Bottom-up parsing approach to modelling local coherence effects

University Utrecht

Kunstmatige Intelligentie

Niek van der Gaag

Supervisor: Jakub Dotlacil

Second Assesor: Sanne Boing

June 2021

Abstract

In the field of psycholinguistics, sentence interpretation has huge discussions due to the level of ambiguity. Human sentence processing occurs incrementally. The central question of grammatical constraints lies with the levels of ambiguity. Several theories and parsing models have been tested to attack distracting effects of wrongly interpreted sentences. Previous Noun Phrase-Verb Phrase sequences have been tested on human reading times and found local syntactic coherence effects. This paper asks if a bottom-up transition-based parsing model can predict activation numbers that compare to human results in order to further support the theory of local coherence effects. We ran 20 sentences in 4 different conditions through a bottom-up parser and used mean activations to showcase the effects of local ambiguity. Our data shows the results are consistent with the hypothesis and show promising results for further research in modelling local coherence effects.

Keywords: sentence processing, local coherence, Good Enough parsing, garden-path sentences, bottom-up parsing, transition-based parsing

Contents

Chapter 1

Introduction

In the world of language processing, sentence interpretation has huge discussions due to the level of ambiguity. The role of grammatical constraints have been a central point of discussion in the past few decades within the field of psycholinguistics. Modern linguistics is often concerned with structures, patterns and theories of those structures. With its close links to cognitive science, proposals of different theories are trying to solve the problem of relating grammars and parsers. How exactly grammar fits into the rest of science has been studied for decades. The role of meaning plays an important part in language processing. Mainly since semantic values have an influence in the structural division of a sentence. The vision on how we approach parsing bottles down to how grammar is conceptualized. The notion of grammar which we use is that of Chomsky (1957): language structures are derivable in a formal system that constitutes a theory of them. Sentence processing is a task which most people perform adequately. Understanding language which is hardly ever failed by most, has been found difficult to theorise due to its extreme complex features and level of ambiguity. There have been a number of proposals that suggest how language is deemed to be processed in order to deal with the problems that occur in sentence ambiguity. Ambiguity occurs when a sentence has two or more possible meanings within that single sentence. Ambiguity lends itself to cognitive psychology and psycholinguistics. Within this field of psycholinguistics we encounter garden path sentences which are a form of an ambiguous sentence. A lot has been studied about garden path sentences, which refers to a grammatically correct sentence that most reader will likely interpret the wrong way. Garden path phenomenon occurs when the parser encounters a temporarily ambiguous sentence. Difficulty arises when a multitude of parses are made because of the grammar and level of ambiguity. Early psycho-lingual theories suggest the reader ranks each parse parallel where the wrong one is often ranked the highest. Later literature found that human sentence processing is known to proceed incrementally. The syntactic and semantic interpretations of a sentence unfolds gradually as readers read the sentence, as opposed to waiting for the entire sentence to finish and then observe an interpretation. Different phenomena are being questioned as to what decisions made early in the sentence have an effect in the parsing later on.

In this paper, we consider the possibility of local syntactic coherence in the input as a theory of syntactic parsing. Earlier findings of local coherence accounts (LCAs) stem from an approach to parsing that we refer to as "Self Organized" parsing (Vosse and Kempen, 2000). Under the version of self-organized parsing Tabor and Hutchins (2003), found that later arriving words can form a syntactic combination that is inconsistent with the global parse. The findings suggest reading times in garden path sentences are higher on the ambiguous verb than in control conditions. We intend to further investigate the role of ambiguity in garden path sentences and look for correlations between sentence parsing and local syntactic coherence effects. Since the effects of local coherence have been previously studied we aim to predict those findings with a bottom-up transition based parsing model to see if the model predictions are consistent with previous human data results. This papers asks how well a bottom-up parsing model can predict the effects of local coherence by previous findings of garden-path examples. We expect to see consistent similarities between the human data results from tabor et al. (2004) and the predictions from the bottom-up transition-based parsing model retrieved from Dotlacil (2021).

1.1 Connection to AI

In the field of Artificial Intelligence, language processing is a widely discussed topic along psycholinguists, cognitive scientists and data scientists. The processing of natural language is necessary when a user wants intelligent systems like robots or algorithms to perform like the instructions plead. In hospitals where intelligent systems and computers are being introduced to function as dialogue based clinical expert systems, it is extremely important to process language both syntactically and semantically. The method for processing language is called Natural Language Processing (NLP). Input and output of NLP systems can be both speech and written text. There are two main components with NLP: Language Understanding and Language Generation.

Natural Language Understanding (NLU) involves the understanding of the language. It involves tasks like mapping a certain input in natural language and try putting that information into useful representations so the intelligent system can use the data. Another important task is analyzing the different aspects of the language. We as humans don't always understand language to the fullest, having intelligent systems that can analyze and conceptualize such languages offers great support in research.

Language Understanding is a bit more challenging than Language Generation. Some of the major difficulties arise due to the level of ambiguity. Because of the structure and form of natural language, there are different levels of ambiguity. Lexical ambiguity is at a base level. For example, it can be about deciding weather a certain word should be treated and seen as a verb or a noun. Syntactic ambiguity is the level we are most interested in. Syntax level ambiguity is about how sentences can be and should be parsed. Reading and parsing in the field of psycholinguistics is a difficult task, therefore there have been multiple researches on how we can model these tasks correctly. A multitude of theories are formed in order to conceptualize the effects of ambiguity levels. The misinterpretations of ambiguous sentences like garden-path sentences are heavily discussed in the field of psycholinguistics. Over the last decades, researches have tried to model the effects and try to conceptualize how we as human readers interpret ambiguous sentences. With the rise of intelligent systems that are better at conceptualizing ambiguous sentences in sentence processing, more models are created to try and predict the outcomes of human parsers. Behaviour of human language processing is being used in order to create parsing models that are able to predict similar effects of human parsing. Local Coherence effects are a prime example of how theories are formed that try to explain the local and global ambiguities effects. Since a lot of research on Local Coherence effects has been conducted previously, we are able to try and see if those effects can be predicted by models that compute cognitive memory abilities from human parsers. Ultimately, this research is proposing a transition-based as well as a bottom-up parsing model that predicts the reading times of human parsers. With those predictions we are able to compare the results with human data where local coherence effects are believed to be present, and see if the model is able to predict those effects as well.

Chapter 2

Informational Background

2.1 Garden-Path

Garden-pathing is a good example of the core problem with ambiguity in sentence processing. Garden-path sentences in psycholinguistics are used to demonstrate when humans read, they process one word at a time. The definition of garden-path stems from being deceived or tricked. One of the more famous examples used for garden-path sentences is shown in the following sentence

1: *The horse raced past the barn fell.*

This sentence is a well-formed grammatically correct English sentence. This sentence could very well be generated by a correct competence grammar. It can't go unnoticed that for most people the sentence feels odd to read and interpret. The process of the comprehension shows an interesting phenomenon. Readers would often act surprised once arrived at the word "fell". When reaching this last word in the sentence, it becomes evident that the substring "The horse (noun phrase) raced (verb) past the barn (prepositional phrase)" was not in fact a full sentence. The true meaning and interpretation of sentence *a* can be shown with an extended version of the sentence

2: *The horse that was raced past the barn fell.*

We can see how in sentence (2) the initial substring is a reduced version of the relative clause. There have been all sorts of methods trying to find the effects of garden-path instances. One of the empirically suggested method is measuring reading times on the critical word, in the case of sentence (1) and (2): "fell". The trickery and magic in garden-path sentences are often compared to the cognitive behaviours of optical illusions. With this in mind, the effect of garden-path sentences are a selection of structural alternatives that are wrongly interpreted.

An important aspect of sentence processing is the fact that it is compositional, meaning the interpretation of a sentence is determined by the way its words can be combined. But we can't gather any information from previous encounters because that would imply we wouldn't be able to parse novel sentences. The interpretation is constructed when one is confronted with a new example of a sentence.

We must point out that sentences like (1) are only ambiguous if we assume incremental interpretation. Frazier and Rayner (1982) show that it is possible to study garden paths using eye tracking movements, but don't offer incremental explanation. Incremental interpretation has been argued before this however. Frazier (1979) and Marslen-Wilson (1973) explain and claim that people find difficulties with garden path sentences and suggest the sentences are analysed incrementally based on garden paths. With that, the empirical results suggest the direct object is analysed initially. Moving to the embedded sentence causes for increased complexity when reading. Since the realisation of the true interpretation happens late into reading the sentence, conscious garden paths which are also described as local ambiguities have the effect of leading a ready towards a wrong analysis. After the wrong interpretation is made it is difficult to recover from this mistake as a reader.

2.2 Local Coherence

Most of the theories on ambiguous parsing include the Garden Path Theory (Frazier, 1979). This theory includes a serial modular parsing model. It suggests that a single parse is constructed by a syntactic module. Semantic and contextual factors influence processing at a later stage and can induce re-analysis of the syntactic parse. Garden-path sentences have been studied intensely well over the last decades. It has brought huge developments and new insights into human sentence parsing and cognitive processing. Tabor et al.(2004) differed from earlier theories and proposed a different structure for local ambiguities. Tabor and colleagues suggested the possibility of a Local Syntactic Coherence structure. They provided evidence that proposes local parsers have a detectable influence on the time course of processing. The difficulty arises in the "merely local coherent" string. It must be said that there is no structural or formal difference between garden-paths and local coherence effects. Most literature would describe local coherence as the appearance where previous syntactic context should rule out the misparse of the locally coherent substring. Sentences that are long enough can contain sequences that form partial or whole sentences on their own if they are viewed in isolation. Meaning they are functionally containing a substring. Take the example sentence

3: *The coach smiled at the player tossed a frisbee by the opposing team.*

In the case of (3), the meaning of the embedded clause diverges from the global content of the entire sentence. In theory, the global content should not be considered a possible interpretation since it is ruled out by the context of the latter words in the sentence. However, previous findings have made several attempts at explaining the effects of local coherence and theorize local syntactic coherence (LSC) affect human sentence processing. This theory questions the global consistency assumption which holds most theories of language processing.

We define local syntactic coherence as such; sequences of two or more words in the text stream which form a phrase that cannot be grammatically unified with the parse of the preceding words. "The player tossed a frisbee" in sentence (3) has an interpretation that is locally coherent. However, if we look at the context of the sentence which favors "the player" as the object of "smiled at", "tossed a frisbee" can only be parsed as a reduced relative clause of "player". Local coherence effects is about finding weather or not readers would still be distracted by the active interpretation of the sentence. Christianson et al. (2017) conducted eye-tracking experiments to compare the effects of garden-path sentences and local coherence effects. Their results showed that local coherence structures elicited signals of reading disruption that arose earlier and lasted longer than that of garden-path comprehension. If we were to compare the ambiguous sentence of 3 to a control case where we have a less distracting local coherence we might see a contrast between the two conditions and obtain an explanation. A control case of (3) could very well be:

4 *The coach smiled at the player who was tossed a frisbee by the opposing team.*

This unreduced version of sentence (3) portrays a version of the ambiguous sentence that is believed to be less difficult to read. "The player tossed the frisbee" in (3) is by most readers interpreted as an active clause, but this finds no parse in sentence (4).

The effects of local coherence are very specific. In order to further investigate and predict these effects, we need to look at previous conducted research and theories, as well as defining what parsing approach is suitable for finding the effects.

Chapter 3

Theoretic Positioning

This chapter covers different viewpoints on local coherence effects. Different literature is discussed to provide a multitude of models that have a different approach to interpret local coherence effects. We further investigate these approaches in the first section and explain how the models of these approaches work. We later explain how the parsers are used from an algorithmic viewpoint in the sections of both bottom-up parsing and transition-based parsing. These two parsing approaches are two efficient ways of tackling the problems of local coherence effects. We later discuss how in our own model, we combine both parsing approaches to form what is our parsing model.

3.1 Earlier models of Local Coherence effects

The appearance of local coherence and it's effects have been measured and predicted by previous studies. Different results and different case studies make it so the theory on what the effects of local syntactic coherence are, still cause for debate. Generally, the difference in theories are because of the effects of the distinction between a reduced and unreduced sentence. Where the reduced sentence is believed to be more ambiguous than the unreduced sentence. There are two general types of explanations if we were to see a distinction between the two cases:
1) Local Coherence Accounts
2) Self-Consistent Parse Accounts
We tend to mainly focus on Local Coherence Accounts since this explanation stems from a parsing approach called "Self-Organized" parsing (Kempen and Vosse, 1989). With this way of parsing if a group of later-arriving words can form a syntactic combination that is inconsistent with the global parse, then this combination will form and compete with the main parse. It is predicted that with

Local Coherence Accounts at (or shortly after) the critical word in the sentence, which is the second verb in the sentence, the reading time slows down. If we then present a control case like an unreduced version of the sentence, the local formation will see very little competition of this form and reading times will be significantly faster. Tabor and Hutchins (2003) describe the Self-Organized model as a parser where reading a word activates a set of lexical anchored tree fragments. Those fragments spread activation to compatible fragments, so that the system stabilizes to a correctly parsed sentence. The class of models that assume self-organized parsing models are unified as 'SOPARSE', which stands for Self-Organized Parse accounts.

Levy (2008) used a noisy-channel model to argue in favor of local coherence effects. Other theories that have been tested used 'clean' tokens to find the optimal solution to the problem of language comprehension. Levy (2008) suggests to use a novel yet simple noisy-channel model of sentence comprehension that uses probabilistic features. Testing language comprehension this way, there is uncertainty about the word-level representation which was thought of as a better way to represent language processing. Levy predicted that larger changes in beliefs results into greater processing difficulty and longer reading times. The model had the capability to read the effect of every word consistently. However, some of its flaws lie within the grammatical parses as it does not capture Good Enough parsing (GE). The theory of *Good Enough* parsing is being described as: People sometimes compute local interpretations which are inconsistent with the overall sentence structure, this indicates that the overall comprehension system can be lazy about computing a more global meaning (Ferreira and Patson., 2007). This is relevant because good enough parsing can be seen as a local coherence effect. It originates from the idea that people construct a representation of a sentence that is just good enough to complete the task (the parse) at hand. Since this representation is locally influenced by the local information (partial parse trees), the effects of the local coherence are visible. Good enough language processing approach emphasizes people's tendency to generate superficial and even inaccurate interpretations of sentences (Ferreira. Lowder., 2016).

Another previous model that conceptualizes local coherence effects has been founded by Bicknell and Levy (2009). Local coherence effects are viewed as resulting from a belief-update process and show that the relevant probabilities in the model are calculable from a probabilistic parser. Incrementally processing a sentence is being viewed as a process of updating one's beliefs. The effects of local coherence are being modelled as the consequences of an update from a bottom-up prior belief to a posterior belief that uses top-down information. The find difficulty in the locally coherent substring because the bottom-up beliefs make strong predictions on the category of the substrings. But this is information is being contradicted

with the top-down information. This model uses the GE parsing idea and looks at substrings of different lengths. The problem is, it cannot integrate information of substrings from different lengths.
There is no generic widely covered and accepted model of human parsing that has been able to implement GE parsing strategy. Instead, The computational models have thus far had to account for local coherence effects indirectly.
In order to possibly solve these problems of previous theories we investigate how a combination of Bottom-up parsing and transition-based parsing can offer information from subtrings that have different lengths.

3.2 Bottom-up Parsing

Bottom-up parsing have been a popular parsing method for long periods of time in the field of AI. It is a parsing strategy that looks at the lowest level of the parsing tree first, then works it way up the tree by using a set of grammatical rules. In contrast to a top-down algorithm, bottom-up's main decision is to select when to use a production rule in order to reduce to string to get the starting symbol. A 'shift-reduce' parser is a bottom-up parsing model in its simplest instance. A shift-reduce algorithm is arguably the most common type of a bottom-up parser. The parser looks at the words in the sentence and combines them into constituents using grammar rules. The two fundamental parsing operations; 'shift' and 'reduce' move the algorithm to the next word in the sentence and combines the found constituents into new constituents respectively. Found categories are pushed onto a 'stack' in order to keep track of all categories. When two categories appear on top of the stack, the algorithm reduces them by using grammatical rules. The only category remaining is the S which presents the parsed sentence. A 'shift-reduce' parser makes a decision with every word in the sentence: either shifting to the next word by determining the category or reducing if categories on the stack match those of a grammatical rule.

There are some known problems with bottom-up parsing in the field of psycholinguistics. It may receive problems when tied to incremental interpretation (Crocker, 1999). We use a bottom-up parsing algorithm since it is arguably the most common parser that we can use to parse phrase structures that are transition-based.

3.3 Transition-based Parsing

Transition-based parsing is a fast and effective approach for dependency parsing. A dependency parser processes an input and predicts a sequence of parsing

actions. Transition-based parsing models have parsing systems that follow decisions made by a classifier, predicting transitions from one state to another. The parser starts with an initial configuration. At every step, the parser asks a guide to choose between one of multiple transitions (actions) into new configurations. The parsing stops if the parser reaches a so called 'terminal configuration'. The parser then returns the dependency tree associated with the terminal configuration. Transition-based parsing approaches have achieved great success for tree parsing in order to build general dependency graphs (Zhang et al., 2016).
In a typical transition-based parsing process, the input is put into a queue and partially built structures are organized by a stack. The parser must decide what transition is the appropriate transition given the configuration. Our parsing algorithm works with 2 databases. The constructed trees S and a stack of words with their POS-tags W.

3.4 The model

This study uses a transition-based bottom-up parsing model that has been trained on data in the Penn Treebank (Marcus et al., 1993). It must be highlighted that the model uses a combination of both a bottom-up parsing algorithm as well as a transition-based parsing algorithm specifically. Using a transition-based bottom-up parsing algorithm, we can create a parser that heavily relies on local context. This is extremely useful since we want to generate information on the local ambiguity. Predicting features that can show effects of local coherence is what is necessary to make claims about local coherence effects. Our model is believed to be able to capture locally coherent reading which distracts from the globally coherent interpretation. Strictly bottom-up parsing was claimed to be unable to model human parsing since it does not allow for incremental interpretation. Strictly speaking, this means nothing can be interpreted until the very end of the sentence because it is only then that the parse is supposedly completed. To overcome this argument, we combine our bottom-up parsing model with that of a transition-based model. The model in question was previously used by Dotlacil and de Haan (2021) to explore the rational theory of memory and human parsing. It is assumed that declarative memory consists of chunks that represent correct parsing steps. Those chunks are collected from the data that was retrieved from the Penn Treebank. Previous findings on syntactic parsing predictions show this class of parsers can be used in computational linguistics and are compatible with parsing rules in memory (Dotlacil and de Haan). Transition based parsing models are compatible with memory structures and can be embedded in 'Adaptive Control of Thought-Rational' (ACT-R), which is a cognitive architecture to model the

structure of our cognitive memory and is one of the best hypothesis about such an architecture (Anderson et al., 2004).
Effects of activation might show spill over effect. Spill over effect implies the activation levels might spill over to the next word in the sentence. This effect only occurs in one direction since a sentence is read from left to right.

Chapter 4

Experiments

We present one set of an experiment that has been split between retrieving the model and processing prior information to conduct analysis on the results.

4.1 Method and Materials

All 20 experimental sentences used to conduct the experiment were pre-processed and prepared for parsing and finding traceable activation levels. The 20 sentences used for parsing are listed in Appendix A and were retrieved from Tabor et al. (2004). Every single sentence involved 4 conditions:
A/R, Ambiguous, Reduced; A/U, Ambiguous, Unreduced; U/R, Unambiguous, Reduced; U/U. Unambiguous, Unreduced.
In total, 80 experimental sentences (4 conditions per item with 20 items in total) were created. Every item was included with a noun phrase in a non-subject position which was modified by a relative clause in passive voice. Relative clauses were either reduced or unreduced. Verbs were either ambiguous or unambiguous.
A bottom-up parsing model retrieved from Dotlacil, J (2021) was used to process and parse all 80 sentences and create activation results. Activations showcase the inverse difficulty level of a certain word on a specific position. Meaning when activation levels are low, the reader is predicted to interpret the word with difficulty hence the reading time is expected to be high. When the activation levels are high, we expect the difficulty level to be easy hence the reading time is expected to be low.

The input sentences were given Part of Speech tags (POS). Every sentence needed to be filled in with the correct POS-tag and a non-existent POS-tag on the target word (position 0) in the sentence. The target word in the sentences used for parsing is referring to the second verb in the sentence where the ambiguity effect is believed to take place.

4.2 Procedure

The sentences were parsed by the parsing model and an overview of all activations on every word was created. Data predictions were retrieved from the model and used to calculate the mean activations to generate relevant data in order to compare results with the human parsed data. All mean activations were calculated for the critical position and the following 3 positions in every sentence. Mean activations for positions 0-3 were retrieved.
The experiment was executed by the programming language Python (1995) and ran on Windows 10 OS.

4.3 Results

Position	**A/R**	**A/U**	**U/R**	**U/U**
0	5.4627985558163	9.2599021225395	5.4627985558163	9.2599021225395
1	4.1579294336339	6.631806738065	6.3946917240753	6.631806738065
2	5.9030789910839	5.6396767631104	5.9097754504768	5.6396767631104
3	5.7706131054706	6.1290606763408	5.8788826423062	6.010269051857

Figure 4.1: table with mean activations for every condition

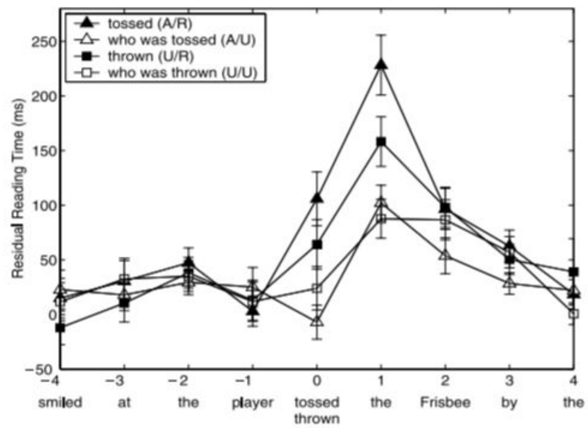

Figure 4.2: Previous results recovered from Tabor et al.(2004)

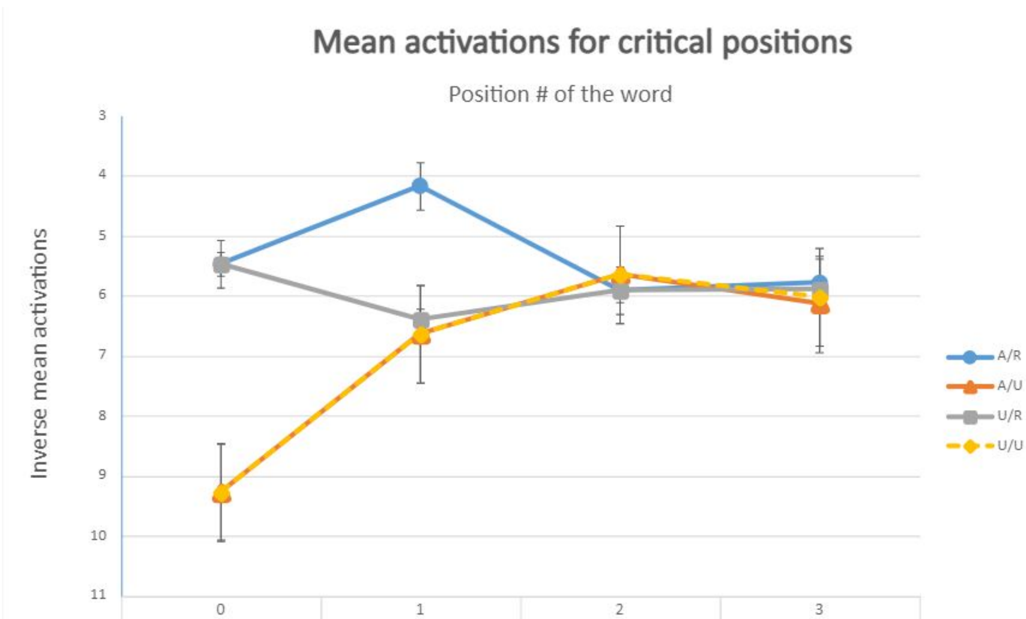

	0	1	2	3
A/R	5.46	4.16	5.9	5.77
A/U	9.26	6.63	5.64	6.13
U/R	5.46	6.39	5.91	5.88
U/U	9.26	6.63	5.64	6.01

Figure 4.3: Mean predicted activations from Experiment 1. Error bars show one standard error around each data point.

The data from 1 item, 4 sentences in total, was not processed properly by the

parser and therefore removed from the analysis.
Figure 4.1 shows the average activations for every condition in all the 4 relevant positions. Figure 4.2 shows the results of human reading times on critical positions for each condition retrieved by Tabor et al.(2004).
Figure 4.3 shows the mean activations of those same findings visualized in a graph. The findings in figure 4.3 are inverted on the Y-axis to match the reading times of figure 4.2.
We analyzed activation numbers by looking at the relative increase or decrease for each condition when moving up positions. A lower activation is equivalent to a higher reading time. A higher activation is equivalent to a lower reading time. It is only at the second verb in each sentence where the distracting analysis was first supported by a combination of words. Therefore this word is being referred to as the "critical" word and it is the position 0 in every sentence.

4.4 Discussion

Results from experiment 1 show consistency with the hypothesis that predictions of the transition-based bottom-up parsing model shows strong similarities with the human retrieved data from Tabor et al. (2004) and thus shows signs of predicting Local Syntactic Coherence effects. Interpretations of these results should be based on the comparisons between the activations and the human reading times. The reason we draw this conclusion is because consistent similarities are shown between the results of figure 4.2 and 4.3. Activations in the unreduced conditions were expected to be higher as results from human data show unreduced conditions have shorter reading times. Both the critical position 0 and the following word after, position 1, show a significant increase in activation when changing to the unreduced condition. Effects of the activations we encounter show some parts of spill over effect. In the A/R vs the U/R condition, spill over effect is believed to be taken place on position 1. This can be explained due to the activation results of position 0. There is believed to be a very small spill over effect on position 0 as well, this effect however is much more difficult to explain.
Other cases of local coherence could be studied to generate and predict more information on local coherence effects. This study only conducted one specific experiment with 20 items. Further research could investigate the role of local coherence by adding more sentences or using different parsing approaches in order to find effects of local ambiguity.

Chapter 5

Conclusion

Natural Language Understanding has been found one of the more difficult tasks to achieve for intelligent systems. Natural Language Processing systems are widely integrated in the field of AI, the understanding for natural language systems is therefore an important goal. Sentence parsing has previously shown massive problems due to parsing constraints. These constraints are caused by levels of ambiguity which need a certain set of interpretations. Earlier literature suggests interpretation is heavily influenced by local ambiguity effects. Those theories were tested with human data reading times. Our goal was to model the process by which local coherence effects emerge as a result of Good Enough parsing. A major finding in this research is that the predictions of a transition-based bottom-up parsing model show positive results towards local coherence effects. A model that has been previously used to show the rational theory of cognition is now showing it's ability to predict local coherence effects, as its results compared to human data show strong similarities. Our conclusive results show how a bottom-up approach to sentence processing is capable of showing local ambiguity effects. Some predictions are more difficult to explain due to spill over effects. These experiments show promising results for future research as theorizing sentence processing gets more advanced.

Chapter 6

Appendix A

Sentences used for parsing

- 1. The kindergartners liked the little girl (who was) brought/chosen a toy by her parents on the first day of Chanukah.
- 2. We saw a movie about an artist (who was) painted/drawn a picture by her father while he was on his deathbed.
- 3. The health officials pounced on a restaurant (which was) sent/flown a shipment of salmon by a company that had failed to comply with refrigeration laws.
- 4. At the dinner party, I met a man (who was) allowed/forbidden the pleasure of eating sweets by his doctor.
- 5. One expects a man (who is) told/forgiven his sins by his own god to have tolerance for weaknesses in others. (Experiment 3: one should respect a man told/forgiven his sins by his own god.)
- 6. An elderly gentleman addressed the woman (who was) offered/given a beer by the hostess.
- 7. Balthazar praised the professor (who was) taught/given Swahili lessons by a graduate student.
- 8. The manager watched a waiter (who was) served/given pea soup by a trainee.
- 9. James entertained the children (who were) dyed/hidden Easter eggs by their teachers.

- 10. The foreman yelled at a carpenter (who was) cut/sawn a board by his buddy.
- 11. The preschool teacher congratulated the little boy (who was) knitted/sewn a hat by his grandmother.
- 12. The janitor chatted with the young man (who was) rented/ shown an apartment by his uncle.
- 13. The nurse admonished a student (who was) nabbed/stolen a muffin by her friends from the dining hall.
- 14. The play centered around an innkeeper (who was) recited/ sung a verse by a travelling monk.
- 15. The hotel owner questioned a guest (who was) brought/ taken a drink by the bellboy.
- 16. The coach smiled at the player (who was) tossed/thrown a frisbee by the opposing team.
- 17. The anthropologist interviewed a woman (who was) knitted/woven a shawl by her mother.
- 18. The FBI questioned a congressman (who was) mailed/ written a letter by the activist.
- 19. The deliveryman teased the accountant (who was) saved/ given a coupon by her boss.
- 20. The prophet spoke of a man (who was) planted/grown a tree by his daughter.

Chapter 7

References

1 Anderson, J., Bothell, D., Byrne, M. (2004). An Integrated Theory of the Mind. Vol. 111, No. 4, 1036–1060. Carnegie Mellon University. doi: 10.1037/0033-295X.111.4.1036

2 Bicknell, K., Levy, R. (2009). A model of local coherence effects in human sentence processing as consequences of updates from bottom-up prior to posterior beliefs. In Proceedings of North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL HLT) 2009 Conference (pp. 665–673). Boulder, CO: Association for Computational Linguistics.

3 Chomsky, N., Structures, S. (1957). Mouton,". Syntactic Structures", The Hague.

4 Christianson, K., Luke, S. G., Hussey, E. K., Wochna, K. L. (2017). Why reread? Evidence from garden-path and local coherence structures. The Quarterly Journal of Experimental Psychology, 70(7), 1380-1405.

5 Crocker, M. W. (1996). Computational Psycholinguistics: An Interdisci plinary Approach to the Study of Language. Studies in Theoretical Psycholinguistics 20. Kluwer Academic Publishers.

6 Crocker, M. W. (1999). Mechanisms for sentence processing. In Language 887 processing, 191–232. Psychology Press Hove, UK.

7 Dotlacil, J., De Haan, P. (2021) Parsing model and a rational theory of memory. Utrecht Institute of Linguistics OTS, Utrecht University, Netherlands. doi: 10.3389/fpsyg.2021.657705

8 Ferreira, F., Christianson, K., Hollingworth, A. (2001). Misinterpretations of garden-path sentences: Implications for models of sentence processing and reanalysis. Journal of Psycholinguistic Research, 30, 3–20.

9 Ferreira, F., Patson, N. D. (2007). The 'Good Enough' approach to language comprehension. Language and Linguistics Compass, 1(1–2), 71–83.

10 Ferreira, F., Lowder, M. (2016). Chapter Six - Prediction, Information Structure, and Good-Enough Language Processing. Editor(s): Brian H. Ross, Psychology of Learning and Motivation. Academic Press. Volume 65. Pages 217-247.

11 Frazier, L. (1979). On comprehending sentences: Syntactic parsing strategies. University of Connecticut.

12 Frazier, L., 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.

13 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: Lawrence Erlbaum Associates.

14 Futrell, Richard, Edward Gibson, Harry J. Tily, Idan Blank, Anastasia Vishnevetsky, Steven T. Piantadosi, and Evelina Fedorenko. 2018. The natural stories corpus. In Proceedings of LREC 2018, Eleventh International Conference on Language Resources and Evaluation, 76–82. Miyazaki, Japan.

15 Kempen, G., Vosse, T. (1989). Incremental syntactic tree formation in human sentence processing: A cognitive architecture based on activation decay and simulated annealing. Connection Science, 1, 273–290.

16 Konieczny, L., Müller, D., Hachmann, W., Schwarzkopf, S., Wolfer, S. (2009). Local syntactic coherence interpretation: Evidence from a visual world study. In N. A. Taatgen, H. van Rijn (Eds.), Proceedings of the 31st annual conference of the Cognitive ScienceSociety (pp. 1133–1138). Austin, TX: Cognitive Science Society.

17 Levy, R. (2008). A noisy-channel model of rational human sentence comprehension under uncertain input. In Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP). Morristown, NJ: Association for Computational Linguistics.

18 Marcus, Mitchell P., Mary Ann Marcinkiewicz, and Beatrice Santorini. 1993. Building a large annotated corpus of english: The penn treebank. Computational Linguistics 19:313–330.

19 Marslen-Wilson, W. (1973). Linguistic Structure and Speech Shadowing at Very Short Latencies. Nature 244, 522–523.
https://doi.org/10.1038/244522a0

20 Morgan, E., Keller, F., Steedman, M. (2010). A Bottom-Up Parsing Model of Local Coherence Effects. Proceedings of the Annual Meeting of the Cognitive Science Society, 32.
Retrieved from https://escholarship.org/uc/item/4v3265ds.

21 Tabor, W., Hutchins, S. (2003). Evidence for self-organized sentence processing: Digging in effects. Submitted manuscript, University of Connecticut, Department of Psychology.
Available at http://www.sp.uconn.edu/ps300vc/papers.html.

22 Tabor, W., Galantucci, B., Richardson, D. (2004). Effects of merely local syntactic coherence on sentence processing. Journal of Memory and Language, Volume 50, Issue 4, Pages 355-370,
doi: https://doi.org/10.1016/j.jml.2004.01.001.

23 Van Rossum, G., Drake Jr, F. L. (1995). Python reference manual. Centrum voor Wiskunde en Informatica Amsterdam.

24 Vosse, T., Kempen, G. (2000). Syntactic structure assembly in human parsing: A computational model based on competitive inhibition and a lexicalist grammar. Cognition, 75, 105–143.

25 Zhang, X., Du, Y., Sun, W., Wan, X. (2016). Transition-based Parsing for Deep Dependency Structures. Institute of Computer Science and Technology, the MOE Key Laboratory of Computational Linguistics, Peking University, Beijing 100871, China. doi:0.1162/COLI a 00252