Generative Linguistics in Wrocław No. 7

Veranika Puhacheuskaya

THE ROLE OF SYNTACTIC AND SEMANTIC CONSTRAINTS IN RELATIVE CLAUSE ATTACHMENT PROCESSING IN RUSSIAN: AN EYE-TRACKING STUDY

TRACKIN

Veranika Puhacheuskaya

The Role of Syntactic and Semantic Constraints in Relative Clause Attachment Processing in Russian: An Eye-Tracking Study

Center for General and Comparative Linguistics
Pracownia Językoznawstwa Ogólnego i Porównawczego

Generative Linguistics in Wrocław (GLiW) series is published by the Center for General and Comparative Linguistics, a unit of the Institute of English Studies at the University of Wrocław.

Address:
Center for General and Comparative Linguistics (CGCL)
ul. Kuźnicza 22
50-138 Wrocław, Poland

Editor of the series — Joanna Błaszczak

Layout and copy-editing — Joanna Błaszczak
Proof-reading — Markus Saint-Pettersen
Cover design — Marcin Orszulak & Markus Saint-Pettersen

ISSN 2084-5723

ISBN 978-83-932687-5-7
Published: December 2020
Publisher: CGCL

About the Series

Generative Linguistics in Wrocław (GLiW) is meant to provide a suitable forum for the presentation and discussion of the Polish research within the field of generative linguistics. We are interested in studies that employ generative methodology to the synchronic or diachronic analysis of phonology, semantics, morphology, and syntax. Apart from that, we express a keen interest in interdisciplinary research that is based on typology, diachrony, and especially experimental methods taken from psycho- or neurolinguistics and applied so as to provide a psycholinguistic reality to purely theoretical research. We believe that the dissemination of ideas is fundamental to any scientific advancement and thus our choice is to publish research studies in the form of e-books, which are available for free on our website.

Joanna Błaszczak
on behalf of the Editorial Board

TABLE OF CONTENTS

1.1. General background 3
1.1.1. Parsing strategies 4
1.1.2. Parsing stages 6
1.1.3. Syntax-lexicon dissociation 6
1.2. Major parsing models 9
1.2.1. Bever's heuristics (1970) 10
1.2.2. Kimball's two-stage model (1973) 11
1.2.3. Frazier's Garden-Path Model (1979; 1987) 12
1.2.4. Frazier and Clifton's Construal Hypothesis (1996) 14
1.2.5. Constraint Satisfaction models 16
1.2.6. Barton and Sanford's "Good Enough" Model (1993) 17
1.2.7. Van Gompell et al's Unrestricted Race Model (2005) 18
1.3 Chapter summary 19
2.1. General background 20
2.2. Russian data 25
2.3. Chapter summary 30
3.1. Research questions 31
3.2. Method 31
3.3. Subjects 34
3.4. Materials and design 35
3.5. Pre-Tests 38
3.5.1. Frequency 38
3.5.2. Plausibility norming 39
3.6. Differences from the previous studies 40
3.7. General predictions 41

3.7.1. Group 1: Syntactic Disambiguation 41
3.7.2. Group 2: Semantic Disambiguation 44
3.8. Apparatus 47
3.9. Procedure 48
3.10. Results 48
3.10.1. Group 1: Syntactic Disambiguation 53
3.10.2. Group 2: Semantic disambiguation 57
3.10.3. Ambiguous condition 61
3.11. Discussion 63

LIST OF ABBREVIATIONS

ACC	Accusative case
ANOVA	Analysis of variance
AmbA	Ambiguous attachment
F	Feminine gender
GEN	Genitive
EA	Early attachment
ERP	Event-related potentials
IA	Interest Area
LA	Late attachment
M	Masculine gender
NVN	Noun-verb-noun strategy
NOM	Nominative case
NP	Noun phrase
PPP	Preliminary Phrase Packager
PREP	Prepositional (locative) case
rANOVA	Repeated measures ANOVA
RC	Relative clause
RP	Relative pronoun
SD	Standard deviation
SPRT	Self-paced reading task
SemEA	Semantic early attachment
SemLA	Semantic late attachment
SynEA	Syntactic early attachment
SynLA	Syntactic late attachment
SSS	Sentence Structure Supervisor

LIST OF TABLES

Table 1. Parsing stages assumed by the Garden-Path Model 12

Table 2. Groups of experimental sentences 38

Table 3. Interest areas 41

Table 4. Predictions by different parsing models 47

Table 5. Accuracy rate 49

Table 6. Attachment preferences by participants' country of origin 51

Table 7. Mean first-pass time (in milliseconds) 52

Table 8. Mean second-pass time (in milliseconds) 52

Table 9. Mean dwell time (in milliseconds) 52

Table 10. Mean number of regressions to areas 52

Table 11. Mean number of regressions out of areas 53

Table 12. rANOVA results for SynEA, SynLA, AmbA 54

Table 13. ANOVA results for SemEA, SemLA, AmbA 58

LIST OF FIGURES

Figure 1. Possible relative clause attachment sites (Frazier and Clifton, 1997, p. 281) 15

Figure 2-3. One-way ANOVA for verb frequencies 39

Figures 4-5. Welch's two-sample *t*-test comparing plausibility of sentences with relative clauses attached to either NP1 or NP2 40

Figure 6. Predictions of the Garden-Path Model 42

Figure 7. Accuracy rate 49

Figures 8-9. Second-pass and dwell time on the relative pronoun 55

Figures 10-11. Regressions to the relative pronoun and dwell time on the verb 56

Figures 12-13. Regressions to NPs by condition 56

Figures 14-15. Compared regressions to NP1 and NP2 57

Figures 16-17. Second-pass and dwell time on the relative pronoun 59

Figure 18. Regressions out of the verb 60

Figures 19-20. Regressions to NPs by condition 60

Figures 21-22. Compared regressions to NP1 and NP2 61

Figures 25-26. Dwell time on the relative clause in the Ambiguous condition depending on the chosen type of closure 63

ABSTRACT

An eye-tracking experiment was conducted to examine relative clause attachment processing in Russian speakers and how it is affected by syntactic and semantic constraints. Attachment was manipulated by either morphological means (gender-marking) or semantic bias. Previous experiments on Russian found a dissociation between early and late measures, which agreed with the Relativized Relevance principle proposed by Frazier (1990) claiming that early attachment preference observed in cross-linguistic studies is determined by rather late processes and disguises effects of the Late Closure strategy (Chernova and Chernigovskaya, 2015). The present experiment showed no evidence in favor of the Late Closure strategy in initial syntactic commitments (first-pass time). Early attachment had a significant reading time advantage in second-pass and dwell time, agreeing with the Constraint Satisfaction accounts. When disambiguation was syntactic (i.e., by morphological means), late attached relative pronouns received the highest number of incoming regressions. When disambiguation was semantic, late attached verbs demonstrated the highest number of outgoing regressions. Regressions to competing NPs showed that NP1 was reread twice more often than NP2, which was not correlated with noun frequencies. Accuracy was significantly higher for syntactically constrained sentences and sentences with early attached relative clause modifiers. Everything seems to confirm early attachment preference in Russian.

Introduction

One thing that is very characteristic of natural human languages—as compared to most context-free programming languages, for example—is their massive ambiguity. Yet people hardly ever notice it. Only very occasionally do we encounter sentences that force us to use our metalinguistic awareness (reflecting on and consciously pondering about speech) to resolve an ambiguity in the input. In written language deprived of prosodic cues there exist at least three levels of ambiguity: lexical (concerned with word meanings and their lexical class), syntactic (concerned with packing words into hierarchical structures), and scope (concerned with semantic scope due to the presence of quantifiers). Ambiguity could also be either temporal (sentences containing strings of words that can be configured in multiple ways but having only one grammatically acceptable structure in the end) or global (sentences that have more than one grammatically acceptable structure). The way ambiguities are processed and resolved provides an important window into how we process written language in general and what strategies our brain uses to interpret it.

Apart from that, people do not possess unlimited time and cognitive resources to interpret linguistic input. Sounds are transient, do not stay around, and thus have to be processed quickly. Written speech interpretation usually imposes less time pressure but is also affected by processing resources. Cognitive processing has both temporal and quantitative limitations, such as the amount of material our working memory can hold simultaneously and the decay rate of this material (Frazier 1979:1). In order to interpret and assign meaning to an acoustic signal, it must remain active in the working memory buffer, which is generally considered to be limited to about seven independent chunks of information (Miller 1956). Moreover, the interpreted signal must also be integrated with the preceding utterance (Traxler 2012:195). The decay rate of the material also depends on the type of information to be processed: information with poor inherent organization will require more cognitive processing (Williams 2009:231) and will decay faster (Frazier 1979:1). This leaves the human brain with a quite narrow time-span during which the incoming information should be structured, interpreted, and integrated in the discourse. And people, as empirical evidence suggests, cope with it excellently and with very little conscious effort.

To explain this relative ease with which people process language so fast and so efficiently, different models of processing have been proposed. Most of them were originally developed on the basis of English data and turned out to make wrong cross-linguistic predictions (e.g., the Garden-Path Model postulating the existence of the universal Late Closure strategy has been consistently proven to be incorrect for many languages, including Russian; see Chapters 1 and 2). Relative clauses are very useful for examining cross-linguistic parsing strategies, because these structures often have comparable syntax. Since the phenomenon first attracted mass attention of psycholinguists in 1988, with the publication of Cuetos and Mitchell's paper on Spanish, a vast amount of research has been conducted on different languages. However, as of today, no agreement has been reached, with many studies showing contradictory results. Being a morphologically rich language, Russian is a perfect material to examine *how* and *when* syntactic and semantic constraints affect processing of adjunct modifiers.

Given the scale that the relative clause attachment research has reached in the last forty years, it is very surprising how few studies on Russian have been conducted so far. None of them have examined syntactically unambiguous clauses with relative pronouns using eye-tracking nor studied what role semantics plays in resolution of global syntactic ambiguities, i.e., how fast it is accessed, how reliably it is used, and how it interacts with attachment preferences. This study is intended to shed some light on these questions.

This work is structured as follows. Chapter 1 reflects on the notion of parsing and tracks the progress in this area from the origination of psycholinguistics as a field (1950s). It discusses the milestones in parsing models, from the earliest modular accounts to more recent constraint-based and shallow processing models. Chapter 2 discusses the phenomenon of relative clause attachment based on cross-linguistic data. It examines the major hypotheses proposed to account for large variance observed in these studies. A separate subsection is devoted to previous works in this area done on Russian and discusses their findings and limitations. Chapter 3 describes an eye-tracking experiment conducted to examine relative clause attachment in Russian and how it is affected by different constraints (syntactic and semantic). It discusses the technology of eye-tracking and the most commonly used measures: how they are affected by different linguistic manipulations and what stages of processing they usually reflect. It also describes the methodology, materials, two pre-tests (frequency norming and plausibility norming), predictions by different models, and results. Chapter 4 provides the summary of this investigation with an overall discussion of the experiment and possibilities for future research.

Chapter 1: Syntactic Processing

The aim of this chapter is to introduce the notion of parsing and what has been discovered about its mechanisms in humans so far. Section 1.1 provides a definition of parsing, elaborating on its incrementality, possible strategies and stages. It presents two most influential groups of parsing models which differ with respect to the number of stages and strategies they assume: modular and constraint satisfaction accounts. The issue of syntax-lexicon dissociation in the brain is examined from different angles (theoretical, psychological, and neurological). Section 1.2 provides an overview of the most influential parsing models proposed so far, and discusses evidence supporting or contradicting them.

1.1. General background

Syntactic processing, also known as *parsing*, refers to the process of organizing incoming strings of words into hierarchical structures (Traxler and Gernsbacher 2011:147) and assigning meaning to them based on the syntactic rules. It is widely accepted that the parser is equipped with grammatical rules of a given language and thus does not construct ungrammatical representations, which would later have to be eliminated by some specialized grammar device[1]. Evidence that people do construct ungrammatical representations would force us to revise all the existing parsing models (Traxler and Gernsbacher 2011:484). Most commonly, syntactic representations that the parser constructs are depicted as syntax trees. On a neurophysiological level, of course, no trees are present: the information about sentence structure is encoded as patterns of neural activity, and it is one of the greatest challenges to establish how the former translates into the latter.

Empirical studies, in particular eye-tracking and brain imaging experiments, suggest that the human parser is highly incremental: it does not wait until all the input has been read to start interpreting[2] it and building a structure, even though that implies making parsing decisions

[1] But see Traxler (2012:483–486) for an interesting overview of several experiments suggesting that people may, in fact, at least consider ungrammatical analyses, especially when the grammatically acceptable analysis is too complicated to comprehend (such as in multiple central embeddings, semantically anomalous sentences, “reversible” passives, etc.).

[2] As a matter of fact, activation of word meaning during reading was consistently shown to be autonomous: even when activating word meaning is task-irrelevant, people find it difficult or even impossible to prevent (see MacLeod, 1991 for a good overview of multiple Stroop experiments exploiting this phenomenon).

in the conditions of partial information. Locally ambiguous sentences (also known as garden-path sentences) in which difficulty is visible immediately when the disagreeing word is encountered, are the most informative for that purpose (see, for example, Rayner et. al. 1983; Mitchell and Holmes 1985; Ferreira and Clifton 1986; Trueswell et al. 1993; 1994; Tanenhaus et al. 1995; Clifton et al. 2003, among many others).

Given the limited resources available to the human parser, it seems only logical for it not to shelve any incoming strings of words and do as much interpretive work as it can to clear the "buffer" and maximize efficiency. Though the debates on incrementality and the exact size of processing chunks are not yet resolved, there is a great deal of evidence suggesting that the human parser operates on a word-by-word basis. This implies that it is impossible for the parser to "lay aside" the current word and wait for subsequent words to shed light on how the current word must be interpreted and attached to the tree. In this respect, it significantly differs from deterministic bottom-up computer parsers that are allowed to look ahead at *n* input symbols before deciding how to parse current symbols, which makes it possible for them to avoid backtracking (for the description of standard deterministic parsers and a possible adaptation of them to natural language, see Abney 1989).

Research on the topic of parsing started as early as the field of psycholinguistics itself—in the 1950-60s. Different parsing models have been mushrooming ever since, trying to address the questions of what governs initial parsing strategies and through how many stages the human parser goes to assign meaning to an utterance. Before we move to major parsing models proposed during around 70 years of psycholinguistic research, let's first address the very notions of a parsing strategy and stage of parsing, and tackle the problem of syntax-lexicon dissociation.

1.1.1. Parsing strategies

The aim of a parsing strategy is to guide the parser in making a decision when the information it needs is not (or not yet) available (Frazier 1979:3). There are currently about 7,000 languages in the world, and the absolute majority of people are speakers of at least one of them. Given such a variety, different configurations of parsing strategies are possible.

First, parsing strategies can be universal and used by all speakers in all languages we examine. As Frazier put it in her 1987 work, "[i]deally we should be able to remove the grammar of English from our theory of sentence processing, plug in the grammar of some other language, and obtain the correct theory of processing of that language. […] And, if it should turn out that language-specific parsing differences do exist, this fact along with the detailed differences will require explanation" (*ibid*.:565). There are several possible explanations for why that would be the case, but the most plausible one is that such strategies would yield processing advantage and increase efficiency, thus representing an optimal way of using human cognitive resources (Cuestos and Mitchell 1988:75). Given both temporal and quantitative limitations the human parser faces, being able to structure the input faster and with less resources would obviously have a huge processing advantage. However, considering the amount of structural variation among languages, such strategies may not be the most effective for *all* languages and for *all* cases, although they should be invoked when dealing with *similar* (or *analogous*) structures in different languages. Thus, if such strategies exist, they may constitute only a part of all strategies that speakers of different languages use. We may call such strategies *processing load driven*.

Alternatively, strategies vary from language to language. Such strategies would be "customized" to suit particular languages or even particular structures within such languages. These strategies would obviously be shaped by past experience with language, being exposure-based. The most plausible motivation for them would be the frequency of syntactic constructions. They may also bring processing advantage by dealing with some input in a particular language in the most efficient way. There is much evidence that this is indeed the case. For example, languages have been shown to differ in cues people use to interpret inconsistent sentences, with German speakers relying heavily on animacy, English speakers on word order, and Italian speakers on noun-verb agreement (see Bates et al. 1984; MacWhinney, Bates, and Kliegl 1984; McDonald 1986; Sokolov 1989; MacWhinney, Osmán-Sági, and Slobin 1991 to name just a few). We may call such strategies *language-driven*.

Strategies may also vary from speaker to speaker. Such strategies would most likely be driven by individual cognitive variance. Evidence suggests that there is, indeed, a significant by-subject variability in performance on syntactic processing tasks (for working memory-related theories, see King and Just 1991; Just and Carpenter 1992; MacDonald, Just, and Carpenter 1992; Caplan and Waters 2002; Swets, Desmet, Hambrick, and Ferreira 2007; for cognitive control and information suppression theories, see Gernsbacher and Faust 1991; Gernsbacher 1993; for perceptual interference theories, see Leech et al. 2007, among others). We may call such strategies *individual-driven*.

Obviously, we may find an intricate combination of all the above strategies. For example, humans may be born with processing load driven strategies that are then shaped by both an exposure to a particular language and our individual differences in cognitive abilities.

1.1.2. Parsing stages

Speaking of stages of parsing, it is reasonable to begin with addressing the topic of brain functional specialization. In a series of seminal papers, Forster (1974; 1976; 1979) proposed an idea that the human language comprehension system consisted of several autonomous processing modules (a lexical processor, a syntactic processor, and a message processor). The whole notion of such a module is, perhaps, best addressed in Fodor's influential monograph *The Modularity of Mind* (1983), which triggered many years of heated debates among cognitive scientists. A module, in his definition, is domain-specific, informationally encapsulated from the background knowledge, fast, mandatory, and has a fixed neural architecture. Despite its somewhat misleading name, the monograph did not assert that the human mind was massively modular. Instead, Fodor argued that only lower-level cognitive processes (which he called "input systems") were modular. The job of such input systems is to take a sensory stimulus and perform basic recognition and description. Upon completing the task, input systems feed their output to higher-level cognitive processes (an example of which is analogous thinking) that, according to Fodor, are nonmodular, since they have access to all information contained within the cognitive system. As far as language is concerned, Fodor hypothesized that language processing was a module encapsulating the grammar and the entire lexicon: "Presumably, the language processing system has access to a grammar of the language that it processes, and a grammar must surely contain a lexicon. What words are in the language is thus one of the things that the language module can plausibly be assumed to know consonant with its modularity" (J. A. Fodor 1985:5). The concept of modularity shaped the parsing research and laid the foundation for many models of parsing, as well as provoked heated discussions on whether there were any submodules (such as syntax) within a big language module.

1.1.3. Syntax-lexicon dissociation

Whether, as Fodor put it, grammar indeed contains lexicon, is a big topic for discussion. At the dawn of modern linguistics as a discipline, syntax and lexicon where generally perceived as two distinct components. Early Chomsky (1965) and theoretical linguists following his framework considered "abstract syntactic rules" and lexical information to be two disjoint subsets. Early psycholinguists also posed a dichotomy between the two, often treating syntax as having a privileged position over lexical information, which is described by a famous saying "syntax proposes, semantics disposes." Forster (1974) proposed the idea of the "autonomy of syntax." However, as more and more linguistic data has been accumulated, this distinction has become more and more blurred. There is little doubt that some abstract syntactic rules exist, because people manage to do some processing[3] of Jabberwocky sentences constructed according to grammar rules of a given languages but from made-up words, which renders them semantically meaningless:

Twas brillig, and the slithy toves

Did Gyre and gimble in the wabe;

All mimsy were the Borogoves,

And the mome raths outgrabe.

Lewis Carroll (1872)

However, in most other cases, the boundary between syntactic and lexical information becomes less and less clear-cut (see Clifton et al. 1984; MacDonald et al. 1994; Trueswell et al. 1994 for frequency effects not only for syntactic constructions but also for lexical items in particular syntactic contexts), although there is still a lot of evidence for a dichotomy between syntactic and other information from psycho- and neurolinguistic studies using a variety of techniques. Eye-tracking and SPRT (self-paced word-by-word reading tasks) experiments reported the so-called "garden-path effect"—a necessity to reanalyze the initially misparsed syntactically ambiguous fragment after the disambiguating word has been encountered (Rayner, Carlson, and Frazier 1983; Ferreira and Clifton 1986; Mitchel 1987[4]; Ferreira and Henderson 1990; Clifton 1993, among others). The reanalysis is visible in the reduced pace of reading or regressive eye movements. Another piece of evidence comes from experiments using the speed-accuracy trade-off technique. For example, McElree and Griffith (1995) used it to examine whether there was a temporal dissociation between structural and other types of information. The participants were presented with sentences with semantic anomalies (1a), subcategorization violations (1b), or syntactic category violations (1c):

[3] See Kako (2006) for grammaticality judgments research and Hahne and Jescheniak (2001) for event-related potentials.

[4] But see Adams, Clifton, and Mitchell (1998) for conflicting evidence.

(1) a. *Some people alarm books.*

b. *Some people agree books.*

c. *Some people rarely books.*

Each sentence was followed by a tone at various time intervals, upon which the participants had to immediately reply whether the sentence made sense. The accuracy of responses after a rapidly following tone was at a chance rate, while the accuracy of responses after a long delay was the same as in untimed tasks, indicating that the asymptote was reached. However, the participants reached an asymptote slower for (1a) than for (1b) or (1c), suggesting that semantic information is processed slower (or is delayed) in comparison with syntactic category or subcategorization information (but see Spivey, Fitneva, Tabor, and Ajmani 2002 for an alternative explanation). Finally, several ERP studies showed the dissociation between syntactic and other types of information (Hagoort, Brown, and Groothusen 1993; Ainsworth-Darnell, Shulman, and Boland 1998; Münte, Heinze, and Mangun 1993; Rösler et al. 1993), in which the size of N400 was larger when it occurred with gender violation on the same word, whereas P600 was unaffected by semantic violations.

The above evidence covers the topic of syntax-lexicon dissociation from the view of modular vs. connectionist debates (that is, whether language and/or its components are separate isolated systems or interact with other cognitive systems). Regarding the issue of localization (whether there are brain areas sensitive to either only linguistic information or just one type of linguistic information but not others), most of neurophysiological evidence suggests that no region in the brain is sensitive to only lexical or only syntactic information, although there are many debates as to the experimental items representing these two opposite poles. Fedorenko, Nieto-Castañon, and Kanwisher (2012) used multi-voxel pattern analyses to examine which information, pure lexical or pure syntactical, is represented in the brain more robustly, and whether some brain regions reliably distinguish between "pure" lexical information (lists of words) and "pure" abstract syntactic information (Jabberwocky sentences) in their pattern of activity. They found that lexical information was represented more robustly. There was a better discrimination between conditions that differed along the lexical dimension (sentences vs. Jabberwocky, and word lists vs. nonword lists) than between conditions that differed along the syntactic dimension (sentences vs. word lists, and Jabberwocky vs. nonword lists). Also, surprisingly, they found that some regions in the inferior frontal gyrus and posterior temporal cortices reliably discriminated between pure lexical and pure syntactical information in their patterns of firing. Fedorenko, Nieto-Castañon, and Kanwisher (2012) reasonably argued that there was a continuum between purely lexical items on one end and abstract syntactic rules on the other, with most linguistic data falling somewhere in between—that is, stored in the lexicon together with syntactic/semantic contexts they frequently occur in. They also concluded that lexical information was a very important source of information guiding initial sentence interpretation.

1.2. Major parsing models

All the above split parsing research into several camps, with more and more models emerging as more and more theoretical and empirical cross-linguistic evidence was obtained. In the 1985 commentary to Fodor's modularity theory, Janet D. Fodor was one of the first to classify parsing models with respect to the amount of interaction they allow. She divided them into "algorithmic" and "detective" (or "heuristic").

Algorithmic models, later known as **two-stage models** or **modular models**, are based on the assumption that either language as a whole or some aspect of it (e.g., syntax) is subserved by a domain-specific module that exists independently of a central store of general knowledge. The flow of information is blocked from both sides: module is cognitively impenetrable and does not reference any other systems in order to perform its task. Janet D. Fodor (1985) referred to the algorithmic account of processing as "deeply unintelligent": "[...] the parsing mechanism is programmed to examine input words sequentially as they are received and to respond to each one in some quite specific way, such as adding certain nodes to a phrase marker in temporary memory" (*ibid.*:8). Because the algorithmic parser operates on a pool of template-based strategies (or responses, as Janet D. Fodor called them) and does not communicate with other systems while making initial decisions, it is normally assumed that such a parser will only construct one analysis at a time. If some additional information appears that makes the initial parse incorrect, then the reanalysis occurs. This account is called *serial*.

Detective models, which are based on the same idea as **interactive models** or **constraint-satisfaction models**, are the opposite of algorithmic models. They claim that almost all cognitive processes are interconnected, and that there is a free exchange of information between them. Detective models draw on a vast number of different clues in order to make the best structural guess. Compared to the algorithmic processor, the detective processor should be much more intelligent and much harder to simulate as a computer program, primarily because its computations are "global." Of course, so much unstructured data would create a massive ambiguity and slow the processor down. It is thus commonly assumed that different clues provide different amount of support to different analyses, and if two constraints conflict, one will take precedence: "Typically, no one clue will be decisive for sentence structure; each must be weighted and integrated. If clues conflict, then one must be allowed to override another, and so forth" (Janet D. Fodor 1985:8). The parser must weigh all the available data and choose the analysis that receives most support. In an interconnected network, everything communicates with everything, and there is constant feedback between different systems and stages of processing, contrary to modular organization. Also very importantly, "[…] what counts as a useful superficial clue to structure is likely to be highly language-relative, suggesting that the success of a detective procedure requires considerable experience with parsing this particular language" (*ibid.*). As discussed above, this has been consistently shown to be the case. These models mostly assume that the parser constructs multiple analyses and ranks them accordingly. If new information supports the candidate that is lower on the list, then the re-ranking occurs. This account is called *ranked parallel*. One big flaw of such models is their poor predictive power: as of today, no complete list of possible constraints has been created, let alone describing how exactly they affect processing.

Let's now discuss the major parsing models proposed over almost 70 years of psycholinguistic research.

1.2.1. Bever's heuristics (1970)

The development of the field started with the assumption that initial sentence processing is governed by template-based heuristics that get the job done fast, although sometimes at the cost of accuracy. One of the first models of parsing was proposed by Bever (1970). The model consisted of over a dozen heuristic strategies that, as he claimed, were a subpart of general cognitive processes and drew on the same principles as general perceptual mechanisms. The evidence for heuristics, Bever argued, came from misparsed sentences, when people constructed analyses that were not grammatically sanctioned. Almost 40 years later, Townsend and Bever (2001) formulated it this way: "a quick and dirty parse is initially elicited. [...] This preliminary analysis uses a variety of surface schemata in conjunction with verb argument and control information to organize an initial hypothesis about meaning" (*ibid.*:163). Bever's work was intended to provide an alternative approach to the then-popular Derivational Theory of

Complexity—the idea that perceptual complexity of a sentence is a function of the number of grammatical rules employed in its derivation (Miller 1962). It was hypothesized that during comprehension the human parser has to "reverse engineer" the input to go back to the initial stage, from which the process of sentence building would normally start during production. In other words, the parser has to remove the applied transformational rules one by one, and the number of such rules is the primary contributor to perceptual complexity of a sentence. Bever, on the contrary, argued that instead of directly applying these complex syntactic algorithms to comprehend a sentence, people used heuristics (which he also called "perceptual strategies") to identify the deep structure of the sentence.

Bever proposed two major groups of speech perception strategies—segmentation strategies (which establish clausal relations in the sentence) and functional labelling strategies (which establish structural relations within clauses). The latter draw on, among other, some sort of semantic and frequency information. Apart from his Strategy A[5], which later became known as the NVN strategy, these principles have been mostly abandoned ever since; however, they had an important influence on the development of the field. Bever's model was one of the earliest, if not the earliest, examples of a detective-style (heuristic) model.

1.2.2. Kimball's two-stage model (1973)

Kimball (1973) examined parsing strategies from the standpoint of sentence acceptability. He argued that human languages differ from computer languages in two important ways. First, compared to context-free programming languages whose grammars are unambiguous[6] and deterministic (yielding a unique parsing tree for each string), human languages are ambiguous, so the parsing model for human languages must differ from that of computer languages and allow for multiple underlying structures. Second, the computer parser has almost unrestricted memory: it can go *n* symbols ahead and decide that the appropriate action is to read in the next symbol. Human short-term memory is, on the contrary, quite limited, so the parser must constantly build trees over input strings for them to be cleared out of the memory (*ibid.*:20). Kimball then proceeds with formulating "six or seven" principles of surface structure parsing. For reasons of space, I will not list them all here. What is relevant now is that this model was

[5] Strategy A: Sequence together any sequence X…Y, in which the members could be related by primary internal structural relations, 'actor, action, object…modifier' (Bever 1970:290).
[6] However, context-free grammars that are deterministic (that is, always unambiguous) only constitute a subclass of grammars in computer languages. Moreover, some classical context-free languages have been shown to be inherently ambiguous (e.g., Flajolet 1987). In fact, many programming languages are ambiguous due to the dangling-else problem and other issues.

an early example of a modular model that states a clear dichotomy between syntactic and other types of information.

1.2.3. Frazier's Garden-Path Model (1979; 1987)

Not long after Bever and Kimball, Frazier (1979; 1987) came up with a more influential two-stage model that shaped the study of parsing mechanisms for more than ten years, before it was substantially revised in 1996 (Frazier and Clifton 1996). The Garden-Path Model separates syntactic processes from other linguistic processes (such as semantic, thematic and discourse processes) and encapsulates them in a module that only communicates with other systems through input and output. It consists of two parsing stages:

Table 1. Parsing stages assumed by the Garden-Path Model

Stage 1	Access to syntactic information only (modular); One candidate initially created (serial).
Stage 2	Access to semantic, thematic, and contextual information; Initial candidate evaluated in context; Revision for incorrect candidates.

The Garden-Path model assumes that, to incorporate new phrases into the preceding tree, the parser uses three principles: Minimal Attachment, Late Closure[7], and Active Filler[8] (a recent addition to the theory). These principles are claimed to apply cross-linguistically.

The Minimal Attachment principle, according to Frazier, means that no potentially unnecessary nodes in the syntactic representation should be postulated. It is immediately obvious that such a principle is heavily theory-dependent: different syntactic theories may assume different number of nodes for the same structures. It also poses the risk of inverse operation: if some structure is preferred over another, we can postulate that it has "fewer nodes", thus being minimal. Frazier provided an explicit psychological motivation for the Minimal Attachment principle: "Minimal attachment analyses will be available earlier than nonminimal ones due to the relative number of phrase structure rules that must be accessed for the two analyses" (Frazier 1987:564). In other words, the parser is able to compute the minimal structure faster, thus yielding a processing advantage. Taking into account working memory

[7] The Minimal Attachment principle is an adaptation of Kimball's Right Association strategy, and Late Closure is a modified Closure strategy.
[8] "Active Filler Hypothesis: When a filler has been identified, rank the option of assigning it to a gap above all other options" (Frazier and Clifton 1989: 95).

limitations and other issues we have discussed above, such a principle looks well-motivated. Judging by our everyday experience, pre-existing rules and templates that do not conflict with each other facilitate the decision-making process and release the cognitive load. The Garden-Path Model is essentially a model of minimizing processing load: universal rules allow the processor to spend on parsing decisions as few cognitive resources as possible. However, there is always a trade-off. By sticking to predefined parsing rules that are context-independent, the parser increases its speed but has to go through the costly process of reanalysis, if the information that comes later (semantic, thematic, etc.) disagrees with the initial parsing decision.

Minimal Attachment was first introduced in the earliest version of the model—the Sausage Machine—proposed in Frazier and Janet D. Fodor (1978). The authors argued that the human parser builds structures for word strings in two steps. First, it takes substrings of roughly six[9] words and assigns lexical and phrasal nodes. Then, it merges the phrasal packages into a complete sentence structure by adding higher non-terminal nodes (*ibid.*:2). They hypothesized that the parser consists of two "subparsers": the Preliminary Phrase Packager (or the Sausage Machine) and the Sentence Structure Supervisor. The Preliminary Phrase Packager, or PPP, they argued, is "short-sighted" (sees only six words at a time) and is in some respects insensitive to the rules of well-formedness. The Sentence Structure Supervisor, or SSS, on the contrary, has a much larger span: it can keep track of long-distance dependencies and long-term structural commitments, and survey the whole sentence structure. Frazier and Fodor claimed that the Minimal Attachment principle was the only principle guiding the structuring of a sentence.

Frazier (1979; 1987) finalized the model by making several important adjustments. First, she got rid of subdivisions inside the parser (PPP and SSS). Second, she added the Late Closure strategy. This strategy states that the parser attaches the incoming words to the lowest node (the currently processed constituent) because it has less cognitive cost. She also proposed the Active Filler strategy, but it did not gain the same popularity as the first two principles.

Not much is known as to what exactly syntactic information is used during the first stage. Frazier assumed that the processor constructed its initial analysis based solely on the lexical categories of words represented in the input (most likely, on a word-by-word basis) and then feeds this input into the syntactic parsing mechanism (Traxler 2012:148). This means that the parser does not know what specific words are represented in the input. Given that originally the model was developed on the data from English, it makes some sense. However, it can be

[9] This number is based on the evidence that human working memory is limited to about seven independent chunks of information (Miller 1956).

that case, number, and gender are also used at this stage. Case is a rather complicated phenomenon, with some instances being more syntactic (like nominative and accusative) and some being more thematic (like instrumental). To my knowledge, there are no studies so far focusing on this very question, so for the rest of this work I will assume that the parser initially employs only word category information. This has direct predictions that will be tested in this work's experiment: if only word category information is used initially, then gender-marked relative pronoun should be first late-attached (following the Late Closure strategy) and then reanalyzed when there is a disagreement, increasing reading time (see Chapter 3, sections 3.4 and 3.7).

Because the Garden-Path Model belongs to processing load driven models and leverages the general cognitive architecture of the human brain, it assumes that the principles hold universally and are independent of individual languages. As Frazier herself put it, "[a]ssuming that the need to structure material quickly is related to restrictions on hum an immediate memory capacity, we might expect all humans to adopt the first available constituent structure analysis. If so, we expect the minimal attachment and late-closure strategies to be universal" (Frazier 1987:564-565).

1.2.4. Frazier and Clifton's Construal Hypothesis (1996)

After multiple studies showed considerable cross-linguistic variation and failure to observe the parsing principles proposed under the Garden-Path Model, Frazier and Clifton (1996) revised the model to account for the new data. The Construal Hypothesis claims that the universal parsing principles only hold for the so-called primary phrases, which include the subject and main predicate of any finite clause as well as complements and obligatory constituents of primary phrases (Frazier and Clifton 1997:279). Frazier and Clifton also claimed that phrases temporarily taken to be primary will be treated as if they were primary phrases. Non-primary phrases, which include RC adjuncts, are not processed using the universal parsing principles. The authors suggested that such RC adjuncts were underspecified phrases and could associate to either the last theta-assigner or its projection:

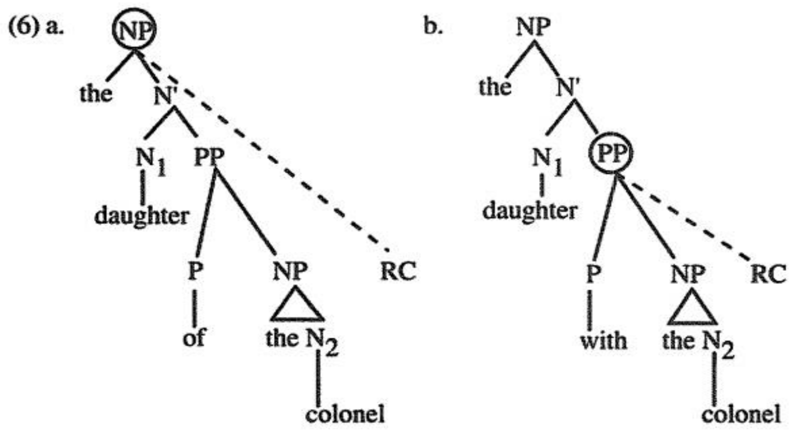

Figure 1. Possible relative clause attachment sites (Frazier and Clifton, 1997: 281)

According to Frazier and Clifton, in cases like (6a) in Figure 1, when the relative clause is encountered, *the daughter* is the last theta-role assigner, so two heads are now available for attachment: the bigger NP headed by *the daughter* and the smaller NP headed by *the colonel*, giving rise to ambiguity. In contrast, in (6b) in Figure 1 the last theta-role assigner is the preposition *with*, so only *the colonel* (being the only head in the theta-role assigner domain, PP) is available for attachment, eliminating ambiguity. They also hypothesize, that, because the alternative Saxon genitive construction is available (*colonel's daughter*), which can be unambiguously used to attach the relative clause to *the daughter*, attachment to *the colonel* is preferred in (6a) but does not exclude the possibility of attachment to *the daughter*[10].

In many Slavic languages and Russian in particular, only one possessive construction is possible (there is no analog of Saxon's genitive). Therefore, if the Construal Hypothesis is on the right track, there should be massive ambiguity and no preferred attachment. Frazier and Clifton (1997) admit they do not have an explanation for what makes the processor choose this or that analysis: "The question of what constitutes sufficient evidence for the parser to commit itself to a particular attachment consistent with an association is one that will take additional research to answer" (*ibid.*:281). They hypothesize that a variety of semantic, pragmatic and lexical factors influences the selection. What exactly those factors are will be explored in detail in Chapter 2.

[10] However, see Mitchell et al. (2000) for conflicting evidence from Dutch. Although both constructions are possible in Dutch, attachment to the first noun was found to be preferred in an "ambiguous set-up" (the Norman genitive).

1.2.5. Constraint Satisfaction models

In 1994, MacDonald and colleagues published an article titled *The Lexical Nature of Syntactic Ambiguity Resolution* that marked a milestone for interactionist accounts of syntactic parsing. The authors argued that lexical and syntactic ambiguities, which were commonly assumed to be resolved by different mechanisms, might actually be resolved by the same mechanism, meaning that people build syntactic structures in mostly the same way as they determine a word's meaning. As the authors themselves put it, "recent types of theorizing eliminate the strong distinction between accessing a meaning and constructing a syntactic representation, which was central to previous accounts" (*ibid.*:682). They suggested that semantic and contextual information might be deployed much earlier, during initial parsing, and that formal syntactic algorithms might not be directly applicable in online processing. The account they proposed was much more integrated and unified than Frazier's Garden-Path Model. It is part of a general class of *Constraint Satisfaction accounts* (MacDonald, Pearlmutter, and Seidenberg 1994; Spivey-Knowlton and Sedivy 1995; Trueswell 1996; Trueswell, Tanenhaus, and Kello 1993).

The main idea behind these accounts is that the human parser uses all information available at the moment to construct multiple analyses and rank them accordingly. When one analysis receives significantly more activation, parsing is easy, but when several structures receive similar activation, processing load increases. This increase in reading times has been observed in resolution of lexical ambiguities, where balanced ambiguous words were read faster if the context made one of its two meanings more appropriate, and imbalanced words were read longer if the context made their subordinate meaning more appropriate, thus making two constraints conflict: frequency and contextual activation (Duffy, Morris, and Rayner 1988). According to Constraint Satisfaction models, resolution of syntactic ambiguities proceeds in the same fashion. When information available later in the sentence disagrees with the top-ranked analysis, a disruption and rearrangement of analyses occurs, and the more activated was the top-ranked analysis, the harder it will be for the parser. What information exactly the parser uses has been a matter of hot debates. In general, Constraint Satisfaction accounts propose three large groups of constraints:

- lexical (semantic and frequency information from individual lexical items) (MacDonald et al. 1994; Trueswell and Tanenhaus 1994);
- structural (word order, frequency of specific constructions) (Hawkins 1995);

- discourse-level (a top-down constraint on processing, forming when the word is integrated with the rest of the utterance) (Altmann and Steedman 1988).

Several studies found that animacy (Trueswell, Tanenhaus, and Garnsey 1994; Mak, Vonk, and Schriefers 2002) and subcategorization preferences (Trueswell, Tanenhaus, and Kello 1993) could also guide initial preference and eliminate difficulty with processing sentences when there is any conflict. As far as semantic plausibility unrelated to animacy goes, no clear evidence that it eliminates difficulty with reduced clauses or affects syntactic ambiguity resolution has been found so far (Rayner, Carlson, and Frazier 1983; McRae, Spivey-Knowlton, and Tanenhaus 1998, among others). Schriefers, Friederici, and Kuhn (1995) conducted an experiment on German examining both plausible and implausible subject and object relatives. They found that implausible relatives were harder to process (meaning that plausibility played a role in processing), but the difference in difficulty between subject and object relatives was unaffected by whether the plausible interpretation supported the more difficult object relative structure or not.

Another evidence comes from the effect of context on further processing. The Garden-Path Model, assuming universal parsing principles driven by minimization of the processing load, predicts that context should not have any influence on subsequent processing. However, multiple experiments demonstrated that this was not the case (see Crain and Steedman 1985; Altmann and Steedman 1988; Zagar, Pynte, and Rativeau 1997; cf. also Mitchell, Corley, and Garnham 1992; Rayner, Garrod, and Perfetti 1992).

1.2.6. Barton and Sanford's "Good Enough" Model (1993)

In a 1993 article, Barton and Sanford demonstrated that human mental representations during speech comprehension were far from being ideal, coherent, and correct—an assumption that for many years had been taken for granted in parsing research. The authors approached the issue mostly from a semantic point of view. They carried out several experiments with sentences like "When an airplane crashes, where should the survivors be buried?" and found that in many cases people had not spotted an anomaly. They proposed the so-called "shallow processing" account of speech comprehension, arguing that our mental representations were very often underspecified and "good enough." Sanford and Sturt (2002) elaborated on the issue even more, digging also into pragmatics and context and demonstrating that many ambiguities (whether scope, lexical or other) may remain unresolved (see also Frazier and Rayner 1990 for evidence of underspecification in comprehension of nouns with multiple sense, like *newspaper*). There is less evidence for syntactic underspecification, and Sanford and Sturt (2002) provided only one example of "shallow parsing" in computational linguistics (automatic generation of indexes for large texts).

1.2.7. Van Gompell et al's Unrestricted Race Model (2005)

Some experiments, however, showed results that disagreed with both interactive and modular models. In an eye-tracking study, Van Gompell et al. (2005) found that globally ambiguous sentences were read faster than their locally ambiguous counterparts:

(2) a. *I read that the bodyguard of the governor retiring after the troubles is very rich. (globally ambiguous)*

b. *I read that the governor of the province retiring after the troubles is very rich. (high attachment)*

c. *I read that the province of the governor retiring after the troubles is very rich. (low attachment)*

d. *I read quite recently that the governor retiring after the troubles is very rich. (syntactically unambiguous)*

(Van Gompell et al. 2005:289)

The constraint-satisfaction models predict that sentences like (2a) should be the hardest to process because of the strongest competition: here both semantics and syntax support two interpretations. Sentences like (2b) should rank second because in English low-attachment preference was found (Cuetos and Mitchell 1988; Carreiras and Clifton 1999). Sentences like (2d) should be the easiest. Surprisingly, Van Gompell et al. found that (2a) was, in fact, easier to process than (2b) and (2c) and did not differ from (2d). Interestingly, an experiment conducted on Russian that compared sentences like (2a) to (2d) found the opposite pattern: syntactically ambiguous sentences with complex noun head were read significantly longer that their unambiguous counterparts with just one noun (Anisimov, Fedorova, and Latanov 2014). To account for this data, van Gompell and colleagues proposed the Unrestricted Race Model. It claims that in the case of balanced syntax-semantics ambiguities like (2a), the processor adopts one of the possible analyses roughly half the time. Given that it never has to reanalyze, no difficulty is experienced further on. This finding is also consistent with Traxler, Pickering, and Clifton (1998) who found that globally ambiguous sentences were read faster than determinately attached sentences. This is an interesting idea that can also be tested by offline measures. If the processor indeed adopts both analyses about half the time in balanced ambiguities, then answers to comprehension questions ("Who is retiring after the troubles?") must be evenly distributed. However, from multiple cross-linguistic studies we know that this is not the case, with early attachment dominating the research (see next chapter for detailed discussion).

1.3 Chapter summary

This chapter provided an overview of the most influential parsing accounts proposed during over seventy years of psycholinguistic research. As we have seen above, there is no conclusive evidence in favor of any of those models. The main division line between them lies in what information is used during the initial parse, whether some of it is more "privileged" over the other, and how much interaction there is between different cognitive systems. Although there is a lot of evidence for at least some dichotomy between syntactic and other information (the so-called "garden-path effect" in eye-tracking and self-paced reading studies; modulated N400 when it occurs with gender violation on the same word, whereas P600 remains unaffected by semantic violations; participants approach the asymptote (performance as in untimed tasks) slower for plausibility violation in comparison with syntactic and subcategorization information), other neurological evidence suggests that no region in the brain is sensitive to only lexical or only syntactic information. Also, many experiments have demonstrated that people used discourse information, semantics, and frequency of not only syntactic constructions but also lexical items in particular syntactic contexts during initial parses, which disagrees with what modular accounts predict. Given that most of those models were initially developed on the basis of English, the evidence is even more mixed when it comes to cross-linguistic studies. It is obvious that more research is needed to verify the main parsing accounts by findings from morphologically rich languages.

Chapter 2: Relative Clause Attachment

The goal of this chapter is to examine the phenomenon of relative clause attachment. It starts with discussing why the said phenomenon has been recently receiving so much attention from psycholinguists and what it could tell us about parsing in different languages. It examines cross-linguistic data and discusses different explanations proposed to account for a great variance we have observed, including the role of semantics in relative clause attachment processing. Section 2.2 provides an overview of experiments that have been conducted on the Russian data and elaborates on their findings and limitations.

2.1. General background

Relative clause attachment ambiguity is a very interesting phenomenon, because as of today no parsing model has managed to account for all cross-linguistic data we have. Moreover, it represents a structure that uses comparable syntactic devices in many languages, which makes it cross-linguistically comparable. The said phenomenon attracted the mass attention of psycholinguists in 1988, with the publication of Cuetos and Mitchell, which provided the first challenge to the claimed universality of the Late Closure strategy. To examine relative clause attachment preferences, the authors conducted several offline and online experiments on Spanish data and found out that Spanish speakers, unlike English speakers, did not prefer late attachment but rather early attachment with comparable materials:

(3) a. *El periodista entrevisto a la hija del coronel que tuvo el accidente.*

b. *The journalist interviewed the daughter of the colonel who had had the accident.*

(Cuetos and Mitchell 1988:77)

What they found exactly was three things. First, Spanish readers preferred early attachment in ambiguous sentences in offline questionnaires. Second, in SPRT, reading times at the point of disambiguation were much longer when the relative clause was forcefully attached to the second noun, compared to when it was attached to the first noun or could attach to both. Third, there was a significant reading time advantage in Spanish speakers when the relative clause was forced to attach to the first noun (by either pragmatics devices or gender mismatches). Carreiras and Clifton (1993; 1999) also confirmed these findings for Spanish.

Today the literature on relative clause attachment is very ample and diverse. Interestingly, though, only English has more or less consistently proven to have late attachment preference (Cuetos and Mitchell 1988; Carreiras and Clifton 1999, among other). All other languages either showed early attachment preference (Italian: Cuetos, Mitchel, and Corley 1996; German: Konieczny et al. 1997; Dutch: Brysbaert and Mitchell 1996; French: Zagar, Pynte, and Rativeau 1997; etc.) or no real preference, with a lot of subject-to-subject variability (Carreiras and Clifton 1993) or variability depending on other factors (such as a preposition type, presence of a determiner, etc. which will be reviewed below). All this data contradicts the claimed universality of the Late Closure strategy.

Frazier and Clifton's (1996) Construal Hypothesis was an attempt to account for such variance. It states that in sentences with non-primary relations there is no initial commitment to any structural analysis. The relative clause is thus assumed to be "associated" to an entire theta-domain and remains unattached until other "range of information yet to be determined" (Carreiras and Clifton 1993: 365) comes into play:

a. Associate a relative clause to the current thematic processing domain.
b. Interpret the relative clause with any grammatically permissible material in the associated domain using structural and semantic/pragmatic information.

(Frazier and Clifton 1996:31-32).

Multiple attempts have been made to elaborate on this topic. Gilboy et al. (1995) formulated the so-called *Referentiality Principle* claiming that restrictive modifiers (such as relative clauses) preferentially attach to those hosts that introduce discourse entities into a discourse model with the help of a determiner. The authors tested three different types of noun phrase complexes: substance and quantity NPs where NP2 did not introduce discourse entities, appearing without a determiner (*a sweater of wool*), noun phrase complexes with different semantic relations between the two NPs (*daughter of X*, *assistant of X*) where both nouns introduced discourse entities and were referential, and complexes with the preposition *with* restricting the theta-domain to NP2 only (*a steak with a sauce*). They found that different factors affect the resolution of ambiguities: the type of a preposition (when the preposition is the last theta-role assigner, no ambiguity should arise), referentiality (if NP2 lacks a determiner, it becomes a less preferable host for attachment), and semantic relations between two nouns in the noun phrase complex.

De Vincenzi and Job (1993; 1995) found early attachment preference in Italian readers with the preposition *of* but late attachment preference with the preposition *with*. Konieczny et al. (1997) found differences between relative clauses and prepositional phrase modifiers: German readers in their experiments preferred to attach the former to the first noun and the latter to the second noun. Interestingly, though, while De Vincenzi and Job (1993; 1995) found differences in offline preferences depending on the type of a preposition, in their self-paced reading experiments they found that Italian readers were initially faster to read adjunct modifiers in both conditions when they semantically modified the second noun phrase rather than the first one.

Frazier (1990) made another attempt to explain the early attachment preference consistently found cross-linguistically. She noticed that sentences used in the Cuetos and Mitchell's (1988) experiment were very long, with disambiguation occurring almost in the end, so by this time higher-level interpretation processes had enough time to intervene. It is thus hard, she argued, to determine whether early attachment preference really reflects initial syntactic commitments or is driven by the second stage of processing. She formulated *the Relativized Relevance* principle: "Other things being equal, e.g., all interpretations are grammatical, informative and appropriate to discourse, preferentially construe a phrase as being relevant to the main assertion of the sentence" (Frazier 1990:321). This principle claims that once the initial structurally determined choice has been made, another mechanism called the thematic processor comes into play very quickly and can trigger a full revision of the analysis that has been made, thus disguising any effects of the Late Closure strategy. If reader's commitments are checked at any point after that, there will be no traces of the initial analysis. Attachment of the new material to the more prominent host (the most salient discourse referent) is preferred, other factors being equal. Thus, for example, in sentences where the first noun is the direct object of the main verb, this noun is the most salient discourse referent, with the other noun being outside of focus, and the relative clause should be attached to it. Unambiguous sentences should still be read faster when they contain late attached modifiers rather than early attached.

Russian is perfectly suited to examine initial attachment preferences because the relative clause could be disambiguated already on the first word (relative pronoun) by means of morphology. The first experiment that used this feature was Brysbaert and Mitchell (1996) who examined early attachment preferences in Dutch. They manipulated the latency of disambiguation (immediate and late). Interestingly, when total reading time was analyzed, they found only a marginally significant main effect of position of disambiguation, which did not interact with that of attachment preference. This suggests that the nature of conflict resolution in sentences with immediate disambiguation is more complicated than the processing of delayed disambiguation, apparently because the preferred analysis does not manage to grow strong enough for the disruption to be visible. Another experiment that used a similar design was Zagar, Pynte, and Rativeau (1997) who examined relative clauses disambiguated by means of gender-marking on the adjective in French. Although they did find a significant effect of attachment in first-pass in one of the two conditions, with longer fixation times on the disambiguating zone during first-pass reading when it was forcefully attached to NP2, it should be noted that the critical region was still separated from the second noun in the complex noun phrase by three words:

(4) a. *Un journaliste aborda l'avocat de la* ***chanteuse*** *qui semblait plus* ***confiante*** *que de raison.*

A journalist approached the barrister [male] of the singer [female] who seemed more confident [feminine gender] than she ought to.

b. *Un journaliste aborda l'avocat de la* ***chanteuse*** *qui est plus* ***confiante*** *que les autres.*

A journalist approached the barrister [male] of the singer [female] who seemed more confident [masculine gender] than the others.

(Zagar, Pynte, and Rativeau 1997:424, my emphasis)

Thus, it can be argued that interpretation processes still had time to enter into play before the disambiguating zone was reached, as Frazier (1990) noticed. It is interesting to see whether the relative clause disambiguated already on the first word would show difference between early attached, late attached, and ambiguous pronouns, and this would be tested in the current experiment.

Mitchell et al. (1995) proposed an exposure-based model called the *Tuning Hypothesis*. It states that preferences of the parser in dealing with ambiguities of any kind are shaped by the person's previous encounters with ambiguities of the same kind. It also predicts that every time the chosen strategy by the parser turns out to be correct, it strengthens the preferred analysis for this type of structures, and the parser would be more likely to choose this strategy again. It also predicts close correspondence between corpus data and behavioral data. In the Spanish corpus, for example, early attachment accounts for 60 percent of cases, whereas in the English corpus, for only 38 percent of cases. However, Brysbaert and Mitchell (1996) found counterevidence: in the Dutch corpus, early attachment accounts for only 29 percent of cases, whereas Dutch speakers have been shown to favor the early attachment strategy. Also, inconsistencies between subjects' responses and corpus analysis have been reported for English (Gibson, Schütze, and Salomon 1996).

The *Predicate Proximity/Recency theory* put forward by Gibson, Schütze, and Salomon (1996) claims that two factors compete while the parser selects a suitable host for RC attachment:

a. The structural proximity of each NP to the head of the entire predicate phrase;
b. The relative distance between the modifier and each NP.

The first principle basically says "attach to the head of a predicate," and the second "attach to the most recent host." The relative weight of these two factors seems to differ across languages, resulting in an observed variance. It also assumes that the weight of the Predicate Proximity factor is enough to strongly outweigh the Recency factors in the majority of languages examined but not in English.

Constraint Satisfaction accounts in general do not focus specifically on RC ambiguity resolution; however, MacDonald, Pearlmutter, and Seidenberg (1994:697-698) proposed one account to capture cross-linguistic data. They suggested that attachment of the modifier to one of the competing NPs is determined mostly by the lexical properties of these NPs. If NP1 "attracts" the modifier more than NP2, then it wins the competition, and the modifier would be attached to NP1. The overall preference observed in a language for either early or late attachment reflects a stronger lexical bias of most NP1 or NP2 tested in the experiments to appear alongside the modifier. Even when the same nouns are used to examine attachment preferences in different languages, these nouns may have different modifier-attracting properties that could tip the scale.

Semantic disambiguation was not studied that extensively as disambiguation by syntactic means; however, there were a couple of experiments in this area. The most relevant study is that of Traxler et al. (1998). They conducted three eye-tracking experiments, in two of which they manipulated semantic plausibility and syntactic agreement in the remaining one (gender). One of their experiments tested sentences like:

(5) a. *The driver of the car with the moustache was pretty cool.*

b. *The car of the driver with the moustache was pretty cool.*

c. *The son of the driver with the moustache was pretty cool.*

(6) a. *The driver of the car that had the moustache was pretty cool.*

b. *The car of the driver that had the moustache was pretty cool.*

c. *The son of the driver that had the moustache was pretty cool.*

(Traxler et al. 1998:563)

They found the clearest effects for the relative clause sentences in the dwell time rather than in first-pass time, same as Carreiras and Clifton (1999). Some previous experiments (e.g., Frazier and Rayner 1982) observed disruption due to syntactic factors (or even some other factors) already in first-pass time or even first fixation time. Also, in Traxler et al. (1998), readers spent less time processing the modifier when both analyses produced a plausible semantic interpretation than when only one analysis produced a plausible semantic interpretation (even when a preferred type of attachment resulted in a plausible interpretation). This contradicts Constraint Satisfaction accounts which claim that increased competition between different possible analyses slows the processor down. However, the authors mentioned a possible confound: Readers, in fact, could have never resolved the ambiguity at all. They might have evaluated both attachment sites, realized that establishing only one host was not possible, and moved on to the next sentences. So as soon as it became apparent that both interpretations were possible, processing stopped. In summary, they found that relative clause modifiers and prepositional modifiers were treated differently, that reading times for sentences with late attached modifiers were just as long as for sentences with early attached modifiers, despite the fact that while answering questions readers preferred to attach modifiers to NP2.

2.2. Russian data

Surprisingly, research on relative clause attachment in Russian is very scarce and mixed. In cross-linguistic discussions in textbooks, Russian is usually claimed to have early attachment preference (mostly referencing Sekerina 2003). Sekerina was apparently the first one to examine attachment preference in Russian by conducting a paper-and-pencil questionnaire requiring judgments about the accessibility of each of two different interpretations of a sentence (on a scale from 0 to 3), like in the example below:

(7)	*Николай*	*хорошо*	*знал*	*сына*	*полковника,*
	Nikolaj.NOM	well	knew	son.ACC	colonel.GEN
	который	*погиб*	*в*	*автомобильной*	
	who.NOM	was killed	in	car	
	катастрофе.				
	accident				

Interpretation 1: 'Nikolaj knew well the son whose father, the colonel, was killed in a car accident.'

Interpretation 2: 'Nikolaj knew well the colonel's son who was killed in a car accident.'

The questionnaire only contained eight experimental sentences (four with a complex NP without a preposition and four with a lexical preposition). She found a significant effect for attachment preference and no effect of preposition (which contradicts the Construal Hypothesis). Sekerina (2003) also conducted a second experiment reported in the same paper using "a whole-sentence reading technique": the subjects had to read a sentence from the computer screen and click the button when finished (with a 9-second timeout for display). Each sentence was followed by a comprehension question. The results agreed with those obtained from the questionnaire: early attachment preference was highly significant, with no effect of preposition. Although she stated that reading time data was collected, she did not report it. On the basis of these findings, Sekerina concluded that Russian, along with Spanish, Italian, French, German, Dutch, and Japanese, was an early attachment language.

Sekerina's (2003) experiments were pioneering and very valuable for the field; however, they had significant limitations. First and foremost, she had only four target items per condition. Even with 30-40 native speakers, this might not be enough to achieve statistically significant results. Moreover, she did not control for animacy: in the prepositionless condition, one item had both an animate and inanimate noun, while in the lexical preposition condition, all nouns were inanimate. In addition, they all were of a very different structure and with complex noun phrases of very different types (e.g., an object and its material like "necklace made of amber" or an object relative to some other object in space like "a room near the corridor"). Such variance coupled with an insufficient number of experimental items may compromise the results. Second, she only got offline measures. Even though she herself labeled the second experiment as "online," it was not online in the current definition of the term: it was not possible to establish first-pass reading times for individual critical areas, probability of regressions, etc.

Anisimov, Fedorova, and Latanov (2014) conducted the first eye-tracking study examining attachment preference in Russian. They compared globally ambiguous sentences with relative clauses (8a) to unambiguous counterparts with only one noun (8b):

(8) a. *В прессе редко упоминали о пианистке солистки, которая часто выезжала на гастроли.*

'The media rarely mentioned the pianist of the soloist who often went on concert tours.'

b. *Студент негромко говорил со своей племянницей, которая ему не нравилась.*

'The student was quietly talking to his niece whom he did not like.'

(Anisimov, Fedorova, and Latanov 2014:525)

Although this was an eye-tracking study, the method of presentation was not perfect: the sentences were split into three fragments located on new lines—the format which is not characteristic of normal reading. They found that "the reading of a fragment in clauses with ambiguity (second strings) slows down by 17 percent, compared to the control sentences without ambiguities" (*ibid.*:528), concluding that ambiguity increases processing load. They also found that the reading time of the first noun in the complex noun phrase in target sentences significantly exceeded the reading time of the second noun (dwell time, without excluding first-pass). However, there was no significant difference in the reading time of the relative pronoun when the subject attached it to the first noun and when they attached it to the second: "the time of reading RPs at early closure (normalized per symbol) proved to be slightly less than at late closure (respectively, 51.9 ± 1.1 against 54.4 ± 1.8 ms per symbol), although such a difference proved to be nonsignificant ($t = -1.20$, $p < 0.232$)" (*ibid.*:527).

Chernova and Chernigovskaya (2015) examined attachment preferences with participle constructions in self-paced reading and eye-tracking experiments:

(8) a. AmbA condition

Svidetel'	*upomjanul*	*naparnika*	*voditelja,*	*pozavčera*
witness	mentioned	workmate.ACC	driver.GEN	yesterday

videvšego	*eto*	*ograblenie.*
having-seen.ACC=GEN	this	robbery

b. LA[11] condition

Svidetel'	*upomjanul*	*o*	*naparnike*	*voditelja,*
Witness	mentioned	about	workmate.PREP	driver.GEN

pozavčera	*videvšego*	*ograblenie.*
yesterday	having-seen.GEN	robbery

c. EA condition

Svidetel'	*upomjanul*	*o*	*naparnike*	*voditelja,*
witness	mentioned	about	workmate.PREP	driver.GEN

pozavčera	*videvšem*	*ograblenie.*
yesterday	having-seen.PREP	robbery

(Chernova and Chernigovskaya 2015:130)

They got quite interesting results in both experiments. For SPRT, they reported "surprisingly many mistakes with the experimental sentences" and that "participants very often ignored the case morphology on the participle" (*ibid.*:131). They did not report the exact accuracy rate for experimental sentences, however, although this is a very important indicator: if the accuracy rate for some condition is at a chance level or even lower, it may compromise online measures, since the participants most likely did not manage to understand the sentence at all (or, in this case, they interpreted all sentences in the same way, thus doing the same amount of processing).

[11] The paper uses the terms "low attachment" (LA) and "high attachment" (HA). For the sake of consistency, they were changed to "late attachment" (LA) and "early attachment" (EA) in this work.

For ambiguous sentences, they stated that in 67.3 percent of cases they were interpreted as EA and concluded that "participants interpreted about two thirds of target sentences as EA paying little attention to case morphology" (*ibid.*:130). Interestingly, though, online measures showed a different picture: reaction times in the participle region were much shorter for LA rather than EA condition. All other measures turned out to be insignificant, and reading times for EA and AmbA sentences virtually coincided in all interest regions. They also reported that AmbA sentences interpreted as LA were read faster than those interpreted as EA.

For the eye-tracking experiment, they reported longer first-pass reading times (that are usually assumed to reflect initial syntactic commitments) in the participle area in EA compared to LA ($p = 0.042$). They did not find any difference in total dwell time on participle across all conditions. During regressions to competing NPs, NP1 was reread twice more often than NP2. Here they also reported an extremely low accuracy rate for the LA condition (38.6 percent), which actually suggests that participants did not understand the sentences. This may also explain why no difference in dwell time was found: if, assuming that the preferred attachment is EA, the participant did not make any extra cognitive work to attach the relative clause to NP2, there should be no difference between EA and LA in this respect. Unfortunately, they did not report dwell time for the whole relative clause, which might have been insightful. If the condition with such an accuracy rate is acceptable at all, the results of their experiment showed no early attachment preference during initial syntactic analysis. They concluded that the results agreed with the Late Closure principle. However, if that is the case, it leaves an open question: why early attached participles had longer first-pass than late attached participles but not ambiguous ones? If the Late Closure is correct, then ambiguous participles should have been equally easily late attached. However, they did not find the above mentioned difference.

Overall, we can argue that to date no conclusive results on attachment preference in Russian have been obtained. All experiments had their own limitations, such as structural variety, insufficient number of experimental items, extremely low accuracy rate, unnatural mode of presentation, or too few measures provided. Early attachment preference in Russian is mostly determined on the basis of offline data (answers to questions about who did what) or of the number of rereadings of the first noun in the noun phrase complex. It is clear that more experiments are needed, which would investigate not only regressions to the noun phrase complex, but also first-pass and dwell time on critical regions.

2.3. Chapter summary

As we have seen, current models fail to fully account for the variance observed in cross-linguistic experiments on relative clause attachment. In the majority of languages examined, either early attachment preference was found or no consistent attachment patterns. The main hypotheses proposed to account for the variance are the Construal Hypothesis, the Tuning Hypothesis, the Relativized Relevance principle, the Predicate Proximity/Recency theory, and the Referentiality principle. Attachment preferences have been also shown to depend on the existence of alternative possessive constructions (e.g., Saxon’s genitive in English), type of preposition, presence of a determiner, and semantic relations between nouns in the noun phrase complex.

As far as Russian is concerned, the evidence is mixed. Sekerina (2003) found a strong early attachment preference in question responses with no effect of a preposition, which contradicts the Construal Hypothesis. Anisimov, Fedorova, and Latanov (2014) also found early attachment preference in question responses and regressions to competing NPs. They also found that ambiguity increased cognitive load (although just in late measures), which is consistent with Constraint Satisfaction accounts. Chernova and Chernigovskaya (2015) found late attachment preference in first-pass on case-marked participles but early attachment preference in question answers and regressions to NPs (however, their accuracy rate for sentences with late attached clauses was unacceptably low), confirming the predictions of the Relativized Relevance principle.

Chapter 3: An Experimental Study

The goal of this chapter is to present the experiment conducted for this study. The chapter provides a detailed description of research questions, the methodology (together with the description of the eye-tracking method in general and its most commonly used measures), subjects, materials and design. It also presents the results of the two conducted pre-tests (frequency and plausibility norming), elaborates on the differences from the previous studies, and discusses the general predictions for the experiment. It ends with presenting the results (both offline and online measures) and a general discussion.

3.1. Research questions

The experiment conducted for this study was designed to answer the following questions:
Q1: Do Russian speakers show a preferred relative clause attachment? If yes, how exactly it is reflected (in early online measures, late online measures, or offline measures)?
Q2: In what fashion and how fast does syntactic and semantic information affect the processing of adjunct modifiers? How different will be syntactically disambiguated sentences from ambiguous sentences, and semantically disambiguated sentences from ambiguous sentences?
Q3: Does syntactic ambiguity itself have processing cost?

3.2. Method

The method chosen for this study was eye-tracking. Eye movements have proven to be extremely informative with respect to moment-to-moment comprehension processes. This informative value, however, rests on two core assumptions explicitly formulated by Just and Carpenter in their early work proposing a model of reading comprehension (1980). The first one is the so-called “eye-mind assumption.” It claims that there is no appreciable delay between what is being fixated and what is being processed. In other words, there is supposed to be a 1:1 mapping between “the eye” and “the mind”: the reader fixates on a word as long as they process it and leaves it right after. This assumption is usually taken for granted by researchers performing eye-tracking[12]. The second assumption is the immediacy principle, which is also usually taken for granted and without an explicit formulation. For Just and Carpenter, it means that each word is interpreted right upon encountering, even in the conditions of partial information when guesses may later turn out to be wrong (*ibid.*:330). To put it differently, the parser adds the incoming words to a tree as soon as possible[13] in order not to accumulate pending elements in the working memory, which would increase cognitive load. In the notion of "interpreting" they included encoding the word, choosing its meaning from the mental dictionary, assigning it to its referent, and integrating the word in the discourse. Now, after accepting that the word is interpreted immediately upon encountering and that eyes fixate on the word as long as it is being processed, the next logical step is to argue that gaze duration on a word is a function of the processing load on the brain. Frazier (1999) distinguished between a weak immediacy principle ("Some interpretation takes place immediately") and a strong immediacy principle ("All interpretations takes place as soon as logically possible, that is, as soon as a possible choice point is encountered") and argued that the latter is incorrect (*ibid.*:35).

Indeed, if, assuming the above principles, one word is fixated longer than another, it means that it requires more cognitive processing (during one or several stages of interpretation). This explains why eye movements are so valuable when assessing linguistic complexity and modelling theories of written language comprehension.

In comparison to word-by-word SPRT in which measured response latencies may be affected by the very pattern of segmentation, eye-tracking experiments investigate a reading process that is as natural as possible. Moreover, the available methods of recording eye movements do not hinder participants' reading rate. The main artificial component is a fixed positioning of the head, although advanced hardware nowadays can remove even this limitation (Traxler and Gernsbacher 2011:613), and (very often) a sentence-by-sentence means of presentation.

Although readers have a subjective experience of a continuous pass over text, eyes do not glide in smooth lines but rather move in a rapid series of jumps (saccades), remaining still[14] in-between (fixations). Meaningful information from the text is only extracted during fixations; the visual system does not register any information picked up by the retina during saccades. On

[12] But see Reichle and Reingold (2013) for interesting findings about a significant amount of parafoveal processing unavoidable during reading.

[13] For more research confirming the incrementality of the parser, see Tanenhaus et. al (1995), Kamide, Altmann, and Haywood (2003), among others.

[14] Technically speaking, even during fixations eyes are not perfectly still but constantly perform miniature movements, such as microsaccades and ocular drifts, controlled by the same mechanisms that generate large saccades (Krauzlis, Goffart, and Hafed 2017).

the average, fixations last around 200-250 ms, with minimum and maximum being 50-100 ms and 500 ms, respectively (Gaskell 2007:327). Saccades are divided into backward (known as regressions) and forward, typically last around 20-40 ms, and span the distance of about 7-9 letter spaces; however, this can vary from one to twenty characters. The percentage of regressions correlate with reading skills: in skilled readers, forward saccades constitute 90 percent of all saccades, and eyes move backwards either to resolve some comprehension difficulty or to correct an error in programming forward saccades (*ibid*.).

In comparison to listening, reading (in natural circumstances) has one very crucial advantage. Readers themselves control the rate of input—that is, they are free to make pauses, reread those parts of a text that are unclear to them, or skip particular words (Just and Carpenter 1980:329). These actions give us valuable information about the level of complexity of the material and the amount of processing load. The duration of both regressions and fixations, in comparison to their size, strongly reflects cognitive processes and correlates with text difficulty: when reading complicated texts, readers tend to make longer fixations and shorter saccades (Traxler and Gernsbacher 2011:615). Not all words are fixated. Short words, extremely frequent words, and words highly predictable from the context tend to be skipped more often; function words (e.g., determiners and prepositions) are skipped more than half of the time (Gaskell 2007:328). Almost all content words are fixated[15]. Importantly, though, just because some words are skipped does not mean that they are not processed. Fisher and Shebilske (1985) conducted an experiment where they removed the frequently skipped words from the text and found that comprehension dropped dramatically. This suggests that people do process words they do not fixate, most likely through parafoveal preview.

As we have seen above, temporal measures are much more informative than spatial when investigating cognitive processes; that is why the majority of commonly used measures are temporal. The most common are as follows:

First fixation duration: the duration of the first fixation in a particular region. It is the earliest measure where we can expect to see some effect of a manipulation, because it corresponds to the first time the reader lands on a region. Longer first fixation duration in one condition relative to another usually suggests that the difficulty was immediate. For example, infrequent words cause longer first fixation durations compared to more frequent words of the same length, because the latter have a high base level of activation and consequently require less additional

15 Rayner et al. (2016) investigated speed-reading and found that, while speed readers made more skips, there was a trade-off between speed and accuracy, and that it was unlikely that readers could double their speed while still being able to understand it.

activation to retrieve them (Just and Carpenter 1980:338). Longer first fixation durations can also be registered when a word disambiguates a sentence toward a dispreferred syntactic analysis (Rayner, Carlson, and Frazier 1983).

Gaze duration: the sum of the durations of all fixations on a word before leaving the word.

First-pass duration: time spent on a region from first entering the region before moving on or looking back. This is a very useful measure for examining early processing: longer first-pass reading time in one condition relative to another usually suggests that the difficulty with this region was immediate.

Regression path duration (also known as *cumulative region reading time*): time spent on a region from first entering the region to first moving the eyes beyond that region to the right, including time spent rereading previous parts of the sentence.

Second-pass duration: the sum of the durations of re-fixations on a region, following the first pass time.

Probability of a regression: the percentage of regressive eye-movements out of a region. This is also called a probability of a first pass regression, because such a regression usually terminates first-pass duration and signals some processing difficulty. However, compared to regression path duration, this measure may not always be informative, because time spent on rereading an earlier portion of text is not taken into account. The reader may regress the same number of times in different conditions but spend longer time rereading an earlier portion of text (Liversedge, Paterson, and Pickering 1998:59).

Total dwell time: the sum of the durations of all fixations on a region including rereadings. Unlike with first-pass duration, longer total dwell time in one condition relative to another usually suggests that the effect of some linguistic manipulation on processing is relatively late.

3.3. Subjects

Thirty-three Russian native speakers (21 females, mean age 24.5, median age 23, range 18-36 years, SD = 5.3) took part in the experiment on a voluntary basis. Distribution by country: 12 from Russia, 10 from Ukraine, 7 from Belarus, 2 from Uzbekistan, 1 from Azerbaijan, and 1 from Kazakhstan. All participants reported that they were born to Russian-speaking parents and acquired Russian as their first language. The data from three participants had to be excluded because of very poor calibration and from another two because of very low overall accuracy (< 65 percent). In the end 28 people were analyzed (19 females, mean age 24.5, median age 23, range 18-36 years, SD = 5.3). All participants had normal or corrected vision. All participants were naïve with respect to the purpose of the study. Each participant gave a written consent to the participation in the experiment. The experiment was conducted in the eye-tracking lab of the Institute of English Studies of the University of Wrocław.

3.4. Materials and design

Stimuli. The current experiment consisted of 75 experimental items and 85 fillers. All experimental sentences had the following structure:

> Schematic: *NP1 — Verb — NP2 — NP3 — Relative Pronoun — Verb — NP4*
>
> Detailed: *NP1 (proper name) — Verb (transitive, perfective) — NP2 (human, animate, accusative) — NP2 (human, animate, genitive) — Relative Pronoun — Verb (transitive, perfective) — NP4 (accusative)*

Fillers were of varied syntactic structure but all contained a relative pronoun to hide the purpose of the experiment. All experimental sentences, including fillers, were followed by a question that forced the participant to choose between two nouns mentioned in the sentence. Correct answers to questions were pseudorandomized (the correct answer could appear on both the right and the left side equally frequently). Apart from that, two different test versions with “mirrored” location of answers were created, with all the subjects randomly assigned to one of the two test versions.

Experimental conditions were as follows:

- Syntactic Early Attachment (SynEA)
- Syntactic Late Attachment (SynLA)
- Ambiguous Attachment (AmbA)
- Semantic Early Attachment (SemEA)
- Semantic Late Attachment (SemLA)

In the *Syntactic* conditions, nouns were of different gender and the relative clause was unambiguously attached to one of them by gender-marking on the relative pronoun. The gender of the noun to which the relative clause attached was always male, in order to eliminate any effects of markedness present in feminine nouns. The relative pronoun immediately followed the noun phrase complex, eliminating the problem with higher-level processes intervention raised by Frazier (1990). It is thus most likely that the first-pass time on the relative pronoun would reflect initial syntactic commitments, since no higher-level interpretation process would have time to kick in.

(9) a. *Сеня окликнул* ***ученика*** *скрипачки,* ***который***

Senia called **student.M** violinist.F **who.M**

уронил ноты. *(Early Attachment)*

dropped.M notes

'Senia called **the student** of the violinist **who** dropped his notes.'

b. *Юра подозвал сиделку* ***пенсионера****,* ***который***

Jura called sitter.F **pensioner.M** **who.M**

пропустил завтрак. *(Late Attachment)*

missed.M breakfast

'Jura called the sitter of **the pensioner who** missed his breakfast.'

In the *Semantic* conditions, nouns were of the same gender but the relative clause was semantically biased on the verb towards one of them. Given that the two nouns in the noun phrase complex were always animate because there was no possibility to counterbalance for animacy in Russian (inanimate nouns always precede animate nouns in noun phrase complexes), semantic plausibility was a less reliable disambiguator than in Traxler, Pickering, and Clifton (1998), where drivers having moustaches were compared to cars having moustaches. The verb was chosen in such a way that it was strongly associated with one of the two nouns (such as "prescribed" and "doctor", or "delivered" and "courier"). This should provide a good constraint, if Constraint Satisfaction models are on the right track, because they emphasize the importance of lexical association (MacDonald, Pearlmutter, and Seidenberg 1994:697-698). Theoretically, it may be possible to see some early effect already on the verb.

(10) a.

Вера	*закутала*	***ребёнка***	*соседа,*	*который*
Vera	wrapped	**child.M**	neighbor.M	who.M
выплюнул	*соску.*		*(Early Attachment)*	
spit-out.M	pacifier			

‘Vera wrapped **the child** of the neighbor who **spit out** the pacifier.’

b.

Вика	*вызвала*	*отчима*	***школьника****,*	*который*
Vika	invited	stepfather.M	**pupil.M**	who.M
прогулял		*уроки.*	*(Late Attachment)*	
skipped.M		lessons		

‘Vika invited the stepfather of **the pupil** who **skipped** school.’

The *Ambiguous* condition served as a control for the above two groups. In this condition, the relative clause was semantically unbiased and syntactically ambiguous. In the statistical analysis, syntactically disambiguated sentences (early and late attached) would be compared to ambiguous sentences, and semantically disambiguated sentences would also be compared to ambiguous sentences. This would keep the disambiguation point constant (relative pronoun in the first group and verb in the second). AmbA should also reveal participants' true attachment preferences that persist in the absence of any constraints (judging by their question answers). For a full list of experimental sentences and fillers, see Appendix 1.

(11)

Катя	*увидела*	*ассистента*	*лектора,*	*который*
Katia	saw	assistant.M	lecturer.M	who.M
обронил	*ключи*			
dropped.M	keys			

‘Katia saw the assistant of the lecturer who dropped the keys.’

Thus, two target groups were identified:

Table 2. Groups of experimental sentences

Group	**Conditions in the group**	**Translated example sentences (* point of disambiguation)**	**Attachment**
Syntactic disambiguation	SynEA	Senia called the student of the violinist **who*** dropped the notes.	Early
	SynLA	Jura called the sitter of the pensioner **who*** missed breakfast.	Late
	AmbA	Katia saw the assistant of the lecturer **who** dropped the keys.	Ambiguous
Semantic disambiguation	SemEA	Vera wrapped the child of the neighbor who **spit-out*** pacifier.	Early
	SemLA	Vika invited the stepfather of the pupil who **skipped*** lessons.	Late
	AmbA	Katia saw the assistant of the lecturer who **dropped** the keys.	Ambiguous

All verbs in the relative clauses consisted of 2-3 syllables (6-9 characters) and nouns of 1-3 syllables (4-7 characters).

3.5. Pre-Tests

3.5.1. Frequency

As shown by Just and Carpenter (1980), the frequency of words affects their first fixation duration: eyes stay longer on less frequent words because they need more activation to be retrieved from the mental dictionary. This experiment manipulated both gender on the relative pronoun and semantic bias on the verb. Since relative pronouns were always the same, there was no need to check their frequency, unlike the frequency of verbs. To balance them, frequency per million words in the Russian National Corpus was used, taken from Lyashevskaya and Sharov (2009). All frequencies are provided in Appendix 2. One-way ANOVA was highly insignificant: $F(4, 70) = 0.652$, $p = 0.617$, which means that if longer reading times on the verbs are found, they could not be attributed to unbalanced frequency and reflect other processes.

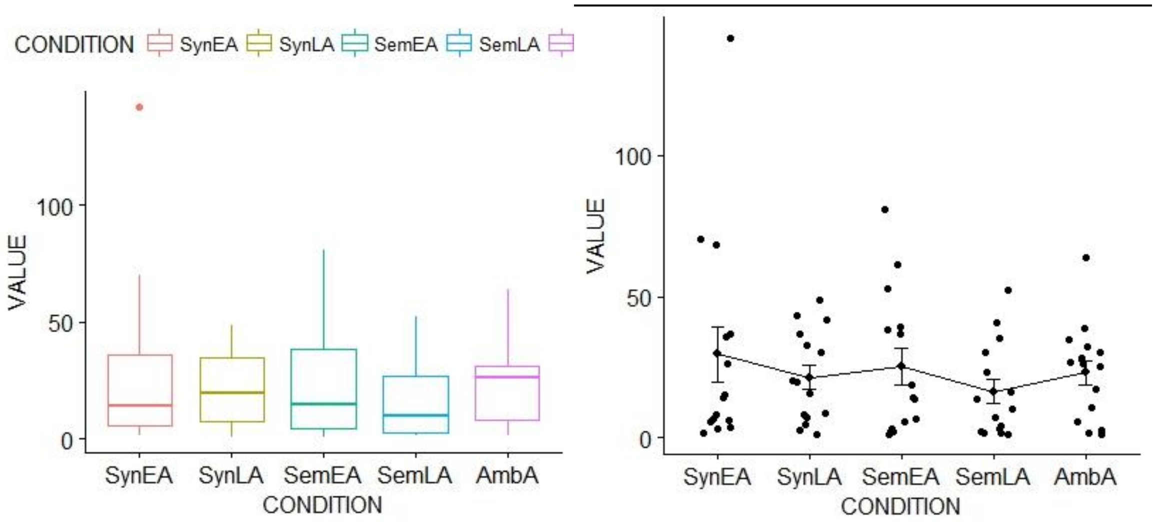

Figure 2-3. One-way ANOVA for verb frequencies

3.5.2. Plausibility norming

AmbA condition was constructed to discover participants' preferences in the absence of any constraints. Thus, to guarantee that in AmbA sentences the relative clause could be plausibly attached to both nouns, a plausibility norming study was carried out. Two separate online questionnaires with 15 sentences were created, which contained a combination of the noun phrase (either NP1 or NP2, distributed 50/50 between two questionnaires) and a relative clause from the AmbA condition. The order of sentences was randomized. One hundred and two native speakers (51 per each questionnaire) who did not take part in the main study assigned them a plausibility score from 1 to 7 (7 = most plausible):

- Questionnaire 1: *Ассистент обронил ключи* ("The assistant dropped the keys") (see (11))
- Questionnaire 2: *Лектор обронил ключи* ("The lecturer dropped the keys") (see (11))

No significant differences were found (according to the Welch's two-sample *t*-test): $F = 0.878$, $p = 0.3742$, (the mean of group "NP1" = 5.59; the mean of group "NP2" = 5.84) (the norming data is presented in Appendix 3).

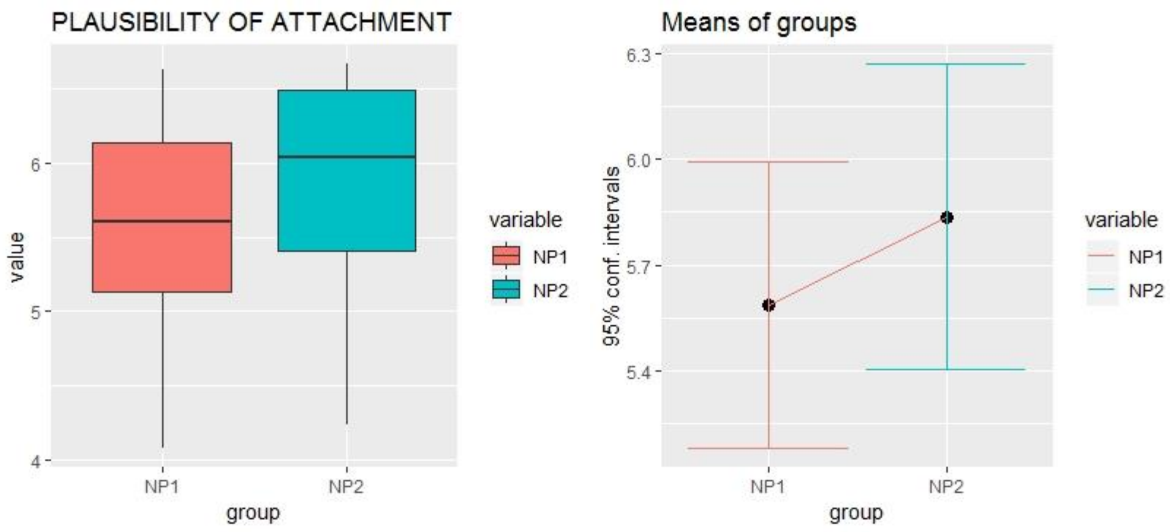

Figures 4-5. Welch's two-sample *t*-test comparing plausibility of sentences with relative clauses attached to either NP1 or NP2

3.6. Differences from the previous studies

The experiment contained several important modifications that made it different from those already conducted on the Russian data.

First and foremost, neither of the experiments on the Russian data manipulated semantic plausibility. Sekerina (2003) manipulated the prepositions in the complex noun phrases; Chernova and Chernigovskaya (2015) manipulated the case on the participles; and Anisimov, Fedorova, and Latanov (2014) only compared sentences with two animate nouns to unambiguous sentences with one inanimate noun.

Second, all verbs, most importantly all verbs in the relative clauses, were telic (in the past perfective form). Many studies either do not report whether verbs were of the same aspect or it is clear from the examples that they were not.

Third, all experimental sentences were of a strict uniform structure.

Also, in comparison to most cross-linguistic experiments on relative clause attachment as well as other types of syntactic ambiguities, semantics here was a less reliable disambiguator. In examples like "evidence examined" (see Rayner, Carlson, and Frazier, 1983), the garden-pathing analysis does not really belong to our world knowledge: evidence simply cannot examine anything. In the Semantic condition in this experiment, the alternative analysis in some cases was more or less plausible:

(12) a. *Vitia complimented the cook of the captain who cooked a broth.*

b. *Zhora shocked the neighbor of the poacher who shot a deer.*

In (12a), although it makes more sense if the cook was the one who cooked a broth, nothing prevented the captain to do it. In (12b), the neighbor could have well been the one who shot a deer, although the poacher fits the description better. Thus, it is interesting to see how fast the participant would interpret both types of sentences and how the answers to questions will be distributed.

The experiment was also as close to the natural reading process as possible: unlike many previous experiments on attachment that segmented the sentences into several strings, which by itself may have produced disruption, the sentences were displayed in full and the participants did not have any time restrictions.

3.7. General predictions

All experimental items were divided into 6 interest areas: the beginning of the sentence (1), NP1 (IA 2), NP2 (IA 3), relative pronoun (IA 4), verb (IA 5), and final noun (IA 6). The beginning of the sentence was not analyzed; all other areas consisted of one word. In general, two critical regions were distinguished as the most important: IA 4 and IA 5 (marked by an asterisk).

Table 3. Interest areas

Condition/IA	**1**	**2**	**3**	**4***	**5***	**6**
Syntactic	Сеня подозвал Senia called	ученика student.M	скрипачки violinist.F	который who.M	уронил dropped.M	ноты notes
Semantic	Вера закутала Vera wrapped	ребёнка child.M	соседа neighbor.M	который who.M	выплюнул spit out.M	соску pacifier
Ambiguous	Катя увидела Katia saw	ассистента assistant.M	лектора lecturer.M	который who.M	обронил dropped.M	ключи keys

Models differs as for their predictions of processing difficulty for different groups of items.

3.7.1. Group 1: Syntactic Disambiguation

The Garden-Path Model predicts that readers, following the Late Closure strategy, would always initially attempt to attach the modifier to NP2 (during the first stage of parsing), then the thematic processor during the second stage should evaluate this analysis for correctness. Reading times should be faster if the parser's initial analysis produces a plausible interpretation and longer when it does not. Thus, this area in SynEA should demonstrate longer first-pass time because the initial analysis of the parser produces an incorrect result. According to this model, the parser is initially guided strictly by word categories, so, when it reaches the relative pronoun in SynEA, it first attaches it to the second noun in the noun phrase complex:

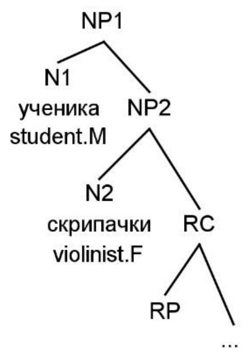

Figure 6. Predictions of the Garden-Path Model

When the parser detects an error and has to do a reanalysis, first-pass reading times and regressions increase (Rayner et. al. 1983; Mitchell and Holmes 1985; Ferreira and Clifton 1986; Trueswell et al. 1993; 1994; Tanenhaus et al. 1995; Clifton et al. 2003). In all other conditions this would not happen: SynLA agrees with the Late Closure strategy, whereas AmbA is ambiguous on the relative pronoun and thus the RP should be attached to NP2 without extra difficulties. Being modular, the Garden-Path Model assumes that the parser always utilizes at least three strategies (Minimal Attachment, Late Closure, and Active Filler) and constructs only one analysis at a time. It is clear that this model predicts no increase in reading time for ambiguous parts of the sentence—it is only the reanalysis that is costly and causes regressions and longer reading times.

We may thus observe:

1) longer reading time in SynEA on IA 4 than in SynLA and AmbA;

2) more regressions in SynEA on IA 4 than in SynLA and AmbA.

Although for IA 5 the Garden-Path Model does not give clear prediction, since the disambiguation point occurred (or did not occur) earlier, it is reasonable to assume that we may see the same effects on the verb as on the relative pronoun. For example, Traxler, Pickering, and Clifton (1998) found the sharpest effect in regressions from the area immediately following the critical (disambiguating) region. We may thus observe:

1) longer reading times and more regressions on IA 5 in SynEA than in SynLA and AmbA.

Since the model predicts that in AmbA participants should attach the modifier to NP2, this should also be reflected in their answers to questions.

Constraint Satisfaction models predict that online processing will reflect offline preferences. Given the previous experiments conducted on Russian (Sekerina's questionnaires (2003) and answers to sentences with ambiguous modifiers in Anisimov, Fedorova, and Latanov 2014 and Chernova and Chernigovskaya 2015) it is reasonable to assume that Russian speakers exhibit early attachment preference for adjunct modifiers, which most likely reflects the higher frequency of early attachment constructions in Russian. Constraint Satisfaction models assume a ranked parallel analysis of parsing (constructing multiple analyses and ranking them according to the strength of various clues). They claim that if different analyses get similar amount of activation, the processing is slowed down. Because in SynLA their preferred analysis competes with the actual attachment, the processing should be more difficult. In AmbA, nothing indicates where the relative pronoun must be attached. Thus, pronouns in AmbA should be harder to process than syntactically disambiguated pronouns in SynEA and SynLA. Thus, the following should be observed:

1) longer reading time and more regressions in SynLA than in SynEA on IA 4;
2) longer reading time and more regressions in AmbA than in SynLA and SynEA on IA 4.

For IA 5, these accounts again predict longer reading times for SynLA than for SynEA. They predict even harder processing in AmbA since all possible analyses received even more activation. We should thus see the same effect for IA 5 as for IA 4 but greater in size.

The Construal Hypothesis treats relative clauses as non-primary phrases and thus assumes that they are not processed using the universal parsing principles. Both noun phrases here lie within the active theta domain while the adjunct modifier is processed, which means that both sites should be evaluated as hosts for the modifier simultaneously while the modifier is being processed. The parser's actions should be as follows: access the current word category, identify that it does not belong to the subject and main predicate of any finite clause as well as complements and obligatory constituents of primary phrases, and then (in case of a relative pronoun) associate it to either the entire theta domain, evaluating both hosts on the way. For AmbA, the Construal Hypothesis does not give any clear predictions on how the final attachment should be chosen since no structural factors that could affect it have been used, such as determiners (which make the noun more discourse-prominent and thus a more preferred host for attachment) and pronouns (which may restrict theta-domain). Russian also does not have

Saxon's genitive which is claimed to play a role in RC attachment resolution. Thus, this model is mute as to what attachment the participants will choose in AmbA. When the relative pronoun is unambiguous, there should be no difference between SynLA and SynEA, since both attachment sites are evaluated simultaneously and syntax clearly dictates the host. We should thus see:

1) no difference between SynEA and SynLA on IA 4 and IA 5.

The Relativized Relevance principle formulated by Frazier (1990) claims that early attachment is determined by rather late processes, disguising the results of the Late Closure strategy. Thus, at some point in the RC in AmbA the thematic processor must come into play and reverse participants' preferences from late attachment to early attachment, resulting in the same pattern in question answers. As far as disambiguated items are concerned, this principle does not predict any differences between them. It also does not say anything about the processing cost of ambiguity, so whether AmbA would be any different from the rest is unclear.
The Unrestricted Race model (and to some extent **the Underspecification model**) predicts lower processing load in AmbA than in both SynEA and SynLA, claiming that in such cases all our interpretations are correct and the parser does not have to reanalyze. It does not predict any differences between SynLA and SynEA, because they are disambiguated immediately and the parser should make no mistakes.

3.7.2. Group 2: Semantic Disambiguation

Since the relative pronoun is ambiguous in all conditions in this group (AmbA, SemEA, SemLA), all models predict no difference in processing of IA 4.
The **Garden-Path Model** predicts that on IA 5 there should be visible difficulty in SemEA: it should first be added to the same late attached branch (because the processor is only guided by word categories in the first step) and then reanalyzed at stage 2, when the processor detects a disagreement between the initially chosen attachment and the most plausible attachment. We may thus observe:

1) longer reading times on IA 5 in SemEA in comparison to SemLA and AmbA.

The Relativized Relevance principle gives mixed predictions. It is hard to say for sure when the thematic processor responsible for early attachment should kick in. In this experiment, there is one ambiguous word between the last noun in the complex noun phrase and the semantically biased verb. In theory, the thematic processor might be able to start its operation. We should thus observe:

1) longer reading times in SemLA than in SemEA.

If the processor did not kick in, however, the opposite should be true (because the parser is assumed to initially followed the Late Closure strategy):

1) longer reading times in SemEA than in SemLA.

Again, it is unclear whether AmbA would be any different from the rest, apart from the fact that the principle predicts clear early attachment preference in offline measures (question answers).

Constraint Satisfaction accounts predict fierce competition on the relative pronoun in all conditions. If reading time reflects offline preferences, then early attachment should be finally ranked the highest. Then, when the verb is reached, we should see a reranking in SemLA, when semantics begins to support another analysis, resulting in longer reading times. AmbA should still be the most difficult, since multiple analyses continue to receive equal amount of support. We should thus observe:

1) longer reading times and more regressions in SemLA than in SemEA;
2) longer reading times and more regressions in AmbA than in SemLA and SemEA.

However, as noticed by Traxler, Bybee, and Pickering (1997), when disambiguation is carried out by grammatical means, the reader's response can be very quick, whereas when the sentence is disambiguated by inferencing or semantic interpretation, the response can be much slower. The eyes may move to the next region and only then come back to the critical region or even earlier regions, which will result in disruption in late measures (second-pass, dwell time, total number of regressions) than in first-pass effects. We may thus expect less clear difference between early attachment and late attachment in the Semantic condition, with longer reading times in general and more regressions.

Also, different models give different predictions for regressions to competing NPs (for both groups, Syntactic Disambiguation and Semantic Disambiguation).

- **IA 2 (NP1) and IA 3 (NP2)**

For disambiguated sentences, we may assume two different scenarios:

1) the participants would regress more to that NP that agrees with the actual relative clause attachment;

2) the participants would regress more to that NP attachment to which they naturally prefer, making more regressions in general in those sentences where there is a disagreement between the actual attachment and their preferred attachment.

In AmbA sentences, we may expect:

1) more frequent regressions to NP2 (**the Garden-Path Model**);

2) more frequent regressions to NP1 (**the Relativized Relevance principle**);

3) more frequent regressions to NP1 or equally frequent regressions to both NPs (**the Constraint Satisfaction models**);

4) equally frequent regressions with no preference (**The Unrestricted Race model)**

Accuracy should also vary according to the preferred attachment (with less preferred analysis being more error prone).
Summary of the predictions is provided overleaf.

Table 4. Predictions by different parsing models

Group 1 (Syntactic disambiguation)			
IA/Model	**Garden-Path Model**	**Constraint Satisfaction models**	**Relativized Relevance principle**
IA 4 (RP) and IA 5 (Verb)	1) longer reading time and more regressions in SynEA than in SynLA and AmbA; 2) no difference between SynLA and AmbA; 3) late attachment preference for AmbA reflected in questions answers and more regressions to NP2.	1) longer reading time and more regressions in SynLA than in SynEA; 2) longer reading time and more regressions in AmbA than in both SynLA and SynEA; 3) more regressions to NP1 in SynLA and SynEA; 4) more regressions to NP1 (or to both NPs) in AmbA.	1) no difference between SynLA and SynEA; 2) early attachment preference for AmbA reflected in questions answers and more regressions to NP1.
Group 2 (Semantic disambiguation)			
IA 4 (RP)	1) no difference	1) no difference	1) no difference
IA 5 (Verb)	1) longer reading time and more regressions in SynEA than in SynLA and AmbA; 2) no difference between SynLA and AmbA; 3) late attachment preference for AmbA reflected in questions answers and more regressions to NP2.	1) longer reading time and more regressions in SemLA than in SemEA; 2) longer reading time and more regressions in AmbA than in both SemLA and SemEA.	1) longer reading time and more regressions in SemLA than in SemEA if the thematic processor managed to start its operation, or longer reading time and more regressions in SemEA than in SemLA if it did not; 2) early attachment preference for AmbA reflected in questions answers and more regressions to NP1.

3.8. Apparatus

The experiment was conducted on an SR Research EyeLink 1000 Plus, with a sampling rate of 2000 Hz. Viewing was binocular but only one dominant eye (determined by a simple test with a kaleidoscope) was tracked. Experimental items were displayed on a 24-inch BenqXL monitor in a single line in black monospaced font on a white background. Participants were seated 61 cm from the screen, their head was immobilized by means of a head and a chin rest to minimize head movements. At this distance, 3.7 characters subtended 1 degree of visual angle. The resolution of the eyetracker was < 0.01º RMS.

3.9. Procedure

All participants were instructed to silently read the sentences and to respond to comprehension questions after every sentence. The experiment started with the nine-point calibration procedure followed by a training session consisting of 8 sentences. The acceptable calibration error was less than 0.5 degrees of visual angle. The main experiment consisted of four blocks with 40 sentences per block, presented in an individually pseudorandomized order. The calibration procedure was repeated before each block. Every sentence was preceded by a drift correction, after which a fixation point (a gray square) appeared at the location of the first letter of the sentence. Once the participant made a stable fixation on this point (or 2,000 ms after the fixation point appeared), the sentence was displayed. After reading the sentence, the participant had to press a button, and a question appeared on the screen. The participant had to choose between two sentence interpretations. Choosing an answer triggered the next sentence to appear. No time limit was imposed in order to make the reading process as natural as possible and not to stress the participants.

3.10. Results

The statistical analysis was done using the SPSS software.

Offline measures

Accuracy

Below is the accuracy rate for all conditions, excluding AmbA, where no correct answer was assumed and which will be reviewed separately.

Table 5. Accuracy rate

Condition	Mean accuracy, percent	SD
Syntactic Early Attachment (SynEA)	99.0	2.4
Syntactic Late Attachment (SynLA)	96.3	5.9
Semantic Early Attachment (SemEA)	97.4	4.5
Semantic Late Attachment (SemLA)	92.4	8.3
Fillers	95.4[16]	4.3

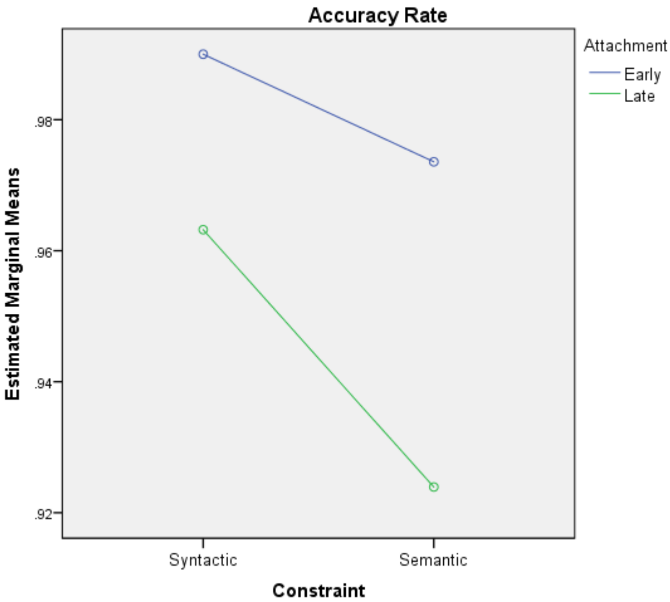

Figure 7. Accuracy rate

The two-way rANOVA (**Constraint** (Syntactic, Semantic) × **Attachment** (Early, Late)) showed main effects of both **Constraint** ($F(1, 27) = 5.857$, $p = 0.023$, $\eta^2 = 0.178$) and **Attachment** ($F(1, 27) = 11.547$, $p = 0.002$, $\eta^2 = 0.300$), with no interaction ($F(1, 27) = 1.916$, $p = 0.178$, $\eta^2 = 0.066$). As expected, the highest accuracy was obtained in SynEA: only 4 participants made one mistake each, the rest had 100 percent correct answers. Also predictably, although still interestingly, second best was SemEA and not SynLA.

SemLA showed the lowest accuracy apparently because initial syntactic commitments in many cases were too strong to be overridden just by semantics, even though the participants did not have any time restrictions and could read both the sentence and the question for as long

[16] The accuracy rate for fillers was affected by several items of the following type:

В клубе стояли автоматы и бильярд, который заинтересовал Пашу.
in club stood slot machines.PL and billiard.SG.M that.SG.M interested.SG.M Pasha.
'In the club there were slot machines and a billiard that Pasha found interesting.'

Even though the morphology on nouns unambiguously indicated to which of them the relative clause should be attached, the participants made surprisingly many mistakes, with only occasional correct answers. This suggests that they did not split the coordinated NP into two different nouns but rather treated it as a single unit, attaching the relative clause to it as a whole ("slot machines and a billiard"). Why this is the case is very interesting, though unfortunately outside the scope of this work.

as they needed. Strictly speaking, questions in SemEA and SemLA could have also been answered logically even without reading the sentence itself, e.g.

(Igor heard out the patient of the oculist who prescribed the glasses)

Who prescribed the glasses?

Patient *Oculist*

However, the participants did not do this, which is reflected in how many mistakes they made. It suggests that they were reading the sentence and were misled by their default syntactic preference.

Interestingly, though, the distribution of answers for AmbA sentences was 51.6 percent (SD 20.6) towards early attachment and 48.4 percent towards late attachment, thus failing to show a preference and disagreeing with all the previous studies. These results, although agreeing with the unrestricted race model, most likely were caused by the experimental design: previous experiments on Russian compared syntactically ambiguous sentences containing noun phrase complexes either to unambiguous sentences containing just one noun (Anisimov, Fedorova, and Latanov 2014) or to syntactically unambiguous sentences containing noun phrase complexes (Chernova and Chernigovskaya 2015), which means that participants did not have to thoroughly consider semantics while establishing preference in syntactically ambiguous sentences. In this experiment, however, in two Semantic conditions the relative clause was biased toward one of the two nouns, being at the same time syntactically ambiguous. This may be the reason why, when reading fully ambiguous sentences with no clues whatsoever (syntactic or semantic), the participants felt lost and replied at random (or were trying to find at least some clue, being aware of the ambiguity).

What is also interesting is that attachment preference in AmbA sentences differed depending on the participant's country of origin. The samples for Uzbekistan, Kazakhstan, and Azerbaijan were too small to be representative; however, the samples for Russia and Ukraine (and Belarus to some extent) were comparable. Unexpectedly, speakers from Russia showed a clear late attachment preference, whereas speakers from Ukraine showed the opposite pattern. The heterogeneity of the population may have affected both offline and online measures in this experiment and is an interesting topic for further investigation.

Table 6. Attachment preferences by participants' country of origin

Country (Number of participants)	Early attachment, %	Late attachment, %	SD
Russia (9)	39	61	18.2
Ukraine (8)	57	43	15.8
Belarus (7)	48	52	20.6
Uzbekistan (2)	74	26	—
Kazakhstan (1)	67	33	—
Azerbaijan (1)	93	7	—

Online measures

Spaces between words were divided in half, with each half added to the adjacent interest area. Outliers (2SD from the mean per interest area per condition per participant) were removed, with the percentage of trials removed not exceeding 1 percent of all trials per each interest area. Sentences to which incorrect responses were provided were excluded from the final analysis (0.01 percent of sentences in SynEA, 0.03 percent of sentences in SynLA, 0.02 percent of sentences in SemEA, and 0.08 percent of sentences in SemLA). Because no correct answer was assumed in AmbA, no items were excluded from the final analysis in this respect. Also, fixations shorter than 70 ms were excluded, because, as shown by Rayner and Pollatsek (1989), readers normally do not extract much information during very short fixations.

The analysis was carried out in the following manner. Section 3.10.1 presents the results of one-way rANOVA for SynEA, SynLA, and AmbA—that is, for Group 1, comparing syntactically disambiguated clauses with fully ambiguous controls with the main factor of **Attachment** (early, late, ambiguous). Section 3.10.2 presents the results of one-way rANOVA for SemEA, SemLA, and AmbA—that is, for Group 2, comparing semantically disambiguated clauses with fully ambiguous controls with the main factor of **Attachment** (early, late, ambiguous). Section 3.10.3 presents some analyses of the AmbA condition separately.

Five measures are reported below: *first-pass time* (summation of the duration of all fixations on the interest area from first entering the area before moving on or looking back), *second-pass time* (summation of the duration of all fixations on the interest area from entering the area second time before moving on or looking back), *dwell time* (summation of the duration of all fixations on the interest area), *regression in count* (number of times the interest area was entered from an interest area with a higher ID), and *regression out count* (number of times the interest area was exited to an interest area with a lower ID before an interest area with a higher ID was fixated). First-pass belongs to early measures and is supposed to reflect initial syntactic commitments, whereas the rest belong to late measures. Below are all the above measures for every interest area, with SD given in the brackets.

Table 7. Mean first-pass time (in milliseconds)

Condition	NP1	NP2	RP	V	NOUN
SynEA	382 (135)	390 (135)	224 (44)	402 (124)	260 (76)
SynLA	319 (103)	414 (131)	219 (49)	413 (144)	259 (91)
SemEA	352 (127)	372 (107)	216 (39)	411 (114)	251 (95)
SemLA	325 (98)	403 (139)	224 (53)	406 (95)	250 (86)
AmbA	365 (109)	399 (157)	215 (37)	401 (131)	270 (107)

Table 8. Mean second-pass time (in milliseconds)

Condition	NP1	NP2	RP	V	NOUN
SynEA	292 (76)	273 (72)	224 (55)	298 (85)	281 (151)
SynLA	277 (62)	271 (64)	234 (51)	279 (65)	229 (86)
SemEA	304 (79)	291 (57)	204 (30)	310 (96)	260 (77)
SemLA	295 (75)	295 (96)	209 (33)	291 (85)	251 (175)
AmbA	308 (65)	290 (86)	220 (44)	287 (98)	228 (78)

Table 9. Mean dwell time (in milliseconds)

Condition	NP1	NP2	RP	V	NOUN
SynEA	1017 (354)	848 (341)	499 (135)	717 (219)	354 (139)
SynLA	899 (475)	928 (397)	557 (174)	767 (258)	349 (168)
SemEA	1043 (416)	926 (284)	492 (126)	759 (224)	324 (146)
SemLA	994 (428)	1052 (395)	520 (161)	744 (259)	323 (132)
AmbA	1239 (446)	1191 (435)	624 (197)	841 (301)	354 (145)

Table 10. Mean number of regressions to areas

Condition	NP1	NP2	RP	V	NOUN
SynEA	1.15 (0.56)	0.42 (0.30)	0.47 (0.07)	0.29 (0.19)	0 (0)
SynLA	0.91 (0.41)	0.55 (0.32)	0.60 (0.06)	0.29 (0.24)	0 (0)
SemEA	1.35 (0.57)	0.51 (0.35)	0.37 (0.07)	0.27 (0.23)	0 (0)
SemLA	1.19 (0.54)	0.68 (0.50)	0.40 (0.05)	0.18 (0.16)	0 (0)
AmbA	1.61 (0.77)	0.80 (0.48)	0.47 (0.06)	0.36 (0.26)	0 (0)

Table 11. Mean number of regressions out of areas

Condition	NP1	NP2	RP	V	NOUN
SynEA	0.11 (0.12)	0.39 (0.29)	0.02 (0.07)	0.25 (0.34)	1.09 (0.30)
SynLA	0.10 (0.12)	0.32 (0.28)	0.01 (0.05)	0.21 (0.24)	1.10 (0.38)
SemEA	0.06 (0.11)	0.43 (0.35)	0.04 (0.05)	0.18 (0.28)	1.09 (0.26)
SemLA	0.05 (0.11)	0.37 (0.39)	0.03 (0.10)	0.30 (0.35)	1.10 (0.33)
AmbA	0.07 (0.12)	0.32 (0.31)	0.01 (0.05)	0.20 (0.28)	1.17 (0.24)

3.10.1. Group 1: Syntactic Disambiguation

Below are the results for syntactically disambiguated items vs. ambiguous controls for areas of interest. All measures were subject to rANOVA with the main factor of Attachment (early, late, ambiguous), with a separate analysis for each IA. $p < 0.05$ is marked with one asterisk, $p < 0.01$ is marked with two asterisks, $p < 0.001$ is marked with three asterisks.

Table 12. rANOVA results for SynEA, SynLA, AmbA

Measure	IA	df	Attachment		
			F	***p***	***η²***
First-pass	4	2, 54	1.006	.372	.036
		2, 42	1.836	.172	.080
	5	2, 54	.432	.651	.016
		2, 42	.261	.771	.012
	6	2, 54	.760	.473	.027
		2, 42	.414	.664	.019
Second-pass	4	2, 54	1.887	.161	.065
		2, 42	2.884	.067	.121
	5	2, 54	.722	.491	.026
		2, 42	.797	.458	.037
	6	2, 54	1.618	.209	.063
		2, 42	1.273	.291	.057
Dwell time	4	2, 54	8.325	.002**	.236
		2, 42	8.154	.001**	.280
	5	2, 54	4.353	.030*	.139
		2, 42	4.640	.015*	.181
	6	2, 54	.058	.913	.002
		2, 42	.132	.877	.006
Regressions in	2	2, 54	25.027	.000***	.481
		2, 42	32.336	.000***	.606
	3	2, 54	12.960	.000***	.324
		2, 42	23.627	.000***	.529
	4	2, 54	4.721	.013*	.149
		2, 42	6.316	.004**	.231
	5	2, 54	2.330	.107	.079
		2, 42	1.593	.215	.071
	6	2, 54	0	0	0
		2, 42	0	0	0
Regressions out	4	2, 54	.116	.891	.004
		2, 42	.691	.507	.032
	5	2, 54	.812	.449	.029
		2, 42	.631	.537	.029
	6	2, 54	1.377	.261	.049
		2, 42	1.505	.234	.067

IA 4 (relative pronoun)

First-pass on the relative pronoun did not show a significant effect ($F_1(2, 54) = 1.006$, $p = 0.372$, $\eta^2 = 0.036$; $F_2(2, 42) = 1.836$, $p = 0.172$, $\eta^2 = 0.080$). These results agree with some other experiments (e.g., Carreiras and Clifton 1999) and highlight the more complicated nature of immediate disambiguation, suggesting that, even if the parser made an error in the initial parse, it was corrected too quickly to be detected.

Second-pass revealed a marginally significant main effect by items ($F_1(2, 54) = 1.887$, $p = 0.161$, $\eta^2 = 0.065$; $F_2(2, 42) = 2.884$, $p = 0.067$, $\eta^2 = 0.121$), with early attached pronouns being read much faster than late attached (see Figure 8).

Dwell time showed a highly significant main effect ($F_1(2, 54) = 8.325$, $p = 0.002$, $\eta^2 = 0.236$; $F_2(2, 42) = 8.154$, $p = 0.001$, $\eta^2 = 0.280$). Participants spent significantly less time on IA 4 in SynEA than in AmbA ($p_1 = 0.001$, $p_2 = 0.001$) and marginally less than in SynLA ($p_1 = 0.062$, $p_2 = 0.123$) (see Figure 9).

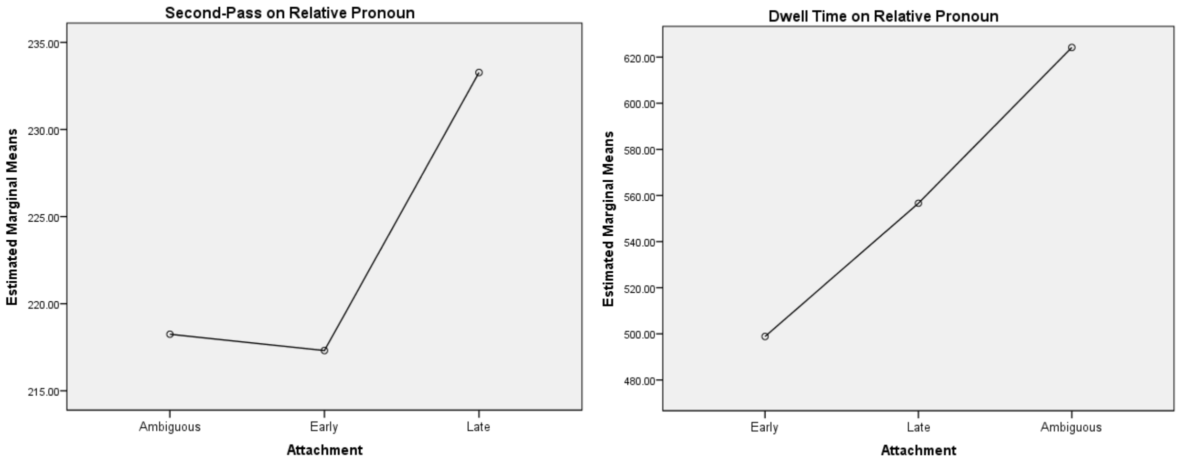

Figures 8-9. Second-pass and dwell time on the relative pronoun

This clearly demonstrates that late attachment was the hardest to process, with difficulty visible as early as during second-pass. No difference between SynLA and AmbA was found in any measure. Early attachment gave a significant reading time advantage when dwell time was analyzed.

Regressions to the relative pronoun showed a highly significant main effect ($F_1(2, 54) = 4.721$, $p = 0.013$, $\eta^2 = 0.149$; $F_2(2, 42) = 6.316$, $p = 0.004$, $\eta^2 = 0.231$), with SynLA receiving significantly more regressions than both SynEA ($p_1 = 0.053$, $p_2 = 0.008$) and AmbA ($p_1 = 0.068$, $p_2 = 0.016$) (see Figure 10).

This confirms early attachment preference in Russian, with late attachment being significantly harder to process and causing more regressions.

IA 5 (verb)

Same as for IA 4, dwell time showed a highly significant main effect ($F_1(2, 54) = 4.353$, $p = 0.030$, $\eta^2 = 0.139$; $F_2(2, 42) = 4.640$, $p = 0.015$, $\eta^2 = 0.181$). Pairwise comparisons showed that SynEA was read significantly faster than AmbA ($p_1 = 0.016$; $p_2 = 0.012$) (see Figure 11).

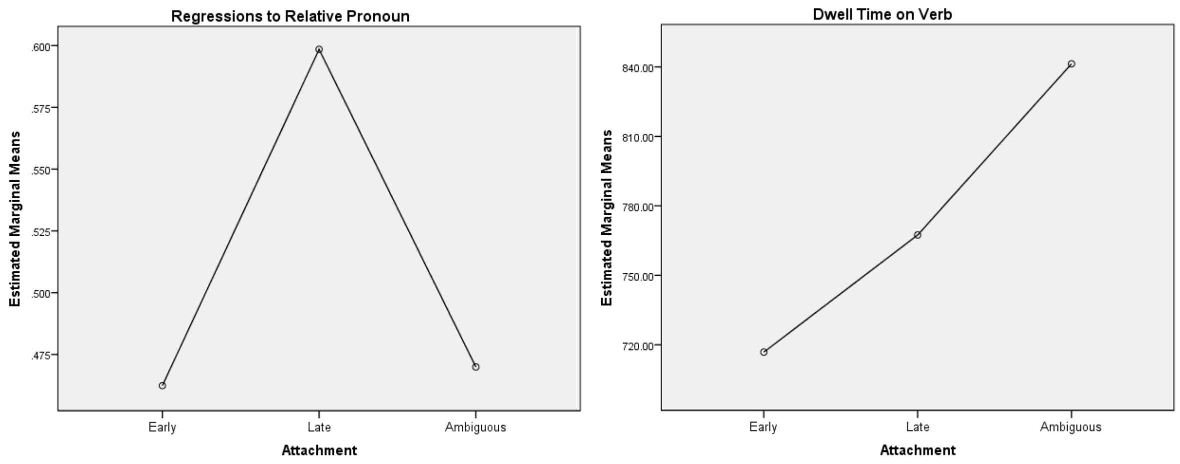

Figures 10-11. Regressions to the relative pronoun and dwell time on the verb

Analyses from the two critical regions (IA 4 and IA 5) showed a very clear picture. For Russian speakers, the processing of early attached relative clauses was much easier than late attached and ambiguously attached relative clauses, without significant difference between the last two.

NPs

Regressions to NP1 showed a very significant main effect ($F_1(2, 54) = 25.027$, $p < 0.0001$, $\eta^2 = 0.481$; $F_1(2, 54) = 32.336$, $p < 0.0001$, $\eta^2 = 0.606$), with a difference in each pair. Regressions to NP2 demonstrated the same pattern (see Figures 12-13).

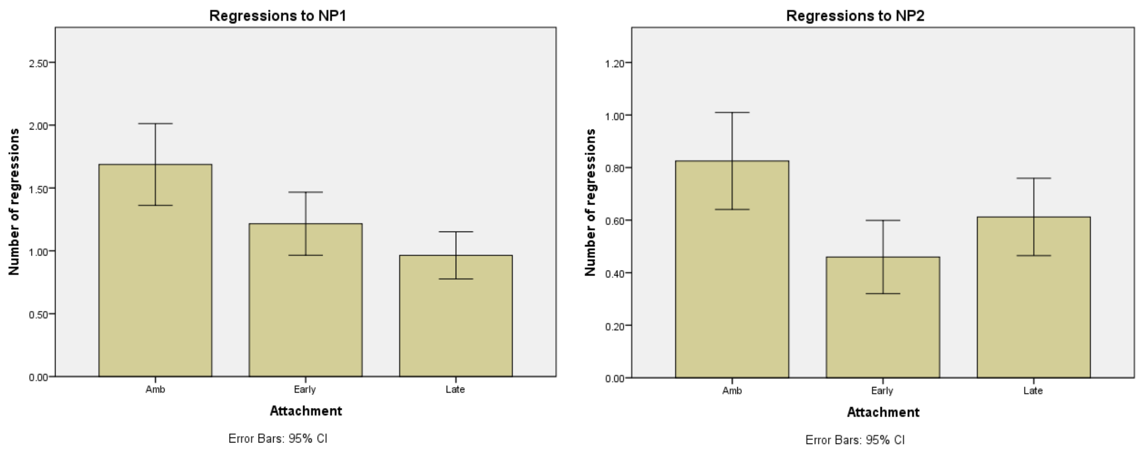

Figures 12-13. Regressions to NPs by condition

This does not reveal a clear picture. In order to compare regressions to NP1 with those to NP2 in general, an independent Student's *t*-test was carried out, with regressions as a dependent variable and NP as an independent variable. The test was highly significant: $t = 7.498$, $df = 138$, $p < 0.0001$. The mean for NP1 was 1.22, the mean for NP2 was 0.59. This suggests that across all three conditions participants regressed to NP1 twice more often than to NP2, regardless of the actual attachment.

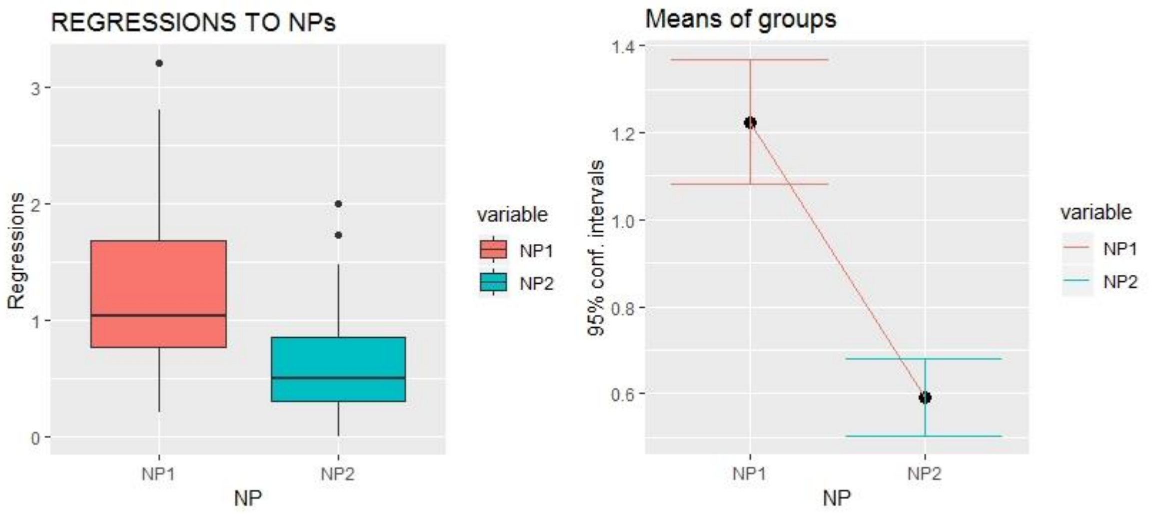

Figures 14-15. Compared regressions to NP1 and NP2

3.10.2. Group 2: Semantic disambiguation

Below are the results for semantically disambiguated items vs. ambiguous controls for areas of interest. All measures were subject to rANOVA with the main factor of Attachment (early, late, ambiguous), with a separate analysis for each IA. $p < 0.05$ is marked with one asterisk, $p < 0.01$ is marked with two asterisks, $p < 0.001$ is marked with three asterisks.

Table 13. ANOVA results for SemEA, SemLA, AmbA

Measure	**IA**	**df**	**Attachment**		
			F	***p***	***η²***
First-pass	4	2, 54	1.473	.240	.052
		2, 42	2.267	.116	.097
	5	2, 54	.284	.754	.010
		2, 42	.083	.920	.004
	6	2, 54	1.511	.230	.053
		2, 42	1.703	.194	.075
Second-pass	4	2, 54	2.319	.108	.079
		2, 42	3.468	.040*	.142
	5	2, 54	.934	.399	.033
		2, 42	.291	.749	.014
	6	2, 54	1.131	.305	.054
		2, 42	.167	.847	.008
Dwell time	4	2, 54	14.988	.000***	.357
		2, 42	15.021	.000***	.417
	5	2, 54	3.821	.028*	.124
		2, 42	3.910	.028*	.157
	6	2, 54	1.501	.232	.053
		2, 42	2.305	.112	.099
Regressions in	2	2, 54	11.887	.000***	.306
		2, 42	11.349	.000***	.351
	3	2, 54	7.035	.005**	.207
		2, 42	11.689	.000***	.358
	4	2, 54	2.618	.082	.088
		2, 42	5.554	.007**	.209
	5	2, 54	8.980	.000***	.250
		2, 42	7.133	.002**	.254
	6	2, 54	0	0	0
		2, 42	0	0	0
Regressions out	4	2, 54	1.277	.287	.045
		2, 42	3.288	.047*	.135
	5	2, 54	3.348	.043*	.110
		2, 42	3.726	.032*	.151
	6	2, 54	1.597	.212	.056
		2, 42	2.015	.146	.088

IA 4 (relative pronoun)

Since in all the above conditions the relative pronoun was ambiguous, no differences in the first-pass were expected. And this turned out to be the case: $F_1(2, 54) = 1.473$, $p = 0.240$, $\eta^2 = 0.052$; $F_2(2, 42) = 2.267$, $p = 0.116$, $\eta^2 = 0.097$.

Surprisingly, however, second-pass revealed a significant main effect by items ($F_1(2, 54) = 2.319$, $p = 0.108$, $\eta^2 = 0.079$; $F_2(2, 42) = 3.468$, $p = 0.040$, $\eta^2 = 0.142$), with SemEA being read significantly faster than AmbA ($p_2 = 0.038$) (see Figure 16). This suggests that second-pass on the early attached relative clause, even when it started with an ambiguous pronoun, was still the easiest.

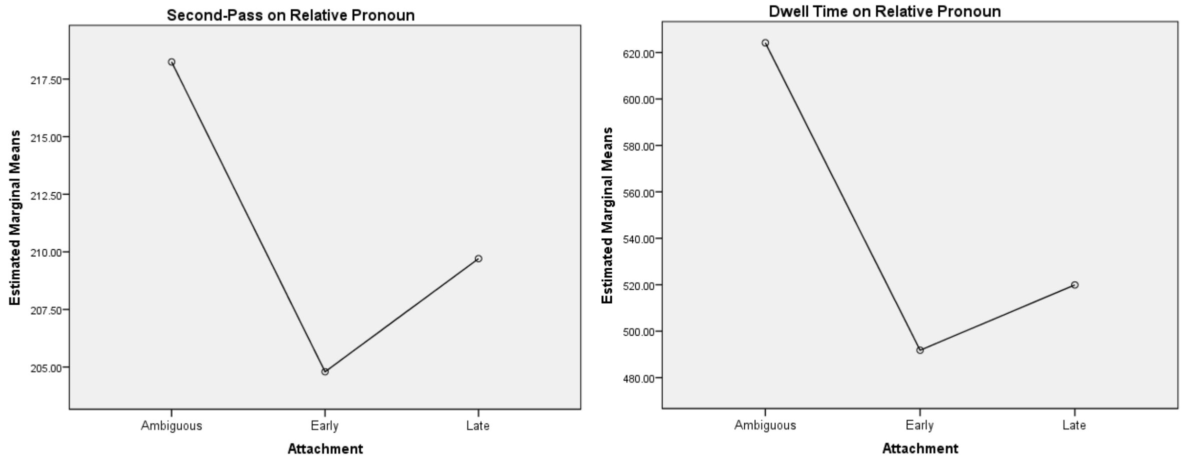

Figures 16-17. Second-pass and dwell time on the relative pronoun

Dwell time showed a highly significant main effect ($F_1(2, 54) = 14.988$, $p < 0.0001$, $\eta^2 = 0.357$; $F_2(2, 42) = 15.021$, $p < 0.0001$, $\eta^2 = 0.417$), with AmbA being read significantly slower than both SemEA (*p2* < 0.0001) and SemLA (*p2* < 0.0001), which is consistent with the predictions of Constraint Satisfaction accounts.

IA 5 (verb)

Dwell time showed a significant main effect ($F_1(2, 54) = 3.821$, $p = 0.028$, $\eta^2 = 0.124$; $F_2(2, 42) = 3.910$, $p = 0.028$, $\eta^2 = 0.157$), with verbs in AmbA being read significantly slower than in SemLA (841 ms vs. 744 ms, $p_1 = 0.049$; $p_2 = 0.036$). This, however, was not because of the relative ease of the latter but rather because participants regressed significantly more often from a late attached verb: regressions out showed a significant main effect ($F_1(2, 54) = 3.348$, $p = 0.043$, $\eta^2 = 0.110$; $F_2(2, 42) = 3.726$, $p = 0.032$, $\eta^2 = 0.151$), with SemLA demonstrating significantly more regressions from them than SemEA ($p_2 = 0.042$) (Figure 18). This suggests the highest processing complexity of SemLA.

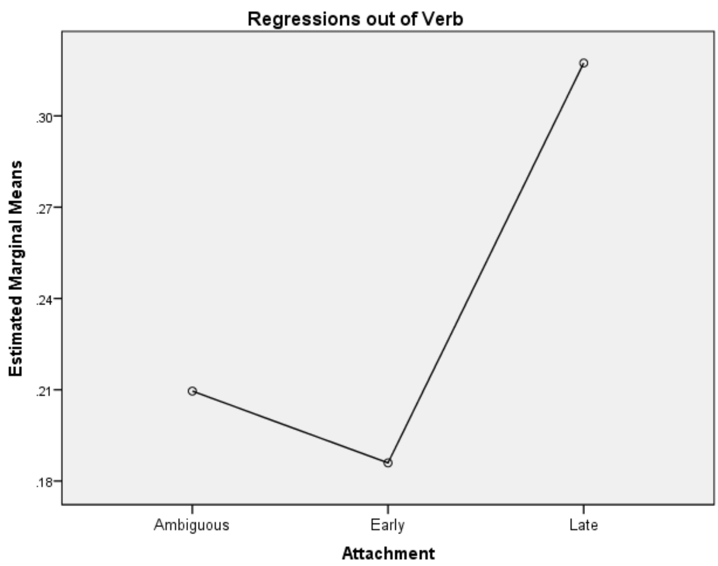

Figure 18. Regressions out of the verb

NPs

Same as for the previous set of items, regressions to competing NPs by condition were not very telling (see Figures 19-20).

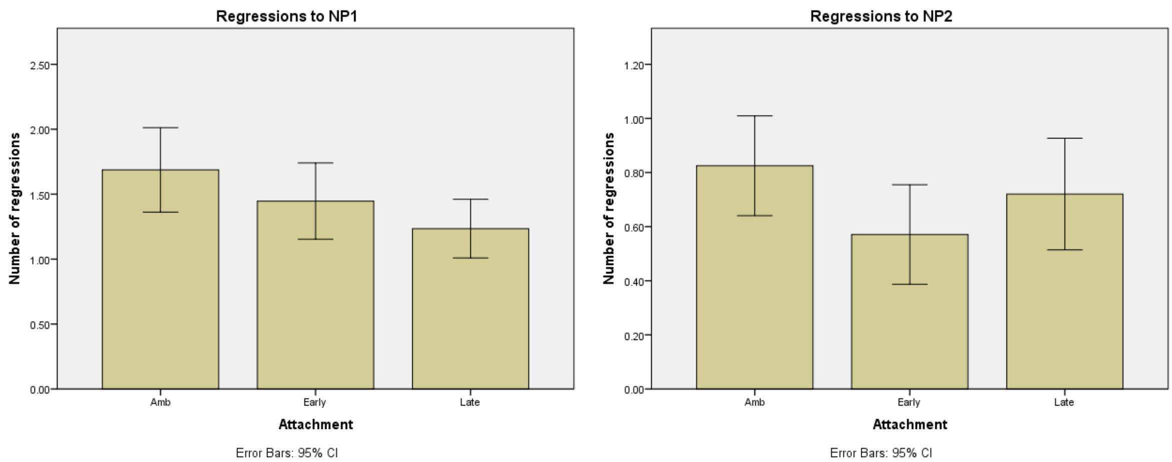

Figures 19-20. Regressions to NPs by condition

In order to reveal participants' true preferences, regressions to NP1 were again compared to regressions to NP2 in general using an independent Student's *t*-test. The test was highly significant: $t = 8.280$, $df = 149$, $p < 0.0001$. The mean for NP1 was 1.38, the mean for NP2 was 0.66. This suggests that across all four conditions participants regressed to NP1 twice more often than to NP2, regardless of the actual attachment (see Figures 21-22).

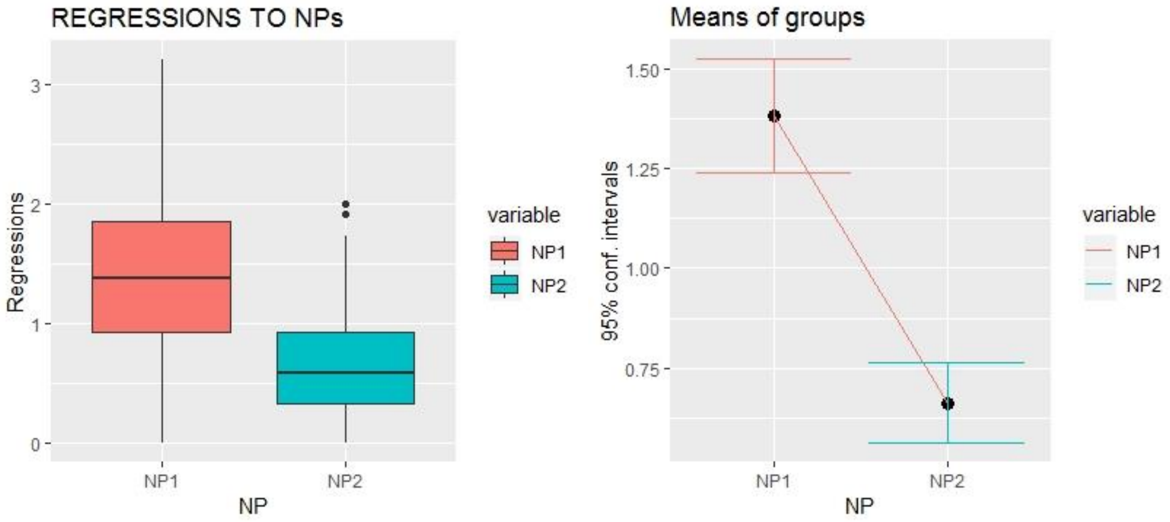

Figures 21-22. Compared regressions to NP1 and NP2

3.10.3. Ambiguous condition

Since the participants did not reliably choose one of the two nouns while answering the questions in the AmbA condition, it was important to check whether they still regressed more to NP1 or NP2, to confirm the hypothesis that the absence of preference in this experiment was caused by its design rather than an actual lack of preference in Russian native speakers.

NP1 and NP2

In order to determine participants' preferences in the absence of any clues, whether syntactic or semantic, regressions to NP1 and NP2 were analyzed with the help of Student's *t*-test. The test turned out to be highly significant: $t = 4.725$, $df = 45$, $p < 0.0001$. The mean for NP1 was 1.61, the mean for NP2 was 0.80. This suggests that participants regressed to NP1 twice more often than to NP2 (see Figures 23-24).

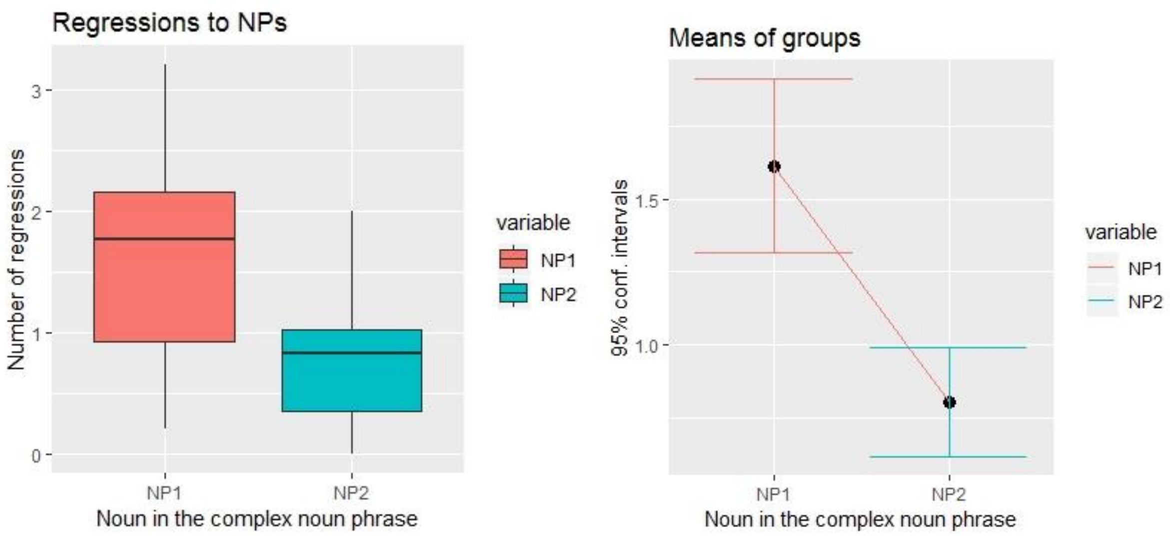

Figures 23-24. Regressions to competing NPs in the Ambiguous condition

Since the participants clearly regressed more to NP1 and NP2 across all conditions, it was also important to exclude other possible reasons other than their natural preferences. The bivariate Pearson correlation test was performed for all conditions, with regressions to nouns and their frequency in the Russian National Corpus (see Appendix 3) as variables. The test did not reveal any statistically significant correlation (the magnitude of the association $r = 0.129$, $p = 0.117$), suggesting that frequency had no effect on regression patterns.

Both Anisimov, Fedorova, and Latanov (2014) and Chernova and Chernigovskaya (2015) conducted an additional analysis to determine whether the total reading time of the relative clause was correlated with the attachment participants finally chose. The same analysis was performed for this experiment using a *t*-test. Even though the mean reading time for clauses when they were interpreted as EA was lower than when they were interpreted as LA (1858 ms vs. 2013 ms), this was not statistically significant ($t = 0.847$, $df = 54$, $p = 0.401$) (see Figures 24-25).

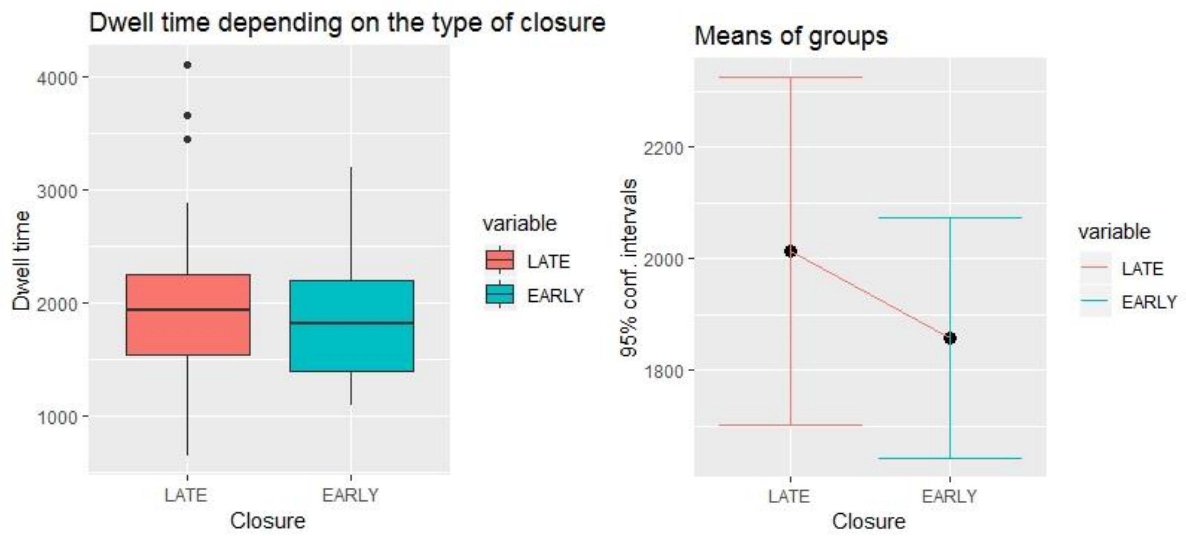

Figures 25-26. Dwell time on the relative clause in the Ambiguous condition depending on the chosen type of closure

3.11. Discussion

As predicted, offline measures (question answers) showed a strong early attachment preference in Russian, with significantly more mistakes in sentences where the relative clause was forcefully attached to the second host. They also revealed a higher complexity of sentences constrained only by semantics, with more errors in them than in syntactically unambiguous sentences. This also agrees with other experiments conducted on Russian, which showed a much higher accuracy for early attached clauses (Chernova and Chernigovskaya 2015). However, in comparison to Chernova and Chernigovskaya, the mean accuracy for items with late attached clauses did not drop below 90 percent in this experiment, regardless of the constraints, suggesting that participants experienced no severe comprehension troubles.

A surprising difference between this experiment and the previous ones conducted on Russian was the distribution of answers in the Ambiguous condition. When answering the questions, participants chose early attachment almost as often as late attachment (51.6 percent vs. 48.4 percent), with no real preference, whereas in the previous experiments they chose early closure significantly more often: 67 percent of times in Anisimov, Fedorova, and Latanov (2014) and 67.3 percent of times in Chernova and Chernigovskaya (2015). However, as indicated before and confirmed by online measures, these results were most likely caused by the experimental design, which contrasted syntactically ambiguous but semantically biased sentences with fully ambiguous sentences. This was not done in the previous experiments on Russian. Unlike in Traxler et al. (1998), when it became apparent to participants that both interpretations were plausible, processing did not stop but rather became more difficult, increasing regressions and resulting in longer dwell times. This is also confirmed by the fact that participants still regressed to NP1 significantly more often (twice more often) than to NP2 in ambiguous sentences.

Interestingly, no significant differences were found in the first-pass time on any area. It poses challenges for all the existing parsing models, unless we assume that, as claimed by Constraint Satisfaction accounts, constraints provided were enough to resolve the ambiguity right from the start. However, assuming this would only create additional problems: why differences in second-pass and dwell time were significant? Why first-pass on the relative pronoun in the Ambiguous condition was not longer than that on unambiguous ones? Most likely, the absence of significance simply reflects a lack of experimental power. It may have also been caused by the fact that when disambiguation was very strong (syntactic), it occurred immediately on the next word following the noun phrase complex, so even if the parser did a mistake or experienced some troubles, it resolved it too quick to be detected. When disambiguation occurred later in the relative clause, giving the initial misparse a chance to grow strong enough for a disruption to be detectable, it was semantic (less reliable), and the effect might have been delayed.

Late measures revealed a much more interesting picture. For syntactically disambiguated items, early attachment had a marginally significant second-pass advantage and a significant total time advantage for both critical areas (gender-marked pronoun and gender-marked verb). This finding is consistent with other studies that found early attachment advantage only in total times (e.g., Carreiras and Clifton 1999). This also seems to agree with Constraint Satisfaction models that predict easier processing when one analysis receives clear activation from multiple constraints and much harder processing when different constraints support different analyses. In late attached relative clauses, participants' general offline preferences (accuracy rate in this experiment and attachment preferences for ambiguous items consistently found for Russian before) were violated, which means that two analyses received similar activation from different constraints, slowing the processor down. Late attached relative pronouns also received the highest number of regressions, additionally confirming the highest processing difficulty.

For semantically disambiguated items, the picture was a bit different. Whereas ambiguous relative pronouns still had a significant reading time advantage when they belonged to early attached clauses, this effect did not show up on the disambiguating verb. Rather, semantically late attached verbs had less time spent on them because of the highest number of regressions out.

Thus, syntactic and semantic constraints had different effects on processing. Whereas there were more regressions *into* syntactically late attached regions, there were more regressions *out of* semantically late attached regions. This is reasonable: given that two nouns in the syntactically disambiguated sentences were of different gender, regressing back from them to disambiguating regions would help establish the agent. In semantically disambiguated sentences, participants reread the relative clause and actively regressed from it to competing NPs to, most likely, reanalyze their initial misparse.

Ambiguity did not have higher processing cost in early measures but produced the longest total time. This seems to agree with Constraint Satisfaction models predicting harder processing for ambiguous fragments, since many syntactic analyses simultaneously receive equal activation, overloading the parser.

Regressions to competing NPs showed that NP1 was reread twice more often than NP2 across all conditions with syntactic and semantic constraints regardless of their actual attachment. The same was true for the unconstrained condition (AmbA). This fully agrees with finding of Chernova and Chernigovskaya (2015) for adjunct modifiers with participles and indicates that there is a clear early attachment preference in Russian in late measures. Also, no significant difference was found in dwell time on the relative clause in AmbA when it was interpreted as EA and when it was interpreted as LA (the value was taken from participants' answers to questions), which agrees with the same absence of significance found by Anisimov, Fedorova, and Latanov (2014) but disagrees with reading time advantage for late attached clauses found in Chernova and Chernigovskaya (2015).

Summarizing the above, we may conclude that the results mostly agree with the predictions of Constraint Satisfaction accounts. The Relativized Relevance principle was not confirmed because the participants did not read sentences containing late attached clauses any faster (in fact, vice versa) and did not reliably choose NP1 when answering the questions to ambiguous sentences. The Late Closure strategy was again disconfirmed, with no reading time advantage for sentences with late attached clauses and no preference for NP2 in question answers. Constraint Satisfaction accounts were right in that ambiguity imposed extra processing cost (at least in late measures) and that sentences with early attached clauses (that is, corresponding to participant's offline preferences) were read faster. However, the findings present problems for all the existing parsing models, because all these results were visible only in late measures.

Conclusions

The aim of this study was to investigate relative clause attachment preferences in Russian and how different constraints—syntactic and semantic—affect them. Several previous experiments conducted on Russian speakers showed very mixed results: an eye-tracking study by Chernova and Chernigovskaya (2015) showed evidence for the Late Closure strategy in the first-pass reading time, which contradicted participants' offline attachment preferences (question answers) in all other experiments, including that of Chernova and Chernigovskaya themselves. The authors also found no differences in late measures (dwell time), whereas many cross-linguistic studies reported them (Traxler et al. 1998; Carreiras and Clifton 1999). This investigation was intended to provide additional evidence to the debate by manipulating two variables: constraints (syntax or semantics) and attachment (ambiguous, early, or late). It was also novel in terms of verb aspect (all verbs used in this experiment were perfective) and heterogeneity of the population (participants were born in different countries but all acquired Russian as their first language).

Chapter 1 of this study provided a background of what had already been discovered about human parsing mechanisms: what quantitative and qualitative limitations they have, how incremental they are, and how many stages they may contain. The question of syntax-lexicon dissociation in the brain was examined based on the results of multiple experiments, from SPRT and eye-tracking to ERP and other neuroimaging techniques. Two huge groups of models were presented: modular accounts positing a dichotomy between different language subcomponents and interactive accounts claiming that there is constant interaction between them. The chapter also provided an overview of the most influential models, such as the Garden-Path Model, Constraint Satisfaction models, the Construal Hypothesis, the Unrestricted Race Model, the Underspecification Model, and other. In summary, the empirical evidence is very inconclusive, with some experiments supporting the idea that the parser operates using pre-existing templates initially and has to revise when its parse turns out to be incorrect, and some contradicting it.

Chapter 2 summarized major results found in numerous cross-linguistic relative clause attachment experiments. Different hypotheses proposed to account for a great variance found in these experiments were reviewed: The Relativized Relevance principle claiming that attachment to the most discourse-prominent NP is preferred and is determined by relatively late discourse-related processes, the Tuning Hypothesis based on the frequency of particular attachment in a given language, the Late Closure principle claiming the universality of late attachment, the Predicate Proximity/Recency principle, Constraint Satisfaction accounts, and other. In summary, cross-linguistic evidence clearly contradicts the Late Closure strategy and shows some support for each of the above principles.

Chapter 3 described the conducted experiment. It was designed to test how quickly and how reliably two different constraints were used: gender that rendered sentences syntactically unambiguous and semantics, which kept them syntactically ambiguous but provided a clue for establishing attachment. Early measures (first-pass) revealed no differences between conditions. This is problematic for all the existing parsing models. No support for the Late Closure strategy was found, unlike in the Chernova and Chernigovskaya (2015) experiment. This suggests that their findings might have been the result of the extremely low accuracy in the late attachment condition (38 percent) and the fact that the participants simply did not understand the sentences. Later measures, however, agreed with the Constraint Satisfaction accounts, revealing easier processing when the relative clause was forcefully attached to the first host (NP1). They also demonstrated the highest processing complexity of unconstrained (ambiguous) sentences.

Syntax and semantics had different effects on processing. There were more regressions out of the RC when semantics supported the less preferred analysis (LA) and more regressions to the RC when syntax supported the less preferred analysis (LA).

This experiment included several modifications that could have affected its results. First, the population was very heterogeneous, with participants divided almost equally between three countries: Russia, Ukraine, and Belarus. Other studies may manipulate only this variable, because differences in attachment depending on the participant's country of origin were found, although the samples were two small to establish reliable correlations. Second, the verbs used in relative clauses were all telic (past tense, perfective aspect). It would be interesting to manipulate only aspect and see how (and whether) it affects participants' attachment preferences. Third, semantics was a less reliable disambiguator compared to some other experiments (e.g., Traxler, Pickering, and Clifton 1998), which was done in order to keep both potential attachment sites animate. Other experiment may, however, manipulate animacy of the hosts (and have stronger semantic plausibility as a result). It would also be interesting to examine effects of semantic plausibility on dispreferred attachment in an ERP experiment.

References

Abney, S. P. (1989). A computational model of human parsing. *Journal of Psycholinguistic Research*, *18*(1), 129–144. doi: 10.1007/BF01069051.

Adams, B. C., Clifton, C., and Mitchell, D. C. (1998). Lexical guidance in sentence processing?. *Psychonomic Bulletin & Review*, *5*(2), 265–270.

Ainsworth-Darnell, K., Shulman, H. G., and Boland, J. E. (1998). Dissociating brain responses to syntactic and semantic anomalies: Evidence from event-related potentials. *Journal of Memory and Language*, *38*(1), 112–130. doi: 10.1006/jmla.1997.2537.

Altmann, G., and Steedman, M. (1988). Interaction with context during human sentence processing. *Cognition*, *30*(3), 191–238.

Anisimov, V. N., Fedorova, O. V., and Latanov, A. V. (2014). Eye movement parameters in reading sentences with syntactic ambiguities in Russian. *Human Physiology*, *40*(5), 521–531. doi: 10.1134/S0362119714040033.

Barton, S. B., and Sanford, A. J. (1993). A case study of anomaly detection: Shallow semantic processing and cohesion establishment. *Memory & Cognition*, *21*(4), 477–487. doi: 10.3758/BF03197179.

Bates, E., MacWhinney, B., Caselli, C., Devescovi, A., Natale, F., and Venza, V. (1984). A cross-linguistic study of the development of sentence interpretation strategies. *Child Development*, *55*(2), 341–354. doi:10.2307/1129947.

Bever, T. G. (1970). The cognitive basis for linguistic structures. *Cognition and the Development of Language*, *279*(362), 1–61.

Brysbaert, M., and Mitchell, D. C. (1996). Modifier attachment in sentence parsing: Evidence from Dutch. *The Quarterly Journal of Experimental Psychology Section A*, *49*(3), 664–695. doi: 10.1080/713755636.

Caplan, D., and Waters, G. (2002). Working memory and connectionist models of parsing: A reply to MacDonald and Christiansen (2002). *Psychological Review*, *109*(1), 66–74. doi: 10.1037/0033-295X.109.1.66.

Carreiras, M., and Clifton, C. (1993). Relative clause interpretation preferences in Spanish and English. *Language and Speech*, *36*(4), 353–372. doi: 10.1177/002383099303600401.

Carreiras, M., and Clifton, C. (1999). Another word on parsing relative clauses: Eyetracking evidence from Spanish and English. *Memory & Cognition*, *27*(5), 826–833.

Chernova, D., and Chernigovskaya, T. V. (2015). Syntactic ambiguity resolution in sentence processing: New evidence from a morphologically rich language. In: G. Airenti, B. G. Bara, G. Sandini (Eds.), *Proceedings of the EuroAsianPacific Joint Conference on Cognitive Science* (pp. 129–133), Torino, September 25-27, 2015. [published at http://ceur-ws.org/Vol-1419/paper0017.pdf]

Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.

Clifton, Jr, C. (1993). Thematic roles in sentence parsing. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale, 47*(2), 222–246. doi: 10.1037/h0078817.

Clifton Jr, C., Frazier, L., and Connine, C. (1984). Lexical expectations in sentence comprehension. *Journal of Verbal Learning and Verbal Behavior*, *23*(6), 696–708.

Clifton Jr, C., Traxler, M. J., Mohamed, M. T., Williams, R. S., Morris, R. K., and Rayner, K. (2003). The use of thematic role information in parsing: Syntactic processing autonomy revisited. *Journal of Memory and Language*, *49*(3), 317–334.

Crain, S., and Steedman, M. (1985). On not being led up the garden-path: The use of context by the psychological parser. In: D. R. Dowty, L. Karttunen, and A. M. N. Zwicky (Eds.), *Natural language parsing* (pp. 320–358). Cambridge, UK: Cambridge University Press.

Cuetos, F., and Mitchell, D. C. (1988). Cross-linguistic differences in parsing: Restrictions on the use of the Late Closure strategy in Spanish. *Cognition*, *30*(1), 73–105. doi: 10.1016/0010-0277(88)90004-2.

Cuetos, F., Mitchell, D. C., and Corley, M. M. B (1996). Parsing in different languages. In: M. Carreiras, J. E. Garcia-Albea, and N. Sebastian-Galles (Eds.), *Language processing in Spanish* (pp. 145–187). Mahwah, NJ: Erlbaum.

De Vincenzi, M., and Job, R. (1993). Some observations on the universality of the late-closure strategy. *Journal of Psycholinguistic Research*, *22*(2), 189–206.

De Vincenzi, M., and Job, R. (1995). An investigation of late closure: The role of syntax, thematic structure, and pragmatics in initial interpretation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*(5), 1303–1321.

Duffy, S. A., Morris, R. K., and Rayner, K. (1988). Lexical ambiguity and fixation times in reading. *Journal of Memory and Language*, *27*(4), 429–446.

Fedorenko, E., Nieto-Castañon, A., and Kanwisher, N. (2012). Lexical and syntactic representations in the brain: An fMRI investigation with multi-voxel pattern analyses. *Neuropsychologia*, *50*(4), 499–513. doi: 10.1016/j.neuropsychologia.2011.09.014.

Ferreira, F., and Clifton, C. (1986). The independence of syntactic processing. *Journal of Memory and Language*, *25*, 348–368.

Ferreira, F., and Henderson, J. M. (1990). Use of verb information in syntactic parsing: Evidence from eye movements and word-by-word self-paced reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*(4), 555–568. doi: 10.1037/0278-7393.16.4.555.

Fisher, D. F., and Shebilske, W. L. (1985). There is more that meets the eye than the eyemind assumption. In: R. Groner, G. W. McConkie, C. Menz (Eds.), *Eye movements and human information processing* (pp. 149–157). North Holland, Amsterdam.

Flajolet, P. (1987). Analytic models and ambiguity of context-free languages. *Theoretical Computer Science*, *49*(2), 283–309. doi: 10.1016/0304-3975(87)90011-9.

Fodor, J. A. (1983). *Modularity of mind: An essay on faculty psychology*. Cambridge, MA: MIT Press.

Fodor, J. A. (1985). Précis of the modularity of mind. *Behavioral and Brain Sciences*, *8*(1), 1–5.

Fodor, J. D. (1985). Module or muddle? *Behavioral and Brain Sciences*, *8*(1), 7–9.

Forster, K. I. (1974). The role of semantic hypotheses in sentence processing. In: F. Bresson and J. Mehler (Eds.), *Current problems in psycholinguistics* (pp. 391–408). Paris: Editions du CRNS.

Forster, K. I. (1976). Accessing the mental lexicon. In: R. J. Wales and E. W. Walker (Eds.), *New approaches to language mechanisms* (pp. 257–287). Amsterdam: North-Holland.

Forster, K. I. (1979). Levels of processing and the structure of the language processor. In: W. Cooper and E. Walker (Eds.), *Sentence processing: Psycholinguistic studies presented to Merrill Garrett* (pp. 27–85). Hillsdale, NJ: Erlbaum.

Frazier, L. (1979). *On comprehending sentences: Syntactic parsing strategies*. PhD dissertation, Indiana University Linguistics Club, University of Connecticut.

Frazier, L. (1987). Sentence processing: A tutorial review. In: M. Coltheart (Ed.), *Attention and performance XII: The psychology of reading* (pp. 559–586). Hillsdale, NJ: Erlbaum.

Frazier, L. (1990). Parsing modifiers: Special purpose routines in the human sentence processing mechanism? In: D.A. Balota, G.B. Flores d'Arcais, and K. Rayner (Eds.), *Comprehension processes in reading*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Frazier, L. (1999). *On sentence interpretation* (Studies in Theoretical Psycholinguistics, Vol. 22). Kluwer Academic Publishers: Springer Science & Business Media. doi: 10.1007/978-94-011-4599-2.

Frazier, L., and Clifton, C. (1989). Successive cyclicity in the grammar and the parser. *Language and Cognitive Processes*, *4*(2), 93–126. doi: 10.1080/01690968908406359

Frazier, L., and Clifton, C. (1996). *Construal*. Cambridge, MA: MIT Press.

Frazier, L., and Clifton, C. (1997). Construal: Overview, motivation, and some new evidence. *Journal of Psycholinguistic Research*, *26*(3), 277–295. doi: 10.1023/A:1025024524133.

Frazier, L., and Fodor J. D. (1978). The sausage machine: a new two-stage parsing model. *Cognition*, *6*(4), 291–325.

Frazier, L., and Rayner, K. (1982). Making and correcting errors during sentence comprehension: Eye movements in the analysis of structurally ambiguous sentences. *Cognitive Psychology*, *14*(2), 178–210. doi: 10.1016/0010-0285(82)90008-1.

Frazier, L., and Rayner, K. (1990). Taking on semantic commitments: Processing multiple meanings vs. multiple senses. *Journal of Memory and Language*, *29*(2), 181–200.

Gaskell, M. G. (2007). *The Oxford handbook of psycholinguistics*. Oxford; New York: Oxford University Press.

Gernsbacher, M. A. (1993). Less skilled readers have less efficient suppression mechanisms. *Psychological Science*, *4*(5), 294–298. doi: 10.1111/j.1467-9280.1993.tb00567.x

Gernsbacher, M. A., and Faust, M. E. (1991). The mechanism of suppression: A component of general comprehension skill. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*(2), 245–262. doi: 10.1037/0278-7393.17.2.245.

Gibson, E., Schütze, C. T., and Salomon, A. (1996). The relationship between the frequency and the processing complexity of linguistic structure. *Journal of Psycholinguistic Research*, *25*(1), 59–92.

Gilboy, E., Sopena, J.-Mm., Clifton, C., and Frazier, L. (1995). Argument structure and association preferences in Spanish and English complex NPs. *Cognition*, *54*(2), 131–167. https://doi.org/10.1016/0010-0277(94)00636-Y

Hagoort, P., Brown, C., and Groothusen, J. (1993). The syntactic positive shift (SPS) as an ERP measure of syntactic processing. *Language and Cognitive Processes*, *8*(4), 439–483. doi: 10.1080/01690969308407585.

Hahne, A., and Jescheniak, J. D. (2001). What's left if the Jabberwock gets the semantics? An ERP investigation into semantic and syntactic processes during auditory sentence comprehension. *Cognitive Brain Research*, *11*(2), 199–212. doi: 10.1016/S0926-6410(00)00071-9.

Hawkins, J. A. (1995). *A performance theory of order and constituency* (Cambridge Studies in Linguistics). Cambridge, UK: Cambridge University Press. doi:10.1017/CBO9780511554285.

Just, M. A., and Carpenter, P. A. (1980). A theory of reading: From eye fixations to comprehension. *Psychological Review*, *87*(4), 329–354.

Just, M. A., and Carpenter, P. A. (1992). A capacity theory of comprehension: Individual differences in working memory. *Psychological Review*, *99*(1), 122–149.

Kako, E. (2006). The semantics of syntactic frames. *Language and Cognitive Processes*, *21*(5), 562–575.

Kamide, Y., Altmann, G. T., and Haywood, S. L. (2003). The time-course of prediction in incremental sentence processing: Evidence from anticipatory eye movements. *Journal of Memory and Language*, *49*(1), 133–156.

Kimball, J. (1973). Seven principles of surface structure parsing in natural language. *Cognition*, *2*(1), 15–47. doi: 10.1016/0010-0277(72)90028-5.

King, J., and Just, M. A. (1991). Individual differences in syntactic processing: The role of working memory. *Journal of Memory and Language*, *30*(5), 580–602. doi: 10.1016/0749-596X(91)90027-H.

Konieczny, L., Hemforth, B., Scheepers, C., and Strube, G. (1997). The role of lexical heads in parsing: Evidence from German. *Language and Cognitive Processes*, *12*(2–3), 307–348. doi: 10.1080/016909697386871.

Krauzlis, R. J., Goffart, L., and Hafed, Z. M. (2017). Neuronal control of fixation and fixational eye movements. *Philosophical Transactions of the Royal Society B: Biological Sciences, 372*(1718), 1–12. doi: 10.1098/rstb.2016.0205.

Leech, R., Aydelott, J., Symons, G., Carnevale, J., and Dick, F. (2007). The development of sentence interpretation: Effects of perceptual, attentional and semantic interference. *Developmental Science, 10*(6), 794–813. doi: 10.1111/j.1467-7687.2007.00628.x.

Liversedge, S. P., Paterson, K. B., and Pickering, M. J. (1998). Chapter 3 – Eye movements and measures of reading time. In: G. Underwood (Ed.), *Eye guidance in reading and scene perception* (pp. 55–75). Amsterdam: Elsevier Science Ltd.

Lyashevskaya, O. N., and Sharov, S.A. (2009). *Chastotnyi slovar' sovremennogo russkogo yazyka (Na matrialakh Natsional'nogo korpusa russkogo yazyka)* [Frequency vocabulary of the Russian language (on the materials of the Russian National Corpus)]. Moscow: Azbukovnik.

MacDonald, M. C., Just, M. A., and Carpenter, P. A. (1992). Working memory constraints on the processing of syntactic ambiguity. *Cognitive Psychology, 24*(1), 56–98.

MacDonald, M. C., Pearlmutter, N. J., and Seidenberg, M. S. (1994). The lexical nature of syntactic ambiguity resolution. *Psychological Review*, *101*(4), 676–703. doi: 10.1037/0033-295X.101.4.676.

MacLeod, C. M. (1991). Half a century of research on the Stroop effect: An integrative review. *Psychological Bulletin*, *109*(2), 163–203.

MacWhinney, B., Bates, E., and Kliegl, R. (1984). Constraint validity and sentence interpretation in English, German, and Italian. *Journal of Verbal Learning and Verbal Behavior, 23*(2), 127–150.

MacWhinney, B., Osmán-Sági, J., and Slobin, D. I. (1991). Sentence comprehension in aphasia in two clear case-marking languages. *Brain and Language*, *41*(2), 234–249. doi: 10.1016/0093-934X(91)90154-S.

Mak, W. M., Vonk, W., and Schriefers, H. (2002). The Influence of animacy on relative clause processing. *Journal of Memory and Language*, *47*(1), 50–68. doi: 10.1006/jmla.2001.2837.

McDonald, J. L. (1986). The development of sentence comprehension strategies in English and Dutch. *Journal of Experimental Child Psychology, 41*(2), 317–335.

McElree, B., and Griffith, T. (1995). Syntactic and thematic processing in sentence comprehension: Evidence for a temporal dissociation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*(1), 134.

McRae, K., Spivey-Knowlton, M. J., and Tanenhaus, M. K. (1998). Modeling the influence of thematic fit (and other constraints) in on-line sentence comprehension. *Journal of Memory and Language, 38*(3), 283–312.

Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, *63*(2), 81.

Miller, G. A. (1962). Some psychological studies of grammar. *American Psychologist*, *17*(11), 748–762. doi: 10.1037/h0044708.

Mitchell, D. C. (1987). Lexical guidance in human parsing: Locus and processing characteristics. In: M. Coltheart (Ed.), *Attention and performance 12: The psychology of reading* (pp. 601–618). Hillsdale, NJ, US: Lawrence Erlbaum Associates, Inc.

Mitchell, D. C., Brysbaert, M., Grondelaers, S., and Swanepoel, P. (2000). Chapter 18 – Modifier attachment in Dutch: Testing aspects of Construal Theory. In: A. Kennedy, R.

Radach, D. Heller, and J. Pynte (Eds.), *Reading as a perceptual process* (pp. 493–516). Oxford: North-Holland. doi: 10.1016/B978-008043642-5/50023-1.

Mitchell, D. C., Corley, M., and Garnham, A. (1992). Effects of context in human sentence parsing: Evidence against a discourse-based proposal mechanism. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*(1), 69–88.

Mitchell, D. C., Cuetos, F., Corley, M. M., and Brysbaert, M. (1995). Exposure-based models of human parsing: Evidence for the use of coarse-grained (nonlexical) statistical records. *Journal of Psycholinguistic Research*, *24*(6), 469–488.

Mitchell, D. C., and Holmes, V. M. (1985). The role of specific information about the verb in parsing sentences with local structural ambiguity. *Journal of Memory and Language*, *24*(5), 542–559.

Münte, T. F., Heinze, H. J., and Mangun, G. R. (1993). Dissociation of brain activity related to syntactic and semantic aspects of language. *Journal of Cognitive Neuroscience*, *5*(3), 335–344.

Rayner, K., Carlson, M., and Frazier, L. (1983). The interaction of syntax and semantics during sentence processing: Eye movements in the analysis of semantically biased sentences. *Journal of Learning and Verbal Behavior*, *22*(3). 358–374. doi: 10.1016/S0022-5371(83)90236-0.

Rayner, K., Garrod, S., and Perfetti, C. A. (1992). Discourse influences during parsing are delayed. *Cognition*, *45*(2), 109–139.

Rayner, K., and Pollatsek, A. (1989). *The psychology of reading.* Englewood Cliffs, NJ: Prentice Hall.

Rayner, K., Schotter, E. R., Masson, M. E., Potter, M. C., and Treiman, R. (2016). So much to read, so little time: How do we read, and can speed reading help? *Psychological Science in the Public Interest*, *17*(1), 4–34.

Reichle, E. D., and Reingold, E. M. (2013). Neurophysiological constraints on the eye-mind link. *Frontiers in Human Neuroscience*, *7*, Article 361, 1–6. doi: 10.3389/fnhum.2013.00361.

Rösler, F., Pütz, P., Friederici, A., and Hahne, A. (1993). Event-related brain potentials while encountering semantic and syntactic constraint violations. *Journal of Cognitive Neuroscience*, *5*(3), 345–362.

Sanford, A. J., and Sturt, P. (2002). Depth of processing in language comprehension: Not noticing the evidence. *Trends in Cognitive Sciences*, *6*(9), 382–386. doi: 10.1016/S1364-6613(02)01958-7.

Schriefers, H., Friederici, A. D., and Kuhn, K. (1995). The processing of locally ambiguous relative clauses in German. *Journal of Memory and Language, 34*(4), 499–520. doi: 10.1006/jmla.1995.1023

Sekerina, I. (2003). The late closure principle in processing of ambiguous Russian sentences. In: P. Kosta and J. Frasek (Eds.), *Current Approaches to Formal Slavic Linguistics: Contributions of the Second European Conference on Formal Description of Slavic Languages (FDSL II) held at Potsdam University, November 20-22, 1997*. Frankfurt/M., Berlin, Bern, Bruxelles, New York, Oxford, Wien: Peter Lang, 205–217.

Sokolov, J. (1989). Constraint validity and the acquisition of Hebrew. In: B. MacWhinney and E. Bates (Eds.), *The crosslinguistic study of sentence processing* (pp. 118–57). Cambridge, UK: Cambridge University Press.

Spivey, M., Fitneva, S., Tabor, W., and Ajmani, S. (2002). The time course of information integration in sentence processing. *The Lexical Basis of Sentence Processing: Formal, Computational and Experimental Issues*, *4*, 207.

Spivey-Knowlton, M., and Sedivy, J. C. (1995). Resolving attachment ambiguities with multiple constraints. *Cognition*, *55*(3), 227–267.

Swets, B., Desmet, T., Hambrick, D. Z., and Ferreira, F. (2007). The role of working memory in syntactic ambiguity resolution: A psychometric approach. *Journal of Experimental Psychology: General*, *136*(1), 64–81. doi: 10.1037/0096-3445.136.1.64.

Tanenhaus, M., Spivey-Knowlton, M., Eberhard, K., and Sedivy, J. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science*, *268*(5217), 1632. doi: 10.1126/science.7777863.

Townsend, D. J., and Bever, T. G. (2001). *Sentence comprehension: The integration of habits and rules*. Cambridge, MA: MIT Press.

Traxler, M. J. (2012). *Introduction to psycholinguistics: Understanding language science*. Chichester, West Sussex: Wiley-Blackwell.

Traxler, M. J., Bybee, M. D., and Pickering, M. J. (1997). Influence of connectives on language comprehension: eye tracking evidence for incremental interpretation. *The Quarterly Journal of Experimental Psychology: Section A*, *50*(3), 481–497.

Traxler, M., and Gernsbacher, M. A. (Eds.). (2011). *Handbook of psycholinguistics*. UK: Academic Press.

Traxler, M. J., Pickering, M. J., and Clifton, C. (1998). Adjunct attachment is not a form of lexical ambiguity resolution. *Journal of Memory and Language*, *39*(4), 558–592. doi: 10.1006/jmla.1998.2600.

Trueswell, J. C. (1996). The role of lexical frequency in syntactic ambiguity resolution. *Journal of Memory and Language*, *35*(4), 566–585.

Trueswell, J. C., and Tanenhaus, M. (1994). Toward a lexical framework of constraint-based syntactic ambiguity resolution. In: C. Clifton, L. Frazier, and K. Rayner (Eds.), *Perspectives in sentence processing* (pp. 155–179). Hillsdale, NJ: Lawrence Erlbaum Associates.

Trueswell, J. C., Tanenhaus, M. K., and Garnsey, S. M. (1994). Semantic influences on parsing: Use of thematic role information in syntactic disambiguation. *Journal of Memory and Language*, *33*, 285–318.

Trueswell, J. C., Tanenhaus, M., and Kello, C. (1993). Verb-specific constraints in sentence processing: Separating effects of lexical preference from garden-paths. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*(3), 528–553. 10.1037//0278-7393.19.3.528.

Van Gompel, R. P. G., Pickering, M. J., Pearson, J., and Liversedge, S. P. (2005). Evidence against competition during syntactic ambiguity resolution. *Journal of Memory and Language*, *52*(2), 284–307. doi: 10.1016/j.jml.2004.11.003.

Williams, D. L. (2009). *Developmental language disorders: Learning, language, and the brain*. San Diego, CA: Plural Publishing.

Zagar, D., Pynte, J., and Rativeau, S. (1997). Evidence for early closure attachment on first pass reading times in French. *The Quarterly Journal of Experimental Psychology Section A*, *50*(2), 421–438.

Appendix A: Test sentences

Ambiguous Attachment (AmbA) condition

1. Катя увидела ассистента лектора, который обронил ключи.
   Katia saw the assistant of the lecturer who dropped the keys.

2. Боря обокрал заместителя аудитора, который выписал штраф.
   Boria robbed the assistant of the auditor who issued a fine.

3. Егор вспомнил охранника олигарха, который оформил заказ.
   Egor recalled the bodyguard of the oligarch who placed an order.

4. Петя выслушал оппонента политика, который одобрил аборты.
   Petia heard out the opponent of the politician who encouraged abortions.

5. Сеня припугнул приятеля студента, который замазал афишу.
   Senia scared the friend of the student who painted out the advertising bill.

6. Лиза привела сверстника подростка, который подобрал щенка.
   Liza brought over the peer of the teenager who picked up the puppy.

7. Оля обманула посредника инвестора, который оценил проект.
   Olia deceived the intermediary of the investor who evaluated the project.

8. Паша обвинил поставщика ресторатора, который свернул бизнес.
   Pasha blamed the supplier of the restaurateur who closed the business.

9. Валя отпустила стажёра инструктора, который подхватил грипп.

Valia dismissed the trainee of the coach who caught a flu.

10. Игорь схватил напарника командира, который спрятал карту.

Igor caught the partner of the captain who hid the map.

11. Миша упомянул попутчика бизнесмена, который провёз виски.

Misha mentioned the companion of the businessman who smuggled whiskey.

12. Люда отыскала секретаря министра, который отложил отпуск.

Luda found the secretary of the minister who postponed the vacation.

13. Юра встретил помощника прокурора, который отключил телефон.

Jura met the helper of the attorney who switched off the phone.

14. Рита допросила коллегу бригадира, который поранил руку.

Rita interrogated the colleague of the supervisor who cut his arm.

15. Лёша подозвал оператора журналиста, который закурил трубку.

Liosha called the cameraman of the journalist who started smoking the pipe.

Syntactic Early Attachment (SynEA) condition

1. Сеня окликнул ученика скрипачки, который уронил ноты.

Senia called the student (masc) of the violinist (fem) who dropped (masc) his notes.

2. Олег впустил менеджера гимнастки, который подделал пропуск.

Oleg let in the manager (masc) of the gymnast (fem) who forged (masc) his pass.

3. Петя разбудил племянника хозяйки, который проспал поезд.

   Petia woke up the nephew (masc) of the hostess (fem) who missed (masc) the train.

4. Соня угостила мальчишку соседки, который покрасил забор.

   Sonia gave food to the boy (masc) of the neighbor (fem) who painted (masc) the fence.

5. Коля оправдал сожителя студентки, который разбил машину.

   Kolia justified the cohabitant (masc) of the student (fem) who broke (masc) the car.

6. Денис обсчитал наследника принцессы, который продал виллу.

   Denis cheated on the heir (masc) of the princess (fem) who sold (masc) the villa.

7. Дима похвалил практиканта сотрудницы, который приютил голубя.

   Dima praised the intern (masc) of the colleague (fem) who sheltered (masc) a pigeon.

8. Ира обыграла квартиранта сокурсницы, который принёс шашки.

   Ira defeated the roommate (masc) of the fellow student (fem) who brought (masc) checkers.

9. Илья разыграл любимчика бабушки, который поломал очки.

   Ilia played a trick on the fair-headed boy (masc) of the grandmother (fem) who broke (masc) his glasses.

10. Ваня высмеял шофёра чемпионки, который спутал педали.

    Vania ridiculed the chauffeur (masc) of champion (fem) who confused (masc) the pedals.

11. Влад перебил ровесника аспирантки, который уточнил график.

    Vlad interrupted the peer (masc) of the postgraduate (fem) who clarified (masc) the schedule.

12. Юра осмотрел первенца стюардессы, который проглотил монету.

Jura examined the first-born (masc) of the stewardess (fem) who swallowed (masc) the coin.

13. Оля обвинила соавтора поэтессы, который нарушил закон.

Olia accused the co-author (masc) of the poetess (fem) who violated (masc) the law.

14. Арина одобрила ухажёра дочери, который приобрёл участок.

Arina approved the suitor (masc) of the daughter (fem) who bought (masc) the lot.

15. Вова сменил партнёра балерины, который порвал костюм.

Vova replaced the partner (masc) of the ballerina (fem) who tore (masc) his suit.

Syntactic Late Attachment (SynLA) condition

1. Артур отчитал сиделку пенсионера, который пропустил завтрак.

Jura criticized the sitter (fem) of the pensioner (masc) who missed (masc) the breakfast.

2. Миша обрадовал бабушку знакомого, который выплатил кредит.

Misha delighted the grandmother (fem) of the acquaintance (masc) who paid back (masc) the loan.

3. Слава успокоил невесту сотрудника, который вывихнул плечо.

Slava calmed down the fiancé (fem) of the colleague (masc) who twisted (masc) his shoulder.

4. Зина утешила девушку сноубордиста, который сломал бедро.

Vasia comforted the girlfriend (fem) of the snowboarder (masc) who broke (masc) his thigh.

5. Олег обольстил любовницу режиссёра, который забросил съёмки.

Oleg seduced the lover (fem) of the director (masc) who abandoned (masc) the shooting.

6. Коля покорил спутницу профессора, который изобрёл вакцину.

   Kolia gain the heart of the companion (fem) of the professor (masc) who invented the vaccine.

7. Лёша проводил подругу спасателя, который покинул страну.

   Liosha accompanied the friend (fem) of the rescuer (masc) who left (masc) the country.

8. Света сыграла медсестру штурмовика, который обстрелял причал.

   Sveta played the nurse (fem) of the stormtrooper (masc) who pelted (masc) the mooring.

9. Антон поздравил супругу альпиниста, который покорил Эверест.

   Anton congratulated the wife (fem) of the alpinist (masc) who climbed (masc) Everest.

10. Вита пожалела помощницу хореографа, который перенёс концерт.

    Vita felt sorry for the assistant (fem) of the choreographer (masc) who postponed (masc) the concert.

11. Костя поцеловал ученицу академика, который возглавил кафедру.

    Kostia kissed the student (fem) of the academician (masc) who led (masc) the chair.

12. Дима разыскал падчерицу москвича, который поджёг дверь.

    Dima found the stepdaughter (fem) of the Moscow resident (masc) who set (masc) the door on fire.

13. Петя поддержал фанатку хоккеиста, который завершил карьеру.

    Petia supported the fan (fem) of the hockey player (fem) who finished (masc) his career.

14. Алиса обняла акушерку главврача, который закатил скандал.

    Alisa embraced the midwife of the chief physician (masc) who made (masc) a scene.

15. Митя допросил разведчицу генерала, который прекратил атаку.

Mitia interrogated the secret service woman (fem) of the general (masc) who stopped (masc) the attack.

Semantic Early Attachment (SemEA) condition

1. Лена встретила хирурга жениха, который вправил вывих.

Lena met the surgeon of the fiancé who set the bone.

2. Вера закутала младенца соседа, который выплюнул соску.

Vera wrapped the infant of the neighbor who spit out the pacifier.

3. Петя навестил адвоката приятеля, который защитил дело.

Petia visited the lawyer of the friend who defended the case.

4. Юра отчитал официанта консула, который обслужил банкет.

Sasha criticized the waiter of the consul who served a banquet.

5. Аня заметила уборщика министра, который вымыл посуду.

Ania noticed the cleaner of the minister who washed the dishes.

6. Юля оценила пианиста монарха, который сыграл сонату.

Julia appraised the pianist of the monarch who played the sonata.

7. Зина описала курьера чиновника, который доставил письмо.

Zina described the courier of the deputy who delivered the letter.

8. Витя похвалил повара капитана, который сварил бульон.

Vitia complimented the cook of the captain who cooked a broth.

9. Коля напугал бухгалтера начальника, который рассчитал налоги.
Kolia scared the accountant of the director who calculated the taxes.

10. Паша нарисовал кузнеца герцога, который подковал лошадь.
Pasha painted the blacksmith of the bishop who shoed a horse.

11. Нина уволила садовника депутата, который посадил яблони.
Nina fired the gardener of the councilman who planted apple trees.

12. Миша рассчитал портного директора, который пришил замок.
Misha dismissed the tailor of the director who sew down a zip fastener.

13. Рома наградил корректора писателя, который исправил ошибки.
Roma rewarded the proofreader of the writer who corrected the mistakes.

14. Лера догнала механика дедушки, который починил бампер.
Lera caught the mechanic of the grandfather who fixed the bumper.

15. Дима подозвал коллегу именинника, который подарил конфеты.
Dima called the colleague of the birthday boy who gave sweets.

Semantic Late Attachment (SemLA) Condition

1. Ксюша вызвала отчима школьника, который прогулял уроки.
Ksiusha called the stepfather of the pupil who skipped lessons.

2. Вика утешила клиента мошенника, который выманил кольцо.
Vika comforted the client of the crook who rooked the ring.

3. Юра освободил заложника террориста, который взорвал метро.

Jura released the hostage of the terrorist who blew up the metro.

4. Лёша успокоил пассажира таксиста, который проехал поворот.

Liosha calmed down the passenger of the taxi driver who drove past the turning.

5. Игорь выслушал пациента окулиста, который прописал очки.

Igor heard out the patient of the oculist who prescribed the glasses.

6. Гриша обрадовал тренера теннисиста, который выиграл кубок.

Grisha delighted the coach of the tennis player who won a cup.

7. Алина заметила ребёнка сантехника, который прочистил трубы.

Alina noticed the child of the plumber who unclogged the pipes.

8. Лена огорчила студента профессора, который запретил прогулы.

Lena upset the student of the professor who forbid truancy.

9. Жора шокировал соседа браконьера, который застрелил оленя.

Zhora shocked the neighbor of the poacher who shot a deer.

10. Света допросила племянника хакера, который взломал сервер.

Sveta questioned the nephew of the hacker who hacked a server.

11. Настя увидела земляка пожарного, который потушил склад.

Nastia saw the countryman of the fireman who extinguished the warehouse.

12. Боря задержал водителя разбойника, который ограбил банк.

Boria detained the driver of the bandit who robbed the bank.

13. Вадим порадовал кузена детектива, который раскрыл дело.
    Vadim delighted the cousin of the detective who solved the case.

14. Маша выручила опекуна картёжника, который проиграл ферму.
    Masha helped the caregiver of the gambler who gambled away the farm.

15. Вера пожалела правнука ветеринара, который усыпил собаку.
    Vera comforted the great-grandson of the veterinary who put down the dog.

Fillers

1. На празднике были конкурсы и караоке, которое раздражало Свету.
2. Статья, которую опубликовал в журнале Денис, шокировала учёных.
3. Сироп, который Соня привезла из Греции, очень понравился Наде.
4. Петя купил торт и билет на спектакль, в котором играла Виолетта.
5. Домик, который облюбовала кошка Алисы, смастерил Коля.
6. Документы, на которых стояла печать Оли, и отчёт проверил Сеня.
7. Фотографии Паши, которые лежали на шкафу, отдала на дачу Катя.
8. Карина открыла чемодан, в котором Женя тайно провёз ракушки.
9. Раскраски, в которых были герои мультфильмов, подарила Наташа.
10. В чулане валялась лампа и погремушка, которая растрогала Настю.
11. Маша спрятала дневник в шуфлядку, в которой хранили украшения.
12. Рюкзак, который привезла Алеся, и кеды забрала старшая сестра.
13. Саша перебрал бумаги, которые доставили вместе с батончиками.
14. В лесу выросли опята и земляника, которую очень любила Карина.
15. Полина распечатала плакат, который Денис сделал на конференцию.
16. Лена открутила со стены полку, на которой стоял аквариум Игоря.
17. За завтраком Артём выпил капучино, который стоял в кружке Лены.
18. Вита пошила слоника и подушки, которые положила на кресло.
19. Сарай и баня, которую построил Андрей, простояли несколько лет.

20. Никита поставил палатку на поляне, на которой осталось кострище.
21. В клубе стояли автоматы и бильярд, который заинтересовал Пашу.
22. Презентация, в которую добавил картинок Дима, всех рассмешила.
23. Оксана украла шкатулку, в которой Маша хранила свои серёжки.
24. Слава поломал вешалку, на которой висел джемпер Наташи.
25. Таня отнесла велосипед в подвал, в котором стоял рояль Жени.
26. В парке были аттракционы и клоун, который очень пугал Игоря.
27. Сладости и кофе, который сделала Вика, оценила вся семья.
28. Марат опрокинул ящик, в котором привезли хрустальные бокалы.
29. Дима вытряхнул коврики, которые лежали в машине Марины.
30. Вера повесила полотенце на батарею, которая стояла у кровати Лёши.
31. В цирке продавали сок и сладкую вату, которая порадовала Мишу.
32. Километраж, который проехал на велосипеде Женя, впечатлил маму.
33. В гостиной лежали индийские сувениры и плед, который удивил Петю.
34. Даша опрокинула горшок с геранью, которую выращивал Егор.
35. Марта израсходовала все румяна, которые заказала себе Арина.
36. Вика кинула сумку на стул, на котором висела куртка Толика.
37. Боря сдал в ломбард серьги и плащ, который недавно купила Ира.
38. Вова поехал рыбачить на лодке, в которой лежали удочки Никиты.
39. Блюдо и кувшин, который разрисовала Люда, отдали знакомой отца.
40. В пещере были сталактиты и ручей, который сфотографировала Аня.
41. Женя перенёс жасмин на балкон, на котором выращивали помидоры.
42. Костя посетил деревню, в которой была древняя белая синагога.
43. Анжела вытряхнула сумку, в которой хранили детские игрушки.
44. Андрей выкинул приставку и картриджи, которые так любил Никита.
45. Люба отдала в детдом пакет, в котором лежали игрушки Вити.
46. Антон решил задачу, которую он нашёл в портфеле Оксаны.
47. Петя перенёс на подоконник вазу, в которой стояли розы Лизы.
48. История, которую отправил на конкурс Егор, растрогала читателей.
49. Джинсы и кофта, которая была из хлопка, приглянулись Кириллу.

50. Марк сложил деньги в конверт, который достал из маминой сумки.

51. Ярослав отмыл машину, на которой Анжелика въехала в грязь.

52. Пейзаж, который нарисовала Лена, и портрет повесили на даче.

53. На сайте было много рекламы, которая очень раздражала Риту.

54. Вера пересыпала гречку в банку, в которой обычно хранили муку.

55. Каролина разбила блюдце и кружку, в которую Миша налил глинтвейн.

56. Халву и статуэтку, на которую повесила серьги Аня, привёз Марат.

57. Егор забрал на почте посылку, в которой пришли сапоги Ксюши.

58. Платье и корсет, который купила мама, очень понравились Вере.

59. В порт приплыли военные корабли, которые впечатлили Марину.

60. Антон упал с крутых ступенек, которые выходили на балкон.

61. Желе и лимонад, который похвалила Катя, приготовил Артур.

62. Саша прибрался в комнате, в которую вчера переехала Марина.

63. Люда вырвала страницы из книжки, которую читал Максим.

64. Костёр, который разжёг Петя, и чай быстро отогрели Лизу.

65. Алеся выставила на аукцион ящик, в котором были вещи Лёши.

66. На свадьбе был тамада и коктейли, которые понравились Мише.

67. Гриша испёк ватрушки по рецепту, который нашёл в журнале.

68. В самолёте разносили бутерброды и кофе, который похвалила Оля.

69. Саша заселилась в отель, в котором прошлым летом работал Артём.

70. Компот, который сварила из сухофруктов Зина, был слишком сладким.

71. В бассейне плавали круг и маска, которая подошла Юре.

72. Саша починил машину ключом, который папа купил на рынке.

73. Кассету, на которой был утренник Марты, и диски отдал Витя.

74. Фильм и интервью, которое сняла Рая, обсуждали несколько дней.

75. Снежана оформила отчёт по шаблону, который лежал на столе.

76. Артём забрался на крышу, с которой был виден старый пивзавод.

77. Шторы, которые висели в комнате Андрея, и скатерть постирала Аня.

78. Костя пошёл в школу в свитере, который Света привезла из Парижа.

79. Конкурс, на который подала заявку Настя, организовал Олег.

80. На конференции выдали бейджи и номер, который огорчил Славу.

81. Хомяка и пуделя, которого подстриг Рома, завела Рита.

82. Кулон и кольцо, которое заказала Зина, украла домработница.

83. Алина спрятала записку в книжку, в которой брат хранил деньги.

84. Платье, которое пошила Полина, и юбку Гриша отвёз в деревню.

85. Рита пересадила розу в горшок, в котором недавно рос папоротник.

Appendix B: Verb frequency

No.	Experimental condition	Russian verb	English translation	Frequency per million words in Russian National Corpus
1	AmbA	обронить (ключи)	drop (keys)	5.8
2	AmbA	выписать (штраф)	issue (fine)	17.2
3	AmbA	оформить (заказ)	place (order)	28.1
4	AmbA	одобрить (аборты)	encourage (abortions)	25.4
5	AmbA	замазать (афишу)	paint out (advertising bill)	2.8
6	AmbA	подобрать (щенка)	pick up (puppy)	34.9
7	AmbA	оценить (проект)	evaluate (project)	64.0
8	AmbA	свернуть (бизнес)	close (business)	30.3
9	AmbA	подхватить (грипп)	catch (flu)	32.5
10	AmbA	спрятать (карту)	hide (map)	38.8
11	AmbA	провезти (виски)	smuggle (whiskey)	1.3
12	AmbA	отложить (отпуск)	postpone (vacation)	26.1
13	AmbA	отключить (телефон)	switch off (phone)	10.6
14	AmbA	поранить (руку)	cut (arm)	1.6
15	AmbA	закурить (трубку)	start to smoke (pipe)	26.6
1	SynEA	уронить (ноты)	drop (notes)	15.3
2	SynEA	подделать (пропуск)	forge (pass)	1.7
3	SynEA	проспать (поезд)	miss (train)	6.9
4	SynEA	покрасить (забор)	paint (fence)	5.9
5	SynEA	разбить (машину)	break (car)	36.7
6	SynEA	продать (виллу)	sell (villa)	68.3
7	SynEA	приютить (голубя)	shelter (pigeon)	3.6
8	SynEA	принести (шашки	bring (checkers)	141.9
9	SynEA	поломать (очки)	break (glasses)	6.1
10	SynEA	спутать (педали)	confuse (pedals)	8.3
11	SynEA	уточнить (график)	clarify (schedule)	26.3
12	SynEA	проглотить (монету)	swallow (coin)	14.1
13	SynEA	нарушить (закон)	violate (law)	35.9
14	SynEA	приобрети (участок)	buy (lot)	70.2
15	SynEA	испачкать (костюм)	stain (costume)	3.3
1	SynLA	пропустить (завтрак)	miss (breakfast)	41.7
2	SynLA	выплатить (кредит)	pay back (loan)	8.9
3	SynLA	вывихнуть (плечо)	twist (shoulder)	1.0
4	SynLA	сломать (ключицу)	break (clavicle)	30.4
5	SynLA	забросить (съёмки)	abandon (shooting)	20.1
6	SynLA	изобрести (вакцину)	invent (vaccine)	15.6
7	SynLA	покинуть (страну)	leave (country)	49.0
8	SynLA	обстрелять (причал)	pelt (mooring)	2.5
9	SynLA	покорить (Эверест)	climb (Everest)	8.3
10	SynLA	перенести (концерт)	postpone (concert)	36.6
11	SynLA	возглавить (кафедру)	lead (chair)	19.8
12	SynLA	поджечь (дверь)	set on fire (door)	7.2
13	SynLA	завершить (карьеру)	finish (career)	32.7
14	SynLA	закатить (скандал)	make (scene)	4.7
15	SynLA	прекратить (атаку)	stop (attack)	43.4

No.	Experimental condition	Russian verb	English translation	Frequency per million words in Russian National Corpus
1	SemEA	вправить (вывих)	set (bone)	1.1
2	SemEA	выплюнуть (соску)	spit out (pacifier)	3.3
3	SemEA	защитить (дело)	defend (case)	39.1
4	SemEA	обслужить (банкет)	serve (banquet)	2.2
5	SemEA	вымыть (посуду)	washed (dishes)	13.6
6	SemEA	сыграть (сонату)	play (sonata)	80.8
7	SemEA	доставить (письмо)	deliver (letter)	37.0
8	SemEA	сварить (бульон)	cook (broth)	14.5
9	SemEA	рассчитать (налоги)	calculate (taxes)	38.1
10	SemEA	подковать (лошадь)	shoe (horse)	2.0
11	SemEA	посадить (яблони)	plant (apple trees)	52.9
12	SemEA	пришить (замок)	sew (zip fastener)	5.9
13	SemEA	исправить (ошибки)	correct (mistakes)	18.5
14	SemEA	починить (бампер)	fix (bumper)	6.6
15	SemEA	подарить (конфеты)	give (sweets)	61.4
1	SemLA	прогулять (уроки)	skip (lessons)	1.3
2	SemLA	выманить (кольцо)	rook (ring)	1.5
3	SemLA	взорвать (метро)	blow up (metro)	13.7
4	SemLA	проехать (поворот)	drive past (turning)	23.3
5	SemLA	прописать (очки)	prescribe (glasses)	16.4
6	SemLA	выиграть (кубок)	win (cup)	52.3

7	SemLA	прочистить (трубы)	unclog (pipes)	1.6
8	SemLA	запретить (прогулы)	forbid (truancy)	35.5
9	SemLA	застрелить (оленя)	shoot (deer)	10.2
10	SemLA	взломать (сервер)	hack (server)	3.0
11	SemLA	потушить (склад)	extinguish (warehouse)	4.4
12	SemLA	ограбить (банк)	rob (bank)	7.4
13	SemLA	раскрыть (кражу)	solve (case)	40.8
14	SemLA	проиграть (ферму)	gamble away (farm)	30.1
15	SemLA	усыпить (собаку)	put down (dog)	2.4

Appendix C: Plausibility pre-test results

Sentence	Mean plausibility of attachment to NP1	Mean plausibility of attachment to NP2
Катя увидела ассистента лектора, который обронил ключи.	6.118	6.392
Боря обокрал заместителя аудитора, который выписал штраф.	4.854	4.235
Егор вспомнил охранника олигарха, который оформил заказ.	4.804	4.765
Петя выслушал оппонента политика, который одобрил аборты.	4.078	5.314
Сеня припугнул приятеля студента, который замазал афишу.	5.608	6.294
Лиза привела сверстника подростка, который подобрал щенка.	5.980	6.667
Оля обманула посредника инвестора, который оценил проект.	5.157	6.608
Паша обвинил поставщика ресторатора, который свернул бизнес.	5.373	6.471
Валя отпустила стажёра инструктора, который подхватил грипп.	6.588	6.510
Игорь схватил напарника командира, который спрятал карту.	6.157	6.039
Миша упомянул попутчика бизнесмена, который провёз виски.	5.882	5.510
Люда отыскала секретаря министра, который отложил отпуск.	5.118	4.784
Юра встретил помощника прокурора, который отключил телефон.	6.314	5.529
Рита допросила коллегу бригадира, который поранил руку.	6.627	6.608
Лёша подозвал оператора журналиста, который закурил трубку.	5.137	5.824

Appendix D: Noun frequency

No.	Condition	NP	Russian word	Frequency per million words in Russian National Corpus
1	AmbA	NP1	ассистент	8.4
2	AmbA	NP1	напарник	9.5
3	AmbA	NP1	попутчик	6.4
4	AmbA	NP1	секретарь	77.0
5	AmbA	NP1	помощник	57.7
6	AmbA	NP1	коллега	30.7
7	AmbA	NP1	оператор	31.9
8	AmbA	NP1	заместитель	99.3
9	AmbA	NP1	охранник	33.3
10	AmbA	NP1	оппонент	16.0
11	AmbA	NP1	приятель	53.9
12	AmbA	NP1	сверстник	10.5
13	AmbA	NP1	посредник	13.1
14	AmbA	NP1	поставщик	24.0
15	AmbA	NP1	стажёр	1.9
1	AmbA	NP2	лектор	12.1
2	AmbA	NP2	командир	110.9
3	AmbA	NP2	бизнесмен	29.5
4	AmbA	NP2	министр	151.4
5	AmbA	NP2	прокурор	57.4
6	AmbA	NP2	бригадир	16.3
7	AmbA	NP2	журналист	7.1
8	AmbA	NP2	аудитор	4.5
9	AmbA	NP2	олигарх	30.5
10	AmbA	NP2	политик	45.5
11	AmbA	NP2	студент	105.4
12	AmbA	NP2	подросток	35.4
13	AmbA	NP2	инвестор	33.0
14	AmbA	NP2	ресторатор	0.8
15	AmbA	NP2	инструктор	13.9
1	SynEA	NP1	ученик	80.4
2	SynEA	NP1	шофёр	42.9
3	SynEA	NP1	ровесник	7.6
4	SynEA	NP1	первенец	3.4
5	SynEA	NP1	соавтор	9.3
6	SynEA	NP1	ухажёр	2.3
7	SynEA	NP1	партнёр	67.1
8	SynEA	NP1	менеджер	28.8
9	SynEA	NP1	племянник	15.8
10	SynEA	NP1	мальчишка	56.0
11	SynEA	NP1	сожитель	1.5
12	SynEA	NP1	наследник	22.6
13	SynEA	NP1	практикант	1.2
14	SynEA	NP1	квартирант	1.6
15	SynEA	NP1	любимчик	1.3
1	SynEA	NP2	скрипачка	0.7
2	SynEA	NP2	чемпионка	3.2
3	SynEA	NP2	аспирантка	3.2

4	SynEA	NP2	стюардесса	4.4
5	SynEA	NP2	поэтесса	5.7
6	SynEA	NP2	дочь	142.4
7	SynEA	NP2	балерина	8.1
8	SynEA	NP2	гимнастка	1.3
9	SynEA	NP2	хозяйка	54.2
10	SynEA	NP2	соседка	28.4
11	SynEA	NP2	студентка	11.2
12	SynEA	NP2	принцесса	15.7
13	SynEA	NP2	сотрудница	7.0
14	SynEA	NP2	сокурсница	0.5
15	SynEA	NP2	бабушка	101.8
1	SynLA	NP2	пенсионер	93.4
2	SynLA	NP2	хореограф	2.8
3	SynLA	NP2	академик	49.2
4	SynLA	NP2	москвич	52.8
5	SynLA	NP2	хоккеист	4.8
6	SynLA	NP2	главврач	4.8
7	SynLA	NP2	генерал	140.0
8	SynLA	NP2	знакомый	71.3
9	SynLA	NP2	сотрудник	135.1
10	SynLA	NP2	сноубордист	2.8
11	SynLA	NP2	режиссёр	103.5
12	SynLA	NP2	профессор	111.8
13	SynLA	NP2	дипломат	19.1
14	SynLA	NP2	штурмовик	5.5
15	SynLA	NP2	альпинист	4.3
1	SynLA	NP1	сиделка	23.9
2	SynLA	NP1	помощница	3.9
3	SynLA	NP1	ученица	7.1
4	SynLA	NP1	падчерица	1.3
5	SynLA	NP1	фанатка	0.5
6	SynLA	NP1	акушерка	1.7
7	SynLA	NP1	разведчица	0.5
8	SynLA	NP1	бабушка	101.8
9	SynLA	NP1	невеста	31.7
10	SynLA	NP1	девушка	213.3
11	SynLA	NP1	любовница	16.5
12	SynLA	NP1	спутница	6.2
13	SynLA	NP1	подруга	79.4
14	SynLA	NP1	медсестра	16.1
15	SynLA	NP1	супруга	24.4
1	SemEA	NP1	хирург	22.1
2	SemEA	NP1	кузнец	6.6
3	SemEA	NP1	садовник	5.4
4	SemEA	NP1	портной	4.5
5	SemEA	NP1	корректор	2.1
6	SemEA	NP1	механик	13.0
7	SemEA	NP1	коллега	80.6
8	SemEA	NP1	младенец	23.6
9	SemEA	NP1	адвокат	49.3
10	SemEA	NP1	официант	1.2
11	SemEA	NP1	уборщик	1.2
12	SemEA	NP1	пианист	10.1
13	SemEA	NP1	курьер	6.8
14	SemEA	NP1	повар	13.3
15	SemEA	NP1	бухгалтер	15.5
1	SemEA	NP2	жених	25.2

2	SemEA	NP2	герцог	6.7
3	SemEA	NP2	депутат	108.7
4	SemEA	NP2	директор	222.2
5	SemEA	NP2	писатель	166.3
6	SemEA	NP2	дедушка	48.8
7	SemEA	NP2	именинник	4.7
8	SemEA	NP2	сосед	97.8
9	SemEA	NP2	приятель	53.9
10	SemEA	NP2	консул	4.8
11	SemEA	NP2	министр	151.4
12	SemEA	NP2	монарх	6.6
13	SemEA	NP2	чиновник	75.9
14	SemEA	NP2	капитан	119.0
15	SemEA	NP2	начальник	198.9
1	SemLA	NP1	отчим	5.5
2	SemLA	NP1	племянник	15.8
3	SemLA	NP1	земляк	13.5
4	SemLA	NP1	водитель	65.9
5	SemLA	NP1	кузен	2.7
6	SemLA	NP1	опекун	3.2
7	SemLA	NP1	правнук	3.0
8	SemLA	NP1	клиент	87.8
9	SemLA	NP1	заложник	19.9
10	SemLA	NP1	пассажир	47.4
11	SemLA	NP1	пациент	35.8
12	SemLA	NP1	тренер	37.1
13	SemLA	NP1	ребёнок	658.3
14	SemLA	NP1	студент	105.4
15	SemLA	NP1	сосед	97.8
1	SemLA	NP2	школьник	27.6
2	SemLA	NP2	хакер	2.7
3	SemLA	NP2	пожарный	6.2
4	SemLA	NP2	разбойник	17.1
5	SemLA	NP2	детектив	10.7
6	SemLA	NP2	картёжник	0.8
7	SemLA	NP2	ветеринар	4.0
8	SemLA	NP2	мошенник	6.4
9	SemLA	NP2	террорист	32.3
10	SemLA	NP2	таксист	8.9
11	SemLA	NP2	окулист	0.8
12	SemLA	NP2	теннисист	3.4
13	SemLA	NP2	сантехник	3.0
14	SemLA	NP2	профессор	111.8
15	SemLA	NP2	браконьер	2.2

ISBN 978-83-932687-5-7