporary ambiguity encountered during Mandarin RC processing, owing to the combination of head-final RCs and the absence of case marking, and also particular patterns of potentially helpful neighboring structures. The complexity of the potential interactions here is quite large, and in the next study, we use a Simple Recurrent Network to explore the effect of these linguistic patterns on relative clause processing.

Study 2 – Simple Recurrent Network

We chose a Simple Recurrent Network (SRN) to model Mandarin RC processing because of its potential to simulate word-by-word human reading times, including in prior studies of relative clause processing (Fitz et al., 2011; MacDonald & Christiansen, 2002; Wells et al., 2009). We trained the model on a mix of relative clauses, helpful neighbor sentences, and competitor sentences (contributing to temporary ambiguities) in proportions based on the Study 1 corpus analysis to investigate how these varied experiences could jointly contribute to relative clause processing. It is important to note that SRNs are not simulating human language comprehension per se but are instead a simulation of a component thought to be a part of comprehension and a factor related to reading times, namely prediction of upcoming input. Thus while terms such as “ambiguity resolution,” “alternative interpretations,” and “garden-path effect” are common in descriptions of human reading times, the model is not adopting any interpretation or calculating meaning. Ambiguity created by the existence of multiple interpretations in human

sentence processing is more adequately termed as “conditional indeterminacy” in the case of an SRN. As the probabilities of grammatical continuations are more varied, the higher the prediction error is for the model. Thus for both model and humans, indeterminacy is costly, and the model is taken to represent one important aspect of the ways in which input can be ambiguous for humans.

Methods

The SRN was a connectionist network using the backpropagation learning algorithm. It contained a context layer in addition to input, output and hidden layers. These layers were connected by trainable weights, except that the context layer directly copied the activations of the hidden layer from the previous time tick. The input pattern (the words of a sentence) is activated in the input layer one word at a time and then propagated onto the hidden layer and the output layer at time step t . The weights are adjusted by comparing the output activations to the desired output. At the next time step $t+1$, the output activation on the hidden layer at time step t is copied to the context layer and then projected backed to the hidden layer to pair with the current input.

Model specification

As shown in Figure 2, the model contained 28 localist units in the input and output layer, representing 27 words, plus one “end-of-sentence” marker. Although the network did not model semantics, animacy was captured distributionally, meaning that units designating “animate” nouns appeared frequently as sentence subjects and rarely as objects, while units designating inanimate nouns had the opposite pattern. The grammatical categories included in this model were the following: Animate Nouns (7 units), Inanimate Nouns (7 units), Transitive Verbs (6

units), Intransitive Verbs (6 units), RC marker DE (1 unit), and End-of-Sentence marker (1 unit). There were 40 units each in the hidden layer and the context layer. The learning rate was set to 0.05, with momentum of 0.9 and the batch size of 1. The simulation was conducted with the software Lens (Version 2.63) (Rohde, 1999).




Figure 2. Architecture of the Simple Recurrent Network used in Study 2

Training

Based on the frequency information obtained from Study 1, we calculated the bigram transitional probabilities from one grammatical category to another, as shown in Table 3. Transitional probabilities were calculated from the probability of occurrence of Y given the previous input X. For example, given the occurrence of a transitive verb (VT), the next word was an animate noun (aniN) 18% of the time, an inanimate noun (inaN) 73% of the time and the start of an object modifying relative clause (objRC) 9% of the time. Similarly, the probability of the occurrence of a subject-modifying SRC with an animate head and intransitive verb is $0.7*0.16*0.11*0.2 = 0.0025$, meaning that there are approximately 25 sentences that contain this type of RC in the 10,000-sentence training corpus.

The model's training faithfully reflected the relative frequencies of RCs found in the corpus, but as with all computational models, it is a simplification of the knowledge that humans bring to the task. These simplifications include the absence of semantics, the limited vocabulary, the omission of other uses of DE (see Appendix A), and having a higher proportion of RCs in the training set than in the corpus as a whole. Some of these features (RC frequency, definitive disambiguation at DE) may make RCs proportionally less difficult for the model than for humans, and other features (absence of fine grained semantics or real world or discourse context, smaller neighborhood effects) could make RCs yield proportionally more error for the model than would be expected based on human reading times. Given that all models necessarily include simplifications, the choices here represent a good starting point with which to examine Mandarin RC processing.

Table 4. Finite state grammar with corpus-based bigram transitional probabilities

S → subNP + VP (0.7) / VP (0.3) (meaning that 70% of sentences had a subject NP and 30% were pro-drop sentences, without a subject NP)	
VP → VI (0.21) / VT + objNP (0.79) (i.e., 21% intransitive verbs and 79% transitive verbs with direct object NPs)	
subNP → aniN(0.5) / inaN (0.34) / subRC (0.16)	objNP → aniN(0.18) / inaN (0.73) / objRC (0.09)
subRC (modifying matrix subject) → SRC_VI (0.11): VI + DE + aniN(0.2) / inaN(0.8) → SRC (0.58): (0.35) VT + aniN + DE + aniN(0.47) / inaN(0.53) (0.65) VT + inaN + DE + aniN(0.87) / inaN(0.13) → ORC (0.31): (0.96) aniN + VT + DE + aniN(0.14) / inaN(0.86) (0.04) inaN + VT + DE + aniN(0) / inaN(1)	objRC (modifying matrix object) → SRC_VI (0.26): VI + DE + aniN(0.31) / inaN(0.69) → SRC (0.47): (0.29) VT + aniN + DE + aniN(0.53) / inaN(0.47) (0.71) VT + inaN + DE + aniN(0.45) / inaN(0.55) → ORC (0.27): (0.84) aniN + VT + DE + aniN(0.16) / inaN(0.84) (0.16) inaN + VT + DE + aniN(0) / inaN(1)

S = sentence; NP = noun phrase; VP = verb phrase; VI = intransitive verb; VT = transitive verb; subNP = subject noun phrase; objNP = object noun phrase; aniN = animate noun; inaN = inanimate noun; subRC = subject-modifying relative clause; objRC = object-modifying relative clause; SRC_VI = subject relative clause with intransitive verb; SRC = subject relative clause with transitive verb; ORC = object relative clause; DE = relative clause marker

A python script was written based on the finite state grammar with the transitional probabilities in Table 4 to produce the 10,000 training sentences. The grammar did not permit sentences with multiple RC embeddings because multiply embedded RCs are very rare in the human corpus. The sentences encompassed the following structures: Transitive/intransitive simple sentences with/without subject, intransitive/transitive subject relative clauses, and transitive object relative clauses. The three types of relative clauses (SRCs with transitive and intransitive verbs and ORCs) could modify either main clause subjects or objects. The script selected words within a given category at random, with equal probability for each word. As with other SRNs, there were very few word units in the model, and all words in a category (such as animate nouns, transitive verbs) were equally frequent as the other words in the category.

Model assessment

The model performance was evaluated via Grammatical Prediction Error (GPE), which has been shown to relate well to behavioral measures such as reading times and grammaticality judgments (Christiansen & Chater, 1999). The measure is based on the idea that, because the current output activation is generated by the model reflecting the context in the previous time ticks, the model should activate what is expected/grammatical and should not activate what is ungrammatical, as defined by the training corpus. The GPE therefore incorporates the concepts of “hits” (sum of activation in grammatical units) and “false alarms” (sum of activation in ungrammatical units), with “misses” (sum of insufficient activation of grammatical units), as described in Christiansen and Chater (1999) and shown below.

$$\text{GPE} = 1 - \frac{\text{hits}}{\text{hits} + \text{false alarms} + \text{misses}}$$

The misses are derived from the difference between the actual output activation from the target activation, based on the transitional probabilities in the training corpus. For example, the target activation of the units representing VT (a transitive verb) at the sentence initial position should sum up to about 0.3 because that value is the sum of all the sentence types with a sentence-initial VT: first, there is 30% chance that the sentence starts as a pro-drop sentence, which is in turn 79% likely to be composed of a transitive verb with an object noun phrase, yielding $0.3 * 0.79 = 0.24$. Second, a sentence-initial VT may start an SRC, with a probability of $0.7 * 0.16 * 0.58 = 0.065$. Combining the probabilities of the two situations ($0.24 + 0.065$), we should obtain a target activation for the sentence-initial VT of around 0.3. If in this example the sum of the total output activation in the VT units is actually 0.26, then the misses should be 0.04. Inside the sentence, the target activation distribution was affected by prior context. For example, following the sentence-initial VT, grammatical continuations were aniN (animate noun), inaN (inanimate noun), VT, and VI (intransitive verb). If the next word was an aniN, the VT+aniN fragment could be the start of a) a pro-drop simple sentence with an animate patient ($0.3 * 0.79 * 0.18 = 0.043$) or b) a subject-modifying SRC ($0.7 * 0.16 * 0.58 * 0.35 = 0.023$). The total probability of VT+aniN occurring in the whole training corpus was $0.043 + 0.023 = 0.066$ but the probability of aniN following VT should be weighted among all four possible continuations. That is, summing up the probabilities of sentence-initial VT+aniN (0.066), VT+inaN ($0.3 * 0.79 * 0.73 + 0.7 * 0.16 * 0.58 * 0.65 = 0.215$), VT+VT ($0.3 * 0.79 * 0.09 * 0.58 = 0.012$), VT+VI ($0.3 * 0.79 * 0.09 * 0.11 = 0.002$) in the corpus, the total probability of sentences starting with an VT should be around 0.30, as calculated above. Among all four legal continuations, the relative proportion of aniN appearing after VT was $0.066 / 0.30 = 0.22$, which should be the summed target

activation for all the units representing aniN. The same procedure applied to the other continuations.

It should be noted that overestimation was also implicitly penalized due to the fact that dislocated activation found in some units means insufficient activation in some other units because the total output activation sums to around 1. The GPE ranges from 0 to 1, with 0 being perfectly accurate in predicting the grammatical categories of the current word based on prior context (that is, 0 error) and 1 being completely incorrect in doing so.

GPEs reflect the prediction error of the next word based on the cumulative context (e.g. the GPE of Word 3 is directly affected by Word 2, which in turn is affected by Word 1), and therefore may implicitly simulate the spillover effect observed in human reading patterns, where reading of one word may be affected by properties of prior input.

Testing

Ten networks with different random initial weights, ranging from 1 to -1 with the mean of 0, were created and trained on the 10,000 sentences in the training corpus. These ten networks simulated the role of “participants” in empirical studies, each of whom may have come from different backgrounds with varying prior experiences and skills. The models were trained on one pass through the training set and were tested using novel test sentences (sentences not contained in the training set) that allowed us to assess the major types of relative clauses that have been investigated in the comprehension literature. There were 16 different types of test sentences, each with 10 tokens, for a total of 160 test sentences. The 16 sentence types were obtained from crossing factors to yield the relative clause types shown in Table 3: two modification positions (main clause subject, object) x relative clause type (SRC vs. ORC) x head noun animacy

(animate, inanimate) and RC noun animacy (animate, inanimate). Intransitive SRCs (also shown in Table 3) were not included in the test set because there are no human reading time data in the literature for this sentence type. The GPE scores were calculated at each word in the critical RC region, the head noun, and the word after the head noun.

Results

In the sections below, we first present results of the model performance at each modifying position (main clause subject or object), noting the general relationships to human data. Then in later sections we present a detailed comparison between the model GPEs and the results of specific experiments. Statistical analyses were conducted with mixed effect models with maximal random effects of participants and items using the lme4 packages in R. Significance values were estimated by likelihood ratio test as suggested in Barr et al. (2013).

Subject-modifying RCs

Figure 3 presents the mean GPEs of SRCs and ORCs at each of five word positions comprising a sentence-initial relative clause, a head noun, and the next word. Statistical analyses were conducted to examine three effects – RC type, RC noun animacy, and head noun animacy – along with their interactions. Note that there are no effects of head noun animacy through the first three words, because this factor does not appear in the sentence until the head noun is reached at the fourth word position. The head animacy effects also illustrate a crucial effect of GPE calculations. Before the head noun was encountered, only the effects of RC type, embedded noun animacy, and their interaction were factored into the analyses from Word 1 to DE. The effect of head animacy and its interactions with other two factors were considered only

at the head and head+1 positions.

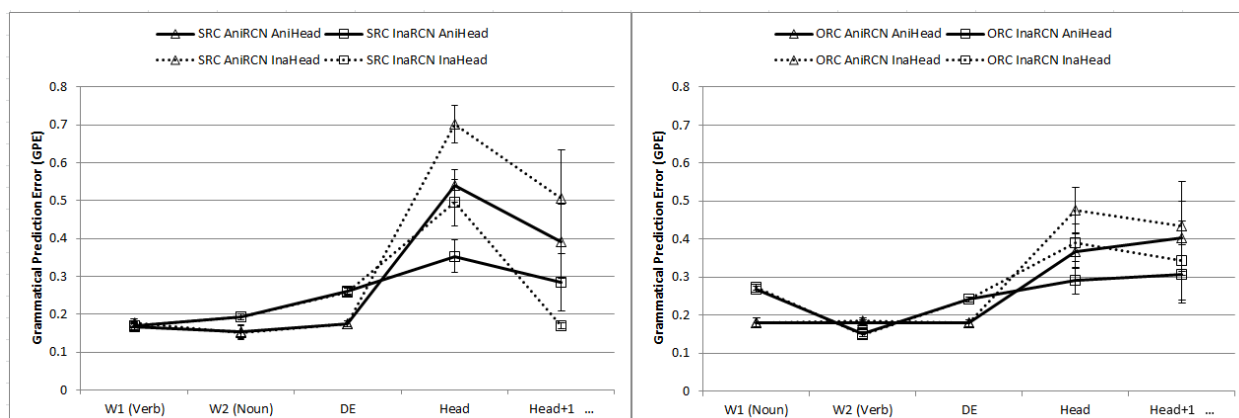


Figure 3. Word-by-word GPEs of subject-modifying SRCs (left panel) and ORCs (right panel) in the critical RC region, the head, and next word after the head

The most important result for these sentences is the contrast in difficulty at the head position for SRCs (left graph) and ORCs (on the right). As we noted in Tables 2-3, SRCs modifying a main clause subject are more ambiguous than ORCs: a sentence-initial verb might be the start of an SRC, but competitor pro-drop structures (simple sentences without an overt grammatical subject) are much more common. Sentence-initial object relatives are comparatively less ambiguous, and the effects these different amounts of ambiguity are clear in the model, with reliably higher GPEs for subject relatives than object relatives at the head position (detailed analyses at each word position are given below). Thus even though sentence-initial SRCs with transitive verbs are more frequent than ORCs at a ratio of almost 2:1, SRCs are substantially harder for the model. This result replicates major findings in the comprehension literature (Hsiao & Gibson, 2003; Gibson & Wu, 2013) and shows how the model's behavior on a particular structure is not simply a reflection of the frequency of that structure. We will elaborate these points after describing the effects at each word position.

Word-by-word analyses.

At W1 (RC object for SRCs, RC verb for ORCs), both main effects of RC type ($\beta=-.009$, $SE=.003$, $t=-3.08$, $p=.002$) and animacy of the RC noun ($\beta=.090$, $SE=.008$, $t=11.69$, $p<.001$) were significant. SRCs were easier than ORCs, and RCs with animate embedded nouns were easier than those with inanimate embedded nouns. An RC type x animacy of RC noun interaction was present ($\beta=-.093$, $SE=.008$, $t=-11.24$, $p<.001$), primarily because ORCs with inanimate RC nouns were the hardest among all conditions.

At W2, the RC type main effect showed an SRC advantage ($\beta=-.031$, $SE=.013$, $t=-2.34$, $p<.001$). The main effect of embedded noun animacy was such that RCs with inanimate embedded nouns were easier than those with animate embedded nouns ($\beta=-.033$, $SE=.003$, $t=-10.70$, $p<.001$). An RC type x RC noun animacy effect existed ($\beta=.074$, $SE=.012$, $t=6.39$, $p<.001$), where SRCs were particularly hard with inanimate RC nouns.

At the RC marker DE, no RC type main effect existed. The RC noun animacy effect remained ($\beta=.063$, $SE=.007$, $t=8.71$, $p<.001$), where RCs with animate embedded nouns were easier. A significant interaction effect of RC type x embedded noun animacy was present ($\beta=.022$, $SE=.007$, $t=3.05$, $p=.008$). SRCs with inanimate RC nouns were particularly hard.

At the head, a main clause transitive or intransitive verb, head animacy was added as a factor in the analysis. All three main effects were significant. The RC type effect showed ORC advantage ($\beta=.017$, $SE=.052$, $t=3.32$, $p=.001$), different from earlier in the RC region. Animacy of the embedded noun stayed significant ($\beta=-.077$, $SE=.028$, $t=-2.78$, $p=.043$), with inanimate embedded nouns being easier. Head animacy was also reliable ($\beta=.011$, $SE=.031$, $t=3.57$,

$p < .001$), where inanimate headed-RCs yielded higher error. Interaction effects were not significant. At the head+1 position, no main effects or interactions were present.

Discussion

Beyond the greater difficulty of SRCs compared to ORCs at the head noun, we observed several other reliable effects in RCs modifying main verb subjects. First, we observed a tendency of early difficulty for RCs with inanimate embedded nouns. Encountering an inanimate noun early in the sentence, a position where animate nouns usually appear as an agent and the grammatical subject, yielded high error. Second, there was an effect of head noun animacy at the head noun, such that animate-headed RCs yielded lower error, reflecting the high frequency of animate nouns being at the main clause subject position in the training corpus. This result has been attested in other corpus data (Wu, 2009) and behavioral studies (Wu et al., 2012).

Third, differences between the two RC types changed over the course of the sentence. ORCs initially yielded higher error rates than SRCs, mostly driven by high error for the (unusual) sentence-initial inanimate noun in ORCs. That is, the performance of the model here shows that simple sentences were initially a competitor (error is high for inanimate nouns sentence-initially, because in the more common simple sentences, sentence-initial nouns are animate). These results are compatible with the results of reading time studies that manipulated noun animacy (Wu et al., 2012). However, other empirical data that found an ORC advantage also exist (e.g. Hsiao & Gibson, 2003; Gibson & Wu, 2013; cf. Lin & Bever, 2006, for non-significant differences). As one reviewer noted, GPEs at the sentence-initial verb in SRCs were no larger than GPEs at the sentence-initial noun in ORCs, despite the fact that sentence-initial nouns are much more common than sentence-initial verbs (a 70:30 ratio). The reason may be attributed to

the prevalence of verb-starting sentences (i.e. pro-drop sentences and SRCs) in the training corpus and the rather low number of simple sentences compared to the realistic statistics in the language. Sentences starting with a transitive verb ($0.3*0.79+0.7*0.16*0.58=30\%$) in the training set occurred almost as often as sentences starting with an animate noun ($0.7*0.5=35\%$). Sentences starting with an inanimate noun, however, were much fewer ($0.7*0.34=23.8\%$) and thus generated higher error. The error triggered by noun-starting sentences as reflected in transitional probabilities averaged across animate and inanimate nouns is 29.5%, even lower than the 30% for transitive verb-starting sentences. The lower errors in SRCs than ORCs may have been the reflection of such subtle difference in the transitional probability. However, at the same time, this may reflect lower sensitivity of the model to sentence-initial variation, in the absence of prior context, or to the comparatively low number of simple sentences to RCs in the training set, which may have affected sentence initial predictions more than in later regions where prior context is available. Future work with a more varied corpus and more training may clarify this result.

At the RC marker DE, the SRC advantage disappeared and became the opposite pattern at the head: ORCs yielded lower error than SRCs. This effect is consistent with many prior studies in Mandarin that found lower reading times for ORCs (e.g. Hsiao & Gibson, 2003; Lin & Garnsey, 2011; Gibson & Wu, 2013). Importantly, this higher error for SRCs over ORCs appears even though SRCs are more frequent: The model's performance does not simply reflect RC frequencies but also the frequency of competitors and neighbor structures. Next, we look at the model performance on object-modifying RCs, which are more constrained by previous context.

Object-modifying RCs

To facilitate comparisons across conditions, all test sentences for object-modifying RCs began with an animate main clause subject noun and a transitive verb. As the sentence unfolded to the main clause object position, possible grammatical continuations were more limited than in the sentence-initial RCs reviewed above, owing to the greater preceding context in object-modifying items. Figure 4 shows the trajectory of GPE values for SRCs and ORCs in the RC region, head, and the end-of-sentence marker.

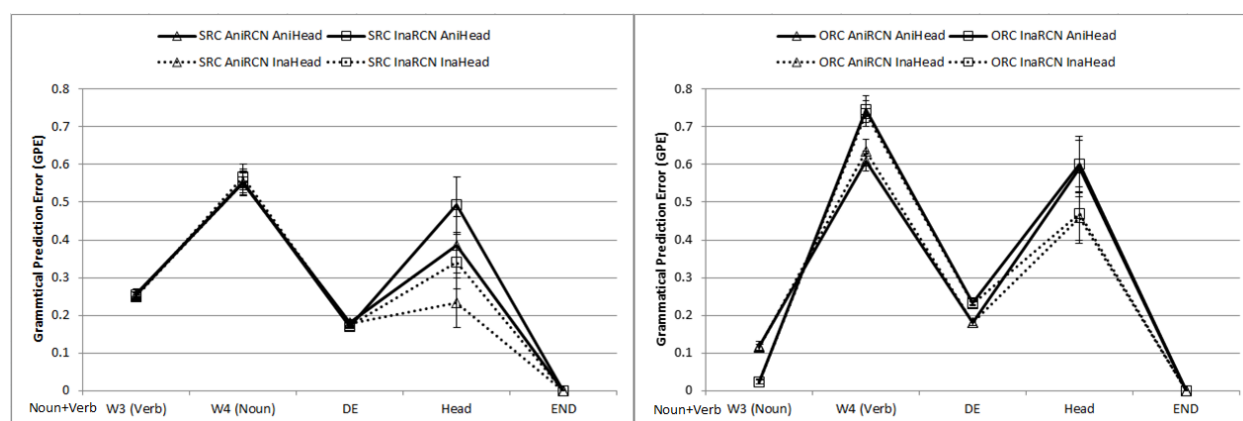


Figure 4. Word-by-word GPEs of object-modifying SRCs (left panel) and ORCs (right panel) in the critical RC region, the head, and next word after the head. All sentences began with a Noun + Verb sequence, not shown in the graphs.

Examination of Figure 4 reveals a general pattern of ORCs being harder than SRCs, in contrast to the subject-modifying relative clauses in Figure 3, for which ORCs were initially slightly harder, and then SRCs were substantially harder after the head noun. This result is also found in the human comprehension literature: Lin & Bever (2006) found longer reading times for object-modifying ORCs than SRCs. Due to the initial unusual word order of two consecutive verbs (main clause verb and RC embedded verb), SRCs are disambiguated early, yielding a small amount of early difficulty but later being easier, because this conjunction of two verbs removes

competitors other than an SRC. By contrast, ORCs remain ambiguous: the first word in ORCs in combination with the matrix clause context (the N V N order) created a strong bias toward a simple sentence, which was hard to revise even at the end.

Word-by-word findings.

Here we call the starting word of an object-modifying RC “W3” (RC verb for SRCs, RC subject for ORCs) because it was preceded by the main clause subject and verb and therefore was the third word in the sentence. At this word, the main effect of RC type indicated that SRCs were significantly harder than ORCs ($\beta=0.137$, $SE=0.017$, $t=7.96$, $p<.001$). Animacy of the RC embedded nouns had a significant effect too ($\beta=-0.094$, $SE=0.006$, $t=-16.30$, $p<.001$), where the inanimate RC nouns induced lower GPEs. The interaction between RC type and embedded noun animacy was significant ($\beta=0.090$, $SE=0.008$, $t=11.16$, $p<.001$). ORCs appeared to have a larger GPE difference between the ones with animate RC nouns (higher) and the ones with inanimate RC nouns (lower), whereas the two kinds of SRCs almost had overlapping GPEs.

At W4 (RC object for SRCs, RC verb for ORCs), the previous significant effects reversed: SRCs were now easier than ORCs ($\beta=-0.072$, $SE=0.024$, $t=-2.98$, $p=.008$) and animate RC nouns instead produced more accurate predictions than inanimate ones ($\beta=.118$, $SE=0.024$, $t=4.97$, $p<.001$). Interaction between the two factors was significant ($\beta=-0.109$, $SE=0.026$, $t=-4.16$, $p<.001$). ORCs with inanimate embedded nouns had particularly high GPEs.

At the RC marker DE, RC type was non-significant. The RC noun animacy effect was still present ($\beta=0.051$, $SE=0.006$, $t=8.32$, $p<.001$). Animate RC nouns again were preferred to inanimate ones. Interaction between the two was significant ($\beta=-0.058$, $SE=0.011$, $t=-5.53$, $p<.001$).

At the head, the main effect of RC type was significant, with GPEs of ORCs higher than those of SRCs ($\beta=-0.20$, $SE=0.098$, $t=-2.08$, $p<.001$). There was also a significant main effect of head noun animacy ($\beta=-0.13$, $SE=0.032$, $t=-4.12$, $p<.001$). Inanimate heads were preferred to animate heads. Other effects and interaction were not significant.

At the end-of-sentence marker, the errors went down to nearly zero for all conditions. No effects were present.

Discussion

Compared to the subject-modifying position, the model performance exhibited an opposite pattern at the object-modifying site: SRCs were easier than ORCs at the head noun. Error rates reflected early effects of low frequency sequences, followed by lower error rates because these low frequency sequences are themselves highly predictive of subsequent input (e.g. the very rare V V sequence in SRCs yields high error, but a V V sequence strongly predicts subsequent input, leading to low error at the next word). Overall, the SRN's performance on object-modifying RCs once again confirmed the model's sensitivity to both the structural and lexical statistics of the target RC structures and those of competing and neighbor structures under the constraining prior context of a main clause subject and verb. The processing advantage of SRCs over ORCs at the head noun resembles the findings in Lin & Bever (2006), which could be a result of the overall higher structural frequency and the lack of competing structures of SRCs at this position. In the next section, we investigate the effect of modifying position by analyzing SRCs and ORCs at both modifying positions together.

Contrasting subject vs. object modification

From the separate analyses above of the two modification sites, we observed almost opposite patterns of GPEs at the head. Prior RC studies have assessed RC difficulty at various regions: words inside the RC, at the head, at the head+1 position, even further down until the end of sentence, or average across the whole sentence. It is very true that the current model showed a dynamic nature in its processing patterns across the RC, the head, and post-head. Due to the head-final order of RCs in Mandarin, the head is considered the first position where processing pattern between SRCs and ORCs is differentiated. The working memory-based account argues that integration of related elements in an RC (filler and gap) occurs at this point. Prior empirical studies of Mandarin RCs typically found robust effect of RC type starting from the head. Therefore, to compare the difficulty between two RC types at two modification positions, we examined the GPEs at the head position. It was not only the position that reflected the accumulated processing difficulty across the RC but also the first position that reflected the head noun animacy effect.

Figure 5 displays the average GPEs of SRCs and ORCs word following the head noun, across both levels of head noun animacy and modification site. It could be observed that ORCs were easier than SRCs when main clause subjects are modified, whereas SRCs were easier than ORCs modifying main clause objects. Animate heads were preferred to inanimate heads as matrix subjects while inanimate heads were preferred to animate heads as matrix objects.

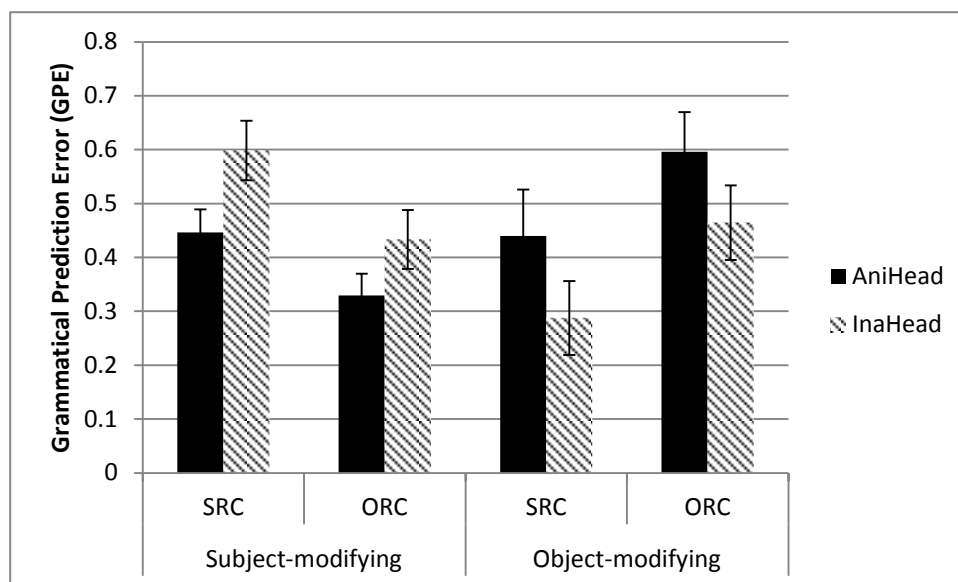


Figure 5. GPEs of SRCs and ORCs at the head noun. AniHead=animate head noun, InaHead=inanimate head noun.

At this word position, even though there were no significant main effects of RC type, head noun animacy, and modifying position, there were significant interactions of modifying position with RC type ($\beta=.37$, $SE=0.043$, $t=8.76$, $p<<.001$) and head noun animacy ($\beta=0.24$, $SE=0.043$, $t=5.61$, $p<.001$). These results suggest that sentential environment (subject vs. object modification) played an important role in the processing the two types of RCs, in combination of the animacy properties of the head. Prior human reading time studies typically have considered only a fraction of the conditions shown, and the conflicting results found in some studies may stem from failures to consider the full pattern of data. Therefore, we examine representative Mandarin RC studies and compare them with our model performance in the next section.

Comparison with previous Mandarin RC reading time studies

In this section we compare the SRN's performance to self-paced reading studies of Mandarin RC processing. As Table 5 shows, there are three major patterns of results in the Mandarin relative clause literature. The first major result is that in studies of subject-modifying

relative clauses, ORCs are easier than SRCs, a reverse of the typical pattern in most other languages. Studies yielding this pattern include Gibson & Wu (2013), Hsiao & Gibson (2003), though only for multiply-embedded sentences. ORC advantage only occurred at the first two words of singly-embedded sentences and Vasishth et al. (in press) report instead an SRC advantage at the head using Hsiao & Gibson's singly-embedded materials), Su et al. (2007), Chen et al. (2008), and Lin & Garnsey (2011). The SRN captured the major result of subject-modifying ORCs being easier than SRCs. Compared to subject-modifying ORCs, subject-modifying SRCs yielded reliably higher error rates in the model at the head noun ($\beta=.17$, $SE=0.052$, $t=3.33$, $p=.003$) and in an average of DE and the head together (as in Gibson & Wu's analysis), the ORC preference was also robust ($\beta=.08$, $SE=0.023$, $t=3.60$, $p=.002$).

Table 5. Comparison of experimental materials and major findings of representative human reading time studies

Major data patterns	Experiment Conditions	Studies	Specific Manipulations	Model Replication?
ORCs easier than SRCs at the head noun and/or nearby words	Only Subject-modifying RCs, all nouns animate	Gibson & Wu (2013)	Supportive context	Yes
		Hsiao & Gibson (2003)	Doubly-embedded RCs (ORC advantage only at pre-DE region in singly-embedded RCs)	
		Su et al. (2007)	Aphasic patients	
		Chen et al. (2008)	Memory spans	
		Lin & Garnsey (2011)	Topicalization	
Non-replication of Hsiao & Gibson (SRCs easier than ORCs) Replication of Gibson & Wu		Vasishth et al. (in press)	Hsiao & Gibson (2003) materials	No
			Gibson & Wu (2013) materials	Yes
Object-modifying SRCs easier than ORCs	Both subject- and object-modifying (all nouns animate)	Lin & Bever (2006)		Yes

No difference between SRCs & ORCs with preferred animacy configuration	Animacy of RC noun and head noun (all subject-modifying)	Wu et al. (2012)	Contrastive animacy of RC noun & head noun	Yes; model shows small differences where Wu et al. find little or no difference
------------------------------------------------------------------------	----------------------------------------------------------	-------------------------	--------------------------------------------	---------------------------------------------------------------------------------

Many of the articles listed in the top row of Table 5 described the ORC advantage as a general tendency of Mandarin sentence processing and working memory limitations, even though their studies investigated only a subset of RCs, namely RCs modifying main clause subjects and containing entirely animate nouns. This brings us to the second and third patterns in Table 6, which reflect the non-universality of the results in the top row. The second major pattern, reported by Lin & Bever (2006), is the effect of modification position. They found that object-modifying RCs were harder than subject-modifying RCs, reflected in reliable reading time differences at many word positions. On first glance, that result does not seem to correspond to the model's performance, but if we examine the model's behavior in the exact sentence types they tested (all animate nouns), the model and human data look much more similar. For this subset of conditions, Lin and Bever's RC-type effect was replicated in our model, with lower GPEs for SRCs compared to ORCs at several word positions (pre-DE: $\beta=-.04$, $SE=.016$, $t=-2.54$, $p=.009$; head: $\beta=-.20$, $SE=0.025$, $t=-8.08$, $p<.001$). We also roughly replicated their main effect of modifying position (subject-modifying easier) at the pre-DE region (W1+W2) ($\beta=-.18$, $SE=.016$, $t=-11.26$, $p<.001$), and also at the head ($\beta=-.22$, $SE=0.025$, $t=-8.81$, $p<.001$). Interactions were significant at these two positions too (pre-DE: $\beta=-.06$, $SE=0.021$, $t=-2.84$, $p=.005$; head: $\beta=.38$, $SE=0.036$, $t=10.51$, $p<.001$). When considering each modifying position separately, Lin and Bever did not find any RC type difference for subject-modifying RCs but a significant difference at DE and head for object-modifying RCs. Our model replicated the strong

effect of SRC advantage at the object-modifying position ($\beta=-.20$, $SE=0.025$, $t=-8.16$, $p<.001$) that Lin & Bever (2006) found.

The third major pattern is the effect of noun animacy on RC processing, exemplified by Wu et al.'s (2012) manipulation of both head and RC noun animacy within subject-modifying relative clauses. Their study did not fully cross head and RC noun animacy and only included contrastive animacy conditions (one noun animate and one inanimate). They found that in the preferred animacy configuration (in which the animate noun was the agent of the RC verb and the inanimate noun the theme), SRCs and ORCs didn't differ in processing difficulty. However, with the dispreferred animacy configuration (inanimate agents, animate patients), ORCs (such as the Mandarin equivalent of *The hiker that the rocks crushed*) were read particularly slowly and an SRC advantage emerged. Our model performance did show such SRC preference for these unusual animacy RCs early at W1 ($\beta=-.09$, $SE=.008$, $t=-11.12$, $p<.001$) and later at DE ($\beta=-.07$, $SE=.007$, $t=-10.04$, $p<.001$), but the ORC became easier at the head ($\beta=.41$, $SE=.053$, $t=7.68$, $p<.001$). For RCs with the preferred animacy configuration, SRCs showed early advantage at W1 ($\beta=-.01$, $SE=.006$, $t=-2.18$, $p=.03$) but later switched to ORC advantage at DE ($\beta=.08$, $SE=.011$, $t=7.25$, $p<.001$) and switched back at the head ($\beta=-.12$, $SE=.050$, $t=-2.48$, $p=.02$) but the effect was rather reduced compared to the disfavored animacy configurations. Furthermore, when considering all possible animacy configurations (i.e. AniRCN+AniHead, InaRCN+AniHead, AniRCN+InaHead, InaRCN+InaHead) rather than only the two that Wu et al. (2012) investigated, we found no difference between the favored configurations, namely animate-headed SRCs and inanimate-headed ORCs, at every word except for W1 ($\beta=-.06$, $SE=.011$, $t=-5.28$, $p<.001$). Thus the model shows comparatively small differences in the same conditions that Wu et al. find little or no difference in reading times.

In sum, the model captured major patterns of comprehension difficulty across several empirical studies, despite the fact that these studies are often thought to conflict with one another. The model's performance suggests that the inconsistencies in the literature are more apparent than real and stem from different experimental materials used in the experiments, which focus on a small subset of relative clause and animacy types. Whereas it is impossible to manipulate all relevant factors within a single self-paced reading experiment, the current SRN model could incorporate 16 types of test sentences in a 2x2x2x2 design. The model data present a more comprehensive picture, in which RC type, modifying position, head noun animacy, and RC noun animacy all have an effect. These data suggest that relative clause difficulty in Mandarin depends on a complex interplay of probabilistic constraints from animacy and other information gleaned from prior experience. Thus contra many claims in the literature (e.g., Hsiao & Gibson, 2003; Gibson & Wu, 2013; Lin & Bever, 2006; Vasishth et al., in press) none of the empirical results warrant broad conclusions about ORCs or SRCs being universally easier or harder.

General discussion

In this paper, we presented a Simple Recurrent Network simulation on the processing of Mandarin relative clauses. We had two related goals. First, we wanted to use the model to investigate issues that are difficult to test in human experiments. We suspected that controversies in the empirical literature stemmed at least in part from complex interactions among a number of factors such that when researchers designed materials tapping different subsets of the factors, different results obtained. Constraints on human studies, such as biases or priming effects that arise when comprehenders encounter many sentences of the same type, typically prevent anything more complex than a 2x2 design in sentence processing studies. By contrast, our 2x2x2x2 design presented in 16 test sentence types showed that the four different factors we

examined (RC type, modifying position, head animacy, and RC noun animacy) interacted in complex ways in the model. The results from the model closely track reading time patterns from a number of human comprehension studies and suggest that there is no overall SRC or ORC preference in Mandarin. Instead, the results strongly depend on which types of relative clauses are contrasted. In relative clauses with the animacy configurations most commonly investigated in human studies to date, ORCs are easier than SRCs when they modify main clause subjects but the reverse is true when they modify main clause objects, but other patterns are obtained with different animacy configurations, which can be seen most clearly in Figure 5. These results suggest that claims for broad categories of relative clause difficulty in Mandarin are premature at best.

Second, we wanted to use an SRN's ability to generalize over similar items (Fitz et al., 2011; MacDonald & Christiansen, 2002) to investigate how competitor and neighbor interpretations affect RC interpretations. Here our corpus analysis showed that RCs in Mandarin can be highly ambiguous for human comprehenders. When the model was trained on the resulting training set, prediction error varied with all four factors investigated here.

An important component of this second goal was our incorporation of animacy information in the corpus analyses and in the model. Human comprehension patterns clearly show the importance of animacy information in relative clause processing (Wu et al., 2012), and the model also captured these animacy effects, despite having no conceptual information that would typically be used to code distinctions between animate and inanimate entities. These results show the power of the sequential information associated with animacy, that sentences containing animate entities have different distributional patterns in the language than those with inanimate nouns. Chang (2009) made a related argument in his examination of cross-linguistic

variation in language production, that learning over the distributional regularities of various animacy configurations in a language is critical for explaining why animacy has different effects in a relatively strict word order language such as English compared with a relatively free word order language such as Japanese—people are learning the sequential information associated with sentences of different types, and the animacy-structure patterns vary as a function of the rigidity of the word order in the language. His results, as well as those of our model, suggest that sequential learning is an important adjunct to conceptual information in accounts of animacy effects in sentence-level language use.

Another important component of our investigation of neighborhood effects is the link between generalization over similar sentences and a topic that initially seems unrelated: the role of working memory in accounts of language comprehension. Because of its generalization over many sentence types, the model shows patterns of difficulty analogous to those in human studies. Critically, it does so in a system in which experience influences the model's computational capacity (and thus its ability to make accurate predictions for upcoming input, MacDonald & Christiansen, 2002). This result contrasts with claims questioning the adequacy of experience-based accounts in RC processing. Levy et al. (in press) and Levy and Gibson (2013) have assessed experience via calculations of a word's surprisal—the conditional probability of that word given prior context. They suggest that surprisal does not correctly predict the full pattern of human relative clause reading times, and they argue that human comprehension difficulty requires supplementing surprisal with an account of human memory burdens, as in Gibson's (1998) Locality theory, in which RC difficulty varies with load in an experience-independent working memory. Given our own results and MacDonald & Christiansen's success in using an SRN to model humans' reading times of English RCs, we think it is premature to reject all

experience-based accounts on the basis of failures of particular implementations (particular surprisal instantiations). First, as Frank (2009) notes, the success of Levy and colleagues' surprisal calculation varies with the richness of the prior input. Thus it may be that a larger or more realistic corpus over which to calculate conditional probabilities would yield a better account of humans' experience and consequently better prediction of reading times. However, we suspect there is a second reason why SRNs can yield different predictions than Levy and colleagues' surprisal results, concerning how context is represented and transformed into predictions. As it has been implemented to date, surprisal is based on an aggregation of past instances, which is used to calculate the conditional probability of upcoming input. By contrast, the SRN is a learning model that compresses and transforms its experience into an internal representation as it learns (Elman, 1990; Frank, 2009; Tabor et al., 1997). As a result, the SRN generalizes over neighboring structures and shows behavior that is not always a sum of instances in the training set, such as when the more frequent subject-modifying SRC sentences yield higher error rates than the rarer subject-modifying ORC sentences. A more direct comparison of SRNs and various instantiations of surprisal should be an important step in better understanding the role of experience in RC processing (see also Frank, 2009). In the meantime, we see no need to complicate our experience-based account with an additional component as in Levy and Gibson's proposal, and indeed we see the success of our SRN as further evidence for the non-independence of experience and computational capacity/working memory (MacDonald & Christiansen, 2002; McClelland & Elman, 1986).

Future directions

The SRN we used was trained on a set of sentences based on realistic human linguistic experiences, gleaned from a detailed corpus analysis. Although other corpus analyses exist for

Mandarin RCs (Hsiao & Gibson, 2003; Pu, 2007; Wu, 2009; Vasishth et al., in press), our study was unique in its large scale (incorporating many non-RC structures), its hand-coding of animacy across both RC and non-RC sentences, and its use in training the SRN. The combination with the SRN is crucial here because RC processing difficulty in both human and model is not simply a function of the frequency of RCs. We consider the following points important steps to further improve the current model.

The claim that Mandarin RC processing is tied to uncertainty of predictions echoes similar claims for other languages (e.g., Gennari & MacDonald, 2008). The SRN offers important opportunities to more stringently test these claims in future research. For example, to test whether simple pro-drop sentences are truly a competitor for certain RCs, the current model (containing pro-drop sentences) can be compared to one trained on a variant of Mandarin without pro-drop (i.e., replacing all pro-drop sentences in the training set with overt subject variants). We are pursuing these and other model comparisons aimed at elucidating the role of neighbors/competitors in the input. In addition, we intend to manipulate the amount of training, as in MacDonald & Christiansen (2002), to further observe the developmental trends in learning and generalizing. Thus SRNs can help trace the sources of processing difficulty using methods that are impossible to use with human comprehenders.

Another future endeavor should involve a more accurate portrayal of the rate of RCs and other related structures in Mandarin. For example, the multiple functions of *DE* (see Appendix A), which has been considered a disambiguating cue in RCs can in actuality create ambiguities depending on the semantic context. Structures involving these other *DE* uses should be considered potential competitors. Another typological feature in Mandarin that might be relevant is the possibility of null object (e.g. He saw (the movie).) with supportive context,

contrastive to null subject considered in the current model. The effects of null object and its combination with null subject are yet to be explored. Of course models necessarily remain simplifications of the entire linguistic experience (and simplifications of many other dimensions of human cognition). Model expansions therefore should not just cover more data but provide new insight into how people weigh multiple probabilistic constraints during sentence interpretation.

An important feature of our training set was its coding of distributional aspects of lexical meaning, as originally demonstrated by Elman (1990). That is, the model had no explicit semantic representations but it came to distinguish animate and inanimate nouns and transitive and intransitive verbs by virtue of their different distributions in the training set. These fine-grained discriminations were crucial for the model's account of animacy effects in RC processing. These and other studies of distributional semantic effects in SRNs (Elman, 1990; St. John & McClelland, 1990) show that there is potential for future work to incorporate other distributional "semantic" effects. There is also potential to make the distributions more natural, thereby capturing some additional ambiguity effects not in the present model. For example, for the current model, all verbs in the training set were either 100% transitive or 100% intransitive, with no optionally transitive verbs such as *eat*, which can occur either with or without a direct object. Optionally transitive verbs add additional indeterminacy, in that a model encountering an optionally transitive verb will be uncertain about an upcoming direct object, an effect which could modulate differences between SRCs (which can have transitive or intransitive verbs) and ORCs (which must have transitive verbs).

Conclusion

The current study confirmed SRN models as a promising tool in modeling human sentence processing, and, in this particular case, appropriate to examine intricate and complicated dynamics of Mandarin RC processing. The architecture of SRNs allows flexibility in modeling multiple effects in a single model, whereas manipulating a large number of factors human studies is nearly impossible. Mandarin is typologically unique in its conjunction of head-final RCs and head-initial SVO basic word, and yet in some sense the model's behavior looks very similar to that of SRN models of English RCs (Fitz et al., 2011; MacDonald & Christiansen, 2002). That is, the patterns of SRC vs. ORC difficulty are wildly different for the two languages, but in both cases, the models are strongly affected by the balance of RCs, competitors and neighbors. The modeling results suggest that rather than arguments for universal Locality (Gibson, 1998) or universal SRC preference (Lin & Bever, 2006), the real universals in human RC processing are exquisite sensitivity to the statistical regularities of across many different types of input.

References

- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255–278.
- Bock, K., & Warren, R. K. (1985). Conceptual accessibility and syntactic structure in sentence formulation. *Cognition*, 21(1), 47-67.
- Chang, F. (2009). Learning to order words: A connectionist model of heavy NP shift and accessibility effects in Japanese and English. *Journal of Memory and Language*, 61, 374-397.
- Chen, B., Ning, A., Bi, H., & Dunlap, S. (2008). Chinese subject-relative clauses are more difficult to process than the object-relative clauses. *Acta Psychologica*, 129(1), 61–65.
- Christiansen, M. H. (1994). *Infinite languages, finite minds: Connectionism, learning and linguistic structure*. (Unpublished doctoral dissertation). University of Edinburgh, Scotland.
- Christiansen, M. H., & Chater, N. (1999). Toward a connectionist model of recursion in human linguistic performance. *Cognitive Science*, 23(2), 157–205.
- Diessel, H., (2004). *The Acquisition of Complex Sentences*. Cambridge: Cambridge University Press, UK.
- Diessel, H. (2007). A construction-based analysis of the acquisition of East Asian relative clauses. *Studies in Second Language Acquisition*, 29, 311-320.

- Diessel, H., & Tomasello, M. (2005). A new look at the acquisition of relative clauses. *Language*, 81, 1-25.
- Elman, J. L. (1990). Finding structure in time. *Cognitive science*, 14(2), 179–211.
- Elman, J. L. (1991). Distributed representations, simple recurrent networks, and grammatical structure. *Machine Learning*, 7(2), 195–225.
- Elman, J.L. (1993). Learning and development in neural networks: The importance of starting small. *Cognition*, 48, 71-99.
- Fitz, H., Chang, F., & Christiansen, M. H. (2011). A connectionist account of the acquisition and processing of relative clauses. In E. Kidd (Ed.), *The Acquisition of Relative Clauses: Processing, Typology and Function*. (Vol 8) (pp. 39–60). Amsterdam: John Benjamins.
- Frank, S. L. (2009). Surprisal-based comparison between a symbolic and a connectionist model of sentence processing. In N.A. Taatgen & H. van Rijn (Eds.), *Proceedings of the 31st Annual Conference of the Cognitive Science Society* (pp. 1139–1144). Austin, TX: Cognitive Science Society.
- Frank, S.L., & Bod, R. (2011). Insensitivity of the human sentence-processing system to hierarchical structure. *Psychological Science*, 22, 829-834.
- Frazier, L. (1987). Syntactic processing: Evidence from Dutch. *Natural Language and Linguistic Theory*, 5(4), 519–559.
- Frauenfelder, U., Segui, J., & Mehler, J. (1980). Monitoring around the relative clause. *Journal of Verbal Learning and Verbal Behavior*, 19(3), 328–337.

- Gennari, S. P., & MacDonald, M. C. (2008). Semantic indeterminacy in object relative clauses. *Journal of memory and language*, 58(2), 161–187.
- Gennari, S. P., & MacDonald, M. C. (2009). Linking production and comprehension processes: The case of relative clauses. *Cognition*, 111(1), 1–23.
- Gibson, E. (1998). Linguistic complexity: Locality of syntactic dependencies. *Cognition*, 68(1), 1–76.
- Gibson, E., Desmet, T., Grodner, D., Watson, D., & Ko, K. (2005). Reading relative clauses in English. *Cognitive Linguistics*, 16(2), 313.
- Gibson, E., & Wu, H.-H. I. (2013). Processing Chinese relative clauses in context. *Language and Cognitive Processes*, 28(1-2), 125-155.
- Haspelmath, M., Dryer, M. S., Gil, D., & Comrie, B. (Eds.) (2005). *The World Atlas of Language Structure*. Oxford: Oxford University Press.
- Hsiao, F., & Gibson, E. (2003). Processing relative clauses in Chinese. *Cognition*, 90(1), 3–27.
- Hundt, M. (2004). Animacy, agentivity, and the spread of the progressive in Modern English. *English Language and Linguistics*, 8(1), 47–69.
- Keenan, E. L. (1985). Relative clauses. In T. Shopen (Ed.), *Language Typology and Syntactic Description, Vol II: Complex Constructions* (pp.141–170). Cambridge: Cambridge University Press.
- Keenan, E. L., & Comrie, B. (1977). Noun phrase accessibility and universal grammar. *Linguistic Inquiry*, 8, 63-99.

- King, J., & Just, M. A. (1991). Individual differences in syntactic processing: The role of working memory. *Journal of Memory and Language*, 30(5), 580–602.
- Kwon, N., Lee, Y., Goron, P. C., Kluender, R., & Polinsky, M. (2010). Cognitive and linguistic factors affecting subject/object asymmetry: An eye-tracking study of pre-nominal relative clauses in Korean. *Language*, 82, 546-582.
- Levy, R. (2008). Expectation-based syntactic comprehension. *Cognition*, 106(3), 1126–1177.
- Levy, R., Fedorenko, E., & Gibson, E. (in press). The syntactic complexity of Russian relative clauses. *Journal of Memory and Language*.
- Levy, R. & Gibson, E. (2013). Surprisal, the PDC, and the primary locus of processing difficulty in relative clauses. *Frontiers in Language Sciences*, 4(229), 1-3.
- Li, C. N. & Thompson, S. A. (1981). *Mandarin Chinese: A Functional Reference Grammar*. Berkeley, California: University of California Press.
- Lin, C. J. C. (2006). *Relative clause processing in typologically distinct languages: A universal parsing account*. (Unpublished doctoral dissertation). University of Arizona, Tucson, AZ, USA.
- Lin, C.-J. C., & Bever, T. G. (2006). Subject preference in the processing of relative clauses in Chinese. In D. Baumer, D. Montero & M. Scanlon (Eds.), *Proceedings of the 25th West Coast Conference on Formal Linguistics* (pp. 254–260). Somerville, MA: Cascadilla Proceedings Project.

- Lin, C.-J.C., & Bever, T.G. (2011). Garden path and the comprehension of head-final relative clauses. In H. Yamashita, Y. Hirose, & J. L. Packard (Eds.), *Processing and Producing Head-final Structures* (pp. 277-297). Springer.
- Lin, Y. B., & Garnsey, M. (2011). Animacy and the resolution of temporary ambiguity in relative clause comprehension in Mandarin. In H. Yamashita, Y. Hirose, & J. L. Packard (Eds.), *Processing and Producing Head-final Structures* (pp. 241-276). Springer.
- MacDonald, M. C., & Christiansen, M. H. (2002). Reassessing working memory: Comment on Just and Carpenter (1992) and Waters and Caplan (1996). *Psychological Review*, *109*(1), 35–54.
- MacDonald, M. C., Pearlmutter, N. J., & Seidenberg, M. S. (1994). Lexical nature of syntactic ambiguity resolution. *Psychological Review*, *101*(4), 676–703.
- Mak, W. M., Vonk, W., & Schriefers, H. (2002). The influence of animacy on relative clause processing. *Journal of Memory and Language*, *47*(1), 50-68.
- Mak, W. M., Vonk, W., & Schriefers, H. (2006). Animacy in processing relative clauses: The hikers that rocks crush. *Journal of Memory and Language*, *54*(4), 466–490.
- McClelland, J. L. & Elman, J. L. (1986). The TRACE Model of Speech Perception. *Cognitive Psychology*, *18*, 1-86.
- Miller, G. A., & Chomsky, N. (1963). Finitary models of language users. In R. D. Luce, R. Bush, & E. Galanter (Eds.), *Handbook of Mathematical Psychology* (pp. 419-492). New York: John Wiley.

- Miyamoto, E. T., & Nakamura, M. (2003). Subject/object asymmetries in the processing of relative clauses in Japanese. In G. Garding, & M. Tsujimura (Eds.). *Proceedings of the 22nd WCCFL* (pp. 342-355). Somerville, MA: Cascadilla Press.
- Pu, M.-M. (2007). The distribution of relative clauses in Chinese discourse. *Discourse Processes*, 43(1), 25–53.
- Ozeki, H. & Shirai, Y. (2007). Does the Noun Phrase Accessibility Hierarchy predict the difficulty order in the acquisition of Japanese relative clauses? *Studies in Second Language Acquisition*, 29, 169-196.
- Real, F., & Christiansen, M. H. (2007). Processing of relative clauses is made easier by frequency of occurrence. *Journal of Memory and Language*, 57(1), 1–23.
- Rohde, D. (1999). LENS: the light, efficient network simulator. *Technical Report CMU-CS-99-164*. Pittsburgh, PA: Carnegie Mellon University.
- Rohde, D. (2005). *TGrep2 Manual*. <http://tedlab.mit.edu/~dr/Tgrep2/tgrep2.pdf>.
- Roland, D., Dick, F., & Elman, J. L. (2007). Frequency of basic English grammatical structures: A corpus analysis. *Journal of Memory and Language*, 57, 348–379.
- Schriefers, H., Friederici, A. D., & Kuhn, K. (1995). The processing of locally ambiguous relative clauses in German. *Journal of Memory and Language*, 34(4), 499-520.
- St John, M. F., & McClelland, J. L. (1990). Learning and applying contextual constraints in sentence comprehension. *Artificial Intelligence*, 46(1), 217–257.

- Su, Y., Lee, S., & Chung, Y. (2007). Asyntactic thematic role assignment by Mandarin aphasics: A test of the Trace-Deletion Hypothesis and the Double Dependency Hypothesis. *Brain and Language*, 101, 1-18.
- Tabor, W., Juliano, C., & Tanenhaus M.K. (1997). Parsing in a dynamical system: An attractor-based account of the interaction of lexical and structural constraints in sentence processing. *Language and Cognitive Processes*. 12(2/3), 211–271.
- Tanenhaus, M.K. & Trueswell, J.C. (1995). Sentence comprehension. In: J.L. Miller & P.D. Eimas (Eds.). *Handbook of Perception and Cognition*. Vol. 11: Speech, language and communication, (pp. 217-262). San Diego, CA: Academic Press.
- Traxler, M. J., Morris, R. K., & Seely, R. E. (2002). Processing subject and object relative clauses: Evidence from eye movements. *Journal of Memory and Language*, 47(1), 69–90.
- Traxler, M. J., Williams, R. S., Blozis, S. A., & Morris, R. K. (2005). Working memory, animacy, and verb class in the processing of relative clauses. *Journal of Memory and Language*, 53(2), 204-224.
- Vasishth, S., Chen, Z., Li, Q. & Kuo, G. (in press). Processing Chinese relative clauses: Evidence for the universal subject preference, *PLoS One*.
- Warren, T., & Gibson, E. (2002). The influence of referential processing on sentence complexity. *Cognition*, 85(1), 79–112.
- Wells, J. B., Christiansen, M. H., Race, D. S., Acheson, D. J., & MacDonald, M. C. (2009). Experience and sentence processing: Statistical learning and relative clause comprehension. *Cognitive psychology*, 58(2), 250–271.

- Wu, F. (2009). *Factors affecting relative clause processing in Mandarin*. (Unpublished doctoral dissertation). University of Southern California, Los Angeles, CA, USA.
- Wu, F., Kaiser, E., & Andersen, E. (2012). Animacy effects in Chinese relative clause processing. *Language and Cognitive Processes*, 27(10), 1489–1524.
- Xue, N., Jiang, Z., Zhong, X., Palmer, M., Xia, F., Chiou, F.-D., & Chang, M. (2010). *Chinese Treebank 7.0*. Philadelphia, PA: Linguistic Data Consortium.
- Yip, V. & Matthews, S. (2007). Relative clauses in Cantonese-English bilingual children: Typological challenges and processing motivations. *Studies in Second Language Acquisition*, 29, 277-300.

Appendix A. **Ambiguity and Competitor Interpretations for Mandarin Relative Clauses**

Due to the head-final feature of Mandarin relative clauses and the fact that the relative pronoun DE appears at the end of a relative clause, Mandarin relative clauses contain many temporary ambiguities. Even at the relative clause marker DE, the structure could still be ambiguous because DE has other functions in Mandarin beyond the relative pronoun. Below, we detail some of these ambiguities—strings that could turn out to be a relative clause of some type but could also turn out to be some other kind of structure. In some cases, these structures may conflict with the relative clause interpretation at one sentence position but may at other times facilitate the RC interpretation (be a neighbor).

Simple Sentences

The dominant Subject-Verb-Object word order in Mandarin creates different amounts of ambiguity for object relative clauses occurring at different main clause positions. For example, an object relative clause modifying a main clause subject begins with the sequence “N-V” (as in 2b) which can be interpreted as a main clause subject and verb (the simple sentence is a temporary competitor), but later regions of the sentence reveal that this initial N-V sequence was instead part of a relative clause. The prevalence of Mandarin Subject-Verb-Object simple sentences could potentially aid the processing of sentences that share word order similarities (Fitz et al., 2011; MacDonald & Christiansen, 2002). Thus the fact that simple transitive sentences and subject-modifying object relatives initially have the same word order creates first an ambiguity, and then when the simple sentence interpretation is removed by additional input, the simple sentences are a helpful neighbor, in that past experience with simple sentences can help interpretation of these relative clauses.

By contrast, simple transitive sentences are not a competitor or neighbor for subject relative clauses modifying a main clause object (as in 3a). These structures contain a sequence of two verbs early in the sentence, which rules out a simple sentence interpretation of the input. Subject relatives might therefore be somewhat difficult early when the unusual sequence of two verbs is encountered, but overall, these sentences are less ambiguous than the object relatives, for which the simple sentence interpretation can persist over more input.

These examples show that generalization and competition from simple sentences manifest differently in the two main clause positions. These patterns are important because the dominant results in the comprehension literature find that subject relatives in subject-modifying position are harder than object relatives (Hsiao & Gibson, 2003; Gibson & Wu, 2013), while the opposite pattern obtains for object-modifying RCs (Lin, 2006; Lin & Bever, 2006).

Pro-drop Sentences

One important competitor structure for relative clauses comes from simple main clause sentences in which the subject NP has been omitted. Mandarin is a pro-drop language (meaning that in context, grammatical subjects are omissible; Li & Thompson, 1981), so that simple sentences such as “VT-Object” exist as complete sentences in the language. These pro-drop sentences are temporary competitors for subject-modifying subject relatives, which have the sequence “VT-Object-DE-Subject”. As more input comes in to disambiguate the structure, the comprehender may recover from the initial difficulty and be aided by this resemblance in word order with pro-drop sentences. Such facilitation effect for subject relatives may be even more pronounced and activated sooner at the main clause object site. The unusual word order of two

verbs in a row after the main clause subject creates short-term difficulty but meanwhile also provides a reliable cue for a subject relative clause and facilitates processing of the rest part.

Subject Relative Clauses with Intransitive Verbs

Subject relatives with intransitive verbs are rarely examined in prior empirical studies. Structurally, they resemble subject relatives with transitive verbs, with the lack of a direct object. Thus, experiences with one structure could be likely transferable to the other.

Other Structures with DE

The surface word order of Mandarin relative clauses is identical to many other structures that also have the particle DE. In addition to serving as a relativizer in a relative clause, DE is also a marker for adjectives and adverbs and also appears in possessives and appositives or simply as a phrase-final particle. With these many functions of DE, which is sometimes considered the disambiguating point for a relative clause, the relative clause structure could still be ambiguous after DE and permits various interpretations; the only definitive disambiguation is discourse context. For example, in the phrase “respect teacher DE parents”, the DE could be interpreted as a relativizer of a subject relative: “the parents that respect the teacher”, or as a possessive marker in a nominalized VP gerund: “respecting the teacher’s parents”. The appositive structure with DE also shares identical word order with a subject relative, such as in “respect teacher DE policy”, which could be either interpreted as “the policy that respects teacher” (RC interpretation) or “the policy about respecting teachers” (appositive interpretation).

Because the SRN in Study 2 will not code for semantics, except for animacy, we focus on only the relativizer use of DE. This represents a simplification compared to the full natural language.

Summarizing from above, we can see that these relevant structures, some of which may well surpass the frequency of relative clauses in comprehenders' linguistic experiences, impose influences on the processing of both the typical subject relative clauses with transitive verbs and object relative clauses, at different main clause site.

Appendix B

Tgrep2 search patterns for structures in Study 1

Simple sentences, transitive	/^IP/<(/^NP-SBJ/!</^NONE-!<<DEC)<(/^VP/<<VV<<(/^NP-OBJ/!</^NONE-!<<DEC))!>>/^IP/
Simple sentences, intransitive	/^IP/<(/^NP-SBJ/!</^NONE-!<<DEC)<(/^VP/<<VV!<<(/^NP-OBJ/))!>>/^IP/
Pro-drop sentences, transitive	/^IP/<(/^NP-SBJ/</^NONE-)<(/^VP/<<VV<<(/^NP-OBJ/!</^NONE-!<<DEC))!>>/^IP/
Pro-drop sentences, intransitive	/^IP/<(/^NP-SBJ/</^NONE-)<(/^VP/<<VV!<<(/^NP-OBJ/))!>>/^IP/
SRC, transitive	/^NP-SBJ/ (or /^NP-OBJ/ for object-modifying)<<(/^IP/<(/^NP-SBJ/</^NONE-)<(/^VP/<<VV<< (/^NP-OBJ/!</^NONE-))\$(DEC../^NP/))
SRC, intransitive	/^NP-SBJ/(or /^NP-OBJ/ for object-modifying)<<(/^IP/<(/^NP-SBJ/</^NONE-)<(/^VP/<<VV!<<(/^NP-OBJ /))\$(DEC../^NP/))
ORC	/^NP-SBJ/(or /^NP-OBJ/ for object-modifying)<<(/^IP/<(/^NP-SBJ/!</^NONE-)<(/^VP/<<VV<<(/^NP-OBJ/</^NONE-))\$(DEC../^NP/))

Chapter 2: Working Memory and Experience: A Connectionist Model of English Relative Clause Processing

A typical finding of language comprehension in English relative clauses is the higher difficulty of object relative clauses (e.g. 1b), compared to subject relative clauses (e.g. 1a).

(1a) Subject Relative Clause: The candidate [who₁ attacked₁ the opponent] won this election.

(1b) Object Relative Clause: The candidate [who₁ the opponent attacked₁] won this election.

The pattern has been replicated in many other languages with post-nominal relative clause structures, since the seminal paper by Miller and Chomsky (1963). The difference in processing difficulty between these two sentence types is often argued to be the result of the limitation of our computational systems. Object relative clauses impose higher working memory demands on the comprehender than do subject relative clauses in that the distance between elements which need to be stored and integrated is greater in object relatives than that in subject relatives, as shown in (1a) and (1b) with subscripts (the head noun “candidate” and the relative clause verb “attack”). It is argued that the dependency structure between related elements demands the parser to keep the partially processed element in active memory for later retrieval and integration (e.g. Gibson, 1998, 2000). Therefore, the longer the dependency distance, the more interruption there is to successfully maintain and retrieve the partially processed element. The dependency structure in an object relative clause seems to involve longer distance between related elements compared to that of a subject relative clause. On the same logic, from the perspective of individual difference, the higher the capacity of working memory in the comprehender, the more likely he or she is to overcome the processing difficulty of object relative clauses (Just & Carpenter, 1992; King & Just, 1991; Waters & Caplan, 1996).

Whether working memory, as modulated by structural complexity, can be viewed as an independent factor on sentence comprehension has been challenged over the years. Evidence shows that, besides syntactic structure, the semantic and discourse properties of the lexical elements in the sentence can contribute significantly in the process of ambiguity resolution (e.g. MacDonald, Pearlmutter & Seidenberg, 1994). The knowledge of syntactic and lexical representation comes from comprehenders' experience of the language's statistics, such as whether certain types of words usually occur in specific types of sentence structure. For example, in English, when the head noun is inanimate in an object relative clause (e.g. replace "the candidate" in 1b with "the policy"), comprehension difficulty is significantly reduced (Traxler, Morris & Seely, 2002), which is a result of comprehenders' conceptual assignment of thematic role being affected by noun animacy. In the real world, agents are more likely animate and the affected are more likely inanimate. This type of animacy effect has also been demonstrated in other languages with varying degree of typological differences from English (Dutch: Mak, Vonk & Schriefers, 2002; 2006; Mandarin: Wu, Kaiser & Andersen, 2012; Hsiao & MacDonald, 2013).

Gennari & MacDonald (2008, 2009) further tested the effect of head noun animacy on how language users would predict a certain sentence structure to follow. They found that when animate head nouns were followed by the relative pronoun and a demonstrative (e.g. the director that the...), participants rarely continued the fragment as a standard object relative clause, where "the director" is the patient (like 1b). Instead, more completions were those in which the "director" served other roles (e.g. the director that the movie was written by, the director that the movie belonged to). Their corpus analyses also showed such patterns, providing evidence that object relative clauses used as materials in psycholinguistic studies contained significant temporary ambiguities, causing comprehension difficulty. They attested this claim by linking the

completion data and corpus data to reading times: the higher completion rate or corpus frequency of the sentence structure, the lower the reading times. That is to say, the more frequently people encounter certain sentence types, or the more they produce them, the more easily people can comprehend them.

Prior research has also investigated non-lexical factors affecting relative clause processing, such as the sentential position of the noun that the relative clause modifies, i.e. whether the relative clause modifies the main clause subject or the main clause object. Several earlier studies found no difference between comprehension of subject-modifying (e.g. 1a and 1b) vs. object-modifying relative clauses (e.g. The reporter ignored the candidate [that the opponent attacked]), using auditory presentation (Baird & Koslick, 1974), phoneme monitoring tasks (Hakes, Evans & Brannon, 1976) or complexity ratings (Gibson & Thomas, 1997). However, there are also studies that reported subject-modifying advantage, such as Holmes (1973), who found that participants could recall more words from subject-modifying relative clauses than object-modifying relatives. More recently, Gibson and colleagues (Gibson, Desmet, Grodner, Watson & Ko, 2005) implemented the on-line self-paced reading measure and found that, on top of higher difficulty of object relative clauses, object-modifying relative clauses were harder than subject modifying relative clauses. They interpreted the result with the information flow hypothesis: old or background information, which is conveyed by restrictive relative clauses, usually occurs early in the sentence, whereas new information that consists of the core assertion of the utterance usually comes later. Therefore, an object-modifying relative clause goes against such information flow and tends to cause comprehension difficulty. This result contradicts syntactic theories in linguistics that argue for advantage in processing right-branching clauses (Miller & Chomsky, 1964; Miller & Isard, 1964) (i.e. the modifier comes after the head,),

compared to the difficulty of center-embedded clauses (i.e. embedding a phrase in another phrase of the same type).

The multitude of factors of relative clause processing examined by previous behavioral studies contributes significantly to our understanding of the complexity of sources of information that the comprehenders draw on. However, whether an inherent working memory limitation exists and how these other different factors, such as extraction type, noun animacy, and sentential modification position, interact among one another is practically untestable in a human experiment. Computational modelling provides a venue to investigate the complicated effects and interactions in language comprehension, and below we review cognitive models of sentence comprehension and computational endeavors to instantiate these models.

Models of sentence comprehension

Models that posit comprehension as a capacity constrained process assume working memory as a separate system from linguistic knowledge. Computational models like this are CC READER (Just & Carpenter, 1992). Instead of endorsing modularity and encapsulation of syntactic processing (Fodor, 1983), Just & Carpenter agree that individual differences in working memory capacity affect the degree to which comprehenders can entertain the non-syntactic information (e.g. noun animacy) during syntactic computation. People with high memory span are more likely to keep both syntactic and non-syntactic information activated. However, criticism from other researchers such as Waters and Caplan (Waters & Caplan, 1996, Caplan & Waters, 1999, 2013) maintains that verbal working memory can be further divided into two pools: one for on-line language comprehension and production and the other for other controlled, conscious, verbally mediated tasks (e.g. recall of a list of unrelated items). Also, the

representation of linguistic knowledge is highly modular. Another resource-based theory is the dependency locality theory (DLT) (Gibson, 1998, 2002), which has also two working memory components: storage and integration. The parser needs to keep track of the partially processed syntactic dependencies, requiring the storage component, and then to connect the newly encountered second element to complete the dependency relationship, involving the integration component. Here, like Waters and Caplan, the storage component is different in nature from other non-linguistic verbal tasks.

The above mentioned models or hypotheses are based on the central idea that working memory is an independent capacity and system from the declarative knowledge of language, such as syntactic rules and vocabulary. An alternative to this rather symbolic approach is the connectionist framework of sentence comprehension. Connectionist models permit processing as the interaction of diverse knowledge sources through updating the weights between connected nodes of units. Model architecture and input are the major influence of how and what the model learns. New patterns can be learned through generalizing from similar patterns from prior encounters. Therefore, from the connectionist point of view, individual differences in processing a sentence involving long distance dependency are a result of difference in experience rather than working memory capacity. That is to say, working memory capacity and linguistic knowledge are inseparable. MacDonald & Christiansen (2002) implemented a type of connectionist neural network typically used for modelling sentence processing called Simple Recurrent Network (SRN). Different from other connectionist models, an SRN has an additional context layer in addition to the input, output and hidden layers. The context layer stores the output of the hidden layer in the previous time tick and pairs it with the current input. This setup resembles human sentence comprehension, where the information of previously encountered elements influences

the interpretation of the new input. The model's task is to predict the next word based on the prior context. MacDonald & Christiansen devised a training corpus that consisted of 5% of relative clauses (2.5% of subject relatives and 2.5% of object relatives) and 95% of a variety of simple sentences. They found that, besides the overall ease of subject relatives compared to object relatives, as the amount of training increased, the error of object relative clauses decreased more dramatically than that of subject relatives. They interpreted the overall advantage of subject relatives as the result of the model rapidly generalizing from the overwhelming number of simple sentences, because of subject relatives' word order: subject-verb-object (noun-verb-noun). Object relatives, on the other hand, have the idiosyncratic word order of object-subject-verb (noun-noun-verb), and could not benefit from generalizing from simple sentences but only from direct exposure to object relatives themselves. A human training study by Wells, Christiansen, Race, Acheson & MacDonald (2009) confirmed this learning pattern. Reading times of object relatives decreased significantly after three training sessions, whereas no improvement was found for subject relatives.

MacDonald & Christiansen (2002) successfully demonstrated that, without a built-in working memory system in the neural network, the model with little training mimicked the performance of humans with low memory span, especially in the difficulty of comprehending object relative clauses. As the amount of training increased, more improvement was seen at processing object relative clauses, to the point that object relatives were not more difficult than subject relatives. This study provides convincing evidence for the inseparability of linguistic experience and knowledge, which gives rise to the individual difference in dealing with linguistic materials. While MacDonald & Christiansen (2002) found the SRNs to be a powerful way of simulating human comprehension behavior, their model was not based on real language statistics,

and the range of structure types they examined was limited. The consequence of this is that it is unclear what factors other than relative clause type are crucial in the processing of relative clauses. Next, we review a recent computational modelling attempt using an SRN to incorporate several important factors while resolving the long-standing controversies of Mandarin relative clause processing.

Mandarin SRN model by Hsiao & MacDonald (2013)

Different from English and many other languages, Mandarin Chinese has a unique combination of word orders in terms of basic simple sentences and relative clauses (Keenan, 1985). The relative clause comes before the modified head noun (e.g. instead of “the candidate that attacked the opponent”, in Mandarin it would be “attacked the opponent the candidate”), the so called “head-final” word order. The basic word order, however, is head-initial: subject-verb-object (SVO). Mandarin thus has become a test ground for various language processing accounts. For example, the resource-based DLT predicts that the head-final relative clause structure causes the dependency distance pattern to flip, opposite from English (e.g. Hsiao & Gibson, 2003; Gibson & Wu, 2013). Therefore, object relatives should be easier than subject relatives. Experience-based theories would hypothesize that, instead of dependency distance, it is the dominant word order of the language SVO that aids the processing of object relative clauses (e.g. Hsiao & MacDonald, 2013). Other theories like structural hierarchy predict subject relatives to be universally easier because of its higher position in the syntactic tree (e.g. Lin & Bever, 2006). The literature has long been divided on which type of relative clause is easier. Findings that support either result exist (e.g. subject relative advantage: Jager, Chen, Ki, Lin & Vasishth, 2015; Lin & Bever, 2006; Vasishth, Chen, Li & Kuo, 2013; object relative advantage: Chen, Ning, Bi & Dunlap, 2008; Su, Lee & Chung, 2007; Gibson & Wu, 2013; Hsiao & Gibson, 2003).

Hsiao & MacDonald (2013) realized that such inconsistency in empirical findings may result from the different manipulations and thus different types of materials used in each experiment. They decided to investigate the distributions in the corpus regarding not only the frequencies of the two relative clause types (subject relatives, object relatives) but also how they vary on the dimension of noun animacy (head noun and relative clause noun) and sentential modification position (main clause subject, main clause object). They then transformed the statistics into the transitional probabilities used to build a finite-state grammar, which generated a 10,000 sentence corpus that an SRN would later be trained on. Indeed, the model exhibited complicated interaction of the three factors. Relative clause type effect was modulated by the sentential position that the clause appeared and also the animacy configuration of the nouns. Specifically, object relative clauses were facilitated by the “neighbor” simple SVO sentences, which comprised most of the corpus, especially at the main clause subject position. Also, like in other languages, animate nouns were better candidates for subjects/agents, and inanimate nouns for objects/themes. The model performance was compared to previous empirical studies and the computational results were comparable to those of humans when manipulations were equated.

The main contribution of Hsiao & MacDonald (2013) lies in the fact that, without built-in grammatical knowledge and a working memory system, the network was able to simulate human-like processing patterns, which are sensitive to a multitude of sources of statistics, be it syntactic or semantic. It also enlightens the field of psycholinguistics of the promising potential of computational modelling, which allows for multiple factors to interact that cannot be done in a behavioral study. Extending this work, we would like to come back to build an improved model of English relative clause processing, based on MacDonald & Christiansen (2002), in the hope of achieving the following: 1) incorporating realistic statistics and the other two factors: noun

animacy, main clause position, in building the training corpus, 2) comparing and investigating the differences in processing patterns between Mandarin and English, as a result of differences in linguistics typology and distributional statistics, 3) demonstrating that SRN can be a valuable tool for modelling human sentence processing.

Like Hsiao & MacDonald (2013), we started in Study 1 by surveying corpus data for statistics of relative clauses and simple sentences, including hand coding the animacy of the all the nouns. In Study 2, we then used the distributional information to develop a finite state grammar, based on which the model's training corpus was generated. The model performance was evaluated and compared to prior human studies in English and also Hsiao & MacDonald's model. We discuss the results' implication on cross-linguistic differences and language processing theories more generally.

Methods

As demonstrated by MacDonald & Christiansen (2002), language comprehension is influenced not only by their direct experience with the target structure but also by their experience with other sentences in the language. They showed through their SRN that processing of English subject relative clauses was facilitated by the highly frequent simple transitive sentences. Both structures share the Noun + Verb + Noun word order with the same thematic role assignment (first noun: agent, second noun: patient/theme). In MacDonald & Christiansen's model, the facilitation came from the overlapping word order of the neighbor structure, with the errors of subject relatives dropping much faster than those of object relatives. Similarly, in the case of head-final relative clauses in Mandarin Chinese, Hsiao & MacDonald (2013) with their SRN attributed the relative ease of object relative clauses compared to subject

relatives at the main clause subject position to be the result of facilitation from the simple transitive sentences (with overt or covert subject), which also have the (noun) + verb + noun sequence, with the first noun being the agent and the second noun being the patient or theme. It is important to note that both MacDonald & Christiansen's and Hsiao & MacDonald's models did not incorporate semantics and the learning was simply based on occurrence of categories of units at specific word positions given the prior word sequence. Different from MacDonald & Christiansen, Hsiao & MacDonald included two additional factors: noun animacy and main clause modification position, which in the Mandarin literature are shown to influence reading times. Although very few prior studies have examined the effect of modification position in English relative clause processing (Gibson, Desmet, Grodner, Watson & Ko, 2005), animacy has been shown extensively and reliably to affect comprehension (Traxler et al., 2002; 2005; Gennari & MacDonald, 2008; 2009). In particular, Gennari and MacDonald (2008) using a sentence completion task showed that English object relative clauses contained more temporary ambiguities as they were disfavored as a continuation given the initial sentence fragment, due to the fact that they had more competing structures that were more frequent.

In this study, we examined comprehenders' experience by surveying the corpus frequency of relative clauses and relevant structures, here namely simple transitive and intransitive sentences, which potentially facilitate or compete with the processing of relative clauses. In addition, we coded for noun animacy and modification position. We then later used the frequency statistics to generate a miniature corpus for training the SRN, which was assumed to produce human-like processing patterns.

Corpus Analysis

The corpus data were based on the Wall Street Journal Treebank 2 (Marcus, Santorini & Marcinkiewicz, 1995), with the statistics of relative clauses directly from Roland, Dick and Elman (2007) and those of simple sentences were extracted by the first author. The corpus contains over 1.6 million words of parsed materials of written texts from the Dow Jones News Service. As a written corpus, it is comparable to the Chinese Treebank 7.0 that Hsiao & MacDonald (2013) used in terms of size (both over a million words) and genre (both are primarily newswire stories). To extract simple sentences, including both transitive and intransitive ones, we used Tgrep2 1.15 (Rohde, 2005) with the search patterns detailed in Appendix A. All the sentences, including both simple sentences we extracted and relative clause sentences we obtained from Roland et al., were hand coded for noun animacy by a native English-speaking research assistant, who was instructed to follow the criteria used also by Hsiao & MacDonald (2013): living entities that possess agency and volition to perform an action were coded animate, and any entities that did not possess these qualities were coded inanimate. Nouns that stood for a group of people (e.g. a company, a school, a country) were considered animate, whereas when these nouns referred to a location they were considered inanimate. When animacy could not be determined from the context, the sentence was excluded from later analysis.

Results

Table 1 displays the token frequencies of simple sentences, both transitive and intransitive, across noun animacy. Table 2 reports the frequencies of relative clauses, including both transitive and intransitive subject relatives and reduced and full object relative clauses (with overt and covert relative pronoun), across noun animacy main clause modification position. It is

important to note that these sentences included those with full noun phrases, proper nouns and pronouns.

Table 1. Token frequencies of simple sentences extracted from Wall Street Journal Treebank 2

	Simple Sentences					
	Animate Subject Nouns			Inanimate Subject Nouns		
	Object Noun Type			Object Noun Type		
Verb Type	Animate	Inanimate	none	Animate	Inanimate	none
Intransitive			2370			3481
Transitive	1036	6292		344	2681	

Table 2. Token frequencies of subject- and object-modifying subject relatives (transitive & intransitive) and object relatives (full & reduced) Wall Street Journal Treebank 2

	Subject-Modifying RCs					
	Animate Head Nouns			Inanimate Head Nouns		
	Relative Clause (RC) Noun Type			Relative Clause (RC) Noun Type		
RC-type	Animate	Inanimate	none	Animate	Inanimate	none
SRC, Intransitive			218			154
SRC, Transitive	69	396		32	285	
ORC, full	13	4		35	10	
ORC, reduced	11	2		123	22	
	Object-Modifying RCs					
	Animate Head Nouns			Inanimate Head Nouns		
	Relative Clause (RC) Noun Type			Relative Clause (RC) Noun Type		
RC-type	Animate	Inanimate	none	Animate	Inanimate	none
SRC, Intransitive			175			435
SRC, Transitive	38	140		48	479	
ORC, full	10	0		142	48	
ORC, reduced	11	4		233	77	

The entire set of extracted sentences was numbered 19418, among which 16204 were simple sentences and 3214 were sentences with relative clauses, with the ratio at around 6:1.

With the corpus data, we made the following observations:

1) Within simple sentences, transitive sentences were more frequent than intransitive sentences. Inanimate nouns occupied the subject position more often than did animate nouns in intransitive sentences (N=3481, 59%), whereas in transitive sentences animate nouns were preferred as the main clause subject (N=7328, 71%) and inanimate nouns were more often the main clause object (N=9873, 87%).

2) Among relative clauses across both modifying positions, SRCs (N=2469, 77% of all RCs), including intransitive and transitive, were more frequent than ORCs, both full and reduced. Transitive SRCs were the most frequent relative clause type (N=1461, 46% of all RCs) and full ORCs were the least frequent (N=262, 8% of all RCs).

3) There were slightly more object-modifying relative clauses (N=1840, 57% of all RCs) than subject-modifying ones (N=1374). At the subject-modifying position, relative clauses with animate heads were not much more frequent (N=713, 52%) than inanimate heads but at the object-modifying position, inanimate-headed relative clauses were dominantly more frequent (N=1462, 79%). This pattern also resembled the observations in simple transitive sentences, where the main clause subjects were usually animate and objects inanimate.

4) Each type of simple sentences outnumbered all types of sentences with relative clauses, except for the simple transitive sentences with inanimate subject and animate object (N=344), a rare animacy configuration. The most frequent simple sentence type was simple transitive

sentences with animate subject nouns and inanimate object nouns, the canonical animacy combination of animate agents acting upon inanimate themes.

5) The animacy distributional patterns in simple sentences were similar to those found in relative clauses. Across two modification positions, SRCs (transitive) with inanimate heads and animate RC nouns and ORCs (both full and reduced) with animate heads and inanimate RC nouns were overall infrequent within each structure type. On the other hand, relative clauses with the canonical animacy configuration were the most frequent within each RC type (except for object-modifying transitive SRCs).

Simple Recurrent Network

Simple Recurrent Network (SRN) has demonstrated by several previous studies to simulate word-by-word human reading time data, some in particular about relative clause processing (Fitz et al., 2011; Hsiao & MacDonald, 2013; MacDonald & Christiansen, 2002; Wells et al., 2009). An SRN has the merits of allowing examination of multiple factors at once, which is improbable in a human behavioral study, and also providing an architecture that generalizes experiences to relevant targets. As both MacDonald & Christiansen (2002) and Hsiao & MacDonald (2013) showed in their SRNs of relative clause processing, the model performance was influenced not only by the direct exposure to the relative clauses per se but also by the simple sentences, which shared word order properties with relative clauses. Hsiao & MacDonald (2013) further investigated the effects of noun animacy and main clause modifying position, which had been manipulated separately in different prior studies that produced inconsistent results. They were able to show that these effects interacted with relative clause type in a complicated manner reflecting statistical frequencies exposed during training. Given the

nature of SRNs being a model that predicts the upcoming output activations, both MacDonald & Christiansen (2002) and Hsiao & MacDonald (2013) successfully simulated model performance that closely tracked human reading times, especially for Hsiao & MacDonald's model, whose training corpus was based on realistic statistics.

Both as SRNs on English relative clause processing, the current model is different from MacDonald & Christiansen's (2002) in that our training corpus, like Hsiao & MacDonald's (2013), was generated based on real corpus data, obtained in the above mentioned corpus analysis. This procedure should enable more precise representation of human's linguistic experience. Additionally, we included noun animacy and main clause modification position, examined in the corpus analysis, as factors that could potentially impact interpretations of relative clause, especially for the effect of head noun animacy as many previous human studies have shown. The current study significantly extended MacDonald & Christiansen's in not only broadening the scope of investigation on the various components causing potential ambiguities but also offering a chance to validate SRNs as a useful tool in simulating human sentence comprehension.

Model specification

The SRN contained an input layer, an output layer and a hidden layer, as well as a context layer. Different from other feedforward networks, the context layer was connected to the hidden layer, from which it copied the activations of the previous time tick. The output activation on the hidden layer at time step t (previous word) was copied to the context layer and then projected back to the hidden layer at time step $t+1$ (the current word) to pair with the current input. The backpropagation learning algorithm was used to adjust weights during training.

The model consisted of 28 localist units in the input and output layers, each representing a word, including one “end-of-sentence” marker. Categories of words were encoded as different units, including various parts of speech and noun animacy, as shown in table 3. Note that although semantics was not modeled by the network, animacy was captured distributionally. For example, units representing animate nouns occurred frequently at the sentence-initial position as the main clause subject, whereas the units representing inanimate nouns mostly followed units denoting transitive verbs at the object position. There were 60 units each in the hidden layer and the context layer. Figure 1 displays the model architecture graphically. The learning rate was 0.05, the momentum was 0.9 and the batch size was 1. Lens 2.63 (Rohde, 1999) was the simulation software used.

Table 3. Categories of words represented by the units and the number of units

Category	Number of units
Animate nouns	7
Inanimate nouns	7
Transitive verbs	6
Intransitive verbs	6
Relative pronoun “that”	1
End of sentence	1

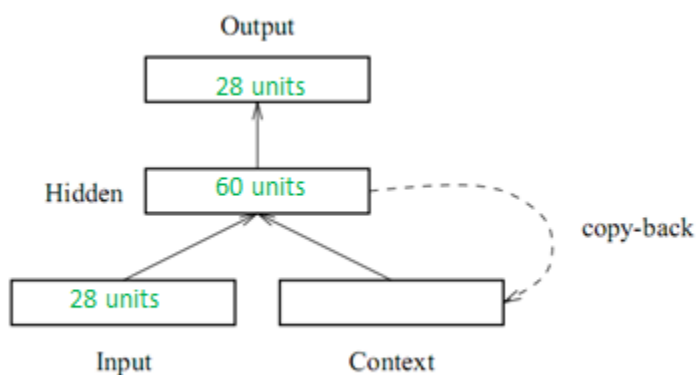


Figure 1. Architecture of the Simple Recurrent Network

Training

The corpus analysis we conducted offered the basis for the network's training set. We calculated the bigram transitional probabilities from one category of words to another, as constrained by the previous input. The transitional probabilities were shown in Table 4. The training corpus of 10,000 sentences was generated based on the finite state grammar in Table 4. The grammar only permitted single embedding, meaning that either the subject or the object position each allows only one relative clause. Also, the grammar did not encode for reduced object relative clauses. All the object relatives were in the full form, with the presence of the relative pronoun "that". However, we aggregated the statistics of both full and reduced object relatives to encode the transitional probabilities for object relatives. The purpose of doing so is to simplify the complexity of the learning task (so that the relative pronoun was a reliable cue across all relative clause types) and preserve the animacy information associated with object relatives as a whole.

The grammar thus allowed the following sentence types: simple transitive/intransitive sentences, transitive/intransitive subject relative clauses, full object relative clauses. The three types of relative clauses could modify either the main clause subjects or objects. A python script was written to code for the finite state grammar and produced 10,000 training sentences containing the above sentence types. Each unit in a given category had equal probability of being selected (at random). For example, unit #1 in the "animate noun" category was chosen equal amount of time to appear in the training sentences as unit #7 in the same category.

Table 4. Finite state grammar with corpus-based bigram transitional probabilities

S → subNP + VP (1)	
When subNP = aniN, VP → VI (0.22) / VT + objNP (0.78) When subNP = inaN, VP → VI (0.50) / VT + objNP (0.50)	
subNP → aniN(0.57) / inaN (0.36) / subRC (0.07)	When subNP = aniN, objNP → aniN (0.12) / inaN (0.74) / objRC(0.14) When subNP = inaN, objNP → aniN (0.1) / inaN (0.77) / objRC (0.14)
subRC (modifying matrix subject) → SRC_VI (0.28): aniN(0.18) + that + VI inaN(0.1) + that + VI → SRC (0.58): aniN (0.38) + that + VT+ aniN(0.15)/ inaN(0.85) inaN (0.2) + that + VT + aniN(0.1)/ inaN(0.9) → ORC (0.14): aniN (0.02) + that + aniN(0.8)/ inaN(0.2) + VT inaN (0.12) + that + aniN(0.83)/ inaN(0.17) + VT	objRC (modifying matrix object) → SRC_VI (0.34): aniN(0.11) + that + VI inaN(0.23) + that + VI → SRC (0.38): aniN (0.11) + that + VT+ aniN(0.21)/ inaN(0.79) inaN (0.27) + that + VT + aniN(0.09)/ inaN(0.91) → ORC (0.28): aniN (0.02) + that + aniN(0.84)/ inaN(0.16) + VT inaN (0.26) + that + aniN(0.75)/ inaN(0.25) + VT

S = sentence; NP = noun phrase; VP = verb phrase; VI = intransitive verb; VT = transitive verb; subNP = subject noun phrase; objNP = object noun phrase; aniN = animate noun; inaN = inanimate noun; subRC = subject-modifying relative clause; objRC= object-modifying relative clause; SRC_VI = subject relative clause with intransitive verb; SRC = subject relative clause with transitive verb; ORC = object relative clause; that = relative clause marker

Model Assessment

The model performance was assessed via the Grammatical Prediction Error (GPE) (Christiansen & Chater, 1999), shown below.

$$GPE = 1 - \frac{\text{hits}}{\text{hits} + \text{false alarms} + \text{misses}}$$

The GPE has been demonstrated to relate well to reading time measures, as in MacDonald & Christiansen (2002) and Hsiao & MacDonald (2013). The concept of GPE relates to signal detection theory, in that the error is calculated based on the success of activating grammatical units (hits), and failure of the model to activate these grammatical units to the

degree specified in the transitional probabilities of the grammar (misses), as well as the failure to suppress the activation of ungrammatical units (false alarms), as learned from the exposure to the training sentences. The GPE ranges from 0 to 1, with 0 being completely accurate in predicting the grammatical categories for the current word given the cumulative context and 1 being completely inaccurate in doing so.

Testing

To simulate the concept of random samples in behavioral studies, ten networks with different random initial weights, ranging from 1 to -1 with the mean of 0 were trained on the same 10,000-sentence training corpus under identical learning parameters. Each network was exposed to 30 epochs of the training sentences and assessed on novel test sentences, which were of 16 sentences types varying on the following three dimensions: two modification positions (main clause subject, object) x relative clause type (transitive SRC vs. full ORC) x head noun animacy (animate, inanimate) and RC noun animacy (animate, inanimate). Note that we did not include intransitive SRCs in our evaluation because the major comparison in the literature is between transitive SRCs and ORCs. Next, we present the networks' GPE performance on the test sentences in the critical relative clause regions.

Results

All the statistical analyses below were mixed effects models with participants (i.e. networks) and items as random effects with random intercepted. The statistical package "lme4" in R was used to fit the models. We treated the effect as significant when the t value reached 2.

Subject-modifying relative clauses

The following graph presents the GPEs of the two relative clause types, subject relatives on the left and object relatives on the right, at the head noun, relative clause region (relative pronoun, relative clause noun and verb) and the main clause verb.

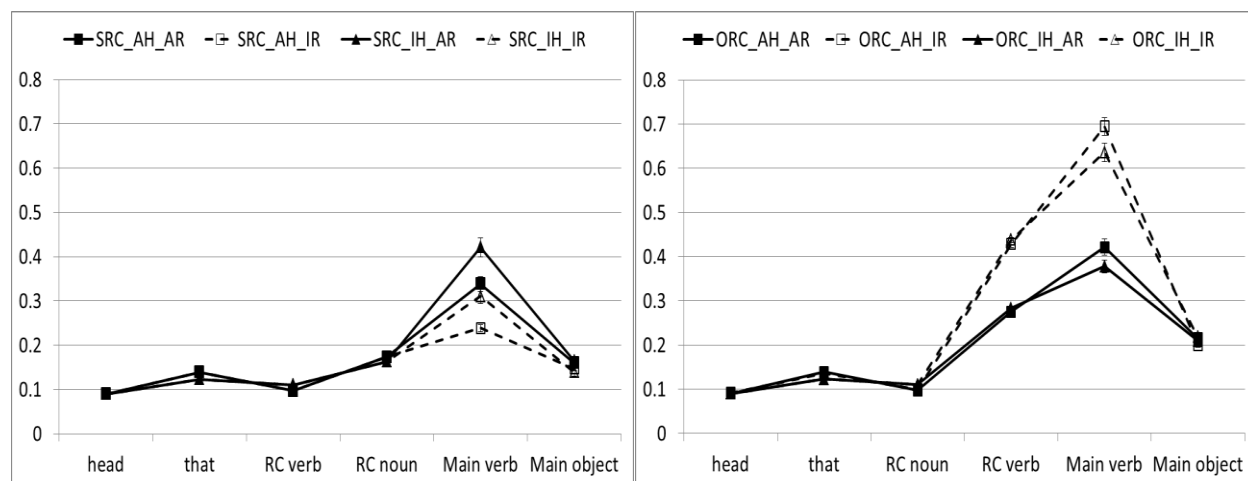


Figure 2. GPEs of the two relative clause types at the subject-modifying position. The GPEs for subject relatives are shown on the left panel and the GPEs for object relatives on the right. (Abbreviations in the legend: AH: animate head, AR: animate RC noun, IH: inanimate head, IR: inanimate RC noun)

Differences in GPEs between the two structure types started to emerge after encountering the word after the relative pronoun “that”. At the relative clause noun for subject relatives and relative clause verb for object relatives, errors started to elevate, possibly due to the reaction of the rarer word sequences of a relative clause, especially for object relatives ($\beta=-0.10$, $SE=0.008$, $t=-12.08$). Object relatives with inanimate relative clause nouns were higher in errors than those with animate relative clause nouns ($\beta=0.15$, $SE=0.007$, $t=24.05$). The same effect persisted to the next word at the main clause verb for object relatives ($\beta=0.25$, $SE=0.034$, $t=7.44$). There was a numerical effect of head noun animacy for object relatives: inanimate-headed ones were easier ($\beta=-0.04$, $SE=0.033$, $t=-1.34$). The head animacy effect was the opposite for subject relatives, where animate-headed ones were lower in error ($\beta=0.08$, $SE=0.032$, $t=2.63$). There was a main effect of relative clause noun animacy: animate nouns incurred higher error than inanimate nouns

($\beta=-0.10$, $SE=0.032$, $t=-3.13$). The overall GPEs were lower at the main verb for subject relatives than object relatives ($\beta=-0.08$, $SE=0.032$, $t=-2.62$). The same relative clause type effect was also found at the next word, the main clause object ($\beta=-0.06$, $SE=0.01$, $t=-5.59$).

Object-modifying relative clauses

Figure 3 displays the GPEs of the two relative clause types at the object-modifying position.

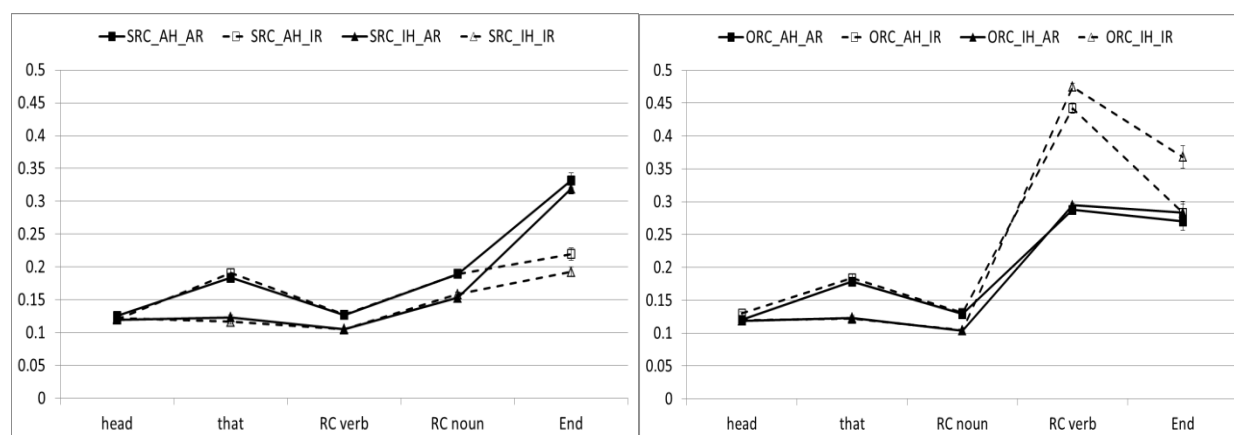


Figure 3. GPEs of the two relative clause types at the object-modifying position. The GPEs for subject relatives are shown on the left panel and the GPEs for object relatives on the right panel.

The preceding words not shown in the graphs are “main clause subject + main clause verb”, where the subject was always animate. Given that the word order of only started to differ between subject relatives and object relatives at the word after “that”, there were no differences in GPEs up until then. There was a clear head noun animacy effect at the word relative clause pronoun, where the model preferred inanimate nouns to be the sentential object ($\beta=-0.06$, $SE=0.007$, $t=-7.36$). The pattern persisted onto the next word ($\beta=-0.03$, $SE=0.002$, $t=-11.10$). The errors for the next word were much lower for subject relatives than those for object relatives ($\beta=-0.10$, $SE=0.010$, $t=-10.11$). Within subject relatives, animate-headed ones remained easier than inanimate-headed ones ($\beta=-0.04$, $SE=0.007$, $t=-4.91$), whereas among object relatives, the ones with animate relative clause nouns (i.e. the agent) were much easier than those with inanimate

relative clause nouns ($\beta=0.15$, $SE=0.012$, $t=13.36$). At the next word, the end of the sentence, the opposite relative clause noun animacy effect was found for subject relatives ($\beta=-0.11$, $SE=0.014$, $t=-7.87$), suggesting that the model preferred inanimate nouns to be the objects/themes.

Comparing the results of both modifying positions, at the critical word where the model finished encountering the relative clause: the main clause verb at the subject-modifying position and the end of sentence for the object-modifying position, the GPEs of subject relatives were overall lower than those of object relatives ($\beta=-0.12$, $SE=0.017$, $t=-6.67$). There was a main effect of sentential modifying position: subject-modifying relative clauses were higher in error than object-modifying relative clauses ($\beta=0.14$, $SE=0.017$, $t=8.16$). Within relative clause type, subject-modifying subject relatives were harder than the object-modifying counterparts ($\beta=0.06$, $SE=0.013$, $t=4.61$) and the same pattern occurs with object relatives ($\beta=0.22$, $SE=0.021$, $t=10.70$).

Discussion

The performance of the model confirmed the well-observed pattern in human studies that subject relatives are easier than object relatives. The pattern was most clear at the subject-modifying position, the site where the majority of previous studies investigated. The network, however, was not equipped with a working memory system or any grammatical knowledge, which humans were presumed to possess. The model learned only through the input made of basic sentences and sentences with types of relative clauses, suggesting that language experience is the major force that helps form learners' knowledge of the language. On the other hand, object relatives causing higher working memory load, as claimed by previous studies (e.g. Gibson, 1998), may not be separate from the fact that object relatives are rarer in frequency and also more uncommon in terms of word order, compared to subject relatives.

The animacy configuration of the nouns in the relative clause also plays a crucial role in the prediction error of the model. Subject relatives were preferred to have animate head nouns and inanimate relative clause objects, whereas object relatives were likely to be easier if the head nouns were inanimate and the relative clause subjects were animate. This pattern agrees with the way we conceptualize the world: agents are most likely alive and with volition, while the entities acted upon are oftentimes immobile or unable to express free will. This is also what previous empirical studies found across languages, regardless of the typological differences (English: e.g. Traxler et al., 2002, 2005; Dutch: Mak et al., 2002, 2006; Mandarin: Hsiao & MacDonald, 2013; Wu et al., 2012). One thing to note is that the model learned this knowledge about animacy through the distribution of animate and inanimate nouns in the sentences, suggesting that language experience consists of an important source where humans acquire knowledge about the world around them.

A less studied factor examined by prior research is the sentential modifying position of the relative clause. The GPE performance of the model shows a subject-modifying relative clauses advantage. The result appears to be consistent with the syntactic theory that argues for the advantage of right-branching structures: subject relatives are easier and subject-modifying relative clauses are easier. A more recent study by Gibson and colleagues (Gibson et al., 2005) using self-paced reading found object-modifying relative clauses to be harder and argue for an information flow theory, which maintains that it is more natural for given or old information to come early in the sentence and new information later in the sentence. Restrictive relative clauses that most prior studies used as materials convey old information, which according to Gibson et al. usually come early in the sentence. However, context or background information was usually not provided in most studies. Therefore, it is difficult to assess whether the restrictive relative clauses

that participants read were treated as given information. Moreover, our corpus analysis shows more object-modifying relative clauses than subject-modifying ones (1840 vs. 1374). Another explanation for our result is that as the sentence progressed, more words were revealed and less ambiguity existed. Due to the head-initial relative clause word order in English, when the model encountered the head noun of an object-modifying relative clause, the sentence was grammatically complete. Any additional words coming after could only be either the end of the sentence or the relative clause pronoun “that”, whereas at the subject-modifying position, many more options were grammatical continuation. This reduction of ambiguity may have contributed to the lower error of object-modifying compared to subject-modifying relative clauses.

General Discussion

Building on the well-known results of subject advantage from prior research, the study consisted of corpus analyses and neural network modelling that examined the processing of English relative clauses, with the attempt to shed light on theories of sentence comprehension, with the working memory-based and experience-based frameworks in mind. First, we investigated the patterns of distribution of relative clauses in the context of a variety of sentence types, varying on the dimensions of the sentential positions and the animacy of the nouns. We found that, as expected, the overall frequency of subject relatives was higher than that of object relatives. Moreover, the animacy configurations in the relative clauses followed the patterns found in simple sentences, where subjects/agents were more likely animate whereas objects/themes were more likely inanimate, meaning that subject relatives generally modified animate entities and object relatives inanimates. The animacy effect interacted with where the relative clauses occurred in the sentence, for example, relative clauses that modified inanimate nouns often appeared in the main clause object position. These distributional patterns were

reflected in the performance of the SRN, the second part of the study, where rarer sentence types were found higher in prediction error at the critical word positions. We replicated the subject relative advantage of numerous prior empirical studies and MacDonald & Christiansen's (2012) modelling result. In addition, we confirmed that animacy and main clause position had crucial modulating effects.

Comparison with Hsiao & MacDonald's (2013) Mandarin SRN

We saw similar distributional patterns in Hsiao & MacDonald (2013) in their Mandarin corpus data. The thematic role preference of agents being animate and themes being inanimate and its interaction with sentential position hold cross-linguistically. However, it appears that in English, with the head-initial structure (i.e. the relative clause follows the modified noun), object-modifying relative clauses outnumbered subject-modifying ones. On the contrary, in Mandarin, with the head-final structure (i.e. the modified noun appears after the relative clause), subject-modifying relative clauses were more frequent than object-modifying ones. Indeed, this difference was also seen in the SRN performance of the two languages. At the critical word regions, English SRN had lower errors for object-modifying relative clauses whereas Mandarin SRN showed lower errors for subject-modifying relatives instead. Another consequence of the word order difference between the two languages is the difficulty of relative clause type. In Mandarin, object relatives have been found in prior human studies to be easier than subject relatives, the opposite from English. Hsiao & MacDonald (2013) also saw this pattern in the Mandarin SRN, majorly in the subject-modifying position. In our English model, we found object relatives to be harder than subject relatives, more so at main clause subject position but also somewhat true at the main clause object position.

We attribute these patterns to the statistical regularities of the word order in English, as did Hsiao & MacDonald (2013) for Mandarin. The head-initial nature of English relative clauses necessarily gives a head-start for English comprehenders, because the sequence of “noun+relative pronoun” (e.g. the girl that/who...), compared to head-final Mandarin relative clauses. Mandarin comprehenders most likely do not realize that they are encountering a relative clause until they read/hear the relative pronoun or even the head noun (because the relative pronoun is a very frequently used homophone for various grammatical usages) (e.g. ... + relative pronoun+ the girl). This word order difference between the two languages in combination with the basic word order of SVO (noun+verb+noun) creates a niche environment for different relative clause types. In English, subject relatives benefit from the similarity to the common word order, whereas object relatives, which have the unusual word order (noun+noun+verb), can only be learned through direct exposure. This is true across main clause modifying positions. In Mandarin, on the contrary, object relatives resemble the basic word order, whereas subject relatives have the uncommon verb+noun+noun order. The advantage of object relatives weakens in the object-modifying position because, due to the fact that the head noun only comes at the end, the parser is likely garden pathed the relative clause noun as the main clause object and has to revise the interpretation when more words come in.

Summing from above, both the English and Mandarin networks exhibit the ability of generalizing from the more frequent neighbor structures: the simple sentences. It suggests that the SRNs achieves these human-like learning results through predicting the most familiar word sequences, instead of calculating the dependency distance between related elements, like the Dependency Locality Theory (DLT) argues. Next, we discuss the difference between our

experience-based approach of sentence comprehension to the capacity-based working memory accounts.

Implications to Working Memory in Sentence Comprehension

Working memory has long been associated with syntactic processing (King & Just, 1999). The cognitive load is high when maintaining stored information while performing computation of new information. Object relative clauses are one such structure that requires the readers to hold and integrate information over long distances. Gibson and colleagues' DLT argues that the head noun information of an object relative needs to be stored across many words until it can be integrated with the relative clause verb (e.g. the boy that the girl kicked) whereas a subject relative does not impose this demand (e.g. the boy that kicked the girl). Our SRN was not equipped with such linguistic knowledge of dependency relationship between nouns and verbs or semantic knowledge about what each word (unit) meant. It simply acquired the sequences where the units of different classes occurred. It was easier for the network to learn two categories of words that co-occurred adjacent to each other and harder to learn when the two categories occurred over long distance. However, with enough exposure, the model was able to learn the latter, as is evident with the similar amount of error between subject relatives and object relatives in MacDonald & Christiansen (2002)'s model and in our current model when those object relatives consisted of nouns with canonical animacy configurations (e.g. inanimate head noun and animate relative subject). The fact that the model did not produce more error at this point for object relatives than for subject relatives suggests that dependency distance cannot single handedly account for comprehension difficulty. Indeed, in behavioral data, low-span participants were found to process syntactic complex or ambiguous structures better after repeated exposure, to the point that the difference is minimal from high-span participants (Long & Prat, 2008).

Experience and its generalization have a significant role in reverting the processing patterns, for which the working memory capacity theories have not yet provided mechanisms to explain.

Many prior pathological studies based on aphasic patients have the assumption that these patients either have reduced working memory capacity (e.g. Rochon et al, 1994) or compromised linguistic knowledge (Grober & Bang, 1995). However, our data suggest that computational capacity and knowledge of language are inseparable. As MacDonald & Christiansen (2002) pointed out, the traditional reading span tasks that were supposed to measure linguistic working memory were language comprehension measures with different task demands. Simply because the reading span tasks were correlated with comprehension ability does not mean that there is a separate construct called working memory capacity. The so-called high-span participants are skilled and experienced readers that are sensitive to probabilistic constraints in the language statistics, and therefore are able to use subtle probabilistic cues to efficiently compute the sentence, whereas low-span readers do not have the sensitivity to detect these different subtleties (Pearlmutter & MacDonald, 1995).

Siding with MacDonald & Christiansen, we do not deny the existence of limits on computational capacity, as some researchers misinterpreted (Conway, 2007). Rather, it is clear that our SRN did have its capacity limits with regard to its architecture that favored local dependency. What we try to advocate is that linguistic working memory is part of the language processing mechanism that can be subject to both innate and external factors, such as the biological architecture and later experience.

Conclusion

The current study extends MacDonald & Christiansen's (2002) SRN on English relative clauses in providing a model that closely simulated human language experience, while also replicating Hsiao & MacDonald's (2013) SRN on Mandarin relative clauses in the intricate complex interactions of various structural and non-structural effects, which produced certain universal and language-specific results. More importantly, along with the two papers, the model replicated crucial human comprehension patterns and in turn offered implications on the roles of verbal working memory and language experience in language processing more broadly.

References

- Baird, R., & Koslick, J.D. (1974). Recall of grammatical relations within clause-containing sentences. *Journal of Psycholinguistic Research*, 3, 165-171.
- Caplan, D., & Waters, G. S. (1999). Verbal working memory and sentence comprehension. *Behavioral and Brain Sciences*, 22, 77-94.
- Caplan, D., & Waters, G. S. (2013). Memory mechanisms supporting syntactic comprehension. *Psychonomic Bulletin and Review*, 20(2), 243–268.
- Chen, B., Ning, A., Bi, H., & Dunlap, S. (2008). Chinese subject-relative clauses are more difficult to process than the object-relative clauses. *Acta Psychologica*, 129(1), 61–65.
- Christiansen, M. H., & Chater, N. (1999). Toward a connectionist model of recursion in human linguistic performance. *Cognitive Science*, 23(2), 157–205.
- Conway, A. R. A., Jarrold, C., Kane, M. J., Miyake, A., & Towse, J. (2007). *Variation in Working Memory*. Oxford University Press.
- Fodor, J. A. (1983). *Modularity of Mind: An Essay on Faculty Psychology*. Cambridge, Mass.: MIT Press.
- Gennari, S. P., & MacDonald, M. C. (2008). Semantic indeterminacy in object relative clauses. *Journal of memory and language*, 58(2), 161–187.
- Gennari, S. P., & MacDonald, M. C. (2009). Linking production and comprehension processes: The case of relative clauses. *Cognition*, 111(1), 1–23.

- Gibson, E. (1998). Linguistic complexity: Locality of syntactic dependencies. *Cognition*, 68(1), 1-76.
- Gibson (2000). The dependency locality theory: A distance-based theory of linguistic complexity. In Miyashita, Y., Marantz, A., & O'Neil, W. (Eds.), *Image, language, brain* (pp. 95-126), Cambridge, MA: MIT Press.
- Gibson, E., Desmet, T., Grodner, D., Watson, D., & Ko, K. (2005). Reading relative clauses in English. *Cognitive Linguistics*, 16(2), 313.
- Gibson, E., & Wu, H.-H. I. (2013). Processing Chinese relative clauses in context. *Language and Cognitive Processes*, 28(1-2), 125-155.
- Gibson, E. & Thomas, J. (1997). The complexity of nested structures in English: Evidence for the syntactic prediction locality theory of linguistic complexity. Unpublished Manuscript.
- Grober, E., & Bang, S. (1995). Sentence comprehension in Alzheimer's disease. *Developmental Neuropsychology*, 11, 95-107.
- Hakes, D. T., Evans, J. S., & Brannon, L. L (1976). Understanding sentences with relative clauses. *Memory & Cognition*, 4(3), 283-290.
- Holmes, V. M. (1973). Order of main and subordinate clauses in sentence perception. *Journal of Verbal Learning and Verbal Behavior*, 12, 285-293.
- Hsiao, F., & Gibson, E. (2003). Processing relative clauses in Chinese. *Cognition*, 90(1), 3-27.
- Just, M. A., & Carpenter, P. A. (1992). A capacity theory of comprehension: Individual differences in working memory. *Psychological Review*, 98, 122-149.

- Jäger, L., Chen, Z., Li, Q., Lin, C.-J. C., & Vasishth, S. (2015). The subject-relative advantage in Chinese: Evidence for expectation-based processing. *Journal of Memory and Language*, 79-80, 97–120. <http://doi.org/10.1016/j.jml.2014.10.005>
- Keenan, E. L. (1985). Relative clauses. In T. Shopen (Ed.), *Language Typology and Syntactic Description, Vol II: Complex Constructions* (pp.141–170). Cambridge: Cambridge University Press.
- King, J., & Just, M. A. (1991). Individual differences in syntactic processing: The role of working memory. *Journal of Memory and Language*, 30(5), 580–602.
- Lin, C.-J. C., & Bever, T. G. (2006). Subject preference in the processing of relative clauses in Chinese. In D. Baumer, D. Montero & M. Scanlon (Eds.), *Proceedings of the 25th West Coast Conference on Formal Linguistics* (pp. 254–260). Somerville, MA: Cascadilla Proceedings Project.
- Long, D. L., & Prat, C. S. (2008). Individual differences in syntactic ambiguity resolution: Readers vary in their use of plausibility information. *Memory & Cognition*, 36(2), 375–391. doi: 10.3758/MC.36.2.375.
- MacDonald, M. C., & Christiansen, M. H. (2002). Reassessing working memory: Comment on Just and Carpenter (1992) and Waters and Caplan (1996). *Psychological Review*, 109(1), 35–54.
- MacDonald, M. C., Pearlmutter, N. J., & Seidenberg, M. S. (1994). Lexical nature of syntactic ambiguity resolution. *Psychological Review*, 101(4), 676–703.

- Mak, W. M., Vonk, W., & Schriefers, H. (2002). The influence of animacy on relative clause processing. *Journal of Memory and Language*, 47(1), 50-68.
- Mak, W. M., Vonk, W., & Schriefers, H. (2006). Animacy in processing relative clauses: The hikers that rocks crush. *Journal of Memory and Language*, 54(4), 466–490.
- Marcus, M., Santorini, B., & Marcinkiewicz, M. A. (1995). Treebank-2 LDC95T7. Web Download. Philadelphia: Linguistic Data Consortium.
- Miller, G.A., & Isard, S. (1964). Free Recall of Self-embedded English Sentences. *Information and Control*, 7, 292--303.
- Pearlmutter, N. J. & MacDonald, M. C. (1995). Individual differences and probabilistic constraints in syntactic ambiguity resolution. *Journal of Memory and Language*, 34, 521-542.
- Rohde, D. (1999). LENS: the light, efficient network simulator. *Technical Report CMU-CS-99-164*. Pittsburgh, PA: Carnegie Mellon University.
- Rohde, D. (2005). *TGrep2 Manual*. <http://tedlab.mit.edu/~dr/Tgrep2/tgrep2.pdf>.
- Roland, D., Dick, F., & Elman, J. L. (2007). Frequency of basic English grammatical structures: A corpus analysis. *Journal of Memory and Language*, 57, 348–379.
- Su, Y., Lee, S., & Chung, Y. (2007). Asyntactic thematic role assignment by Mandarin aphasics: A test of the Trace-Deletion Hypothesis and the Double Dependency Hypothesis. *Brain and Language*, 101, 1-18.

- Traxler, M. J., Morris, R. K., & Seely, R. E. (2002). Processing subject and object relative clauses: Evidence from eye movements. *Journal of Memory and Language*, 47(1), 69–90.
- Vasishth, S., Chen, Z., Li, Q. & Kuo, G. (2013). Processing Chinese relative clauses: Evidence for the universal subject preference, *PLoS One*.
- Waters, G., & Caplan, D. (1996). The capacity theory of sentence comprehension: Critique of Just and Carpenter (1992). *Psychological Review*, 103, 761-772.
- Wells, J. B., Christiansen, M. H., Race, D. S., Acheson, D. J., & MacDonald, M. C. (2009). Experience and sentence processing: Statistical learning and relative clause comprehension. *Cognitive psychology*, 58(2), 250–271.
- Wu, F., Kaiser, E., & Andersen, E. (2012). Animacy effects in Chinese relative clause processing. *Language and Cognitive Processes*, 27(10), 1489–1524.

Chapter 3: Production Predicts Comprehension: Animacy Effects in Mandarin Relative Clause Processing

In order to test theories of how people comprehend language, sentence processing researchers often measure comprehension difficulty via some combination of time to comprehend, such as reading time, and accuracy, as measured by responses to comprehension questions. Because differences in time/accuracy across stimulus conditions cannot by themselves indicate why one sentence type is harder than another, patterns of difficulty are given different interpretations within different theoretical accounts of comprehension processes. Recent sentence processing accounts have been divided between experience-based approaches, including both constraint-based accounts (MacDonald, Pearlmutter & Seidenberg, 1994; Tanenhaus & Trueswell, 1995) and surprisal (Hale, 2001; Levy, 2008), and inherent-difficulty accounts (Gibson, 1998; 2000). Briefly, experience-based approaches suggest that comprehension difficulty arises when the sentence input does not match past experience, and inherent-difficulty accounts suggest that some sentences are harder than others because they are computationally more difficult, such as requiring more working memory. Relative clauses have played an outsized role in generating evidence for both of these theoretical positions. For example, the reduced passive relative clauses in (1) show the influence of comprehenders' use of past experiences with verb-noun thematic role combinations and/or experience with events in the world (e.g., Trueswell, Tanenhaus & Garnsey, 1994): Sentence (1a) is more difficult than (1b) because in comprehenders' past experience, animate nouns such as *girl* are plausible and common agents of verbs such as *kissed*, leading to misanalysis of the sentence, while *trophy* in (1b) is not a good agent of *kiss*, allowing the comprehender to rapidly settle on the correct analysis, in which the trophy is the object of *kissed*.

(1a) The girl kissed by the woman was....

(1b) The trophy kissed by the woman was....

The prime example of relative clauses within the inherent-limitation approaches has been the contrast between Subject Relative Clauses as in sentence (2a) and Object Relative Clauses as in (2b). A large number of studies have found that object relatives (2b, in which the head of the relative clause *girl* is the object of the relative clause verb *kissed*) are harder than subject relatives (2a, where the head *woman* is the subject of *kissed*) in English and many other languages (see Gennari & MacDonald, 2008; O’Grady, 2011, for review). Because these sentences have been assumed to be unambiguous (see MacDonald, 2013, for discussion), the difficulty contrast has generally been thought to owe to inherent complexity differences between the sentences, which affect comprehenders’ inherent computational limits. Gibson (1998), for example, argued that object relatives create a higher working memory load than do subject relatives, because the object relatives involve longer distance relationships between dependent elements, such as between a verb and its subject or object (see also Gordon et al., 2004; Grodner & Gibson, 2005; Hawkins, 1994; King & Just, 1991, for related memory-based approaches).

(2a) Subject Relative Clause: The girl [who kissed the woman] was....

(2b) Object Relative Clause: The girl [who the woman kissed] was...

Recently, several lines of research have questioned this division between experience and memory limitation accounts. First, some researchers have suggested that a full account of

processing difficulty of the (2a-b) contrast requires both memory and experience approaches (e.g., Fedorenko, Woodbury & Gibson, 2013; Staub, 2010). Second, in contrast to this view, MacDonald and Christiansen (2002; also Wells, Christiansen, Race, Acheson & MacDonald, 2009) argued that the memory-experience distinction is untenable, because the amount of memory/computational capacity required to interpret input is itself dependent on experience. The logical argument here is that there cannot be two independent factors if the factors are not actually independent; practically, it is quite difficult to determine the relative contribution of memory limitations independent of experience. Relatedly, Gennari and MacDonald (2008) argued that English relative object relatives actually contain significant ambiguities. This claim is important because lack of ambiguity in relative clauses has been a key component of arguments for inherent complexity accounts—we can attribute processing time differences to inherent complexity differences precisely because ambiguity has been excluded as a possible explanation of the results (MacDonald, 2013). Gennari and MacDonald assessed the degree of ambiguity in relative clauses with a “gated” sentence completion task (McRae, Spivey-Knowlton & Tanenhaus, 1998), in which participants receive a sentence fragment, as in a portion of (2b): *The girl who the...* and complete the sentence. Gennari and MacDonald found that these putatively unambiguous object relative clauses in fact received many completions that were not object relatives, and that the animacy of the head noun (*girl*) strongly determined the rate of object relative completions. Participants readily completed inanimate-headed fragments (*The trophy that the...* as object relatives), but the rate was much lower for animate heads. Importantly, completion patterns for a particular sentence reliably predicted reading times for that sentence. These results suggest that as with the syntactic ambiguities in (1a-b), comprehension difficulty in relative clauses also seems to be linked to animacy-related

ambiguities. Comprehenders' past experiences of pairings of animate and inanimate nouns modifying relative clauses, as revealed by completion patterns, yield different expectations for upcoming structure as the sentence unfolds, and different patterns of comprehension difficulty, as revealed by reading times.

These comprehension results point to an important role for language production processes in understanding language comprehension, particularly in understanding why comprehenders' prior experience and expectations have the character they do. Gennari, Mirkovic and MacDonald (2012) hypothesized that the head animacy-relative clause relationships owe to particular biases within the language production system that promote certain kinds of lexico-syntactic pairings over others (the Production Distribution Comprehension account—PDC; Gennari & MacDonald, 2009; MacDonald, 2013; Montag & MacDonald, 2014, in press). These biases, which we consider in more detail below, conspire to discourage speakers' use of object relatives for animate heads as in (2b). Speakers instead convey the same meaning with passive relative clauses (2c), especially bare passives (without the agent), as in (2d) *The girl who was kissed*. For inanimates, however, the pressure to produce passives is much weaker, and speakers produce a mix of object relatives and passives.

(2c) Passive Relative Clause: The girl [who was kissed by the woman]

(2d) Bare Passive Relative Clause: The girl [who was kissed]

These results: the production patterns and they being significant predictors of the reading times, argue for a strong link between production, statistical distributions in the language, and comprehension of relative clauses. However, it is also clear that the situation is quite complicated

and open to alternative interpretations, e.g., the motivations for speakers' alternative choices, the role of working memory limitations (Gibson, 1998), and the degree to which comprehension or production data require a combination of memory and experience approaches (Fedorenko et al., 2013; Staub, 2010). Extension to other languages would provide an important test of these alternative claims, not only because of the inherent value of replication, but also because the structure of relative clauses differs widely across languages, and because relative clause comprehension has been intensely studied in a number of different languages. After English, perhaps the language for which relative clause processing has received the most scrutiny is Mandarin (e.g., Hsiao & Gibson, 2003; Hsiao & MacDonald, 2013; Jäger, Chen, Li, Lin & Vasishth, 2015; Lin & Bever, 2006, Vasishth, Chen, Li & Guo, 2013). As we will describe below, Mandarin is intensively studied because its relative clause structure is typologically unique, its patterns of comprehension difficulty are markedly different from those found in English and most other languages, and it has been used as a central test case for distance-based memory accounts of relative clause processing. Our goal in this paper is to investigate the production-comprehension linkages in Mandarin relative clauses, from speakers' choices during production, the resulting distributions in the language, comprehenders' expectations for relative clause resolution as they are in the midst of a sentence, and comprehension difficulty, as measured by reading times. We begin with a description of Mandarin relative clauses and the comprehension patterns that have placed them at the forefront of theoretical accounts of relative clause processing.

Production-Comprehension Links in Mandarin

Mandarin has a head-final relative clause structure, meaning that the head noun follows the relative clause, as in the examples in (3). It has an obligatory relativizer DE (similar to *that* in English), but as these examples show, DE does not arrive until the end of the relative clause.

(3a) Subject Relative: [攻击 竞争者 的] 候选人 赢了 这场 选举。

[e₁ attack opponent DE] candidate₁ won this election

The candidate who attacked the opponent won this election.

(3b) Object Relative: [竞争者 攻击 的] 候选人 赢了 这场 选举。

[opponent attack e₁ DE] candidate₁ won this election

The candidate who the opponent attacked won this election.

Japanese and Korean also have this head-final relative clause structure, but these two languages have a rich system of case marking on nouns that presumably aids relative clause processing. Case marked languages usually allow pro-drop (i.e. dropping of subjects or objects), as the case leaves traces for the grammatical role of the omitted noun. When the case is also omitted, pro-drop gives rise to ambiguities in relative clauses processing as the relative clause gap can be mistaken as the dropped noun (see Kwon, Lee, Gordon, Kluender & Polinsky, 2010, for the case of Korean relative clauses). While case-marking is nonexistent in Mandarin, pro-drop is also common and therefore induces similar ambiguities in relative clauses. Hsiao & MacDonald (2013) and Jäger et al. (2015) describe the local ambiguities created by pro-drop in processing relative clauses, as we will also discuss in the next paragraphs. Mandarin is also different from Japanese and Korean in that it has the dominant word order of subject-verb-object (SVO) in main clauses, like English and many European languages, but with the head-final

relative clause structure that is absent in these languages. This combination of SVO basic word order, head-final relative clause structure, and the absence of case marking is attested in the world's languages only in Sino-Tibetan languages such as Bai, and other Chinese languages like Mandarin, Cantonese, and Hakka (Haspelmath, Dryer, Gil, & Comrie, 2005; Keenan, 1985). Thus production and comprehension of relative clauses in Mandarin, and similarities or differences compared to other languages, could be extremely informative about universal vs. language-specific aspects of sentence processing.

There have been relatively few studies of relative clause production in Mandarin, but there has been extensive comprehension research, most of which has examined the comprehension difficulty subject vs. object relatives, as in (3a-b). Results to date have been inconsistent, with some studies finding the typical cross-linguistic pattern of subject relatives being easier than object relatives (e.g. Jäger et al., 2015; Lin & Bever, 2006; Vasishth et al., 2013), and others finding the opposite result (e.g., Hsiao & Gibson, 2003, Gibson & Wu, 2013). This reversal of the dominant cross-linguistic pattern finds a clear interpretation in inherent-difficulty/memory accounts. Hsiao and Gibson argued that subject relatives (e.g. 3a) were more difficult than object relatives (e.g. 3b) because they have higher storage and integration costs, owing to the longer distance between dependencies (see co-indexed filler and gap in the examples) in subject relatives: there are more intervening words between the filler and the gap and thus more new discourse referents and incomplete dependencies in a subject relative than in an object relative. Support for this memory-based view comes from studies in which the added difficulty of subject relatives was higher in participants with lower working memory span (Chen, Ning, Bi & Dunlap, 2008) and patients with aphasia (Su, Li & Chung, 2007). Several studies

manipulating other factors, such as relative clause topicalization (Lin & Garnsey, 2011) and context (Gibson & Wu, 2013), also found a similar object relative advantage.

Still other studies have suggested that the pattern of comprehension difficulty is a function of whether the head noun is the subject of the main clause (as in 3a-b) or the object in the main clause (e.g. *The journalist respected the candidate who...*) (Lin & Bever, 2006; Yang, Johnson & Gordon, 2008). The main clause position effects are related to ambiguity resolution (Hsiao & MacDonald, 2013). A consequence of the head-final structure is that relative clauses are highly ambiguous in Mandarin, as the comprehender who encounters the initial words of the sentences in (3) does not know whether these words begin a relative clause or, (much more common) a simple main clause (Hsiao & MacDonald, 2013; Jäger et al., 2015). For example, the *attack opponent...* sequence at the start of the subject relative in (3a) could instead be the common null-subject pro-drop structure in Mandarin, meaning that someone attacked the opponent. Similarly the object relative in (3b) begins *opponent attack...*, which could be the beginning of a simple sentence in which the opponent is attacking something. Likewise at the main clause object, the relative clause noun of an object relative follows the Noun + Verb sequence of the main clause, leading the comprehender to interpret it as a complete SVO simple sentence at this point. These local structural ambiguities created by the word order of relative clauses and simple sentences affect reading times (for ambiguity and competitor interpretations of relative clauses in Mandarin, see Appendix A in Hsiao & MacDonald, 2013, and Jäger et al., 2015). Hsiao & MacDonald (2013) suggested that the degree of ambiguity in subject and object relative clauses does not only vary with structural factors such as relative clause type and main clause position but also other semantic factors, such as noun animacy. They argued that is possible to reconcile many conflicting results by examining a broader range of lexico-syntactic

conditions than are manipulated in any individual study. Using a neural network model, they were able to show that the complex syntactic and semantic interaction was reflected in the performance of the model, which nicely matched to previous human reading time data. These results argue against any purely structural account of relative clause comprehension such as dependency distance, because the effects of distance actually reverse when considering a broader range of factors.

Hsiao & MacDonald (2013) also presented corpus data relevant to language production. Their results show that rates of production of several relative clause types also are intimately tied to animacy of the head noun. One obvious factor affecting producers' implicit structure choices is the relative accessibility of animate vs. inanimate heads, and many production researchers have observed that animate entities tend to be placed early in the sentence and/or in grammatically prominent positions, such as sentence subject, which yields a passive structure (Bock, 1987; Bock & Warren, 1985; Tanaka, Branigan, McLean, & Pickering, 2011). However, Gennari et al. (2012; MacDonald, 2013) argued that accessibility of animate heads did not explain the whole pattern of results, and that relative clause production choices were also shaped by a bias that MacDonald (2013) called Reduce Interference. Production requires maintaining the utterance plan in memory before it is executed, and semantically similar words held in memory interfere with each other, increasing production difficulty (Smith & Wheeldon, 2004). Gennari et al. observed that in cases where both participants in an event were animate (e.g. a boy splashing a girl), not only were there more passive relative clauses as in (2c), but there were more bare (agentless) passives like (2d). Gennari et al. linked these effects to semantic similarity between the two nouns in the utterance, not simply animacy, and they hypothesized that another motivation for using a passive with an animate patient was to reduce this memory interference—

the passive allows demotion of the agent to a syntactically minor *by* phrase or omission altogether, reducing this similarity-based interference.

It is interesting to consider these claims with respect to Mandarin. These strong animacy effects on relative clause production have typically been conducted in European languages with head-first relative clauses, where the animate/inanimate head noun precedes the relative clause. However, in Mandarin, the relative clause precedes the head noun, and so it is unclear whether this late-appearing element will influence the nature of the earlier relative clause. Hsiao and MacDonald (2013) did see animacy effects in relative clauses in their corpus data, but they did not examine passive relatives (2c-d). The current study addresses this question and investigates the full chain of relationships from production, distribution to comprehension in Mandarin relative clauses. Experiment 1 adopted Gennari et al. (2012)'s picture description paradigm to investigate speakers' choices of relative clause form through manipulating head noun animacy. We then compared the results to corpus data to identify distributional patterns in speech and text. To foreshadow our findings, we found differences between spoken and written production, and we used our and others' corpus and production findings to construct felicitous written stimuli for a gated completion task in Experiment 2 and a sentence comprehension study in Experiment 3. We thus aim to relate comprehension processes in Experiment 3 to comprehenders' expectations, as measured by sentence completion data in Experiment 2, which in turn stem from distributions in the language that owe to speakers' production choices, as in Experiment 1.

The relative clause contrast we investigate in these studies is the contrast between object relatives (3b, also 4a below) and the passive relative counterpart, (4b), which is a type of subject relative in which the verb has passive voice. There are several reasons for this choice. First, the two alternatives are closely matched in meaning, and this is why the object vs. passive relative

was the contrast of choice in Gennari and MacDonald's (2008; 2009) study of the production-comprehension link. In particular, they argued that this choice was preferable to the common practice of comparing object and (active voice) subject relatives as in (3a-b), as these sentence confound sentence structure and meaning. However, we do acknowledge the discourse pragmatic and stylistic differences in the usage of actives vs. passives, making these two structures not entirely equated in meaning. For example, bare passives are sometimes used when the identity of the agent is intended to be hidden. Also, in Mandarin, passives have the connotation of infliction and damage, as we will discuss in detail in the section of corpus analysis. Nevertheless, compared to the typical contrast of subject vs. object relative, the object vs. passive relative comparison allows us to examine producers' choice between two alternatives, which are similar in meaning. With meaning well matched, we can better investigate the effects of structure and its relationship to head noun animacy. Second, Mandarin passive relative clauses do not differ from object relatives in word order; their only difference is a clause-initial passive marker *BEI*, as seen in (4b). Despite the identical word order, passives (which are syntactically subject relatives) have longer dependency distance than object relatives, as shown with the subscripts in (4a-b). This is a very useful property because previous memory-based accounts of relative clause comprehension in both Mandarin and English, have pointed to differences in dependency distance in explaining patterns of processing difficulty in Mandarin object and (active voice) subject relatives (Gibson, 1998; Hsiao & Gibson, 2003). By choosing a contrast with identical word order but different dependency distance, we can test the dependency effect while contrasting it with our own experience-based account, in which comprehension behavior reflects the distributional patterns generated by production choices that are strongly tied to lexical properties of the utterance—head noun animacy.

(4a) Object Relative: [竞争者 攻击 的] 候选人 赢了 这场 选举

[opponent attack e₁ DE] candidate₁ won this election

The candidate who the opponent attacked won this election.

(4b) Passive relative: [被 竞争者 攻击 的] 候选人 赢了 这场 选举

[e₁ BEI opponent attack DE] candidate₁ won this election

The candidate who was attacked by the opponent won this election.

The comparison between the two structures also allows us to examine the production biases proposed by MacDonald (2013). Due to the fact that passives are essentially subject relatives in structure, the head noun fills in the syntactic gap at the prominent relative clause subject position. Gennari et al. (2012) argued that, with their manipulation of head noun animacy, conceptually salient animate nouns were usually assigned the most syntactically prominent function – the subject function. Additionally, like in English, passives in Mandarin allow agent omission (i.e. dropping “candidate” in 4b), which also allows us to test the hypothesis (Gennari et al. 2012; MacDonald, 2013) that producers will produce more agentless passives when both nouns are animate than when one is inanimate.

Experiment 1 – Picture Description Task

This experiment investigated whether noun animacy affected production choices of alternative relative clause forms in Mandarin: active vs. passive relative clauses. We adopted Gennari et al. (2012)’s visual scene description paradigm and used a majority of their materials. Some new experimental pictures were developed to replace ones that did not work well in

Mandarin (e.g. because different verbs are used for actions on animate vs. inanimate entities; see Appendix A for the full set of test pictures). Despite the difference in head noun direction in Mandarin relative clauses compared to languages that have shown effects of head noun animacy on relative clause production (English, Spanish, Serbian; Gennari et al., 2012), we hypothesized that noun animacy would also play a critical role in Mandarin: even though the head noun appears after the relative clause, animate entities would still elicit significantly more passives than inanimate objects, as in other languages. We reasoned that even though the head noun is produced late, speakers must plan it before they plan its modifying relative clause, which should therefore yield effects of head noun animacy. If so, this result could have implications for the degree of incremental processing in production (the degree that planning precedes execution of an utterance). Some researchers have argued for very little advance planning (Kempen and Hoenkamp, 1987; Levelt, 1989; de Smedt, 1996), while others have argued that, depending on speaker needs, more advance planning may occur (F. Ferreira & Swets, 2002; Meyer, 1996). The head-final nature of Mandarin relative clauses would seem to be an excellent example of where more advance planning is needed, and thus we expect animacy effects.

Because the active and passive relative clauses have identical word order in Mandarin (distinguished only by the clause-initial passive marker BEI in the passive relative clause), a preference for animate passives cannot be attributed to preferences for a certain surface word order (as arises in some production environments, e.g. Tanaka, et al., 2011) but rather a preference for a particular grammatical role assignment (e.g. Bock & Warren, 1985): animate entities are more salient and usually occupy the grammatical subject position, and by using passives speakers could align the animate nouns as subject in the relative clause. However, as in Spanish and English, we do not expect that head noun animacy/accessibility is the only

motivation for passive choices, and we hypothesize that when both nouns are animate, participants will be more likely to produce passives in which they omit the agent compared to conditions with one animate and one inanimate noun (Gennari et al., 2012).

Methods

Participants

Thirty-two native speakers of Mandarin Chinese participated in the experiment for course credit in an introductory psychology class at University of Wisconsin-Madison. For this and all experiments reported here, all participants indicated that they had been born and educated in Mainland China or Taiwan and spoke Mandarin Chinese as their dominant language. Most were freshmen or sophomores who did not arrive in the US until the start of their university attendance. All phases of the experiment, including introductory description of the study, obtaining informed consent, and instructions for the task, were conducted in Mandarin with a native Mandarin-speaking experimenter.

Materials and Design

Twenty pictures (most from Gennari et al., 2012, and the rest created specifically for the current Mandarin task, see Appendix A) were used as the experimental stimuli. Each picture showed an action being performed by two human agents. One of these two human agents acted upon an animate patient (e.g. kissing a girl) and the other acted upon an inanimate theme (e.g. kissing a trophy). Also in each picture were additional human characters and objects that were similar to the patient/theme, such as one or more additional girls and trophies in the kissing picture. These similar items were designed to encourage spoken responses with noun modification, especially relative clause modifiers (e.g. “the trophy” doesn’t distinguish among

the two trophies in a picture, and so speakers need to add additional modification to clarify which trophy they are talking about). See Figure 1 for an example experimental picture. Thirty filler pictures were also created. They mostly depicted a scene where a single main event was happening; for example, some people bowling in a bowling alley or several people playing poker.

Two audio questions were recorded by a native Mandarin speaker for each experimental picture. One of the questions targeted the animate patient and the other the inanimate object. For example, the questions recorded for the picture in Figure 1 were “Who has black hair?” and “What is yellow?”, for the animate and inanimate patient/theme of *kiss*, respectively. The structure of the questions and the nature of the pictures (i.e. more than more one girl and one trophy) together encouraged relative clause responses, which specifically distinguish one entity from others (i.e., animate: *the girl that a woman is kissing/the girl that is being kissed by a woman*; inanimate: *the trophy that a girl is kissing/the trophy that is being kissed by a girl*). One question was recorded for each filler item. Filler questions contained a variety of structures and were not designed to elicit relative clauses. Two scripts were prepared to counterbalance the assignment of questions to experimental pictures, so that each participant heard an equal number of animate-targeting and inanimate-targeting questions (10 for each condition) but was asked only one question for each picture. In each script, experimental trials interleaved with fillers in such a way that no more than two experimental pictures occurred in a row.



Figure 1. An example experimental item for the production task in Experiment 1. The target items were the girl and the trophy being kissed. Notice that there were additional entities that resembled the target (e.g. two trophies, one yellow and one orange), making it felicitous to use relative clauses to answer the question (e.g. what is yellow?).

Procedure

Pre-training Task

A pre-training task was administered before the main production experiment. Participants were shown pictures depicting individual actions that would appear later in the actual task. These pictures were cropped from the experimental items. See Figure 2 for an example. Two seconds after the presentation of the picture, the verb that described the action (i.e. *qin* kiss) appeared beneath the picture. Participants were asked to read the verb aloud. The purpose of the pre-training task was to ensure that participants could easily identify the pictured action and would use the designated verb in their production in the actual test so that responses were comparable across participants. Each participant viewed two pictures associated with a verb: the action performed on the animate entity and the inanimate object. This way, participants would not be biased on the animacy of the targets they would see in the main task. The filler pictures were also cropped from the original filler pictures but the accompanying words were nouns instead of verbs. Fillers were included in pre-training so as not to call attention to the experimental items.



Figure 2. An example picture that described the action “kiss” in the pre-training task

Main Task

Upon completion of the pre-training task, participants received instructions about the main task. Participants were told that they would see pictures and answer questions about them. To prompt production of relative clauses, the participants were told that there were multiple similar characters and objects in the pictures and their answers would be shown to a later group of people to guide their recognition of specific characters or objects in the pictures. Participants heard a cover story indicating that describing elements in terms of color or screen position wouldn't help the later group identify the pictured characters, and participants were encouraged to respond by describing the actions that the pictured entities were involved in. Because modification with an action description is accomplished with relative clauses, participants were therefore encouraged to produce relative clauses without having these structures explicitly mentioned in the instructions.

The task began with three practice trials. At the beginning of each trial, a color picture like Figure 1 appeared. Following a three-second delay allowing the participants to inspect the picture, a pre-recorded question was played through speakers. Each participant received ten trials with questions asking about animate patients (such as “Who has black hair?”) and ten about

inanimate themes (such as “What is yellow?”); only one question was presented for each picture. Participants’ spoken responses to the questions were digitally recorded for later analysis. The picture remained on the screen throughout the trial. The picture was removed by a keypress after the participant finished answering, and a new trial began immediately afterwards.

Coding and Results

All responses that did not contain relative clauses were excluded. This process removed 16% of animate trials and 21% of inanimate trials. This level of non-relative clause responses is similar to that in previous studies (Gennari et al., 2012; Montag & MacDonald, 2014) and is not surprising, as participants were never explicitly instructed to produce relative clauses. Excluded trials included ones in which participants did not notice the additional characters or objects that resembled the target and answered only with a noun (e.g. trophy), and ones in which they violated the instructions to avoid color or positional information, producing responses such as *the yellow trophy*, or *the trophy on the right*.

The participants’ relative clause responses were coded as Passive or Active. Analyses in all experiments here were conducted with linear mixed-effects models, using the function `lmer` in the `lme4` package in the statistical software R. A logit mixed effects model was fitted for the data with Animacy as the fixed factor and Subjects and Items as random factors with random intercepts and slopes. As in languages with head first relative clauses (English, Spanish, and Serbian; Gennari et al., 2012; Montag & MacDonald, 2014), animacy significantly influenced participants’ choice of relative clause structures ($\beta = -0.243$, $SE = 0.066$, $t = -3.68$). An overwhelming 98% of responses describing animate entities were passives, while 75% of responses in the inanimate condition were passives, shown in Figure 3.

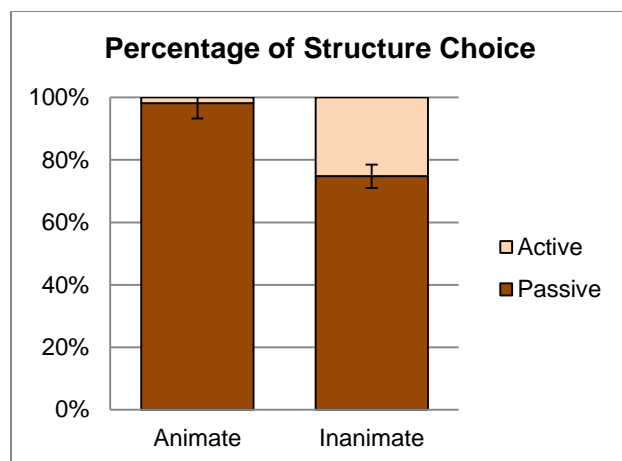


Figure 3. The percentage of structural choice as either in the active or passive relative clause forms when describing animate or inanimate entities. The error bars indicate standard error of means.

We also examined, within passive responses, the rate at which participants omitted the agent of the action in bare passives like [*BEI kiss DE*] *trophy* vs. full passives with the agent noun, e.g. [*BEI girl kiss DE*] *trophy*. Animacy of the head noun significantly affected whether the agent nouns were omitted ($\beta = -0.10$, $SE = 0.04$, $t = -2.67$). 40% of agents were dropped in the animate condition compared with the 31% in the inanimate condition, as shown in Figure 4. This result replicates findings in English and Spanish (Gennari et al., 2012; Montag & MacDonald, 2014).

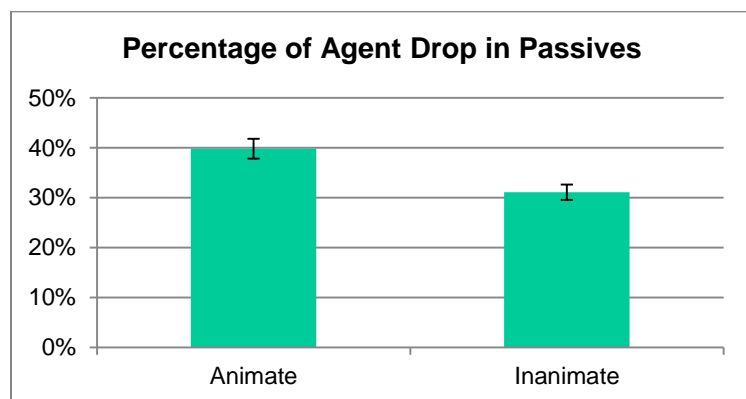


Figure 4. Percentage of agent-drop among passive relative clause responses as a function of animacy of the head noun. The error bars indicate standard error of means.

Discussion

In this picture description task, Mandarin speakers showed an overwhelming passive preference in describing an animate entity, comparable to data from English speakers that Gennari et al. (2012) and Montag and MacDonald (2014) report, and higher than the rates for Spanish and Serbian speakers reported in Gennari et al. The 75% passive rate for inanimates was also very high; English, Spanish, and Serbian speakers in a similar task produced 50% or fewer passives (Gennari et al., 2012, Montag & MacDonald, 2014). Thus Mandarin speakers find passive relative clauses a very viable option for describing both animate and inanimate entities in this spoken production task.

While this study did not investigate factors (other than animacy) that could affect passive rates, we can speculate why it is such a favored option in Mandarin. The relatively high rate of passivization in both the animate and inanimate conditions may be attributable to an interaction among demands of on-line spoken production, the relative clause word order in Mandarin, and noun animacy. To passivize a relative clause in Mandarin, one only needs to add a passivizer *BEI* at the beginning of the phrase. Thus the choice to passivize must be made early, but the realization should be quite easy, involving only a single high-frequency morpheme. The early position of *BEI* and its optional use may allow it to function similarly to how the relative pronoun and complementizer *that* function in English—speakers use these forms in part to provide more time to plan their utterance, especially when they encountered production difficulty (Ferreira & Firato, 2002; Jaeger, 2005; Jaeger & Wasow, 2005; Race & MacDonald, 2003). We are not claiming that added planning time from the addition of phrase-initial *BEI* is the sole

reason of passive use in Mandarin, but future work should investigate whether this feature of passive formation in Mandarin contributes to high rates of passive in this task.

Another incentive for passivization in Mandarin is that while both agents and patients must be mentioned in object relatives, as in English, the agent of the action can be omitted in passives. Gennari et al. (2012) argued that the option to omit the agent was a force behind some passive use, and that the option to omit the agent is beneficial under conditions of similarity-based interference, where both the agent and patient are semantically similar (e.g. *girl* and *woman*). The Mandarin results are consistent with these claims for relative clauses and also consistent with results in simple main clauses in Mandarin, in which agent omission is higher under conditions of similarity-based interference (Hsiao, Gao, & MacDonald, 2014). As one reviewer noted, bare passives may have occurred due to the pragmatic environment set up by the task, where no more than one agent of similar traits were doing the same action (e.g. only one woman is kissing the a girl and only one girl is kissing a trophy, in Figure1). If there were multiple similar agents acting upon similar entities (e.g. two women kissing different girls, two girls kissing different trophies), then dropping of the agents might not have been as frequent. We agree with that the speculation that when only one agent of the kind did the action, specifying only the action but not the agent would be sufficient. However, the fact that we did observe a difference of agent omission between the animate and inanimate conditions suggests that speakers were affected by the semantic similarity between the agent and the patient/theme. Finer grained measures, such as similarity rating scores, will be required to statistically confirm the similarity-based interference effect.

The strong effects of animacy in this study are notable because of the special features of Mandarin relative clauses. First, owing to the head-final word order of Mandarin relative clauses,

the animacy effects in this study show that a late word (the head noun) affects the earliest word in the relative clause. This result that argues against the most radical accounts of incremental processing (Kempen and Hoenkamp, 1987; Levelt, 1989; de Smedt, 1996) because it suggests that speakers must have planned at least some aspects of the head noun before uttering the first words of the relative clause. Second, as we've noted, the active and passive forms are identical in words and word order except for the use of *BEI*. Thus in contrast to many other languages, the active/passive choice here has minimal consequences for the surface form of the utterance, and yet speakers are strongly affected by head noun animacy in their choice of these two forms.

We next investigate the consequences of these relationships for the distributional statistics of the language, using a corpus of written Mandarin. We chose a written corpus because in many languages, the use of passives can differ in written and spoken genres. A corpus analysis also allows us to look beyond the particular materials of our production study. The pictures we used in the production task always involved transitive physical actions that often (though not always) implied adverse outcomes or infliction, which is the connotation embedded in the Mandarin passive morphemes, including *BEI*. Analysis of a written corpus would also allow us to examine the effect of the specific discourse context created for the current production task. Since our test of comprehension will be reading times, we will use the written corpus to investigate passive and active relative clause use in texts.

Corpus Analysis

The aim of this study was to investigate the relationship between head noun animacy and relative clause structure in Mandarin. While there are other previous corpus analyses of Mandarin relative clauses (Hsiao & Gibson, 2003; Hsiao & MacDonald, 2013; Pu, 2007; Wu,

2009), they have not fully considered these lexico-syntactic pairings. The strong animacy effects observed in Experiment 1 suggest that corpus analyses considering only overall frequencies of relative clause types will not adequately capture the important animacy-structure correlations that comprehenders are likely to use in comprehension. That is, even though the animate head noun follows the relative clause, we expect that there are correlations between animacy and structure type that help comprehenders interpret relative clauses and anticipate the animacy of an upcoming head noun. Wu (2009) and Hsiao & MacDonald (2013) are the two corpus analyses on Mandarin relative clauses that included noun animacy as a factor, but neither one fully investigates the conditions relevant to the production data in Experiment 1.

Methods

Hsiao and MacDonald (2013) investigated several kinds of active relative clauses in the Chinese Penn Treebank 7.0 (Xue et al., 2010), but they did not investigate passive relatives. For this study we built on Hsiao and MacDonald's coding and extracted passive relative clauses from the same source. The Chinese Penn Treebank is a parsed corpus consisting of more than one million words in more than 50,000 sentences, with various sources from the news genre, such as newswire data, broadcast materials, web text, magazine articles, from Mainland China and Hong Kong. As in Hsiao & MacDonald, we used TGrep2 1.15 (Rohde, 2005) to extract passive relative clauses, including both full passives (with agent) and agentless forms (see Appendix B for search patterns). We hand coded the head noun animacy in each relative clause, following the criteria used by Hsiao & MacDonald (2013): animate entities refer to living beings with agency and volition to perform an action, whereas inanimate entities refer to those without this property (Hundt, 2004). Coding was performed by a Mandarin-speaking research assistant.

Results

The frequencies of objective and passive relatives coded for the animacy of their head nouns are reported in Table 1. The first row of the table, showing rates of object relatives, indicates that this structure is much more frequent with inanimate head nouns than animate ones, a result that is consistent with prior corpus results in English (Gennari & MacDonald, 2009). If comprehenders are sensitive to these contingencies, then they should tend to expect inanimate heads for object relative clauses (assuming that the lexical content makes that plausible, and also assuming that comprehenders know that they are encountering an object relative).

Table 1. Frequency counts of object relatives and passive relatives in the Chinese Penn Treebank

7.0

	Animate head noun	Inanimate head noun	Total
Object relatives	21	288	309
Full Passive relatives	10	7	17
Bare Passive relatives	10	14	24
Total	41	309	350

The table also shows striking differences to the picture description responses in Exp.1, in that the rate of passives in the written texts is far fewer than in our spoken production task (given the low counts of passives in the corpus, we cannot evaluate the animacy effect here). Roland, Dick & Elman (2007) also reported a frequency disparity of passive relatives between English speech and text, but in the opposite pattern: passives are rarer in the spoken data than in written corpora. They attribute the difference to discourse register: one of the functions of passives is to de-emphasize the agent, which is less necessary in conversation when the speakers and the speakers' actions are the focus of discussion. Discourse register may also be a plausible cause for

the difference found in speech and text in Mandarin but with a different language-specific story. Passive markers, including the most frequently used BEI and other rarer markers that have been grammaticalized to serve the purpose (e.g. *ai, zao, ran, gei, jiao*), and most of them have a negative connotation related to infliction (Chao, 1968). In their corpus analysis, Xiao, McEnery & Qian (2006) found passives in general are less frequent in spoken data than in text, which initially seems to conflict the Chinese Treebank results here. However, Xiao et al. found further differences among genres within the written corpus, with passives being least frequent in newspaper, news editorials and official documents. Because the Chinese Treebank consists of texts from news sources, the low rate of passives in the Chinese Treebank may be due a tradition of low passive use in this genre. Xiao et al. suggest that the high rate of passives in English written texts, especially in formal documents, is due to the function of passives to mark objectivity, whereas passives in Chinese do not have this function. They further noted that the passive marker BEI is used to denote a predominantly negative tone in speech and it does not show a negative prosody frequently in official documents. This potentially accounts for the high rate of passivization in our production data in Experiment 1, where the materials mostly involved physical infliction and damage.

These results suggest that genre has a large effect on passive use in Mandarin. The spoken production task in Experiment 1 appears to have created conditions (adversative events) for which passives are highly felicitous, whereas the Chinese Treebank appears to have selected genres in which passive use is quite rare. Both the spoken production data and the corpus data point to important correlations between head noun animacy and relative clause structure, but neither of these studies is assessing passive use in a discourse environment that would be useful for a comprehension experiment—written materials for which passive relative clauses would be

a viable option. We therefore developed new stimuli for a written sentence comprehension study. These stimuli crossed relative clause type (object relative, passive relative) and head noun animacy, including a preceding context to make the sentence plausible. We examined interpretation of these stimuli first in a gated sentence completion task in Experiment 2 and then in a self-paced reading task in Experiment 3.

Experiment 2 – Gated Sentence Completion Task

In this experiment, we investigated comprehenders' syntactic and semantic interpretations at various points in Mandarin relative clauses using a gated sentence completion task similar to the one in Gennari and MacDonald (2008). An important result in their study was that the animacy of the head of the relative clause affected the degree to which participants adopted an object relative clause interpretation of the sentence input that followed the head noun, which in turn predicted comprehenders' reading times for sentences that turned out to be object relative clauses. Our goal here is similar, in that we aim to investigate how material early in a potential relative clause influences interpretation and thus comprehension difficulty downstream. However, the differences in word order between English and Mandarin do not allow us to examine exactly the same factors. English relative clauses are head-first and thus naturally permit an investigation of early-arriving head noun properties (such as animacy) on the subsequent relative clause. Mandarin has head-final relative clauses, and so head noun animacy is not a predictor but potentially something that is predicted by other, earlier linguistic material.

Another consequence of the head-final relative clause structure in Mandarin is ambiguity. Comprehenders do not necessarily know that the structure they are perceiving is a relative clause. In particular, simple sentence structures are a much more common alternative interpretation for

the Noun + Verb word order that initiates an object relative clause (Hsiao & MacDonald, 2013). The relativizer DE provides strong support for a relative clause interpretation (though other interpretations of DE exists, see Hsiao & MacDonald, 2013, for details), and our choice of gates was designed to identify the interpretations that dominate before and after encountering DE. We investigated whether comprehenders consider a relative clause interpretation of the input before the relativizer DE, and whether the sentence initial passive marker BEI influences expectations for a relative clause. BEI may promote a relative clause interpretation because it functions as a light verb, and a sentence starting with a verb is uncommon, as it is only seen in pro-drop sentences, gerunds or subject relatives, including passive relative clauses. At the point where DE makes a relative clause highly likely, we investigated whether completions favor an animate or inanimate head noun. Experiment 1's production data and corpus results both showed that object relatives are more common with inanimate head nouns, whereas passives were more common with animate heads in Experiment 1, and we expect these relationships between active/passive and head animacy will arise in our completions.

Methods

Participants

56 native speakers of Mandarin Chinese on the campus of University of Wisconsin-Madison participated in exchange for credit in an introductory psychology course, composed of mostly freshmen and sophomores who arrived in the US within the past two years. All of them read and wrote in Simplified Chinese, which was the written script adopted in this experiment. No participants participated in any of other experiments reported here.

Materials and Design

Because relative clauses are much more felicitous in some discourse contexts than others (in Mandarin and in other languages), we developed our sentence fragments for the completion task to include an introductory context phrase for the upcoming relative clause that made additional relative clause modification felicitous. As it is effectively impossible to develop a discourse context that is equally appropriate for an animate and an inanimate head noun modified by a relative clause, we developed two contexts for each fragment, one that supported relative clause modification of an animate head, and one that supported a relative clause modification of an inanimate head. A fragment that could introduce a relative clause immediately followed this context. This potential relative clause was presented either in active form or in passive form, via the insertion of the passive marker BEI clause-initially. The fragments were truncated at two different gates for completions. At the Early Gate, fragments had a “(BEI) Noun Verb” sequence and ended before the appearance of a relativizer DE. The Late Gate was identical but contained one more word, the relativizer DE, which strongly promotes a relative clause. As Table 2 shows, these manipulations yielded a 2x2x2 design, with Context Type (animate head-oriented vs. inanimate head-oriented) and Structure Type (active vs. passive) as within-subject factors and Gate (Early vs. Late) as a between-subject factor. We chose to make Gate a between-subject factor in order not to introduce priming effect, as participants exposed to both Early and Late Gate would potentially be primed to complete the Early Gate fragments with relative clauses. A total of 28 items were created, counterbalanced in eight scripts (see Appendix C for the full list of experimental items adapted from Exp.3).

Table 2. Example items for the sentence completion task

Condition		Example	
Animate head-oriented context 因为不愿意引起轩然大波， Unwilling to cause an uproar,	Early Gate	Active	媒体批评_____。° media criticized_____.
		Passive	被媒体批评_____。° BEI media criticize_____.
	Late Gate	Active	媒体批评的_____。° media criticized DE _____.
		Passive	被媒体批评的_____。° BEI media criticize DE _____.
Inanimate head-oriented context 虽然剧情沉闷了无新意， Although the plot was dull and unoriginal,	Early Gate	Active	媒体批评_____。° media criticized_____.
		Passive	被媒体批评_____。° BEI media criticize_____.
	Late Gate	Active	媒体批评的_____。° media criticized DE _____.
		Passive	被媒体批评的_____。° BEI media criticize DE _____.

Note: Early vs. Late Gate was manipulated between subjects, whereas passive and context type were manipulated within-subjects.

Procedure

The experiment was conducted entirely in Mandarin with a native-speaking experimenter. Participants were randomly assigned to the Early or Late Gate condition and were given a paper survey that listed the sentences to be completed. The active and passive sentence fragments were semi-randomly ordered so that no more than three of the same sentence type (i.e. by head animacy orientation and by active/passive voice) appeared in a row. The same randomized order was applied to each script and to every participant (i.e. all eight scripts had the same item order and all participants who were assigned to a certain script received the same order of stimuli). No time limit or minimal length of a sentence were imposed. Participants were encouraged to continue the sentences until they felt the sentences were complete in the semantic and grammatical sense, and they were told that there were no right answers. Participants' completions were coded by research assistants who were native speakers of Mandarin.

Completions were coded for whether they adopted a relative clause interpretation or some other interpretation of the (BEI)+Noun+Verb sequence, and in the case of relative clause interpretations, whether they were completed with an animate or inanimate head noun. The codings were double checked by a second coder to ensure accuracy.

Results

The results of percentage of completion at the Early Gate and the Late Gate, given either the object relative or passive fragments, are presented in Figure 5.

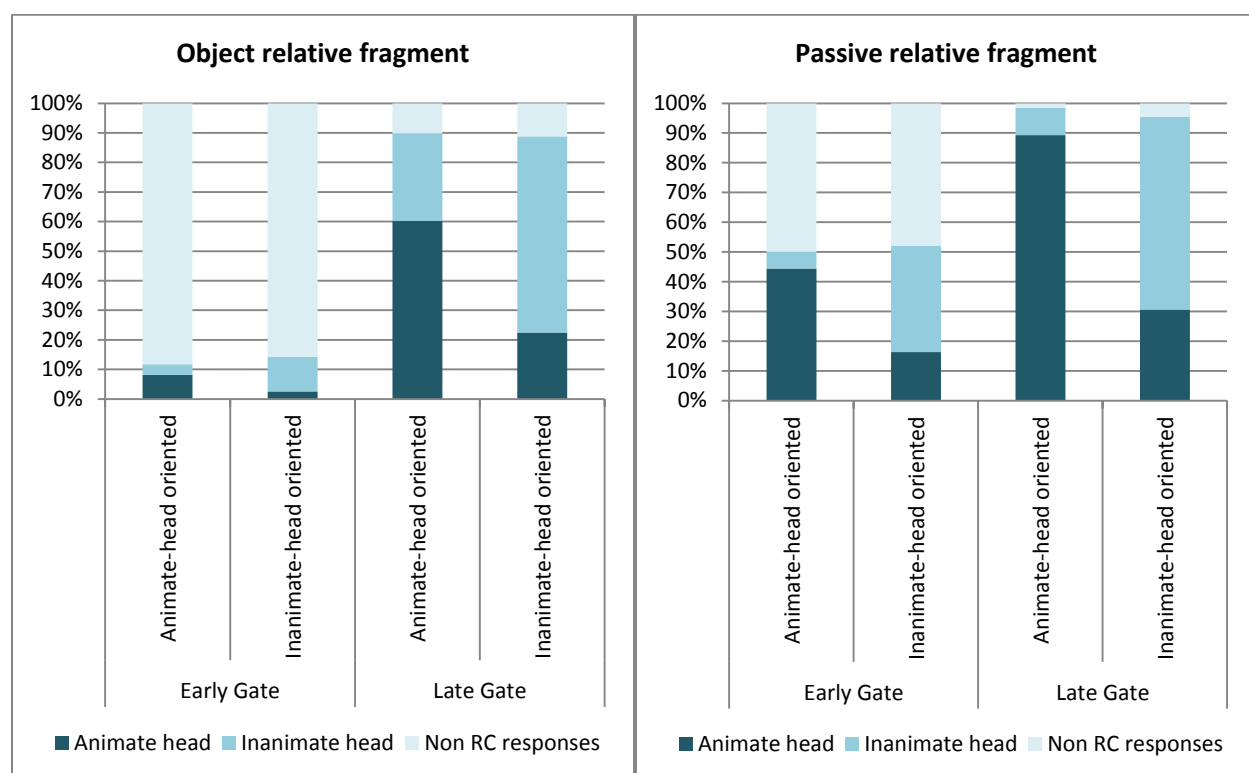


Figure 5. Percentage of completions with 1) relative clauses with animate heads, 2) relative clauses with inanimate heads, and 3) non-relative clause responses, given either the object relative fragment (Noun + Verb + DE) (left panel) or the passive relative clause fragment (BEI + Noun + Verb + DE) (right panel) at Early Gate (before DE appeared) and Late Gate (after DE appeared).

Early Gate

Without the relativizer DE at this gate, completions indicate whether alternative interpretations existed before a clear relative clause cue appeared. There was a large effect of structure on participants' completions: for passive fragments, 51% of the completions were relative clauses (30% completed with animate heads, 21% with inanimate heads, 49% with non-relative clauses) compared to only 13% in the active condition (5% completed with animate heads, 8% with inanimate heads, 87% with non-relative clauses). Most continuations for the active fragments were a simple sentential object, making the sentence a simple Subject-Verb-Object sentence. The high proportion of relative clause completions in the passive condition suggested that the passive marker BEI provided a strong incentive for participants to continue the sentence with a relative clause structure. Among all the valid active relative clause completions, around 41% were completed with an animate head noun, compared to 60% for passive relative clause completions. The context phrases appeared to constrain the animacy of the head nouns for those completions with relative clauses. Among the very few relative clause completions for object relative fragments, animate head-oriented context elicited 70% of animate head completions and inanimate head-oriented context elicited 82% of inanimate head completions. For passive fragments, the pattern was 89% and 68% for the two conditions. The relatively higher constraining power of the animate context (and lower constraint of the inanimate context) for the passives seems to reflect the preference of using passives to modify animate nouns, as discovered in Experiment 1.

Late Gate

At this gate, the relativizer DE was provided in the sentence fragment, and many more relative clause completions were expected here. However, DE has other uses beyond a relativizer, and so responses were not 100% relative clause completions at this gate. There were a few non-relative clause responses resulted from apparent misreading of the fragment (e.g. mistaking DE as another homophone), but most non-relative clause completions reflected alternative interpretations of DE. The majority of the responses were relative clause continuations, with 89% in the active frame (41% completed with animate heads, 48% with inanimate heads, 11% with non-relative clauses) and 97% in the passive frame (60% completed with animate heads, 37% with inanimate heads, 3% with non-relative clauses), again suggesting that BEI is a strong indication for a relative clause and the active fragment, even with DE present, still elicited alternative non-relative clause structure. 62% of the responses among all valid responses in the passive condition were completed with animate head nouns, whereas 46% were animate head in the active condition. In the active frame, the animate head-oriented context elicited 67% of relative clause completions with animate heads, whereas inanimate context yielded 75% of inanimate heads completions. For the passive frame, the percentage was 91% for animate context and 68% for inanimate context. These patterns resemble those of the Early Gate. The passive fragments produced more conforming responses corresponding to context phrases targeted toward animate nouns.

Discussion

This study yielded several important results. First, the completions at the Early Gate are a testament to the extensive ambiguity of relative clauses in Mandarin, with most participants interpreting the relative clause fragments as main clauses. Second, the difference between the

passive and object relative fragments showed that, even though relative clauses are head-final and the (largely) disambiguating relativizer arrives late, the early portions of a relative clause affected comprehenders' expectations for upcoming material. In particular, the result at the Early Gate showed that the passive marker BEI was a strong cue for a relative clause even before the relativizer DE. This result is important because main clause completions are fully grammatical at the early gate, and relative clauses are much less frequent than main clauses, and yet, in the presence of *BEI*, there were nonetheless a number of relative clause completions. This result suggests that in an appropriate discourse context in written text, passive relative clauses are a viable option for Mandarin speakers. Third, the presence of BEI not only promoted a relative clause interpretation, but it also more often led participants to expect a relative clause that modified an animate head noun. This result suggests that comprehenders generate expectations about the pairing of passives and animate heads, consistent with the patterns in the production study.

The next experiment investigates the consequences of these expectations, using the Experiment 2 materials (expanded to be complete sentences) in a comprehension task. Common animacy-structure pairings should result in easier comprehension, as measured by reading times, than uncommon pairings, as in Gennari and MacDonald (2008; 2009). Alternatively, if comprehension difficulty is affected by dependency distance (Gibson, 1998; 2000), then passive relatives, which have a longer dependency distance, should be harder than object relatives.

Experiment 3 – Comprehension task

In this experiment, we investigate the relation between the interpretation preferences of the gated sentence completions in Experiment 2 and the reading times in sentence

comprehension. We hypothesized that the degree of ambiguity revealed in the completion study would positively correlate with the time spent on reading the type of sentence. For example, we expect that the high rate of passive relative clause completions compared to the active completion rate in Exp. 2 at Early Gate would predict shorter reading times for the early part of passive relative clauses, and that the high rate of animate head continuations after the passive relative fragments at Late Gate would predict shorter reading times for passive relative clauses after the head nouns were revealed to be animate. That is, the more likely an interpretation at specific relative clause positions, the less likely it is for alternatives to compete for activation and therefore easier to comprehend when the interpretation matches the intended reading. To evaluate this possibility, we regressed the matching completion rate of each item at a certain gate (e.g. the Late Gate completion rate of a relative clause with an animate head after a passive fragment) onto the reading times of specific word positions in the corresponding item (e.g. the reading times at or after the animate head noun of passive relative clauses).

Methods

Participants

Forty native speakers of Mandarin Chinese enrolled in Introductory Psychology courses at the University of Wisconsin-Madison participated for course credit. Most of them were freshmen or sophomores and lived in the United States for less than two years at the time of the experiment. All of them read Simplified Chinese. Ten participants were assigned to each of the four counterbalanced scripts.

Materials and Design

Two factors were manipulated: the structure of the relative clause (active, passive) and the animacy of the head noun (animate, inanimate), yielding a 2x2 within-subject experimental

design. As in Experiment 2, each stimulus sentence began with a context phrase that created an expectation for an upcoming noun phrase and its corresponding animacy properties. Following the context, the critical sentence always had the sequence: [relative clause] + head-noun + adverb + matrix verb..., followed by one or more additional words to make a sensible sentence. The relative clause always modified the main clause subject, and this subject (head) noun was always immediately followed by an adverb, then the main clause verb. The purpose of including an adverb was to lengthen the number of words that were comparable across items. Examples of the four conditions are in Table 3. As these examples show, the only difference between the active and passive relative clauses was the presence/absence of the passive marker BEI.

There were 28 experimental items in total, 7 in each condition. 50 filler sentences were also created, containing various syntactic structures. Four scripts were created to counterbalance the assignment of items to conditions. A comprehension question was constructed for each test and filler sentence. Half of the questions for the experimental sentences were focused on the head noun and the other half on the relative clause-internal noun. The distribution of yes and no as answers was also counterbalanced across both test and filler items (see Appendix C for the full list).

Table 3. Example items in the comprehension task.

Condition	Example
Animate-headed Active	因为 不愿意 引起 轩然大波, [媒体 批评的] 歌手 低调地 不做任何回应。 Because unwilling cause uproar, [media criticize DE] singer low-key NEG make any comment Unwilling to cause an uproar, the singer that the media criticized made no comments in a low-key way.
Animate-headed Passive	因为 不愿意 引起 轩然大波, [被 媒体 批评的] 歌手 低调地 不做任何回应。 Because unwilling cause uproar, [BEI media criticize DE] singer low-key NEG make any comment Unwilling to cause an uproar, the singer that was criticized by the media offered no

	comments in a low-key way.
Inanimate -headed Active	虽然 剧情 沉闷 了无新意, [媒体 批评 的] 电影 意外地 获得 大奖。 Although plot dull unoriginal, [media criticize DE] movie surprisingly won major award Although the plot was dull and unoriginal, the movie that the media criticized surprisingly won a major award.
Inanimate -headed Passive	虽然 剧情 沉闷 了无新意, [被 媒体 批评 的] 电影 意外地 获得 大奖。 Although plot dull unoriginal, [BEI media criticize DE] movie surprisingly won major award Although the plot was dull and unoriginal, the movie that was criticized by the media surprisingly won a major award.

Note: The sentences in Mandarin are marked with word boundaries by spaces here. Word-by-word glosses are provided here, along with full-sentence translation at the bottom in each cell. Relative clauses are marked with brackets here but were not shown to participants.

Procedure

Sentences were presented with the software Linger (Rohde, 2001), which allowed self-paced word-by-word reading with a moving window display. The stimuli were presented in a quasi-random order so that no more than two experimental items appeared adjacent to one another. Each participant received a different randomized order.

Participants were instructed to read for comprehension and answer a question after each sentence. Feedback was given for incorrect responses to the comprehension question, whereas correct answers would trigger the next trial with no feedback. Six practice trials were administered before the actual experiment. Participants could choose to take a short break after finishing half of the experiment.

Results and Discussion

Mixed effects models were applied for statistical analysis, with Structure and Animacy being fixed factors and subjects and items being modeled for random intercepts. The R package *lme4* was used for modeling and the library language was used for significance testing.

Comprehension Question Accuracy

The mean comprehension accuracy was high overall: 93.93% for experimental items and 91.21% for fillers. There was no difference in accuracy and comprehension question response times among the four conditions (animate active: 95.36%, 1802 ms; inanimate active: 91.58%, 1713 ms; animate passive: 96.42%, 1607 ms; inanimate passive: 92.36%, 1900 ms), suggesting that participants' attention and item difficulty were comparable across conditions. Due to the high overall accuracy, no data were removed.

Word-by-word Reading Times

The analysis was conducted on the reading times in a seven-word critical region: BEI (present only in the passive condition), the relative clause subject, the relative clause verb, DE, the head noun, the post-head adverb and the matrix verb.

Reading times, including those of the filler trials, underwent trimming if beyond 2.5 standard deviation of the participant's mean or longer than 4,000 ms, affecting 2.6% of all data points. These reading time data were analyzed following the steps for a two-stage data treatment described in Jaeger's Lab Blog (<http://hlplab.wordpress.com/2008/01/23/modeling-self-paced-reading-data-effects-of-word-length-word-position-spill-over-etc/>). First, a log transformation was applied to normalize the distribution. The log-transformed reading times were further regressed on a series of predictors to control for extraneous effects, with random intercepts for each participant: difference between fillers and test items ($\beta = -0.016$, $SE = 0.003$, $t = -5.57$), word length ($\beta = 0.052$, $SE = 0.002$, $t = 26.70$), the log-transformed position of the trial in the list ($\beta = -0.108$, $SE = 0.002$, $t = -70.23$), and the restricted cubic spline of the word position in the sentence (we followed the suggestion of one reviewer by equating the position of BEI with that of the following word in order to avoid stripping away the potential effect of filler-gap distance) (β

$\beta_0=0.016$, $SE=0.003$, $t=5.43$; $\beta_1=-0.040$, $SE=0.017$, $t=-2.04$; $\beta_2=0.030$, $SE=0.066$, $t=0.45$; $\beta_3=0.145$, $SE=0.077$, $t=1.89$). The residual log reading times were computed from the above model and served as the dependent variable for the analyses in the next phase. Figure 6 shows on the two panels, separated by animacy, the mean residual log reading times of active and passive relatives at each word position in the 7-word region.

The factors of interest (Animacy, Structure) and their interaction were applied to the second phase of the analyses. Even though the animate vs. inanimate head noun contrast does not appear until word 5, Head noun animacy is included as a factor throughout the analyses because the preceding discourse context varied in setting up a plausible animate or inanimate head noun. Because the word BEI only existed in the passive condition, no structure effect was modeled at this position. Spillover effects of the three words preceding a given word n were also included in all analyses: $n-1$, $n-2$, and $n-3$. The full details of statistical results, including p values, can be found in Appendix D. The following section will only describe the effects of Animacy, Structure and their interaction.

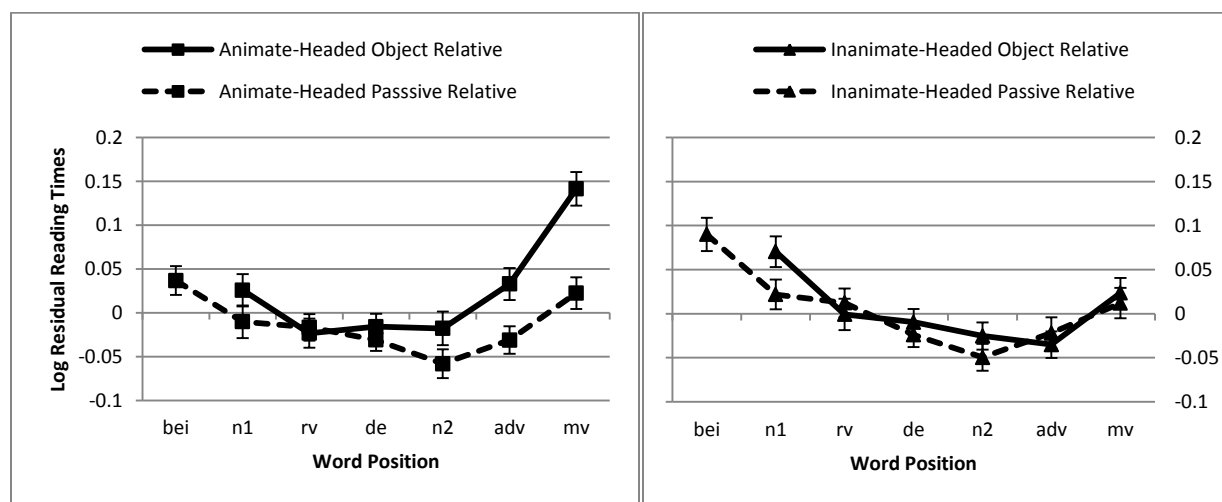


Figure 6. The mean log residual reading times of active and passive relative clauses at each word position, with the left panel showing animate-headed relative clauses and the right panel inanimate-headed relative clauses. *bei* = passive marker (in passive condition only), *n1* = first

noun of relative clause (the agent of the action), *rv* = relative clause verb, *de* = relativizer, *n2* = head noun, *adv* = adverb, *mv* = main verb

Early word positions. Significant animacy effect was found at BEI ($\beta = 0.055$, $SE = 0.023$, $t = 2.36$), showing context clauses oriented at animate heads created more expectation for passives than did context clauses oriented at inanimate heads. The reading times for the next four words (up to the head noun) did not show any main effects of animacy, structure, or the interaction. However, a simple comparison within each level of animacy indicated significant difference in reading times between object and passive relatives at the relative clause subject (*n1*) position: passives were read faster than actives in the inanimate condition ($\beta = -0.054$, $SE = 0.023$, $t = -2.41$) but not in the animate condition. A numerical but non-significant difference was observed between animate-headed object relatives and passive relatives ($p = 0.13$).

Later word positions. Figure 6 shows that at the adverb and the main verb, one condition - the object relatives with animate heads, is more difficult than the other three. Specifically, at the adverb position, there was a reliable interaction of animacy and structure ($\beta = 0.077$, $SE = 0.031$, $t = 2.52$), such that animate-headed object relatives were read reliably more slowly than animate-headed passives ($\beta = -0.062$, $SE = 0.022$, $t = -2.82$), and no difference at this position for the inanimate-headed sentences. Similar patterns were found in the last word region, the main verb (*mv*). There was again a reliable interaction ($\beta = 0.097$, $SE = 0.033$, $t = 2.97$), such that object relative sentences with animate heads were read significantly more slowly than animate-headed passive relatives ($\beta = -0.115$, $SE = 0.024$, $t = -4.84$).

In summary, both animacy and structure had clear effects on comprehension difficulty. Beginning with a marginal effect at the head noun and continuing with reliable effects in the following two words, active object relative sentences with animate heads were the hardest among

the four conditions. As a reviewer correctly pointed out, the effect of context type could contribute to the difficulty. Evident in the sentence completion task, the context phrase effectively constrained the animacy of the head noun. When the participants read the animate head-oriented context, they formed a strong expectation for an animate main clause subject. Given that the relative clause-internal noun (the agent) was always animate and it occurred as the first word after the context phrase, the animate agent in an object relative would be easily mistaken as the main clause subject, which would later have to be reassigned to the animate head noun. This could potentially account for the high reading times of animate-headed object relatives after the head noun, as participants were having difficulty revising the interpretation. However, by simply adding a BEI in front of the agent noun, people quickly settled on a relative clause reading. The results suggest that the passive marker significantly reduced ambiguity and provided strong cue, especially when the modified nouns were animate, matching what we found in the previous two experiments. Similar interactions between structure and animacy have been observed in the head-initial English, in spite of the striking cross-linguistic differences. Gennari and MacDonald (2008; 2009) found in their self-paced reading data that active relative clauses with animate heads were also the hardest to comprehend. They used regression analyses to link expectations, as measured by their completion studies, to reading times, and we pursue similar analyses below.

Regression of Sentence Completion Data onto Reading Times

To investigate the effect of expectations (completions) on reading times from Experiment 3, we conducted regression analyses to relate the mean completion patterns for each item to the mean reading times for that item. As noted in Experiment 2, the presence of DE can strongly alter the interpretation of the sentence, especially the active object relative clause structure,

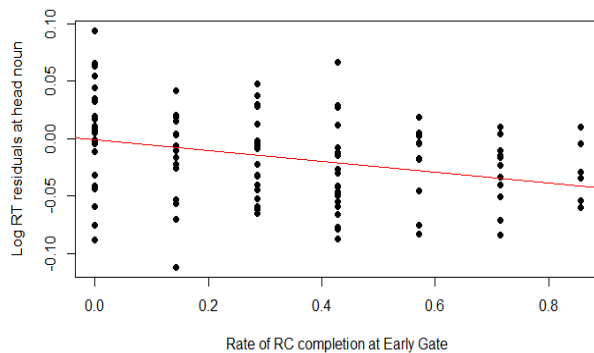
because those sentences are consistent with a simple main clause interpretation before DE is encountered. Because there were very few relative clause completions at the Early Gate, we did not include expectations for a particular head noun animacy at this position and instead regressed only the average rate of relative clause completions for each item. At the Late Gate, after DE had been encountered, relative clause completions predominated, and at this position we used as a predictor the average rate of relative clause completions with the matching head noun animacy specified by the context phrase (identical to the head noun animacy of the expanded sentences used in the comprehension study).

The statistical model was built with the log reading time residuals of the corresponding item obtained in the self-paced reading study as the dependent measure. The fixed effects were 1) the percentage of relative clause completions at the Early Gate and 2) the percentage of relative clause completions at Late Gate that had the same head noun animacy as in the corresponding comprehension item. We ran a separate model for each word starting from DE, which strongly disambiguated the stimuli in favor of a relative clause interpretation. Only the Early Gate completions were used as a predictor in the model for the reading times at DE because the Late Gate was later than DE. Subjects and items are random effects with random intercepts because the fixed effects here did not vary across subjects and items.

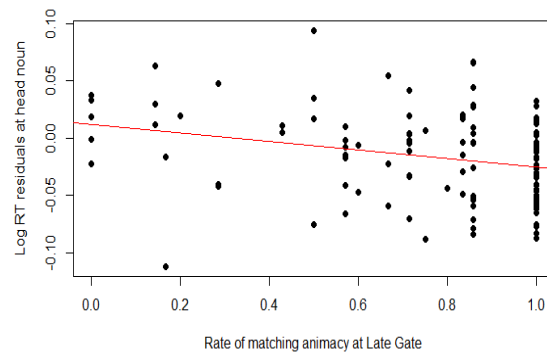
No effect of Early Gate completions were found at DE. However, at the head noun position, both Early and Late Gate completions significantly predicted log RT residuals in a negative trend: higher rates of relative clause completions at the Early Gate and higher rates of the head noun completions with matching animacy at the Late Gate predicted lower reading times at the head noun position in the corresponding item (Early Gate: $\beta = -0.076$, $SE = 0.032$, $t = -2.41$; Late Gate: $\beta = -0.066$, $SE = 0.031$, $t = -2.13$), as presented in panels A and B of Figure 7. The

Early Gate effect faded at the post-head adverb but Late Gate persisted to be significant (Late Gate: $\beta = -0.130$, $SE = 0.032$, $t = -4.08$) (panel C). The same pattern was found in the next word, the matrix verb (Late Gate: $\beta = -0.187$, $SE = 0.035$, $t = -5.38$) (panel D). Since the Early Gate assessed whether participants would continue the sentence with any relative clause, it is not surprising that it was only sensitive to the reading times up to the point of the critical head noun. The Late Gate assessed the specific rates of relative clause completions with matching head noun animacy and structure configuration of the item in the comprehension task and therefore was strongly predictive of reading difficulty at the head noun and the following regions.

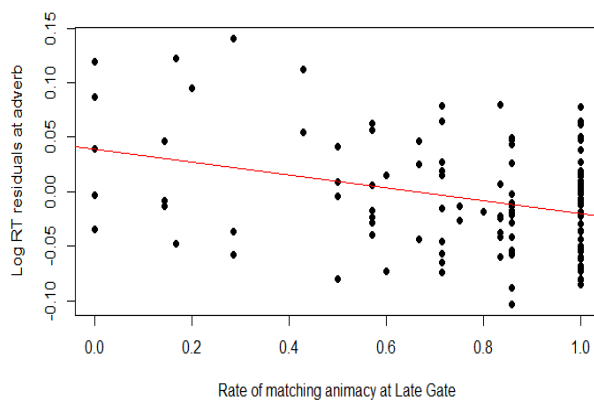
(A)



(B)



(C)



(D)




Figure 7. The four panels show significant regression results of log residual reading times at various word positions being predicted by rate of completion at either the Early Gate (on the rate of sentence frame completed as a relative clause) or Late Gate (rate of relative clause completions with the matching head noun) for each item. The following explains each graph with the format of the word position of the reading times and the gate of the completions. (A): head noun, Early Gate, (B): head noun, Late Gate, (C): adverb, Late Gate, (D): main verb, Late Gate.

Discussion

The results of the comprehension study and the links to sentence completions strongly argue for comprehenders' sensitivity to fine grained information about the lexico-syntactic properties of relative clauses in Mandarin. Beginning at the head noun, the completion data reliably predicted comprehenders' reading times. This head noun position is notable because it is the first position where effects of expectations for head animacy can be tested. Whereas the interaction of structure and animacy was only marginal in the reading experiment at the head noun, the Early and Late Gate completions were reliable predictors, suggesting that comprehenders are guided by finer-grained information relevant to individual sentences.

Importantly, the reading time data show no evidence for an effect of dependency distance (Gibson, 1998; 2000). Passive relatives have a longer dependency distance than the object relatives and thus should be more difficult than the object relatives, but the passives were read at the same rate or more quickly than the object relatives. Thus we see a clear effect of experience on comprehension, as shown by the relationship between the completions and reading times, but dependency distance does not appear to add additional explanatory value to these results.

General Discussion

We examined the link between production and comprehension of relative clauses in Mandarin, a language with unique word order and controversial previous comprehension results.

We approached the topic with the attempt to evaluate two schools of processing accounts: one explains comprehension difficulty through syntactic complexity which induces demands on memory maintenance (e.g. Gibson, 1998; Hsiao & Gibson 2003), and the other is an experience-based account that relates comprehension to learning through language statistics, which Gennari and MacDonald (2008, 2009) trace to production choices that stem from the nature of grammatical encoding in language production. Our results from a picture description task, corpus analysis, a sentence completion task, and a self-paced reading task show a clear linkage between what speakers produce and what comprehenders expect and interpret in comprehension, modulated by the interaction between structure and the lexico-semantic factor of animacy, instead of an effect of syntactic complexity. These results have implications for theories of both production and comprehension.

Implications for Production

Experiment 1 showed clear effects of head noun animacy on speakers' choice of relative clause form, replicating results using a similar task and materials in several typologically diverse languages with head-first relative clauses (Gennari et al., 2012; Montag & MacDonald, 2014; in press). These results argue for simultaneous planning of the head noun and relative clause despite the fact that the head noun comes later in the utterance, and thus they argue against extreme forms of incrementality in which words are uttered very rapidly after they are planned. Theories that support the incremental production process base their argument on the non-serial nature of utterance planning (Kempen & Hoenkamp, 1987; van Nice & Dietrich, 2003; de Smedt, 1996). Instead of fully specifying the conceptual message before forming syntactic structure and then onto articulation, a speaker does not only start producing a sentence after having worked out how to say the entire sentence. Rather, the production process is incremental or piecemeal,

meaning that the surface form is being uttered while conceptualization and syntactic formulation are still ongoing. This language production model has another implication, namely the reduction of processing resources. The easy parts get uttered while the processor deals with the harder portions. Incrementality then facilitates fluent production by dispersing the difficulty to other parts of the sentence, as demonstrated by the case of passivization in English (e.g. Bock, 1986). The more radical view of this incrementality account is held in such a way that “the first conceptualized referent will continue onward as the first-lexicalized and, ultimately, as the first in word order” (van Nice & Dietrich, 2003). Our results, however, do not support such theories. In our picture description task, the modified head noun had to be identified in the first place in order for the participants to formulate a modifying relative clause, but due to the head-final structure in Mandarin, the head noun was last in the word order, not the first. The animacy effect of the head noun on structure choice is clear evidence against radical incrementality. Moreover, the fact that object and passive relatives in Mandarin have identical word order and passivization requires mere addition of the high-frequency morpheme BEI argues against the logic of radical incrementality because even though uttering the very accessible passive marker can unconditionally lengthen the planning time of later elements, people chose not to do so all the time. The fact that participants produced passives, and some omitted agents, as a function of head noun animacy suggests that both the nouns in relative clause and the head noun were planned in advance.

Perhaps because both the relative clause and head noun are planned together, the Experiment 1 results also suggest that similarity-based interference affects speakers’ implicit structure choice. Gennari et al. (2012) argued that passive relative clauses not only allowed animate head nouns to be in subject position, but they also allowed speakers to reduce memory

interference from two semantically similar entities (an animate agent and patient). That is, accessibility of the head noun predicts more passives with animate than inanimate heads, but the Reduce Interference bias (MacDonald, 2013) predicts specifically that agentless passives will be more common when the sentence nouns are semantically similar, because passives are a vehicle for agent omission. Agents can be omitted in passive relative clauses but not object relative clauses in Mandarin, and we saw, consistent with results in English and Spanish (Gennari et al., 2012; Montag & MacDonald, 2014; in press), that agentless (bare) passives were more common with animate than inanimate heads. These results are also consistent with similarity-based interference and agent-omission findings in simple main clauses in Mandarin (Hsiao et al., 2014). Previous accounts of agent omission in Mandarin have suggested that it has simply a discourse origin, where agents are omitted when agent information is given, unimportant or unknown (e.g. Shei, 2014). Discourse factors are of course important in production choices, but our results show that they are not the whole story. In our task, agents were never mentioned in the discourse and were always pictured, and yet they were omitted more when agents acted on animate than inanimate entities.

In sum, these results suggest that several factors conspire to promote different sentence choices. The similarity-based interference effects are particularly interesting in that they suggest that in at least some cases, planning depends not simply on a single accessible noun (the head noun) but on interactions between sentence participants. This work echoes other findings suggesting that speakers are weighing several factors (Stallings & MacDonald, 2011), and they argue for a fairly broad scope of planning. There is no evidence that producers avoid longer dependencies (the passive relative) in favor of the shorter dependency object relatives, in contrast

to one recent hypothesis that dependency distance affects production difficulty (Scontras, Badecker, Shank, Lim & Fedorenko, 2015).

Implications for Comprehension-Production Linkages

Whereas many studies have compared comprehension of object relative clauses and subject relative clauses (2a-b), that comparison confounds sentence meaning and sentence structure (Gennari & MacDonald, 2008; MacDonald, Montag & Gennari, in press). We have instead compared object relatives and the meaning-matched passive relatives (a form of subject relative). In English, this object vs. passive relative comparison yields different word orders, but in Mandarin, the two sentence types are identical except for the passive marker BEI. The result is a comparison that is well matched in both word order and meaning, and the results are very clear. Passives are longer in dependency distance than object relatives, yet they are easier to comprehend for animate heads, and no more difficult than object relatives for inanimate heads. This result is not consistent with dependency distance. Instead, we see these results as most straightforwardly following from the Production-Distribution-Comprehension Account (Gennari & MacDonald, 2009; MacDonald, 2013): production processes favoring a) accessible animate nouns in prominent grammatical positions and b) utterance forms that reduce similarity-based interference conspire to yield more passive relatives for animate heads than inanimate heads. The result is a robust asymmetry in the distributional regularities of the language; perceivers learn the statistics of these distributions and interpret subsequent input with expectations that this input will conform to past experience. Comprehension difficulty, as measured by reading times, is low when the input matches these expectations and higher when it does not.

Distributional statistics provide basis for comprehenders' expectations. However, the frequency space, say in different modalities or genre, may be different. We saw from the

relatively different distributional patterns of active vs. passive relatives in the picture description task and the corpus analysis, suggesting difference of structural usage between speech and text. As Roland et al. (2007) point out, producers or comprehenders do not rely on difference in probabilistic grammars for written and spoken language but rather the difference comes from the difference in language use in different discourse context. To be more exact, the distribution discrepancy between our spoken production and written corpus data can be accounted for with the difference in what people normally do in those discourse situations to achieve their communicative goal. Language users are sensitive to the specific lexical and structural use paired with the certain context. Our corpus analysis supports this argument by showing that although written language usually has higher rate of passives overall, the specific newswire genre that we examined in fact typically had few of them. We also saw the evidence through completions and comprehension data that shared identical content and discourse context that people track statistics closely in a fine-grained constraints-satisfaction manner.

Relationship to Other Accounts: Ambiguity, Mixed Memory-Experience Accounts, Surprisal

Many previous Mandarin relative clause comprehension studies have compared object relative and active subject relatives, which also have a longer dependency distance than object relatives, yielding the prediction that subject relatives would be harder. Several studies have found an advantage for object relative clauses in Mandarin over active subject relatives (e.g. Gibson & Wu, 2013; Hsiao & Gibson, 2003; Li & Garnsey, 2011; Su et al., 2007), consistent with dependency distance accounts. Other studies, however, have found the opposite result—shorter reading times for the active subject relatives than the object relatives (e.g. Lin & Bever, 2006; Vasishth et al., 2013). The lack of consensus in the Mandarin literature was examined by Hsiao & MacDonald (2013) in a computational model that combined multiple structural and

semantic factors important to Mandarin, guided by corpus statistics. They found that the head-final feature of Mandarin interacted with the dominant SVO main clause word order and noun animacy regularities in the language, such that object relatives were easier when they modified the main clause subject but were more difficult when they modified the main clause object. Moreover, animate nouns were preferred as subjects in general (in both simple sentences and relative clauses) and inanimate nouns were preferred as objects, consistent with the lexico-syntactic statistics of Mandarin and other languages. The computational network was able to produce learning results that corresponded to the patterns of prior empirical studies with matching manipulations, without a built-in function of working memory. The results also suggest that the inconsistency across studies is due to the complex combination of relevant factors, which are too numerous to include in a single human experiment but can be examined simultaneously in tests of the computational model. The present studies complement the modeling work by investigating additional conditions, passive relatives with animate and inanimate heads. Our results are compatible with the modeling results in confirming extensive ambiguities in the materials and strong animacy effects in relative clause processing.

While this study is targeted to directly address the Dependency Locality Theory, one of the memory accounts of sentence comprehension, the results can offer potential insight for other memory-based theories. With storage and integration cost being the focus of the Dependency Locality Theory, memory decay is considered the major force behind comprehension difficulty. More intervening information between the filler and gap, hence more time lapses in between, makes it harder to retrieve the target. This type of theories naturally entails that memory capacity predicts how much information can be held actively, which in turn determines the difficulty of retrieval of the right information. Another type of memory accounts emphasizes

the effect of memory interference. Studies have shown that the low working memory capacity is related to inability to attend to goals and maintain on-task thought (Engle, 2002; McVay & Kane, 2011), which are correlated with comprehension deficits. Recent or similar items encountered between the antecedent and the retrieval site can cause comprehension difficulty. Such theories consider the similarity of lexical and contextual representation of the intervening information and the recall target (Gordon et al., 2002; Lewis & Vasishth, 2005; Van Dyke & MacElree, 2006). If the similarity is low, the memory representation/retrieval cue of the target is unique enough from distractors so the comprehension of the dependency relationship should be more effective. In terms of our self-paced reading results, the difficulty of animate-headed active relative clauses may have reflected the effect of such similarity-based retrieval interference, but the comparatively low reading times of both animate- and inanimate headed passives do not seem to support such account.

To this point we have noted that results are not compatible with the Dependency distance account (and thus also not compatible with an account combining memory limitation effects and experience-based effects), but this work also makes a broader point about the applicability of the account to relative clauses. As it is well known that ambiguity affects processing difficulty and that ambiguity resolution processes are strongly affected by experience, it is unclear how to investigate the effect of dependency distance in sentences that also have ambiguities (MacDonald, 2013). Recently Jäger et al. (2015) have made this argument for Mandarin relative clauses, specifically that it is impossible to contrast experience and inherent limitation (dependency distance) accounts with the Mandarin relative clause materials that have been used to date, because they have too much ambiguity. They embedded relative clauses in other clauses in a way that made them nearly free of ambiguities and found evidence for experience-driven

processing. Their observations about the effects of ambiguity and arguments for experience-based processing are in many respects compatible with our own approach.

Jäger et al.'s (2015) perspective differs from ours in a crucial respect that is instructive in distinguishing our own approach from other experience-based approaches, particularly Surprisal (Hale, 2001; Levy, 2008). Surprisal is the degree to which a word is unexpected in a sentence, but the relevant grain over which expectations are developed can take many forms. Jäger et al. view experience in terms of what Demberg & Keller (2008) call “unlexicalized surprisal”, calculating expectations over grammatical categories rather than individual words, so that expectations here essentially are for a particular sentence structure. By contrast, we have emphasized a much finer grain of experience-based expectations, in that we have shown how head noun animacy, and indeed the lexical and discourse properties of individual sentences, affect comprehenders' expectations (as measured by completions) and reading times. Consistent with this view, Gennari et al. (2008) regressed different aspects of their English completion data on comprehenders' reading times to investigate the grains at which expectations appeared to be forming. They found that expectations for a particular thematic role (which incorporates syntactic and lexical properties of nouns and verbs) was a good predictor of reading times, as were predictions for the exact word in some regions, but sentence structure did not account for additional variance.

A second way in which our account differs from Jäger et al.'s (2015) is in the calculation of the relevant experiences generating expectations. The option that Jäger et al. adopt, which is also adopted in other surprisal work (see Frank, 2009, for discussion), is a straightforward assessment of structure frequency, with more frequent structures (that is, frequent sequences of grammatical categories) being more expected than lower frequency structures. By contrast, we

have emphasized that the “neighborhoods” and sub-regularities in sentence structures affect processing difficulty (Fitz, Chang, & Christiansen, 2011; Hsiao & MacDonald, 2013; MacDonald & Christiansen, 2002), so that the difficulty of a structure, such as a relative clause, is dependent not only on the frequency of that structure, but also on the extent to which it shares properties with other structures in the language. Computational models, such as the simple recurrent networks (SRNs) used by Fitz et al., MacDonald and Christiansen, and Hsiao and MacDonald, are an important component in investigating this hypothesis, because the model, trained on a range of sentence types, learns the neighborhood structure and other regularities, so that processing difficulty of a given sentence depends on more than simply the frequency of that sentence type in the input. For example, some more frequent structures yield higher error rates in Hsiao & MacDonald’s (2013) model (and longer reading times for human comprehenders) than some rarer structures (e.g. subject-modifying subject relative clauses are more frequent than object relatives in the corpus but still harder to process). Whereas previous researchers might have used evidence of this frequency-difficulty asymmetry as evidence against experience accounts, our view is that it is evidence only against simpler versions of those approaches, and evidence for neighborhood effects in sentence comprehension. Indeed, in his comparison of surprisal and SRNs in accounting for comprehension difficulty, Frank (2009) argued that though unlexicalized surprisal provides a better language model (that is, it is a more veridical calculation of the frequency of structures in a corpus), the SRNs better account for patterns of human comprehension difficulty: both humans and the models are affected by complex neighborhoods and subregularities. Because these neighborhood effects also affect production as well as comprehension (Fitz et al., 2011), we expect that this approach will be important in continuing to explore comprehension-production links in Mandarin and other languages.

References

- Acheson, D. J., & MacDonald, M. C. (2009). Verbal working memory and language production: Common approaches to the serial ordering of verbal information. *Psychological Bulletin*, 135(1), 50-68. doi:10.1037/a0014411
- Acheson, D. J., & MacDonald, M. C. (2011). The rhymes that the reader perused confused the meaning: Phonological effects during on-line sentence comprehension. *Journal of Memory and Language*, 65, 193-207.
- Allum, P. H., & Wheeldon, L. R. (2007). Planning scope in spoken sentence production: The role of grammatical units. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33(4), 791-810. doi:10.1037/0278-7393.33.4.791
- Almor, A. & Nair, V. A. (2007). The form of referential expressions in discourse. *Language and Linguistics Compass*, 1, 84–99.
- Anderson, J. R. (1983). Retrieval of Information from Long-Term Memory. *Science, New Series*, 220, 25-30
- Arnold, J. E. (2010). How speakers refer: The role of accessibility. *Language and Linguistics Compass*, 4, 187–203. doi:10.1111/j.1749-818X.2010.00193.x
- Arnold, J. E., Eisenband, J. G., Brown-Schmidt, S., & Trueswell, J. C. (2000). The rapid use of gender information: Evidence of the time course for pronoun resolution from eyetracking. *Cognition*, 76, B13-B26. doi:10.1016/S0010-0277(00)00073-1

- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255-278.
- Biberauer, T., Holmberg, A., Roberts, I., & Sheehan, M. (2009). *Parametric variation: Null subjects in minimalist theory*. Cambridge University Press.
- Bock, J. (1986). Meaning, sound, and syntax: Lexical priming in sentence production. *Journal Of Experimental Psychology: Learning, Memory, and Cognition*, 12(4), 575-586.
doi:10.1037/0278-7393.12.4.575
- Bock, J., & Warren, R. K. (1985). Conceptual accessibility and syntactic structure in sentence formulation. *Cognition*, 21(1), 47-67. doi:10.1016/0010-0277(85)90023-X
- Bock, K., Loebell, H., & Morey, R. (1992). From conceptual roles to structural relations: Bridging the syntactic cleft. *Psychological Review*, 99(1), 150-171. doi:10.1037/0033-295X.99.1.150
- Camacho, José. 2013a. *Null subjects*. Cambridge: Cambridge University Press.
- Christianson, K., & Ferreira, F. (2005). Conceptual accessibility and sentence production in a free word order language (Odawa). *Cognition*, 98, 105-135.
- Damian, M. F., & Bowers, J. S. (2003). Locus of semantic interference in picture-word interference tasks. *Psychonomic Bulletin & Review*, 10(1), 111-117.
doi:10.3758/BF03196474

de Smedt, K. (1996). Computational models of incremental grammatical encoding.

Computational psycholinguistics: AI and connectionist models of human language processing, 279-307. Philadelphia, PA US: Taylor & Francis.

Dell, G. S., & Reich, P. A. (1981). Stages in sentence production: An analysis of speech error

data. *Journal of verbal learning and verbal behavior*, 20(6), 611-629

Ferreira, F., & Swets, B. (2002). How incremental is language production? Evidence from the production of utterances requiring the computation of arithmetic sums. *Journal of Memory and Language*, 46, 57-84. doi:10.1006/jmla.2001.2797

Fukumura, K., & van Gompel, R. G. (2011). The effect of animacy on the choice of referring expression. *Language and Cognitive Processes*, 26(10), 1472-1504.

doi:10.1080/01690965.2010.506444

Fukumura, K., van Gompel, R. G., Harley, T., & Pickering, M. J. (2011). How does similarity-based interference affect the choice of referring expression? *Journal of Memory and*

Language, 65(3), 331-344. doi:10.1016/j.jml.2011.06.001

Gennari, S. P., Mirković, J., & MacDonald, M. C. (2012). Animacy and competition in relative

clause production: A cross-linguistic investigation. *Cognitive Psychology*, 65(2), 141-176.

doi:10.1016/j.cogpsych.2012.03.002

Gibson, E. (1998). Linguistic complexity: Locality of syntactic dependencies. *Cognition*, 68, 1-

76.

- Gibson, E. (2000). The dependency locality theory: A distance-based theory of linguistic complexity. In Miyashita, Y., Marantz, A., & O'Neil, W. (Eds.), *Image, Language, Brain* (pp. 95-126), Cambridge, MA: MIT Press.
- Glaser, W. R., & Döngelhoff, F. (1984). The time course of picture-word interference. *Journal of Experimental Psychology: Human Perception and Performance*, 10(5), 640-654.
doi:10.1037/0096-1523.10.5.640
- Gordon, P. C., Grosz, B. J., & Gilliom, L. A. (1993). Pronouns, names, and the centering of attention in discourse. *Cognitive Science*, 17(3), 311-347.
- Gordon, P. C., & Hendrick, R. (1997). Intuitive knowledge of linguistic co-reference. *Cognition*, 62(3), 325-370.
- Hsiao, Y. and MacDonald, M (2013). Experience and generalization in a connectionist model of Mandarin Chinese relative clause processing. *Frontiers in Psychology*, 2013(4):767. doi: 10.3389/fpsyg.2013.00767
- Jäger, L., Chen, Z., Li, Q., Lin, C.-J. C., & Vasishth, S. (2015). The subject-relative advantage in Chinese: Evidence for expectation-based processing. *Journal of Memory and Language*, 79-80, 97–120. <http://doi.org/10.1016/j.jml.2014.10.005>
- Jaeger, TF (2013). Production preferences cannot be understood without reference to communication. *Frontiers in Language Sciences*, 230. doi:10.3389/fpsyg.2013.00230
- Jaeger TF, Furth K and Hilliard C (2012). Incremental phonological encoding during unscripted sentence production. *Frontiers in Psychology* 3:481. doi: 10.3389/fpsyg.2012.00481

- Janssen, N., and Caramazza, A. (2009). Grammatical and phonological influences on word order. *Psychological Science*, 20, 1262–1268.
- Kempen, G., & Hoenkamp, E. (1987). An incremental procedural grammar for sentence formulation. *Cognitive Science: A Multidisciplinary Journal*, 11(2), 201-258.
doi:10.1207/s15516709cog1102_5
- Kennison, S. M., & Gordon, P. C. (1997). Comprehending referential expressions during reading: Evidence from eye tracking. *Discourse Processes*, 24(2-3), 229-252.
- Kurumada, C. & Jaeger, T.F. (2013). Communicatively efficient language production and case-marker omission in Japanese. *Proceedings of the 36th Annual Meeting of the Cognitive Science Society*
- Ledoux, K., Camblin, C. C., Swaab, T. Y., & Gordon, P. C. (2006). Reading Words in Discourse: The Modulation of Lexical Priming Effects by Message-Level Context. *Behavioral and Cognitive Neuroscience Reviews*, 5(3), 107–127. doi:10.1177/1534582306289573
- Levelt, W. M. (1989). *Speaking: From intention to articulation*. Cambridge, MA US: The MIT Press.
- Li, C.N. & Thompson, S.A. (1981). *Mandarin Chinese: A Functional Reference Grammar*. Berkeley: University of California Press.
- MacDonald, M. C. (2013a). How language production shapes language form and comprehension. *Frontiers in Psychology*, 4, doi:10.3389/fpsyg.2013.00226
- MacDonald, M. C. (2013b). Production is at the left edge of the PDC but still central: response to commentaries. *Frontiers in Psychology*, 4, doi:10.3389/fpsyg.2013.00227

- Meyer, A. S., & Schriefers, H. (1991). Phonological facilitation in picture-word interference experiments: Effects of stimulus onset asynchrony and types of interfering stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17(6), 1146.
- Montag, J. L., & MacDonald, M. C. (2014). Visual salience modulates sentence choice in relative clause production. *Language and Speech*, 57, 163-180.
- Neely, J. H. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. *Basic processes in reading: Visual word recognition*, 11, 264-336.
- Rosinski, R. R., Golinkoff, R. M., & Kukish, K. S. (1975). Automatic semantic processing in a picture-word interference task. *Child Development*, 46(1), 247-253. doi:10.2307/1128859
- Schriefers, H., Meyer, A. S., & Levelt, W. J. (1990). Exploring the time course of lexical access in language production: Picture-word interference studies. *Journal of memory and language*, 29(1), 86-102.
- Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of experimental psychology*, 18(6), 643.
- Tanaka, M., Branigan, H. P., & Pickering, M. J. (2005). *The role of animacy in Japanese sentence production*. In CUNY Sentence Processing Conference.
- Van Dyke, J. A., Johns, C. L., & Kukona, A. (2014). Low working memory capacity is only spuriously related to poor reading comprehension. *Cognition*, 131(3), 373-403.

- Wagner, V., Jescheniak, J. D., & Schriefers, H. (2010). On the flexibility of grammatical advance planning during sentence production: Effects of cognitive load on multiple lexical access. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 36(2), 423-440. doi:10.1037/a0018619
- Wang, L. (1987). *Xian Dai Han Yu Yu Fa* [Modern Chinese Grammar]. Taiwan: Lan Deng Wen Hua
- Wang, Q. (2007). An Investigation on the Semantics and Discoursal & Socio-pragmatic Functions of Personal Pronouns in Mandarin Chinese (Master's thesis). Retrieved from <http://www.airitilibrary.com/Publication/alDetailedMesh1?DocID=U0021-0712200716141159#References>
- Wilshire, C. E. (1998). Serial order in phonological encoding: an exploration of the word onset effect using laboratory-induced errors. *Cognition*, 68(2), 143-166.
- Xue, N., Jiang, Z., Zhong, X., Palmler, M., Xia, F., Chiou, F., Chang., M. (2010). Chinese Treebank 7.0. Linguistic Data Consortium, Philadelphia.
- Yang, C.-L., Gordon, P. C., Hendrick, R., & Wu, J. T. (1999). Comprehension of referring expressions in Chinese. *Language and Cognitive Processes*, 14, 715–743.

Conclusions

The three studies in this dissertation examined how the comprehension behavior can be attributed to learning from language statistics, which for a large part can be traced back to the production biases and constraints. The findings altogether suggest that computational capacity limitation, which has long tradition in linguistics and psychology to be considered the cause of humans' inability to process complex sentence structure, is inseparable from experience. I implemented neural networks to simulate processing of Mandarin and English relative clauses, and found that the models generated prediction errors reflecting the corpus statistics that they were trained on. The performance was comparable to human reading time data although no working memory system was built in. I later used production tasks to demonstrate that distributional statistics were constituted by production choices given the specific discourse environment and other constraints and demands of the task, which significantly predicted comprehension patterns. Next, I summarize the findings of each study and discuss their implications on language comprehension and production theories.

In Chapter 1, I presented an SRN on Mandarin relative clause processing, a topic that has long been controversial because of the language's unique typological features and mixed results from human studies. Based on large-scale corpus analyses and hand coding of various sentential and lexical features, the model received training that resembled human language experience. The model performance was affected by the statistical patterns of these multiple factors occurring not only in the target relative clause structure but also in other neighboring structures, such as simple sentences with overt and covert subject nouns. The error patterns of the model was compared to human reading times obtained through a self-paced reading task and also other previous empirical comprehension studies. I found that when manipulated factors were matched with the

ones of the model, the GPE measures of the model showed matching patterns to the human data. The results of the study strongly argue for a constraint satisfaction account for human comprehension, which tracks closely the various sources of information. Instead of a clear-cut dichotomy of processing advantage of a subject or object relative clauses, Mandarin relative clause processing involves complex interactions of these factors that produce different results under different circumstances.

Chapter 2 discusses English relative clause processing using the same methods as in Chapter 1. The well-known advantage of subject relative clauses in human comprehension was replicated with the SRN but this advantage was modulated by animacy of the nouns involved in the relative clauses. As with the Mandarin SRN, the animacy effect showed in the form of agent roles preferring animate nouns and themes inanimate nouns, which was also evident in the corpus data. The cross-linguistic similarity of the animacy effect suggests that semantic factors that are based on world knowledge are likely to manifest as universal, regardless of typological differences across languages. The unfolding of ambiguity, however, was indeed influenced jointly by typological traits, such as word order. Given that English relative clauses are head-initial and Mandarin relative clauses are head-final, as well as the fact that both languages have SVO basic word order, the relative amount of ambiguity differs in the time course of relative clauses unfolding. Again, together with the Mandarin SRN, this English SRN argue for the role of language statistics on comprehension and the parser's ability to generalize predictions from related structures to the target structure.

Chapter 3 links the processes of comprehension to those of production through a series of tasks: on-line spoken production, off-line written production, corpus analyses and self-paced reading task. The picture description task extended the paradigm of Gennari, Mirkovic &

MacDonald (2012) to the language of Mandarin Chinese. Despite the differences in word order, Mandarin was like the many head-initial languages that they investigated and exhibited higher rate of passive relatives when the modified nouns were animate than when they were inanimate. However, the proportion of passives was relatively high compared to other languages, potentially affected by the morphosyntax of passivization in Mandarin, which requires only a simple, high frequency morpheme BEI. Passives in the written corpus were instead quite rare, suggesting that the distributional statistics of passives may vary depending on modality (speech/ written) or specific task demands. The later sentence completion and self-paced reading studies were based on the same materials, reliably showing the relationship of rate of the targeted type of completions and reading times. The results offer support for the Production-Distribution-Comprehension account.

Next, I discuss the implications and future directions of these studies.

Implications on Comprehension Theories

The theme that weaves the three studies together is the reliable tracking of complex statistics of the comprehension system. On top of the biological architecture of the memory system, comprehension difficulty arises with differences in language experience. An important affirmation of the previous literature through this dissertation is the constraints-satisfaction nature of comprehension. Multiple factors contribute to how predictions of the upcoming elements are formed, for example, structural frequency, lexical-semantic property, discourse context, and so on. This dissertation offers an example through computational modeling to examine how these factors, usually investigated separately in human studies, interact. Although language experience has now been well established as a variable that needs to be considered in

sentence processing, my approach distinguishes from other theories that incorporate the role of experience. For example, Levy et al. (Levy et al., 2012, Levy & Gibson, 2013) advocate a mixed model of working memory and experience. They state that comprehenders maintain expectations based on both dependency locality and probabilistic statistics. As seen in Chapter 1 and 2, I demonstrated that dependency relationship was learned by the models through exposure, instead of any built-in knowledge. Also, Levy's experience-based account, which is termed Surprisal, differs from what is entailed by the learning results of my models. In both SRNs, prediction errors of relative clauses were influenced by other neighboring structures, in addition to other non-syntactic properties like animacy. Although the Surprisal measure is not limited to syntax, the calculation of conditional probability does not incorporate generalization from other structures.

However, like Levy and colleagues, I do argue for the importance of cross-linguistic data that they demonstrate statistical differences in grammatical properties of different languages, and that they help disentangle processing theories. For example, the syntax-centered theories that predict a universal structural asymmetry, namely subject relative clause advantage, fails to predict the complicated processing pattern in Mandarin, as shown in Chapter 1 through humans and machine. A future direction of the current approach of the dissertation is to simulate languages varying on more grammatical dimensions other than word order and argument dropping (in the Mandarin model, we included null subject sentences), such as morphological richness. For example, Japanese and Korean are languages with head-final relative clauses, like Mandarin, but they differ from Mandarin and English in having SOV dominant word order. One other feature of the languages is their case marking system, resulting in frequent omission of elements in the sentence when the discourse context allows. This combination of syntactic

features may likely result in different effect of generalization from simple sentences to relative clauses.

Implications on Production Theories

Like comprehension, production is affected by multiple factors. Some of these factors contribute differently across languages, whereas some others provide a rather universal influence. In Chapter 3, I show that the passive relative clauses were produced more often when the modified entities were animate than when they were inanimate. This same pattern was found also in other languages, such as English, Spanish, Serbian (Gennari et al., 2012), Japanese and Korean (Montag, 2014), although these languages differ in many ways grammatically. However, it appears that the degree of the effect and the exact proportion of passive relatives were affected by how frequent passives in main clauses were used (Montag, 2009). Passives in Korean, for example, are very rare in general. Also, Montag (2015) found that English-speaking children at 8 and 12 years of age produced different proportions of passive relatives. The difference was attributed to the different exposure to text containing the passive structures. Children's books targeting older children had more passive sentences than those for younger children, and this difference was reflected in their spoken production. Therefore, when speakers lack experience in practicing some structure, they find it difficult to execute. This finding echoes my modeling results in Chapter 1 and 2 regarding generalizing statistics from neighboring structures and also the overall conclusions made that language users produce and comprehend according to long-term learning of statistics.

Common for speakers across languages is the strategizing of reduction of production difficulty and maximization of fluency. It is often manifested as producing the easier bits first. In

Chapter 3, for example, the picture description task generated high proportion of passive relatives, compared to other languages mentioned above. One explanation is that the participants indiscriminatorily added the passive marker BEI to the beginning of every response, given that the word orders of actives and passives were the same and adding a simple morpheme to initiate the sentence was cognitively cheap.

My other work published as Hsiao, Gao and MacDonald (2014) investigated null subject simple sentences in Mandarin. We found that participants strategically omitted subject nouns when the upcoming object nouns were of the same animacy, i.e. semantically similar. To reduce similarity-based interference, Mandarin speakers avoided producing the subject even when the competitor object noun was further downstream, suggesting that production involves strategic control of advance planning rather than being radically incremental (Ferreira & Swets, 2002; Wagner et al., 2010). This advance planning to avoid semantic competition was also seen in the description task in Chapter 3, where there were more dropping of the agent noun (occurs before the head noun due to the head-final word order) in passive relative clauses when both the agent and the patient were animate than when they were of different animacy. The result adds to the literature on the use of referential expressions in face of similarity interference (see Arnold, 2010 for a review).

Other implications on production include the issue of whether production choices are driven by the need of the comprehender, or the cognitive pressure on the speakers themselves. Some accounts maintain that production preferences are affected by communicative biases. This type of theories is under the general term of “audience design”: speakers produce their utterances based on the estimates of their interlocutors’ beliefs. For example, speakers take their interlocutors’ perspective (e.g. Heller, Gorman & Tanenhouse, 2012), they tend to avoid

ambiguity (e.g. Arnold, Wasow, Asudeh & Alrenga, 2004), and they try to transmit the information efficiently to the listeners by making it uniformly complex (e.g. Jaeger, 2010). There has been evidence claimed to support audience design in a way that speakers do not produce certain things in order to alleviate production difficulty. For example, if fillers such as “um” or “uh” are considered to function as a strategy to buy extra planning time, then not producing anything would be the most economical solution to reducing planning pressure. However, people rarely replace these fillers with pauses and fillers are usually followed by disfluencies (Clark & Fox Tree, 2002). Therefore, production choices should be understood from the perspective of how people achieve efficient communication, mostly through consideration of the comprehenders’ needs.

This dissertation investigates how speakers make choices to reduce production difficulty on their side. Although that no real interlocutors were present at the time of the picture description experiment in Chapter 3, participants were given a cover story about a later group expected to guess what the referred entity was relying on their responses. The fact that there was an effect of animacy on the production of actives/passives, and on the rate of agent dropping among passives, indicate that participants did make these choices with regard to the semantic properties of the lexical items and the flexibility that the two structures offered. Otherwise, the responses would have always been meticulously detailed and specific without any omission, to ensure perfect understanding on the listeners’ side. I argue that although audience design is an important part of what people choose to say and how to say it, because we usually speak to someone instead of talking to ourselves, the computational demands on the speakers is one of the crucial, if not the most crucial, driving force that determine speakers choices.

Implications on Connecting Production and Comprehension

While the production and comprehension are usually studied as two separate processes in psycholinguistics, this dissertation contributes to the understanding of how they are connected. Chapter 1 and 2 provide evidence that comprehension is the reflection of learning from statistical distributions of the language and Chapter 3 links language distributions to production choices. This chain of reasoning from production to distribution to comprehension (PDC) (MacDonald, 2013) offers a unique perspective on language processing theories. Different from a typical frequency-based account of comprehension, PDC provides a way to examine how these statistics originate. In addition, as language users are both producers and comprehenders, production choices are also affected by what they hear and read. The more frequently exposed sentence types are usually better rehearsed and therefore easier to produce.

This argument is related to whether language as a cognitive ability that defines human beings is innate. The knowledge of grammatical rules is considered by some researchers too complex to have evolved naturally but can still be acquired by children with seemingly inadequate amount of exposure (e.g. Chomsky, 1965; Hauser, Chomsky & Fitch, 2002). Along with the PDC account, I argue that language processing is rather like other cognitive tasks that involve non-linguistic elements with regularities in sequencing linearly and hierarchically, for example, perceptual and motor learning tasks (Reber, 1967, 1988) and action planning (Allen et al., 2010). The SRNs in this dissertation also demonstrate that long dependency relations that characterize complex linguistic structures are learnable without an innate set of grammatical rules. It might not even be too different from our morning rituals of coffee or pancake making (Botvinick & Plaut, 2004).

The dissertation contributes to the understanding of the role of learning through language statistics as manifested in comprehension and further about how these distributional patterns are formed through production choices. I demonstrated that language users trace fine grained statistical regularities, and processing difficulty is modulated by experience with these statistical patterns, instead of a steady, inherent computational limitation driven by syntactic complexity. The distributional patterns of the language are largely determined by speakers' production choices, made under the pressure of planning and execution. Besides of goal of getting the message across, production is usually conducted in a way that minimizes effort and maximizes fluency, given what flexibility is offered by the language. These production choices thus become the statistics that comprehenders based on when they hear and read. Through studying production and comprehension jointly, we can gain a more complete picture about both processes and what people bring to bear when they carry out these two tasks.

References

- Allen, K., Ibara, S., Seymour, A., Cordova, N., & Botvinick, M. (2010). Abstract Structural Representations of Goal-Directed Behavior. *Psychological Science, 21*(10), 1518–1524.
- Arnold, J. E. (2010). How speakers refer: The role of accessibility. *Language and Linguistics Compass, 4*, 187–203. doi:10.1111/j.1749-818X.2010.00193.x
- Arnold, J.E., Wasow, T., Asudeh, A., & Alrenga, P. (2004). Avoiding attachment ambiguities: The role of constituent ordering. *Journal of Memory and Language, 51*, 55-70.
- Botvinick, M., & Plaut, D. C. (2004). Doing Without Schema Hierarchies: A Recurrent Connectionist Approach to Normal and Impaired Routine Sequential Action. *Psychological Review, 111*(2), 395–429.
- Chomsky, N. (1965). *Aspects of the Theory of Syntax*. Cambridge, Mass.: MIT Press.
- Clark, H. H., & Fox Tree, J. E. (2002) Using uh and um in spontaneous speaking. *Cognition, 84*, 73-111.
- Ferreira, F., & Swets, B. (2002). How incremental is language production? Evidence from the production of utterances requiring the computation of arithmetic sums. *Journal of Memory and Language, 46*, 57–84. doi:10.1006/jmla.2001.2797
- Gennari, S. P., Mirković, J., & MacDonald, M. C. (2012). Animacy and competition in relative clause production: A cross-linguistic investigation. *Cognitive Psychology, 65*(2), 141–176. doi:10.1016/j.cogpsych.2012.03.002
- Heller, D, Gorman, K. S., Tanenhaus, M. K. (2012). To name or to describe: Shared knowledge affects referential form. *Topics in Cognitive Science. 4*:290–305.

- Hauser, M., Chomsky, N., & Fitch, T. 2002. The faculty of language: what is it, who has it, and how did it evolve. *Science* 198. 1569-79.
- Hsiao, Y., Gao, Y., & MacDonald, M. C. (2014). Agent-patient similarity affects sentence structure in language production: evidence from subject omissions in Mandarin. *Frontiers in Psychology*, 5, 1015. doi:10.3389/fpsyg.2014.01015
- Jaeger, T. F. (2010). Redundancy and reduction: Speakers manage syntactic information density. *Cognitive Psychology*, 61, 23-62.
- Levy, R., Fedorenko, E., Breen, M., & Gibson, T. (2012). The processing of extraposed structures in English. *Cognition*, 12, 12–36.
- Levy, R., & Gibson, E. (2013). Surprisal, the PDC, and the primary locus of processing difficulty in relative clauses. *Frontiers in Language Sciences*, 4(229), 1–3.
- MacDonald, M. C. (2013). How language production shapes language form and comprehension. *Frontiers in Psychology*, 4. doi:10.3389/fpsyg.2013.00226
- Montag, J. L., & MacDonald, M. C. (2015). Text exposure predicts spoken production of complex sentences in eight and twelve year old children and adults. *Journal of Experimental Psychology: General*, 144, 447-468.
- Montag, J. L. & MacDonald, M. C. (2009). Word order doesn't matter: Relative clause production in English and Japanese. In N.A. Taatgen & H. van Rijn (Eds.), *Proceedings of the 31th Annual Meeting of the Cognitive Science Society* (pp. 2594-2599). *Cognitive Science Society*.

Reber, A. S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning and Verbal Behaviour*, 6, 855-863.