

Dependency Structures and Lexicalized Grammars

Marco Kuhlmann

Dissertation zur Erlangung des Grades
des Doktors der Ingenieurwissenschaften
der Naturwissenschaftlich-Technischen Fakultäten
der Universität des Saarlandes

Saarbrücken  2007

Dissertation zur Erlangung des Grades
des Doktors der Ingenieurwissenschaften
der Naturwissenschaftlich-Technischen Fakultäten
der Universität des Saarlandes,

eingereicht von

Marco Kuhlmann, B. Sc., M. Sc.,

geboren am 16. Oktober 1977 in Rinteln

BERICHTERSTATTER:

Prof. Dr. Gert Smolka

Prof. Dr. Manfred Pinkal

Prof. Aravind K. Joshi

DEKAN:

Prof. Dr.-Ing. Thorsten Herfet

PRÜFUNGS AUSSCHUSS:

Prof. Dr. Raimund Seidel

Prof. Dr. Gert Smolka

Prof. Dr. Manfred Pinkal

Dr. Tilman Becker

TAG DES KOLLOQUIUMS:

24. September 2007

Textfassung vom 2. November 2007

Copyright © 2007 Marco Kuhlmann

MATHEMATIKSATZ: Johannes Küster, typoma, Holzkirchen

SCHRIFTEN: Minion, Cronos, typoma® MnMath

Abstract

In this dissertation, I show that both the generative capacity and the parsing complexity of lexicalized grammar formalisms are systematically related to structural properties of the *dependency structures* that these formalisms can induce. Dependency structures model the syntactic dependencies among the words of a sentence. I identify three empirically relevant classes of dependency structures, and show how they can be characterized both in terms of restrictions on the relation between dependency and word-order and within an algebraic framework. In the second part of the dissertation, I develop natural notions of automata and grammars for dependency structures, show how these yield infinite hierarchies of ever more powerful *dependency languages*, and classify several grammar formalisms with respect to the languages in these hierarchies that they are able to characterize. These results provide fundamental insights into the relation between dependency structures and lexicalized grammars.

Acknowledgements

This dissertation reports the major results of roughly four years of research. Among the many people who have contributed to it, there are some to whom I owe a particularly large debt of gratitude.

My first thanks goes to my supervisor, Gert Smolka. He granted me the freedom to develop my own ideas, encouraged me to strive for simplicity in their presentation, and provided guidance and advice. I also thank Manfred Pinkal and Aravind Joshi for accepting to give their expert opinion on this dissertation, and to Raimund Seidel and Tilman Becker for agreeing to join my examination committee.

Guido Tack, as my office-mate, had to suffer from my meanderings on a large number of sometimes errant ideas related to this dissertation. I thank him for the patience he had with me, for his valuable feedback, and most of all, for his friendship. My student, now colleague Mathias Möhl bore with me in endless discussions, and made essential contributions to this work. Joachim Niehren introduced me to tree automata and the algebraic perspective on formal languages. Without him, the machinery used in this dissertation would be far less canonical. Thanks also go to my coauthors: Mathias, Manuel Bodirsky, Robert Grabowski, and Joakim Nivre. Joachim, Bernhard Fisseni, Sebastian Maneth, and Martin Plátek provided me with the opportunity to present preliminary results of my work to their research groups.

The research reported in this dissertation was carried out in the context of the CHORUS project, which is part of the Collaborative Research Centre 378, ‘Resource-adaptive cognitive processes’, and funded by the Deutsche Forschungsgemeinschaft, DFG. Throughout my doctoral studies, I have also been associated to the International Research Training Group 715, ‘Language technology and cognitive systems’, likewise funded by the DFG. In both contexts, I have been privileged to work and study together with Ralph Debusmann, Alexander Koller, and Stefan Thater. Alexander deserves a special thanks for his constant enthusiasm about our research, which has been a source of motivation for me ever since I first met him as a freshman.

During the last five years, my desk was in the Programming Systems Lab at Saarland University. I thank the other members of the Lab for creating a pleasant working environment—both the current members and those who have worked here during my time at the Lab: Gert, Guido, Mathias, Ralph, Mark Kaminski, Sandra Neumann, and Jan Schwinghammer; Joachim, Ondřej Bojar, Thorsten Brunklaus, Denys Duchier, Leif Kornstaedt, Didier Le Botlan, Tim Priesnitz, Andreas Rossberg, Gábor Szokoli, Lutz Straßburger, and Ann van de Veire. Thanks also go to our student assistants Robert Künnemann and Christophe Boutter.

Beyond the people I have already mentioned, there were many more who have listened patiently, answered thoroughly, and been helpful during the composition of this dissertation in all kinds of ways. I specifically want to mention Carlos Areces, Jason Baldrige, Achim Blumensath, Cem Bozsahin, Paul Braffort, Matthias Buch-Kromann, Balder ten Cate, Lukas Champollion, Joan Chen-Main, David Chiang, Michael Daniels, Arno Eigenwillig, Jason Eisner, Phillippe Flajolet, Mary Harper, Günter Hotz, Laura Kallmeyer, Anna Kasprzik, Stephan Kepser, Geert-Jan Kruijff, Ivana Kruijff-Korbayová, Markéta Lopatková, Detmar Meurers, Jens Michaelis, Uwe Mönnich, Timo von Oertzen, Gerald Penn, Fernando Pereira, Giorgio Satta, Yves Schabes, Tatjana Scheffler, Ingo Schröder, Sabine Schulte im Walde, Stuart Shieber, Oliver Suhre, Martin Volk, David Weir, Anssi Yli-Jyrä, and Zdeněk Žabokrtský. Gert, Mathias, Joachim, Ralph, Stefan, and Jiří have all read drafts of parts of this dissertation, and given valuable feedback.

The final phase of the write-up made me appreciate how extremely fortunate I am to be surrounded by a large number of supportive friends and colleagues, such as Colin Bannard, Chris Callison-Burch, Danielle Matthews, Ulrike and Sebastian Padó, Monika Schwarz, and Markus Zacharski. A very special thanks goes to Lars Charbonnier, for his friendship and support during the last ten years.

Finally, I want to express my love and gratitude to my families: to Eva, Richard, Karin, Per, Britta, and Gunnar; to Renate, Eckhard, Maike, Jochen, Christian, Silke, Hannah, and Tilman; to my aunts, uncles, and cousins; to my grandmothers Meta and Waltraud; to my siblings Marina, Ulrike, and Christoph, to my parents Heidrun and Klaus-Gerd, and most of all to my wife, Antje. Thank you for who you are, and for accepting me for who I am.

Saarbrücken, November 2007

Marco Kuhlmann

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Overview of the dissertation	7
2	Preliminaries	13
3	Projective dependency structures	19
3.1	Projectivity	19
3.2	Algebraic framework	22
3.3	Algorithmic problems	30
3.4	Empirical evaluation	32
4	Dependency structures of bounded degree	35
4.1	The block-degree measure	35
4.2	Algebraic framework	39
4.3	Algorithmic problems	45
4.4	Empirical evaluation	50
5	Dependency structures without crossings	51
5.1	Weakly non-projective dependency structures	51
5.2	Well-nested dependency structures	55
5.3	Empirical evaluation	61
6	Structures and grammars	63
6.1	Context-free grammars	63
6.2	Linear Context-Free Rewriting Systems	68
6.3	Coupled Context-Free Grammars	75
6.4	Tree Adjoining Grammar	79
7	Regular dependency languages	83
7.1	Regular sets of dependency structures	83
7.2	Pumping lemmata	91
7.3	Constant growth	97

CONTENTS

8	Generative capacity and parsing complexity	103
8.1	Projection of string languages	103
8.2	String languages and structural properties	107
8.3	Parsing complexity	113
9	Conclusion	121
	Bibliography	125

CHAPTER 1 | Introduction

The central thesis of this dissertation is, that the generative capacity and the parsing complexity of lexicalized grammar formalisms are systematically related to structural properties of the dependency graphs that these formalisms can induce. To back up this thesis, I develop the formal theory of *dependency structures*, show how combining these structures with a regular means of composition yields infinite hierarchies of ever more powerful *dependency languages*, and classify several grammar formalisms with respect to the languages in these hierarchies that they are able to characterize.

1.1 | Motivation

The principal goal of this dissertation is to connect two research areas in computational linguistics: the search for grammar formalisms with adequate generative capacity, and the search for dependency representations with well-balanced structural flexibility. In this section, I review these two research areas, state my principal goal more precisely, and motivate the specific research questions that I want to address.

1.1.1 | Grammar formalisms and generative capacity

Grammar formalisms are mathematical devices that are developed to give explicit descriptions of linguistic theories. Following Pullum and Scholz (2005), we can classify them into two broad categories: *generative frameworks*, which consist of rules that specify how linguistic expressions are composed, and *model-theoretic frameworks*, in which linguistic expressions are characterized as the denotations of logical formulae. In this dissertation, we focus on generative formalisms, even though many of the results that we present also have direct implications for model-theoretic frameworks.

One of the fundamental questions that we can ask about a grammar formalism is, whether it adequately models natural language. In generative frameworks, we can answer this question by studying the *generative capacity* of the formalism: if we interpret grammars as generators of sets of linguistic structures (such as strings, trees, or predicate-argument structures), then we can call a grammar formalism adequate, if it allows us to write grammars that generate exactly those structures that we consider relevant for the description of natural language. The outcome of such an assessment depends on the type of linguistic structures that we base it on: a formalism may be adequate with respect to one type, but inadequate with respect to another. Another fundamental question therefore is, on which type of linguistic expressions should we base our measure of generative capacity?

1 INTRODUCTION

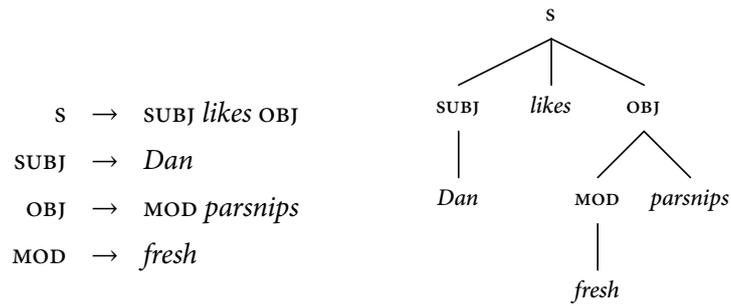


Figure 1.1: A context-free grammar and a parse tree generated by this grammar

To focus the discussion, let us consider the case of context-free grammars. For these grammars, there are two standard measures of generative capacity: we can interpret them as generators of sets of strings (*string-generative capacity*), or as generators of sets of parse trees (*tree-generative capacity*). Figure 1.1 shows a toy context-free grammar together with a parse tree for a simple English sentence. Strings and parse trees are closely related. In particular, for each string generated by a context-free grammar, there is at least one parse tree from which this string can be recovered by reading the leaves of the tree from left to right.

It is widely accepted today that context-free grammars are not adequate for the description of natural language. Independently of each other, Huybregts (1984) and Shieber (1985) showed that certain constructions in Swiss German require grammar formalisms that adequately model these constructions to generate the so-called copy language, which is beyond the string-generative capacity (and, *a fortiori*, the tree-generative capacity) of context-free grammars. If we accept this analysis, then we must conclude that context-free grammars are too weak, and that we should look out for more powerful formalisms. Unfortunately, the first class of formal languages in the Chomsky hierarchy that *does* contain the copy language, the class of *context-sensitive languages*, is too big a leap: it contains many languages that are considered to be beyond human capacity (such as the set of all prime numbers), and while context-free grammars can be parsed in polynomial time, the parsing problem of context-sensitive grammars is PSPACE-complete. For such problems, it is widely suspected that they cannot be solved in (deterministic or non-deterministic) polynomial time.

In search of a class of grammars that extends context-free grammar by the minimal amount of generative power that is needed to account for natural language, several so-called *mildly context-sensitive grammar formalisms* have been developed; perhaps best-known among these is Tree Adjoining Grammar (TAG) (Joshi and Schabes, 1997). The class of string languages generated by TAGs contains the copy language, but unlike general context-sensitive grammars, TAGs can be parsed in polynomial time. More important here than their increased string-generative capacity however, is their stronger power as generators of *dependency structures*.

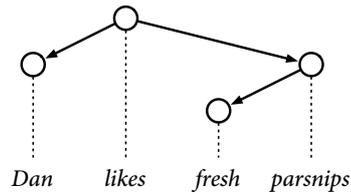


Figure 1.2: A dependency structure

1.1.2 | Dependency representations and structural flexibility

Syntactic representations based on word-to-word dependencies have a long tradition in descriptive linguistics; for a historical overview, refer to Kruijff, 2005. In recent years, dependency representations have also become increasingly used in computational tasks, such as information extraction (Culotta and Sorensen, 2004), machine translation (Quirk et al., 2005), and parsing (Nivre et al., 2007). The basic assumptions behind the notion of dependency are summarized in the following sentences from the seminal work of Tesnière (1959, ch. 1, §§ 2–4, and ch. 2, §§ 1–2):¹

The sentence is an organized whole; its constituent parts are the words. Every word that functions as part of a sentence is no longer isolated as in the dictionary: the mind perceives connections between the word and its neighbours; the totality of these connections forms the scaffolding of the sentence. The structural connections establish relations of dependency among the words. Each such connection in principle links a superior term and an inferior term. The superior term receives the name governor (régissant); the inferior term receives the name dependent (subordonné).

We can represent the dependency relations among the words of a sentence as a graph. More specifically, the *dependency structure* for a sentence $\vec{w} = w_1 \cdots w_n$ is the directed graph on the set of positions of \vec{w} that contains an edge $i \rightarrow j$ if and only if w_j depends on w_i . In this way, just like strings and parse trees, dependency structures can capture information about certain aspects of the linguistic structure of a sentence. As an example, consider Figure 1.2. In this graph, the edge between the word *likes* and the word *Dan* may be understood to encode the syntactic information that *Dan* is the subject of *likes*; similarly, the edge between *likes* and *parsnips* may be interpreted as saying that *parsnips* is the direct object. When visualizing dependency structures, we represent (occurrences of) words by circles, and dependencies among them by arrows: the source of an arrow marks the governor of the corresponding dependency, the target marks the dependent. Furthermore, following Hays (1964), we use dotted lines (and call them *projection lines*) to indicate the left-to-right ordering of the words in the sentence. Note that these lines do not belong to the graph structure as such.

¹ The English translation (by myself) is based on Tesnière (1980).

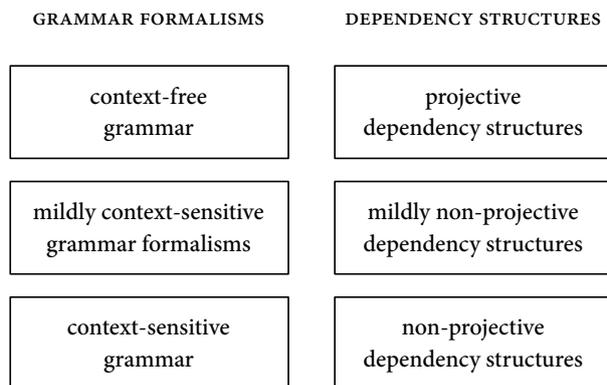


Figure 1.3: *The big picture*

With the concept of a dependency structure at hand, we can express linguistic universals in terms of *structural constraints* on graphs. The most widely used such constraint is to require the dependencies to form a tree. This requirement captures the stipulations that no word should depend on itself, not even transitively, that each word should have at most one governor, and that a dependency analysis should cover all the words in the sentence. The structure in Figure 1.2 satisfies the treeness constraint. More generally, all dependency structures that we consider in this dissertation form trees.

Another well-known constraint on dependency structures is *projectivity*. In contrast to treeness, which imposes restrictions on dependency as such, projectivity concerns the relation between dependency and the left-to-right order of the words in the sentence. Specifically, it requires each dependency subtree to cover a contiguous region of the sentence. Projectivity is interesting because the close relation between dependency and word order that it enforces can be exploited in parsing algorithms (Eisner and Satta, 1999). However, in recent literature, there is a growing interest in *non-projective dependency structures*, in which a subtree may be spread out over a discontinuous region of the sentence. Such representations naturally arise in the syntactic analysis of linguistic phenomena such as extraction, topicalization and extraposition; they are particularly frequent in the analysis of languages with flexible word order, such as Czech (Holan et al., 1998; Veselá et al., 2004). Unfortunately, most formal results on non-projectivity are rather discouraging. In particular, non-projective dependency parsing is NP-complete (Neuhaus and Bröker, 1997; McDonald and Satta, 2007).

In search of a balance between the benefit of more expressivity and the penalty of increased processing complexity, several authors have proposed structural constraints that relax the projectivity restriction, but at the same time ensure that the resulting classes of graphs are computationally well-behaved (Yli-Jyrä, 2003; Nivre, 2006a; Havelka, 2007a). Such constraints identify classes of what we may call *mildly non-projective dependency structures* (see Figure 1.3).

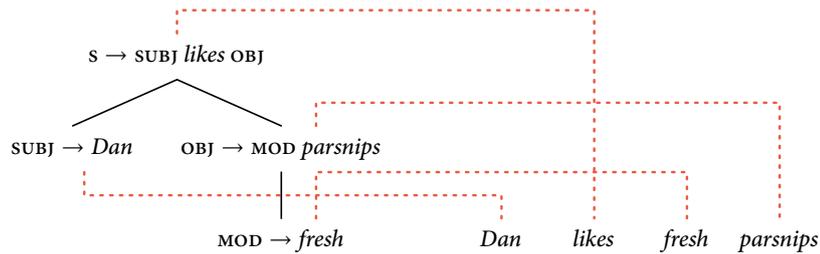


Figure 1.4: Lexicalized derivations induce dependency structures

1.1.3 | Lexicalized grammars induce dependency structures

Let us return to the discussion of generative grammars. An interesting property of the context-free grammar in Figure 1.1 is that it is *lexicalized*: every production of the grammar contains exactly one terminal symbol, the *anchor* of that production (Schabes et al., 1988). Lexicalized grammars play a significant role in contemporary linguistic theories and practical applications. Crucially, every lexicalized grammar can be understood as a generator for sets of dependency graphs, in the following sense.

Consider a derivation d of a terminal string \vec{a} by means of a context-free grammar. A *derivation tree* for d is a tree in which the nodes are labelled with (occurrences of) the productions used in d , and the edges indicate how these productions were combined. To give an example, the left half of Figure 1.4 shows the unique derivation tree of our example grammar. If the underlying grammar is lexicalized, then there is a one-to-one correspondence between the nodes in the derivation tree and the positions in the derived string \vec{a} : every production that participates in the derivation contributes exactly one terminal symbol to this string. If we now order the nodes of the derivation tree according to the string positions of their corresponding anchors, then what we get is a dependency structure. For our example, this procedure results in the structure depicted in Figure 1.2. We say that this dependency structure is *induced* by the derivation d . The notion of induction identifies the dependency relation with the derivation relation, which is sensible for grammars in which all syntactic dependencies are specified locally within the elementary structures of the grammar, such as Tree Adjoining Grammar (Frank, 2002; but see Candito and Kahane, 1998).

Not all practically relevant dependency structures can be induced by derivations in lexicalized context-free grammars. A classic example is provided by the structural difference between the verb-argument dependencies in German and Dutch subordinate clauses, as shown in Figure 1.5: context-free grammar can only characterize the ‘nested’, projective dependencies of German (left), but not the ‘cross-serial’, non-projective assignments of Dutch (right). However, these structures can be induced by lexicalized TAG (Joshi, 1985). The principal goal of this dissertation is to make the relations between grammars and the dependency structures that they induce precise.

1 INTRODUCTION

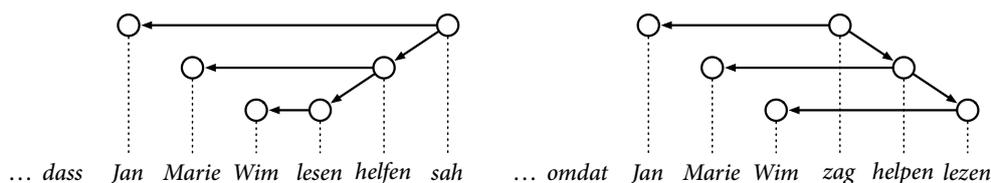


Figure 1.5: Nested and cross-serial dependencies

1.1.4 | Questions asked

The main research question of this dissertation is,

Which grammars induce which sets of dependency structures?

An answer to this question is interesting for at least two reasons. First, it allows us to use dependency structures as the basis of an alternative measure for the generative capacity of a grammar formalism. This is attractive, as dependency structures are more informative than strings, but less formalism-specific and more intuitively accessible than parse trees (cf. Kallmeyer, 2006). Second, an answer to the question allows us to tap the rich resource of formal results about grammar formalisms and transfer them to work on dependency representations. In particular, it could allow us to import the expertise in developing parsing algorithms for lexicalized grammar formalisms into the field of dependency parsing (cf. McDonald and Satta, 2007).

In spite of the apparent connection between the generative capacity of a grammar formalism and the structural properties of the dependency graphs that this formalism can induce, there have been only few results that link the two research areas. A fundamental reason for the lack of such bridging results is that, while structural constraints on dependency graphs are *internal* properties in the sense that they concern the nodes of the graph and their connections, grammars take an *external* perspective on the objects that they manipulate—the internal structure of an object is determined by the internal structure of its constituent parts and the operations that are used to combine them. An example for the difference between the two views is given by the different perspectives on trees that we find in graph theory and universal algebra. In graph theory, a tree is a special graph with an internal structure that meets certain constraints; in algebra, trees are abstract objects that can be composed and decomposed using a certain set of operations. One of the central technical questions that we need to answer in order to connect grammars and structures is, how classes of dependency graphs can be given an algebraic structure. An important result into this direction is Gaifman’s (1965) proof that the string languages that are generated by his projective dependency formalism are exactly the context-free languages. One of the technical objectives of this dissertation is to generalize this result from projective structures and context-free grammars to mildly non-projective dependency structures and mildly context-sensitive grammar formalisms.

1.2 | Overview of the dissertation

This dissertation consists of two parts. In the first part, we develop an algebraic framework within which lexicalized grammars can be compared based on the structural properties of the dependency graphs that they induce. In the second part, we derive a natural notion of regular sets of dependency structures, and use it to study the connection between structural properties such as projectivity on the one hand, and language-theoretic properties such as string-generative capacity and parsing complexity on the other.

1.2.1 | Dependency structures

In the first part of the dissertation, we study dependency structures. These structures clearly separate two relations: the dependency relation, which we call *governance*, and the total order on the nodes of the graph, which we call *precedence*. We discuss three interesting classes of mildly non-projective dependency structures, compare them to other classes in the literature, and evaluate their practical relevance using data from *dependency treebanks*.

Structural constraints

The first two classes of dependency structures that we consider in this dissertation have been studied before. *Projective dependency structures* (Chapter 3), as already mentioned, are characterized by the structural constraint that each subtree must form an interval with respect to the total order on the nodes. As an example, consider the dependency structure depicted in Figure 1.6a: each of the subtrees forms an interval with respect to the precedence relation. In *dependency structures of bounded degree* (Chapter 4), the projectivity constraint is relaxed in such a way that dependency subtrees can be distributed over more than one, but still a fixed number of intervals. For example, in the structure depicted in Figure 1.6c, both the subtree rooted at the node 2 and the subtree rooted at the node 3 span two intervals. We call the maximal number of intervals per subtree the *block-degree* of the structure, and use it to quantify the non-projectivity of dependency graphs.

The third class of dependency structures that we investigate, the class of *well-nested dependency structures* (Chapter 5), is original to this work. Well-nestedness is the structural constraint that pairs of disjoint dependency subtrees must not cross, meaning that there must not be nodes i_1, i_2 in the first subtree and nodes j_1, j_2 in the second such that $i_1 < j_1 < i_2 < j_2$. The dependency structure depicted in Figure 1.6e is well-nested, while the structure depicted in Figure 1.6c is not. Well-nested dependency structures are closely related to several other combinatorial structures, such as non-crossing partitions and Dyck languages. We discuss an empirical evaluation that shows that they are also practically relevant: virtually all dependency analyses in two large and widely-used dependency treebanks obey the well-nestedness constraint.

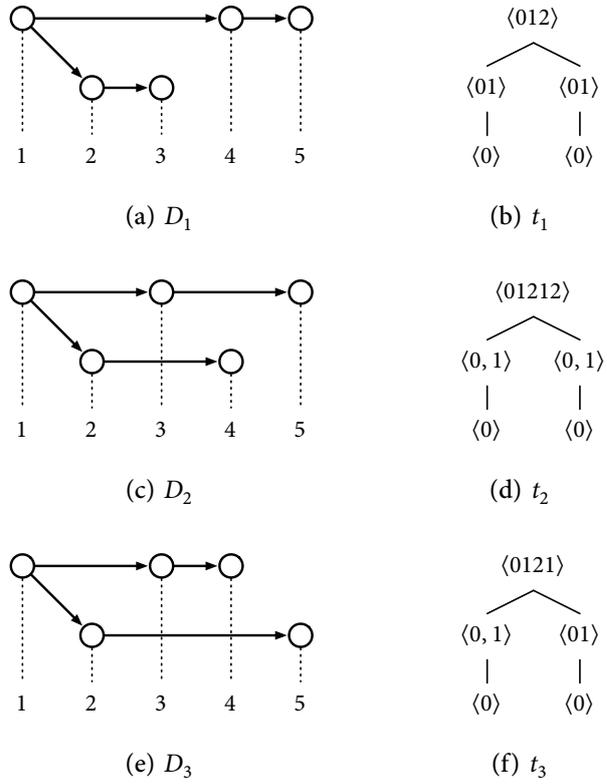


Figure 1.6: A zoo of dependency structures, and their corresponding terms

Algebraic framework

One of the major contributions of this dissertation is an algebraic framework in which projective, block-restricted and well-nested dependency structures can be understood as the outcome of compositional processes. Under this view, structural constraints do not apply to a fully specified dependency graph, but are inherent in the composition operations by which the graph is constructed; this provides a bridge between dependency structures and generative grammar. We formalize the algebraic framework in two steps. In the first step, we show that dependency structures can be encoded into terms over a certain signature of *order annotations* in such a way that the three different classes of dependency structures discussed above stand in one-to-one correspondence with terms over specific subsets of this signature. In the second step, we define the concept of a *dependency algebra*. In these algebras, order annotations are interpreted as composition operations on dependency structures. We prove that each dependency algebra is isomorphic to the corresponding term algebra, which means that the composition of dependency structures can be freely simulated by the usual composition operations on terms, such as substitution.

To give an intuition for the algebraic framework, the right half of Figure 1.6 shows the terms corresponding to the dependency structures in the left half. Each order annotation in these terms encodes node-specific information about the precedence relation. As an example, the symbol $\langle 0, 1 \rangle$ in Figure 1.6d represents the information that the corresponding subtree in Figure 1.6c consists of two intervals (the two components of the tuple $\langle 0, 1 \rangle$), with the root node (represented by the symbol 0) situated in the left interval, and the subtree rooted at the first child (represented by the symbol 1) in the right interval. Under this encoding, the block-degree measure corresponds to the maximal number of components per tuple, and the well-nestedness condition corresponds to the absence of certain ‘forbidden substrings’ in the individual order annotations, such as the substring 1212 in the term in Figure 1.6d.

Structures and grammars

In Chapter 6, we apply the algebraic framework to classify the dependency structures induced by various lexicalized grammar formalisms. The key to this classification is the insight that the notion of induction can be formalized as the interpretation of the derivations of a grammar in a suitable dependency algebra. Based on this formalization, we can generalize Gaifman’s (1965) result that projective dependency structures correspond to lexicalized context-free grammars into the realm of the mildly context-sensitive: the classes of block-restricted dependency structures correspond to Linear Context-Free Rewriting Systems (Vijay-Shanker et al., 1987; Weir, 1988), the classes of well-nested block-restricted structures correspond to Coupled Context-Free Grammar (Hotz and Pitsch, 1996). As a special case, the class of well-nested dependency structures with a block-degree of at most 2 is characteristic for derivations in Lexicalized Tree Adjoining Grammar (Joshi and Schabes, 1997).

1.2.2 | Dependency languages

In the second part of the dissertation, we lift our results from individual dependency structures to sets of such structures, or *dependency languages*. The key to this transfer is the formal concept of regular sets of dependency structures (Chapter 7), which we define as the recognizable subsets of dependency algebras (Mezei and Wright, 1967). From this definition, we obtain natural notions of automata and grammars on the basis of which we can reason about the language-theoretic properties of regular dependency languages.

Automata and grammars

Given the isomorphism between dependency algebras and term algebras, we can derive a natural automaton model for dependency structures from the concept of a *tree automaton* (Thatcher and Wright, 1968). This method in fact applies to all kinds of data structures that are constructible using a finite set of operations; for example, successful applications of the approach have previously led to stepwise tree automata for the data model of XML (Carme et al., 2004) and feature automata for unranked unordered trees (Niehren and Podelski, 1993). From the notion of an automaton, we are led to the concept of a *regular dependency grammar*. By and large, grammars and automata are two sides of the same coin: we get a grammar from an automaton by interpreting the transition rules of the automaton as a directed rewriting system. Using regular dependency grammars, we show a powerful pumping lemma for regular dependency languages, and prove that these languages are semilinear (Parikh, 1966), which is also characteristic for languages generated by mildly context-sensitive grammar formalisms.

String-generative capacity and parsing complexity

In the last technical chapter of the dissertation (Chapter 8), we investigate the connections between structural constraints, string-generative capacity, and parsing complexity. We show how the block-degree measure gives rise to an infinite hierarchy of ever more powerful string languages, and how enforcing the well-nestedness of the underlying dependency structures leads to a proper decrease of string-generative power on nearly all levels of this hierarchy. In proving these results, we see how string languages can ‘enforce’ the presence of structures with certain properties in the corresponding dependency languages. As an example, for every natural number k , we identify a string language $L(k)$ that requires every regular set of dependency structures with block-degree at most k that projects $L(k)$ to contain structures that are not well-nested. Finally, we show that both the block-degree measure and the well-nestedness condition have direct implications for the parsing complexity of regular dependency languages. We prove that, while the parsing problem of regular dependency languages is polynomial in the length of the input string, the problem in which we take the grammar to be part of the input is NP-complete. Interestingly, for well-nested dependency languages, parsing is polynomial even with the size of the grammar taken into account.

1.2.3 | Contributions

In summary, the dissertation makes two main contributions:

- (1) an algebraic theory of mildly non-projective dependency structures and regular sets of such structures (dependency languages), and
- (2) a classification of mildly context-sensitive, lexicalized grammar formalisms in terms of the dependency structures that these formalisms induce.

The algebraic theory complements previous work on dependency representations in that it enables us to link structural constraints such as projectivity, block-degree and well-nestedness to language-theoretic properties such as string-generative capacity and parsing complexity. The classification of grammar formalisms in terms of their ability to induce dependency structures yields a new, practically useful measure of generative capacity. Both results provide fundamental insights into the relation between dependency structures and lexicalized grammars.

1.2.4 | Published results

Some of the results reported in this dissertation have already been published in joint work with coauthors, in particular with my student, now colleague Mathias Möhl. For the present text, I have extended and streamlined the presentation of our results.

Well-nested Drawings as Models of Syntactic Structure. Tenth Conference on Formal Grammar and Ninth Meeting on Mathematics of Language. Edinburgh, UK, 2005. Joint work with Manuel Bodirsky and Mathias Möhl.

Lexicalised Configuration Grammars. Second International Workshop on Constraint Solving and Language Processing. Sitges, Spain, 2005. Joint work with Robert Grabowski and Mathias Möhl.

Extended Cross-Serial Dependencies in Tree Adjoining Grammars. Eighth International Workshop on Tree Adjoining Grammars and Related Formalisms (TAG+8). Sydney, Australia, 2006. Joint work with Mathias Möhl.

Mildly Non-Projective Dependency Structures. 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics (COLING-ACL), Main Conference Poster Sessions. Sydney, Australia, 2006. Joint work with Joakim Nivre.

Mildly Context-Sensitive Dependency Languages. 45th Annual Meeting of the Association for Computational Linguistics (ACL). Prague, Czech Republic, 2007. Joint work with Mathias Möhl.

The String-Generative Capacity of Regular Dependency Languages. Twelfth Conference on Formal Grammar. Dublin, Ireland, 2007. Joint work with Mathias Möhl.

CHAPTER 2 | Preliminaries

In this chapter, I provide a compact review of the basic terminology and notation that we will use in this dissertation. I invite the reader to browse through these preliminaries and return to them whenever some notation or terminology is unclear.

Our formal toolkit is drawn from four main sources: The terminology for mathematical structures and the relations between these structures come from universal algebra (Denecke and Wismath, 2001). The specific formalization of dependency structures takes an order-theoretic perspective (Davey and Priestley, 2001), but also alludes to graph theory (Diestel, 2005). To describe and manipulate structures and sets of structures, we make use of terms and term languages, and of the usual operations defined on them (Gécseg and Steinby, 1997). Note that we use the word ‘term’ for the syntactic object, and the word ‘tree’ when referring to the order-theoretic and graph-theoretic structures.

Basic notations

We write \mathbb{N} for the set of positive natural numbers, and define $\mathbb{N}_0 := \mathbb{N} \cup \{0\}$. Given a natural number $n \in \mathbb{N}$, we write $[n]$ to refer to the set $\{m \in \mathbb{N} \mid m \leq n\}$. Note that by this definition, $[0] = \emptyset$. We put $[n]_0 := [n] \cup \{0\}$. For a set A , we write $|A|$ for the cardinality of A , and $\mathfrak{P}(A)$ for the power set of A .

We use the notations A^* and A^+ to refer to the sets of all and all non-empty strings over the set A , respectively. We treat strings as vectors: the notation a_i refers to the i th element of the string \vec{a} . The length of a string \vec{a} is denoted by $|\vec{a}|$; the empty string is denoted by ε . The concatenation of two strings \vec{x} and \vec{y} is written as $\vec{x}\vec{y}$; only where this could create confusion, we use the alternative notation $\vec{x} \cdot \vec{y}$. An *alphabet* is a finite, non-empty set of *symbols*.

Indexed sets and sorted sets

Let I be a non-empty set. An *I -indexed set* is a total function with domain I . We use the notation $\langle x_i \mid i \in I \rangle$ to refer to an I -indexed set, where x_i denotes the image of i . We freely identify the set of all indexed sets with index set $[n]$, $n \in \mathbb{N}$, with the set of all n -tuples, and with the set of all strings of length n . An *I -indexed family* is an I -indexed set with a set-valued codomain. For indexed families, the usual set-theoretic operation are defined index-wise. In particular, if A and B both are I -indexed families, then $A \times B = \langle A_i \times B_i \mid i \in I \rangle$.

2 PRELIMINARIES

Let \mathcal{S} be a non-empty collection of *sorts*. An \mathcal{S} -sorted set consists of a non-empty set A and a *type assignment* $\text{type}_A: A \rightarrow \mathcal{S}^+$. We write \mathcal{S}_A for the collection of sorts that underlies A . When the sorted set under consideration is irrelevant or clear from the context, we write $a: \vec{s}$ instead of $\text{type}_A(a) = \vec{s}$; this is to be read as ‘ a has type \vec{s} in A ’. The set of all elements of A with type \vec{s} is denoted by $A_{\vec{s}}$. In the following, let A and B be sorted sets. We write $\langle A, B \rangle$ for the sorted set $A \times B$ in which $\langle a, b \rangle: \text{type}_A(a)$, for all $a \in A$, $b \in B$. For an element $a: \vec{s}s$ with $\vec{s} \in \mathcal{S}^*$ and $s \in \mathcal{S}$, the length of \vec{s} is called the *rank* of a , and is denoted by $\text{rank}_A(a)$. If $|\mathcal{S}| = 1$, the type of an element $a \in A$ is uniquely determined by its rank; in this case, the set A is called a *ranked set*, and the set of all elements with rank k is denoted by A_k . We freely treat \mathcal{S} -indexed sets A as \mathcal{S} -sorted sets by stipulating that $\text{type}_A(a) = s$ if and only if $a \in A_s$.

Structures

We now define the notion of a mathematical structure. Let Σ be a sorted set, now called a *signature*. A (concrete) Σ -structure is a pair

$$\mathfrak{A} = (\langle A_s^{\mathfrak{A}} \mid s \in \mathcal{S}_{\Sigma} \rangle, \langle R_{\sigma}^{\mathfrak{A}} \mid \sigma \in \Sigma \rangle),$$

where the first component is an \mathcal{S}_{Σ} -indexed family of non-empty sets, called the *domains* of \mathfrak{A} , and the second component is a Σ -indexed family of relations over the domains such that $R_{\sigma}^{\mathfrak{A}} \subseteq A_{s_1}^{\mathfrak{A}} \times \cdots \times A_{s_n}^{\mathfrak{A}}$, for every symbol $\sigma: s_1 \cdots s_n$. We use the notation $\text{dom}(\mathfrak{A})$ to refer to the domains of \mathfrak{A} . For structures with small signatures over a single sort, we use the compact notation $(A_s^{\mathfrak{A}}; R_1^{\mathfrak{A}}, \dots, R_n^{\mathfrak{A}})$, leaving the signature implicit. A structure is *finite*, if both its signature and all of its domains are finite sets.

Given two Σ -structures \mathfrak{A} and \mathfrak{B} , a *homomorphism* from \mathfrak{A} to \mathfrak{B} is an indexed family $\langle h_s \mid s \in \mathcal{S}_{\Sigma} \rangle$ in which, for each given sort $s \in \mathcal{S}_{\Sigma}$, the object h_s is a total function $h_s: \text{dom}(\mathfrak{A})_s \rightarrow \text{dom}(\mathfrak{B})_s$ with the property that

$$(a_1, \dots, a_n) \in R_{\sigma}^{\mathfrak{A}} \implies (h_s(a_1), \dots, h_s(a_n)) \in R_{\sigma}^{\mathfrak{B}},$$

for every $\sigma: s_1 \cdots s_n$, $a_i \in \text{dom}(\mathfrak{A})_{s_i}$, and $i \in [n]$. The notation $h: \mathfrak{A} \rightarrow \mathfrak{B}$ refers to a homomorphism between Σ -structures \mathfrak{A} and \mathfrak{B} , treating it as a single mapping rather than as an indexed family of mappings. Furthermore, to avoid subscript clutter, for $s \in \mathcal{S}_{\Sigma}$ and $a \in \text{dom}(\mathfrak{A})_s$, we write $h(a)$ rather than $h_s(a)$ for the image of a under the homomorphism h , assuming that h keeps type discipline. A homomorphism is called *epi* or an *epimorphism*, if every member function is injective; it is called *mono* or a *monomorphism*, if every member function is surjective; it is called *iso* or an *isomorphism*, if every member function is bijective. We do not distinguish between structures \mathfrak{A} and \mathfrak{B} for which there exists an isomorphism $h: \mathfrak{A} \rightarrow \mathfrak{B}$. Specifically, a Σ -structure is an equivalence class of concrete Σ -structures modulo isomorphism.

Ordered sets

An *ordered set* is a structure with a single binary relation that is reflexive, transitive, and anti-symmetric. Let $\mathfrak{R} = (A; \preceq)$ be an ordered set, and let $a, b \in A$. We write $a < b$ to assert that $a \preceq b$, and $a \neq b$. We say that a *immediately precedes* b (with respect to \preceq) if $a \preceq b$ and there is no element $c \in A$ such that $a \preceq c \preceq b$. The *up-closure* and the *down-closure* of a are defined as $\lfloor a \rfloor := \{ c \in A \mid a \preceq c \}$ and $\lceil a \rceil := \{ c \in A \mid c \preceq a \}$, respectively. For a given subset $B \subseteq A$, we say that \mathfrak{R} is *total* on B , if $a \preceq b$ or $b \preceq a$ holds for all $a, b \in B$. The structure \mathfrak{R} is called a *chain*, if it is total on A ; it is called a *forest*, if it is total on all down-closures; it is called a *tree*, if it is a forest and additionally contains an element r , the *root node* of \mathfrak{R} , with the property that $\lfloor r \rfloor = A$.

Observe that what we call ‘forests’ and ‘trees’ are the reflexive-transitive closures of the corresponding objects from graph theory. The elements of the domains of trees are called *nodes*. We use the symbols u, v and w for variables that range over nodes. Let $T = (V; \triangleleft)$ be a tree. If $u \triangleleft v$, we say that u *dominates* v . We write $u \rightarrow v$ if u immediately precedes v with respect to dominance; this relation corresponds to the edge relation in the formalization of trees as special directed graphs. We use the standard genealogical terminology to refer to relations between nodes in a tree: If $u \rightarrow v$, then we say that v is a *child* of u , and, symmetrically, that u is the *parent* of v . Distinct children of the same node are called *siblings*. We use the term *yield* as a synonym for the down-closure of u ; notice that $u \triangleleft v$ if and only if $v \in \lfloor u \rfloor$. The set of *descendants* and *ancestors* of u are defined as $\lceil u \rceil - \{u\}$ and $\lfloor u \rfloor - \{u\}$, respectively. Two nodes u and v are *disjoint*, if $\lfloor u \rfloor \cap \lfloor v \rfloor = \emptyset$. Each pair v, w of disjoint nodes has a greatest common ancestor u ; for this situation, we write $v \perp_u w$.

For chains, we define the notion of an *interval*: the interval with endpoints a and b is the set $[a, b] := (\lfloor a \rfloor \cap \lceil b \rceil) \cup (\lceil a \rceil \cap \lfloor b \rfloor)$. We also put $(a, b) := [a, b] - \{a, b\}$. A set is *convex*, if it is an interval.

Dependency structures

We now define *dependency structures*. Dependency structures are trees with a total order on their nodes. They differ from *ordered trees*, where the order relation is only defined on the children of each node, but not on the set of all nodes of the tree.

Definition A *dependency structure* is a structure D with two binary relations: one relation forms a tree, the second relation forms a chain on $\text{dom}(D)$. 201

The tree relation of a dependency structure is called *governance*, the total order is called *precedence*. Just like in trees, the elements of the domains of dependency structures are called *nodes*. Given a dependency structure D , for nodes $u, v \in \text{dom}(D)$, we write $u \trianglelefteq v$ to assert that u *governs* v , and $u \preceq v$ to assert that u *precedes* v in D . To talk about dependency structures, we import all terminology for trees and chains.

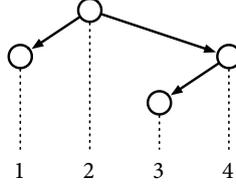


Figure 2.1: A dependency structure

202

Example Figure 2.1 shows how we visualize dependency structures. A picture of a dependency structure contains the nodes of the structure (drawn as circles), *edges* (drawn as pointed arrows), and *projection lines* (drawn as dotted lines). Specifically, we draw an edge between two nodes u and v if and only if $u \rightarrow v$. The nodes are ordered from left-to-right; we place u before v if and only if $u < v$. The projection lines are used to make the left-to-right order more explicit. \dashv

When discussing algorithms that operate on dependency structures, we assume a concrete representation of these structures as a collection of objects, where each object u has access to at least the object representing its parent node, $parent[u]$, and its position in the precedence order among all the nodes of the structure, $pos[u]$. We also make the (reasonable) assumption that both attributes can be accessed in constant time. With this representation in mind, it is straightforward that the following auxiliary mappings can be constructed in time linear in the size of the structure: a mapping $children[u]$ that maps each object u to the set of objects representing the children of u ; a mapping $node[i]$ that maps each position i (where i ranges over the size of the structure) to the object u for which $pos[u] = i$. Similarly, it is straightforward that the following iterations over the nodes in the structure can be supported in linear time: top-down, bottom-up, left-to-right.

Algebraic structures and terms

A Σ -structure $\mathfrak{A} = (\langle A_s^{\mathfrak{A}} \mid s \in \mathcal{S}_{\Sigma} \rangle, \langle f_{\sigma}^{\mathfrak{A}} \mid \sigma \in \Sigma \rangle)$ is called *algebraic* or an *algebra*, if for every symbol $\sigma: s_1 \cdots s_m s$, the relation $f_{\sigma}^{\mathfrak{A}}$ is a total function, meaning that for every $i \in [m]$ and $a_i \in A_{s_i}^{\mathfrak{A}}$, there is exactly one $a \in A_s^{\mathfrak{A}}$ such that $(a_1, \dots, a_m, a) \in f_{\sigma}^{\mathfrak{A}}$. In the context of algebras, we use the notation $\sigma: s_1 \times \cdots \times s_m \rightarrow s$ instead of $\sigma: s_1 \cdots s_m s$, and call m the *arity* of the symbol σ and the corresponding function f_{σ} .

Let Σ be a sorted set. The set of *terms* over Σ is the smallest \mathcal{S}_{Σ} -indexed family T_{Σ} such that if $\sigma: s_1 \times \cdots \times s_m \rightarrow s$ and $t_i \in T_{\Sigma, s_i}$ for all $i \in [m]$, then $\sigma(t_1, \dots, t_m) \in T_{\Sigma, s}$. Let $t \in T_{\Sigma}$ be a term. The set of *nodes* of t , $nod(t)$, is the subset of \mathbb{N}^* that is defined by the equation $nod(\sigma(t_1, \dots, t_m)) := \{\varepsilon\} \cup \{i \cdot u \mid i \in [m] \wedge u \in nod(t_i)\}$. The empty string ε represents the root node of t , and the string $i \cdot u$ represents the i th child of the node u . The *subterm* of t at node u is denoted by t/u , the *substitution* of the

term s at node u in t is denoted by $t[u \leftarrow s]$, and the *label* of t at node u is denoted by $t(u)$. We also put $\text{alph}(t) := \{t(u) \mid u \in \text{nod}(t)\}$, which thereby denotes the set of all labels in t . A *context* over Σ is a pair (t, u) , where $t \in T_\Sigma$ is a term, and u is a leaf node in t . We write C_Σ for the set of all contexts over Σ , and make free use of all term-related concepts even for contexts. Given a context $(c, u) \in C_\Sigma$ and a term $t \in T_{\Sigma, s}$ with $s = \text{type}_\Sigma(c(u))$, we write $c \cdot t$ for the term obtained as $c[u \leftarrow t]$.

The *term algebra* over Σ is the algebra \mathfrak{A} in which $\text{dom}(\mathfrak{A}) = T_\Sigma$, and in which each function $f_\sigma^{\mathfrak{A}}$ is interpreted as a term constructor in the obvious way. We use the notation T_Σ for both the set of terms over Σ and the term algebra over Σ . For every signature Σ and every Σ -algebra \mathfrak{A} , there is a uniquely determined homomorphism $\llbracket \cdot \rrbracket_{\mathfrak{A}}: T_\Sigma \rightarrow \mathfrak{A}$ that evaluates terms in T_Σ as values in \mathfrak{A} .

CHAPTER 3 | Projective dependency structures

The exact relation between dependency and word order is a major point of debate in dependency theory. Over the years, various authors have made proposals for formal constraints to restrict this relation. In this chapter, we study the best-known of these proposals, *projectivity*.

Structure of the chapter. We start by reviewing three standard characterizations of projectivity (Section 3.1), and then introduce a new, algebraic characterization (Section 3.2). This gives rise to an efficient algorithm to test whether a given dependency structure is projective (Section 3.3). We use this algorithm to evaluate the practical relevance of projectivity on data from three dependency treebanks (Section 3.4).

3.1 | Projectivity

Figure 3.1 shows pictures of five dependency structures. One of these, Figure 3.1d, is different from the others in that it displays *crossing edges*—the edge $1 \rightarrow 3$ crosses the projection line of node 2. The projectivity condition is often summarized in the slogan that ‘it disallows dependency structures with pictures that contain crossing edges’. This is a nice mnemonic, but whether a dependency edge crosses a projection line or not of course mainly is a matter of how we *draw* dependency structures, not a property of the structures themselves. For example, Figure 3.1d can be re-drawn without crossing edges if node 2 is moved to a vertical position below node 1 (see Figure 3.2d), while Figure 3.1b and Figure 3.1c exhibit crossing edges when modified such that node 2 is positioned above the root node (see Figure 3.2b and Figure 3.2c). It is clear then that to obtain a precise characterization of projectivity, we need a definition that formalizes the idea of crossing edges without referring to differences in visualization.

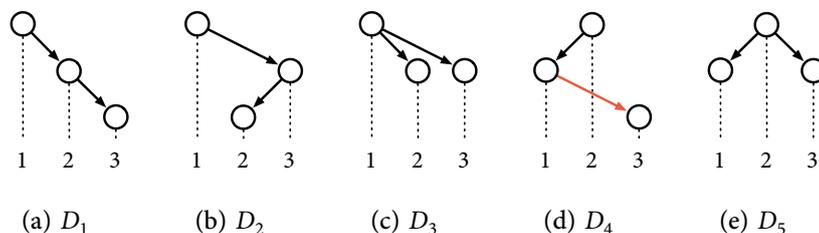


Figure 3.1: Five (of nine) dependency structures with three nodes

3 PROJECTIVE DEPENDENCY STRUCTURES

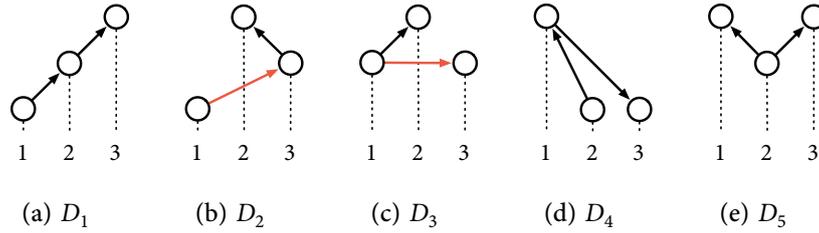


Figure 3.2: Alternative pictures for the dependency structures from Figure 3.1

3.1.1 | Projectivity in the sense of Harper and Hays

A crucial difference between the pictures in Figure 3.1 and the alternative versions in Figure 3.2 is that in the former, all tree edges point downwards, while in the latter, some of them also point upwards. Let us call a picture of a dependency structure *canonical*, if the vertical position of each node is chosen according to its level of depth in the tree, with the root node taking the highest position.

301

Example Figure 3.3 shows canonical and non-canonical pictures of the dependency structures D_2 and D_4 . The horizontal lines visualize the depth levels. ←

In canonical pictures of dependency structures, all tree edges point to the next level of depth. As a consequence, an edge $u \rightarrow v$ can cross the projection line of a node w only if the vertical position of u is the same as or below the vertical position of w . To ban crossing edges in canonical pictures, it is sufficient then to require the node u to govern the node w ; this guarantees that the vertical position of u is strictly above the vertical position of w . The requirement is made formal in the following implication, attributed to Kenneth Harper and David Hays (Marcus, 1967):

$$u \rightarrow v \wedge w \in (u, v) \implies u \preceq w \tag{3.1}$$

We introduce some useful terminology: Let u, v, w be witnesses for the premises of the implication (3.1). We then say that the edge $u \rightarrow v$ *covers* the node w , and call an edge that covers a node but does not govern it a *non-projective edge*. A dependency structure is projective if and only if it does not contain non-projective edges.

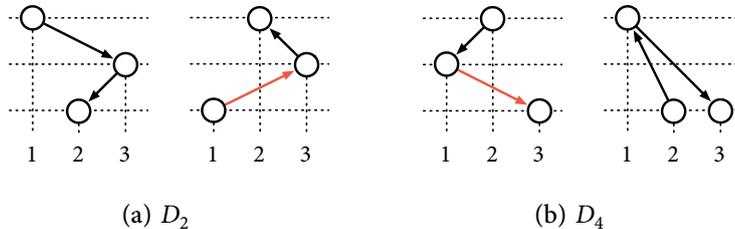


Figure 3.3: Canonical and non-canonical pictures of two dependency structures

Example The edge $1 \rightarrow 3$ in Figure 3.1d and Figure 3.2d is non-projective, as it covers the node 2, but does not govern it. All other edges depicted in Figure 3.1 and Figure 3.2 are projective. 302

3.1.2 | Projectivity in the sense of Lecerf and Ihm

The characterization of projectivity in the sense of Harper and Hays links projectivity to edges. The second characterization that we consider, attributed to Yves Lecerf and Peter Ihm (Marcus, 1967), anchors projectivity at paths:

$$u \triangleleft v \wedge w \in (u, v) \implies u \triangleleft w \quad (3.2)$$

Note that the only difference between this requirement and (3.1) is the first premise of the implication: projectivity in the sense of Lecerf and Ihm requires not only every edge, but every (directed) path from a node u to a node v to cover only nodes w that are governed by u . Since every path consists of a finite sequence of edges, the characterizations of projectivity in the sense of Harper and Hays and in the sense of Lecerf and Ihm are fully equivalent (Marcus, 1967, chapter 6, Theorem 10).

3.1.3 | Projectivity in the sense of Fitalov

We now present a third characterization of projectivity. This characterization formalizes the observation that in a projective dependency analysis, a word and its (transitive) dependents form a contiguous substring of the full sentence. It is usually attributed to Sergey Fitalov (Marcus, 1967).

$$u \triangleleft v_1 \wedge u \triangleleft v_2 \wedge w \in (v_1, v_2) \implies u \triangleleft w \quad (3.3)$$

This condition is equivalent to the preceding two (Marcus, 1967, chapter 6, Theorem 11). Using our terminology for chains, we can rephrase it more succinctly as follows:

Definition A dependency structure D is called *projective*, if every yield in D is convex with respect to precedence. 303

This is the formulation that we adopt as our formal definition of projectivity. We write \mathcal{D}_1 to refer to the class of all projective dependency structures.

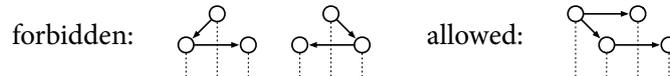
Example (continued) In the dependency structure shown in Figure 3.1d, the yield of the node 1 (the set $\{1, 3\}$) does not form a convex set: the node 2 is missing from it. All other yields in Figure 3.1 are convex. Therefore, the only non-projective dependency structure in Figure 3.1 is structure D_4 . 304

3.1.4 | Related work

⌘ While the fundamental intuitions behind projectivity are already inherent in work on machine translation from the 1950s, the characterizations in the sense of Harper and Hays and in the sense of Lecerf and Ihm appear to be the first formal definitions of the condition; they were both published in 1960. Marcus (1967) collects and compares several definitions of projectivity that circulated in the second half of the 1960s. In particular, he proves the equivalence of the characterizations of projectivity in the senses of Harper and Hays, Lecerf and Ihm, and Fitalov.

⌘ There are several other equivalent characterizations of projectivity; we only name two here. The following formulation is due to Robinson (1970, p. 260); it is sometimes referred to as the ‘adjacency principle’, a term that appears to have been coined by Hudson (1984, p. 98): ‘If A depends directly on B and some element C intervenes between them (in linear order of string), then C depends directly on A or on B or on some other intervening element.’ Havelka (2005) presents an original edge-centric characterization of projectivity based on the difference between the depth of the dependent node of a dependency edge and the material covered by it, and uses it as the basis for an algorithm to test whether a given dependency structure is projective.

⌘ Veselá et al. (2004) propose a characterization of projectivity based on ‘forbidden elementary configurations’ in a dependency structure, but the condition that they define still allows some non-projective structures:



⌘ The characterization of projectivity in terms of convex yields sheds some light on the relation between dependency grammar and phrase-structure grammar: If one accepts that yields reconstruct the notion of constituents that is familiar from phrase-structure grammars, then the projectivity condition amounts to the standard requirement that a constituent should be contiguous. In this sense, projective dependency structures are closely related to standard phrase-structure trees. This correspondence was first investigated by Hays (1960). The survey by Dikovsky and Modina (2000, section 3.2) summarizes some of the formal results obtained since then.

3.2 | Algebraic framework

We have characterized projectivity as a relational property of dependency structures. The immediate value of this characterization is that it is empirically transparent: from a canonical picture of a given dependency structure, we can immediately *see* whether it is projective.

```

PRE-ORDER-COLLECT( $u$ )
1   $L \leftarrow \text{NIL}$ 
2   $L \leftarrow L \cdot [u]$ 
3  foreach  $v$  in  $\text{children}[u]$ 
4      do  $L \leftarrow L \cdot \text{PRE-ORDER-COLLECT}(v)$ 
5  return  $L$ 

POST-ORDER-COLLECT( $u$ )
1   $L \leftarrow \text{NIL}$ 
2  foreach  $v$  in  $\text{children}[u]$ 
3      do  $L \leftarrow L \cdot \text{POST-ORDER-COLLECT}(v)$ 
4   $L \leftarrow L \cdot [u]$ 
5  return  $L$ 

```

Table 3.1: Pre-order and post-order traversal of a children-ordered tree

What is not clear at this point is how the projectivity constraint can be fitted into the framework of generative grammar, where a dependency structure is not given directly, but specified as the outcome of a derivational process. In this section, we clarify this issue: We equip the class of projective dependency structures with a set of algebraic operations. Each application of an operation can be understood as the application of a single production of a generative grammar. Our setup guarantees that all operations only yield projective structures, and that all projective structures can be decomposed into elementary operations. In this way, every projective dependency structure can be understood as the outcome of a complete derivation in a generative grammar with a suitable set of algebraic operations.

3.2.1 | Tree traversal strategies

To convey the basic intuitions behind our algebraic setting, we start this section by looking at *tree traversal strategies*. A *tree traversal* is the process of systematically visiting all nodes of a tree. Two well-known strategies for tree traversal are *pre-order traversal* and *post-order traversal* of children-ordered trees. For the sake of concreteness, let us assume that a children-ordered tree is represented as a collection of nodes, where each node u is annotated with a list $\text{children}[u]$ of its children. We can then specify procedures to collect the nodes of a tree as in Table 3.1.¹ The result of a call to $\text{PRE-ORDER-COLLECT}(u)$ or $\text{POST-ORDER-COLLECT}(u)$ is a list of the nodes in the tree rooted at the node u ; each node of the tree occurs in this list exactly once.

¹ The format of our pseudo-code follows Cormen et al. (2001). We write $[x]$ for the singleton list that contains the element x , and $L_1 \cdot L_2$ for the concatenation of two lists L_1 and L_2 .

3 PROJECTIVE DEPENDENCY STRUCTURES

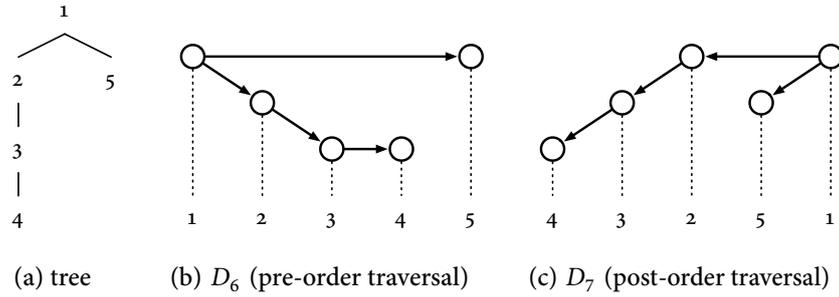


Figure 3.4: Dependency structures obtained by tree traversals of a children-ordered tree

Both PRE-ORDER-COLLECT and POST-ORDER-COLLECT extend the orders among the children of each node into a global order, defined on all nodes of the tree. When we impose this global order on the nodes of the tree that was traversed, we obtain a dependency structure. This observation is formalized in the following definitions.

305 **Definition** Let T be a tree. A *linearization* of T is a list of nodes of T in which each node occurs exactly once. The dependency structure *induced* by a linearization \vec{u} of T is the structure in which the governance relation is isomorphic to T , and the precedence relation is isomorphic to \vec{u} . \dashv

306 **Example** Figure 3.4a shows a children-ordered tree. The pre-order linearization of this tree yields the node sequence 12345. When we impose this order on the nodes in the tree, we obtain the dependency structure D_6 shown in Figure 3.4b. In contrast, the post-order linearization of the children-ordered tree yields the node sequence 43251; this induces the dependency structure D_7 shown in Figure 3.4c. \dashv

We now sketch our plan for the remainder of this section: Our first goal is to find a traversal strategy and a class of order-annotated trees that fully characterize the class of projective dependency structures—traversals should only induce projective structures, and every projective structure should be inducible by some traversal. In a second step, we formalize this framework by regarding the set of all order annotations as an algebraic signature, order-annotated trees as terms over this signature, and tree traversal as the evaluation of these terms in an algebra over dependency structures.

With this roadmap in mind, let us see how far pre-order and post-order traversal take us. One property that both strategies have in common is that they visit the nodes in each subtree as a contiguous block. As a consequence, every dependency structure that is induced by pre-order or post-order traversal is projective. On the other hand, not every projective dependency structure can be obtained as the pre-order or post-order interpretation of a children-ordered tree. Specifically, all structures induced by pre-order traversal are *monotonic* in the sense that $u \triangleleft v$ implies that $u \leq v$, while all structures obtained by post-order traversal are *anti-monotonic*. The fundamental

```

TREELET-ORDER-COLLECT( $u$ )
1   $L \leftarrow \text{NIL}$ 
2  foreach  $v$  in  $\text{order}[u]$ 
3      do if  $v = u$ 
4          then  $L \leftarrow L \cdot [u]$ 
5          else  $L \leftarrow L \cdot \text{TREELET-ORDER-COLLECT}(v)$ 
6  return  $L$ 

```

Table 3.2: Traversal of a treelet-ordered tree

reason for these specific properties is that in both pre-order and post-order traversal, the position of a node relative to its children is hardwired: it is not specified in the order annotations, but in the traversal strategy itself.

3.2.2 | Traversal of treelet-ordered trees

To overcome the restrictions of pre-order and post-order traversal, we include the position of a node relative to its children in the order annotations, and make the traversal strategy sensitive to this information. Let us call the local tree formed by a node u and its children (if there are any) the *treelet* rooted at u , and let us say that a tree is *treelet-ordered*, if each of its nodes is annotated with a total order on the nodes in the treelet rooted at that node. Table 3.2 gives the pseudo-code of a procedure that traverses a treelet-ordered tree and returns a list of its nodes. We assume that each node u in the tree is annotated with a list $\text{order}[u]$ that contains the nodes in the treelet rooted at u in the intended order.

Example Figure 3.5a visualizes a treelet-ordered tree; the sequences at the ends of the dotted lines represent the annotated orders. The traversal of this tree according to the procedure in Table 3.2 yields the sequence 24315. When we impose this order on the nodes in the tree, we obtain the dependency structure D_8 shown in Figure 3.5b. Note that this structure is neither monotonic nor anti-monotonic. 307

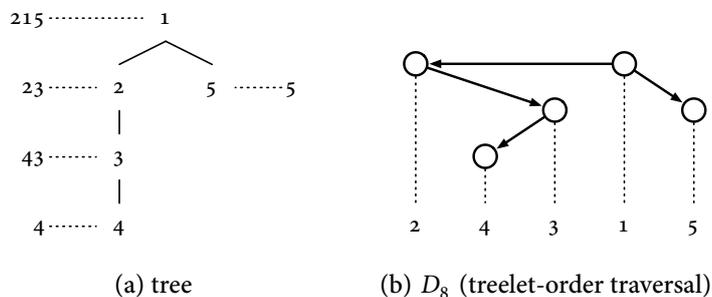


Figure 3.5: A treelet-ordered tree and its corresponding dependency structure

3 PROJECTIVE DEPENDENCY STRUCTURES

We now show that treelet-ordered trees and our procedure for their traversal are expressive enough to fully characterize the class of projective dependency structures. What is more, distinct treelet-ordered trees induce distinct structures. In our proofs, we use two functions on treelet-ordered trees: the function lin that maps each tree to its linearization according to `TREELET-ORDER-COLLECT`, and the function dep that maps each tree T to the dependency structure induced by $\text{lin}(T)$. We then have

$$\begin{aligned} u \preceq v \text{ in } \text{dep}(T) & \quad \text{if and only if} & \quad u \text{ dominates } v \text{ in } T, \text{ and} \\ u \leq v \text{ in } \text{dep}(T) & \quad \text{if and only if} & \quad u \text{ precedes } v \text{ in } \text{lin}(T). \end{aligned}$$

The next three lemmata show that the function dep forms a bijection between the set of all treelet-ordered trees and the set of all projective dependency structures.

308 **Lemma** Let T be a treelet-ordered tree. Then $\text{dep}(T)$ is projective. ←

Proof. Let T be a treelet-ordered tree, and let u be the root node of T . We show that every yield of T is convex with respect to the total order on the nodes of T that is represented by the linearization $\text{lin}(T)$. Our proof proceeds by induction on the depth d of T .

- * Assume that $d = 0$. In this case, the node u is the only node in T , and we have $\text{order}[u] = [u]$. Therefore, $\text{lin}(T) = [u]$, and $\text{dep}(T)$ is the trivial dependency structure. The yield $[u]$, like all singleton sets, is convex with respect to $\text{lin}(T)$.
- * Assume that $d > 0$. In this case, the tree T can be decomposed into the node u and the collection of subtrees rooted at the children of u . Let $w \neq u$ be a node in T , and let v be the uniquely determined child of u that dominates w . By the induction hypothesis, we may assume that the yield $[w]$ is convex with respect to the linearization that was computed by the recursive call `TREELET-ORDER-COLLECT`(v). The result $\text{lin}(T)$ of the call `TREELET-ORDER-COLLECT`(u) is a concatenation of these linearizations and the singleton list $[u]$; thus, the yield $[w]$ is convex even with respect to $\text{lin}(T)$. The yield $[u]$, being the set of all nodes in T , is trivially convex with respect to $\text{lin}(T)$. ←

For the next lemma, we introduce an important auxiliary concept.

309 **Definition** Let D be a dependency structure, and let u be a node in D . The set of *constituents* of u is defined as $\mathcal{C}(u) := \{\{u\}\} \cup \{\{v\} \mid u \rightarrow v\}$. ←

310 **Lemma** For every projective dependency structure D , there is a treelet-ordered tree T such that $D = \text{dep}(T)$. ←

Proof. Let D be a projective dependency structure with root node u . We show how to construct a treelet-ordered tree T such that $D = \text{lin}(T)$. Note that the tree structure

underlying T is isomorphic to the tree structure underlying D . Therefore, it suffices to show that we can find appropriate order annotations for the nodes in T such that $\text{lin}(T)$ corresponds to the precedence relation in D . We proceed by induction on the depth d of the tree structure underlying D .

- * Assume that $d = 0$. In this case, we have only one choice to assign an order annotation to u , $\text{order}[u] = [u]$. With this annotation, we indeed have $\text{lin}(T) = [u]$.
- * Assume that $d > 0$. In this case, the node u has an out-degree of $n > 0$. The set $\mathcal{C}(u)$ of constituents of u forms a partition of the yield $[u]$. Furthermore, every constituent is convex with respect to the order underlying D : the set $\{u\}$ trivially so, and each set $[v]$ because the structure D is projective. We can also verify that for every constituent $C \in \mathcal{C}(u)$, the restriction $D|_C$ forms a projective dependency structure on C . Thus, by the induction hypothesis, we may assume that for every child v of u , we can annotate the subtree T/v such that $\text{dep}(T/v) = D|_{[v]}$. What remains to be shown is that we can annotate u such that the call $\text{TREELET-ORDER-COLLECT}(u)$ arranges the constituents $\mathcal{C}(u)$ according to their relative precedence in D .

Generic procedure to extract the order annotations from a dependency structure.

To construct the order annotation for the node u , let $\pi: [u] \rightarrow [u]$ be the function that maps u to itself and every other node $v \in [u]$ to the uniquely determined child of u that governs v . Now, let L be the list of all nodes in D in the order of their precedence, and let L' be the list obtained from L by replacing each node w with the node $\pi(w)$ if $w \in [u]$, and the symbol \square if $w \notin [u]$. Finally, let $\text{order}[u]$ be the list obtained from L' by collapsing all adjacent occurrences of the same symbol into a single occurrence, and removing all leading and trailing symbols \square . For this order annotation of u , we can verify that the call $\text{TREELET-ORDER-TRAVERSAL}(u)$ returns the constituents $\mathcal{C}(u)$ in the order that they have in D . \dashv

Example (continued) For the dependency structure D_7 shown in Figure 3.5b, the construction described in the proof yields 311

for node 1:	$L = 24315$	$L' = 22215$	$\text{order}[1] = 215,$
for node 2:	$L = 24315$	$L' = 233 \square \square$	$\text{order}[2] = 23,$
for node 3:	$L = 24315$	$L' = \square 43 \square \square$	$\text{order}[3] = 43,$
for node 4:	$L = 24315$	$L' = \square 4 \square \square \square$	$\text{order}[4] = 4,$
for node 5:	$L = 24315$	$L' = \square \square \square \square 5$	$\text{order}[5] = 5.$

Note that these are the order annotations shown in Figure 3.5a. \dashv

Lemma For every projective dependency structure D , there is at most one treelet-ordered tree T such that $\text{dep}(T) = D$. 312 \dashv

3 PROJECTIVE DEPENDENCY STRUCTURES

Proof. Let D be a projective dependency structure, and let T be a treelet-ordered tree such that $\text{dep}(T) = D$. Now let T' be another treelet-ordered tree, distinct from T , and consider the dependency structure $\text{dep}(T')$. We distinguish two cases: If the tree structures underlying T and T' are non-isomorphic, then $\text{dep}(T)$ and $\text{dep}(T')$ are non-isomorphic as well. Otherwise, the tree T' differs from T with respect to some order annotation. Then the call to `TREELET-ORDER-COLLECT` returns a different order for T than for T' ; hence, $\text{dep}(T)$ and $\text{dep}(T')$ are non-isomorphic. \dashv

3.2.3 | Order annotations

We now translate our framework into the language of terms: we regard the list-based order annotations that we used in `TREELET-ORDER-COLLECT` as a ranked set Ω , and treelet-ordered trees as terms over this set. This allows us to reinterpret the function dep as a bijection between T_Ω and the class of projective dependency structures.

While our list-based order annotations were sequences of (pointers to) concrete nodes in a treelet-ordered tree, the ranked set Ω should be defined independently of any particular term over this set. Therefore, we add a layer of indirection: each order annotation in Ω refers to nodes not directly, but by *names* for these nodes; these names are then resolved given the term structure. Specifically, let T be a treelet-ordered tree with root node u . We need two auxiliary sequences: the vector $\vec{v} = v_1 \cdots v_m$ obtained from $\text{order}[u]$ by removing the node u , and the string \vec{i} obtained from $\text{order}[u]$ by replacing every child of u by its position in \vec{v} , and u itself by the symbol 0. The vector \vec{v} orders the children of u ; this order will become the left-to-right order on the children of u in our term representation. The string \vec{i} provides an ‘abstract’ order annotation that makes use of node names rather than concrete nodes: the name 0 denotes the root node of the treelet rooted at u , a name $i \in [m]$ denotes the i th node in the sequence \vec{v} . The term $t(T)$ corresponding to T is then defined recursively as

$$t(T) := \langle \vec{i} \rangle (t(T/v_1), \dots, t(T/v_m)).$$

In this definition, the string $\langle \vec{i} \rangle$ is understood as a term constructor of rank m . We write Ω_m for the set of all such constructors, and put $\Omega := \bigcup_{m \in \mathbb{N}} \Omega_m$. Every term over Ω encodes a treelet-ordered tree in the way that we have just described, and every such tree can be encoded into a term. In this way, we can view the function dep as a bijection $\text{dep}: T_\Omega \rightarrow \mathcal{D}_1$ in the obvious way. We put $\text{term} := \text{dep}^{-1}$.

313 **Example** Figure 3.6 shows the term for the treelet-ordered tree from Figure 3.5a. \dashv

3.2.4 | Dependency algebras

Using the ranked set Ω and the bijection $\text{dep}: T_\Omega \rightarrow \mathcal{D}_1$ between terms over Ω and projective dependency structures, we now give the set \mathcal{D}_1 an algebraic structure: with every order annotation, we associate an operation on projective dependency structures.

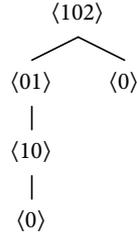


Figure 3.6: The term for the treelet-ordered tree from Figure 3.5a

Definition Let $m \in \mathbb{N}_0$, and let $\omega \in \Omega_m$ be an order annotation. The *composition operation* corresponding to ω is the function $f_\omega: \mathcal{D}_1^m \rightarrow \mathcal{D}_1$ defined as 314

$$f_\omega(D_1, \dots, D_m) := \text{dep}(\omega(\text{term}(D_1), \dots, \text{term}(D_m))). \quad \dashv$$

Composition operations are well-defined because the function dep is bijective. Each composition operation f_ω simulates a single step of the treelet-order traversal: given a sequence of argument structures, it returns the dependency structure that is obtained by taking the disjoint union of the arguments, adding a new root node, and arranging the nodes of the arguments and the root node in the order specified by ω .

Example Figure 3.7 shows some examples for the results of composition operations. 315 The composition of zero arguments, $f_{\langle 0 \rangle}$, is the trivial dependency structure with one node (Figure 3.7a). Starting from this structure, more and more complex dependency structures can be built (Figures 3.7b–3.7d). \dashv

We now have everything we need to define our algebraic setting. In the following definition, we use the function dep lifted to sets in the obvious way.

Definition Let $\Sigma \subseteq \Omega$ be a finite set of order annotations. The *dependency algebra* over Σ is the Σ -algebra that has $\text{dep}(T_\Sigma)$ as its carrier set, and interprets each symbol $\omega \in \Sigma$ by the composition operation corresponding to ω . 316 \dashv

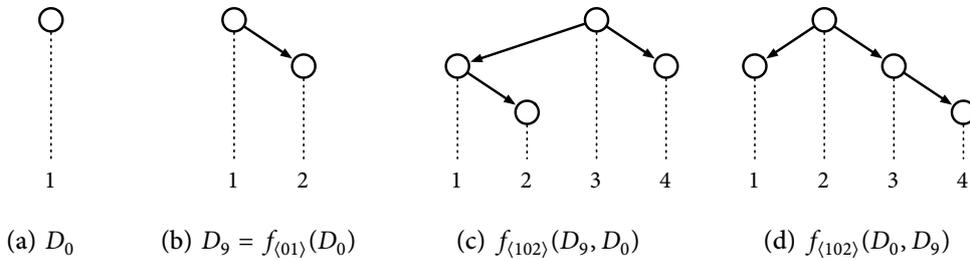


Figure 3.7: Examples for composition operations

By definition, dependency algebras are isomorphic to term algebras:

317 **Theorem** Let $\Sigma \subseteq \Omega$ be a finite set of order annotations. Then the dependency algebra over Σ is isomorphic to the term algebra over Σ , T_Σ . ←

Proof. Let \mathcal{D} be the dependency algebra over Σ . The restriction h of dep to the set of all terms over Σ is a bijection between T_Σ and the set $\text{dep}(T_\Sigma)$, the carrier of \mathcal{D} . Furthermore, from the definition of the composition operations we see that h forms a Σ -homomorphism between the term algebra T_Σ and \mathcal{D} :

$$h(\omega(t_1, \dots, t_m)) = h(\omega(h^{-1}(h(t_1)), \dots, h^{-1}(h(t_m)))) = f_\omega(h(t_1), \dots, h(t_m)).$$

Hence, T_Σ and \mathcal{D} are isomorphic. ←

One convenient consequence of the isomorphism between dependency algebras and their corresponding term algebras is that we can make use of all the terminology and notations available for terms when reasoning about dependency structures.

3.3 | Algorithmic problems

We now address three algorithmic problems associated to projectivity: the problems of encoding a projective dependency structure into its corresponding term, the symmetric problem of decoding a term into a dependency structure, and the problem of deciding whether a given dependency structure is projective.

3.3.1 | Encoding and decoding

The encoding problem for projective dependency structures is to compute, for a given dependency structure D , the term $\text{term}(D)$. Since the tree relation of a dependency structure and its corresponding term are isomorphic, the crucial task when encoding a projective structure into a term is to extract the order annotations for the nodes of the structure. A naïve procedure to solve this task is inherent in our proof of the result that the function dep is onto (page 27). This procedure can be implemented to run in time $O(n^2)$, where n is the number of nodes in D . Each order annotation reflects the restriction of the precedence relation to the nodes in the treelet rooted at u . Consequently, each list $\text{order}[u]$ contains the nodes in the treelet rooted at u in the order in which they appear in D . It is not hard to see that we can populate all of these lists in a single iteration over the nodes in the order of their precedence. Pseudo-code for the procedure is given in Table 3.3. Assuming that all elementary operations on D take constant time, and that an iteration takes time $O(n)$, extracting the order annotations and hence encoding can be done in time $O(n)$ as well.

318 **Lemma** Let D be a dependency structure with n nodes. Then the term $\text{term}(D)$ can be computed in time $O(n)$. ←

```

EXTRACT-ORDER-ANNOTATIONS( $D$ )
1  foreach  $u$  in  $D$ 
2      do  $order[u] \leftarrow \text{NIL}$ 
3  foreach  $u$  in  $D$  ▷ in the order of their precedence in  $D$ 
4      do if  $parent[u] \neq \text{UNDEFINED}$  ▷  $u$  is an inner node
5          then  $order[parent[u]] \leftarrow order[parent[u]] \cdot [u]$ 
6           $order[u] \leftarrow order[u] \cdot [u]$ 

```

Table 3.3: Extracting the order annotations for a projective dependency structure

The problem of decoding a term $t \in T_\Omega$ into its corresponding dependency structure is solved by the tree-traversal procedure that we gave in Table 3.2. Assuming that all elementary operations of that procedure take constant time, it is clear that the full precedence relation can be constructed in time linear in the size of the input term.

Lemma Let $t \in T_\Omega$ be a term with n nodes. Then the projective dependency structure $\text{dep}(t)$ can be computed in time $O(n)$. 319 \dashv

3.3.2 | Testing whether a dependency structure is projective

The isomorphism between projective dependency structures and terms over Ω gives rise to a simple and efficient algorithm for testing whether a given structure is projective. Note that nothing in our encoding procedure hinges on the input structure being projective. At the same time, only for projective structures this encoding produces terms that can be decoded back into the original structures. Therefore, the following algorithm is a correct test for projectivity of a given input structure D : encode D into the term $\text{term}(D)$, decode $\text{term}(D)$ into the dependency structure $D' := \text{dep}(\text{term}(D))$, and test whether D' and D are isomorphic. This test will succeed if and only if D is projective. Since encoding and decoding are linear-time operations, and since checking that two dependency structures are isomorphic is a linear-time operations as well, we obtain the following result:

Lemma Let D be a dependency structure with n nodes. The question whether D is projective can be decided in time $O(n)$. 320 \dashv

3.3.3 | Related work

✎ Havelka (2005) presents two algorithms for testing whether a given dependency structure is projective. The first algorithm, very much like ours, makes use of the one-to-one correspondence between projective dependency structures and treelet-ordered trees. The second algorithm searches for certain types of non-projective dependency edges. Both algorithms run in linear time.

3 PROJECTIVE DEPENDENCY STRUCTURES

✎ Another projectivity test is proposed by Möhl (2006). It uses a post-order traversal of the input dependency structure to compute, for each node u , a bit vector representing the yield of u , and afterwards checks whether this bit vector represents a convex set. The number of bit vector operations used by this procedure is linear in the size of the input structure. It is difficult however to compare this machine-dependent measure with the asymptotic runtime that we have given for our algorithm.

3.4 | Empirical evaluation

In this section, we evaluate the practical relevance of the projectivity condition. Should it turn out that all interesting dependency structures of natural language utterances are projective, then that result would indicate that theories that do not obey the projectivity restriction fail to reflect a deeper truth about the nature of dependency. Of course, we cannot hope to ever have access to ‘all interesting dependency structures’. However, we can estimate the empirical adequacy of projectivity by looking at representative samples of practically relevant data.

3.4.1 | The projectivity hypothesis

Before we describe our experimental setup, we take a brief look at the historical assessment of projectivity as a constraint on dependency analyses.

Early work on formal dependency grammar shows conviction that projectivity has the status of a linguistic universal. To witness, Marcus (1967, p. 230) cites Lecerf, who claimed that ‘almost 100 percent of French strings are projective. The same seems to be true for German, Italian, Danish, and other languages’. This rather radical *projectivity hypothesis* is disputable even without empirical evaluation. In particular, one should note that projectivity is a property of theory-specific *analyses* of sentences, not of the sentences themselves. Consequently, not ‘almost 100 percent of French strings’, but at most all of their dependency analyses can be projective. This fundamental flaw of the argument may have been varnished over by the supremacy in the 1960s of dependency grammar formalisms that embraced projectivity as a central grammatical principle (Hays, 1964; Gaifman, 1965): there simply was no dependency grammar beyond ‘projective dependency grammar’. In the linguistic schools of Eastern Europe, where the objects of linguistic description are languages with a word order far less rigid than English, the status of projectivity as a linguistic universal was early mistrusted (see e.g. Kunze, 1972; Pericliev and Ilarionov, 1986; Dikovskiy and Modina, 2000). This assessment eventually became accepted even in the Western literature, and today, ‘most theoretical formulations of dependency grammar regard projectivity as the norm, but also recognize the need for non-projective representations of certain linguistic constructions, e.g., long-distance dependencies’ (Nivre, 2006b).

	DDT		PDT 1.0		PDT 2.0	
projective	3 730	84.95%	56 168	76.85%	52 805	77.02%
non-projective	661	15.05%	16 920	23.15%	15 757	22.98%
TOTAL	4 391	100.00%	73 088	100.00%	68 562	100.00%

Table 3.4: The number of projective dependency structures in three treebanks

With the availability of large corpora of dependency analyses, *dependency treebanks*, we are able today to complement theoretical considerations concerning projectivity by collecting data on its *practical* relevance: the data that we are evaluating forms the basis for many current applications that build on dependency-based representations, and the degree of projectivity in this data may have direct consequences for the design of these applications. Furthermore, under the assumption that the treebank data forms a representative sample of the set of useful dependency structures, these data also provide an indirect evaluation of the empirical adequacy of projectivity.

3.4.2 | Experimental setup

Our experiments are based on data from the Prague Dependency Treebank (PDT) (Hajič et al., 2001, 2006) and the Danish Dependency Treebank (DDT) (Kromann, 2003). The PDT was used in two versions: version 1.0 contains 1.5M, version 2.0 contains 1.9M tokens of newspaper text. Sentences in the PDT are annotated in three layers according to the theoretical framework of Functional Generative Description (Hajičová et al., 2004). Our experiments concern only the analytical layer, and are based on the dedicated training section of the treebank. The DDT comprises 100k words of text selected from the Danish PAROLE corpus, with annotation of primary and secondary dependencies based on Discontinuous Grammar (Kromann, 2005). Only primary dependencies are considered in the experiments, which are based on the pseudo-randomized training portion of the treebank.² A total number of 19 analyses in the DDT were excluded because they contained annotation errors.

3.4.3 | Results and discussion

The results of our experiments are given in Table 3.4; we report the number and percentage of structures in each data set that satisfy or violate the projectivity condition.

Under the assumption that the three treebanks constitute a representative sample of the set of practically relevant dependency structures, our experiments clearly show that non-projectivity cannot be ignored without also ignoring a significant portion of

² Since the DDT does not have a dedicated training section, it is custom practice to create such a section by splitting the entire data into blocks of 10 analyses each, and keeping blocks 1 to 8 for training.

3 PROJECTIVE DEPENDENCY STRUCTURES

real-world data. For the DDT, we see that about 15% of all analyses are non-projective; for the PDT, the number is even higher, around 23% in both versions of the treebank. Neither theoretical frameworks nor practical applications that are confined to projective analyses can account for these analyses, and hence cannot achieve perfect recall even as an ideal goal. In a qualification of this interpretation, one should note that projectivity fares much better under an evaluation metric that is based on the set of individual edges, rather than on the set of complete analyses: less than 2% of the edges in the PDT data, and just around 1% of the edges in the DDT data are non-projective (McDonald et al., 2005; Nivre and Nilsson, 2005).

3.4.4 | Related work

✎ Our experiments confirm the findings of recent studies on data-driven parsing of non-projective dependency grammar (McDonald et al., 2005; Nivre and Nilsson, 2005). They are particularly similar in vein to a study presented by Nivre (2006a). Nivre’s main objective was to evaluate, how large a proportion of the structures found in the DDT and the PDT can be parsed using several restricted versions of the ‘Fundamental Algorithm’ for dependency parsing (Covington, 2001). Using a version of that algorithm that only recognizes projective structures, and employing the treebanks as oracles to resolve ambiguities, Nivre effectively tested for projectivity. For the PDT part of the data, our results are identical to his, which in turn agree with counts previously reported by Zeman (2004, p. 95). The minor deviation between our results and Nivre’s for the DDT part of the data is explained by the 19 analyses that we excluded because they contained annotation errors. Havelka (2007a) provides data on the frequency of non-projective structures in data sets for Arabic, Bulgarian, Czech, Danish, Dutch, German, Japanese, Portuguese, Slovene, Spanish, Swedish, and Turkish. Notice however, that some of these data sets are no dependency treebanks, but result from the automatic conversion of treebanks that were originally annotated using constituent structures.

✎ A qualitative rather than quantitative approach towards the evaluation of projectivity was taken by Pericliev and Ilarionov (1986). They used a hand-written dependency grammar for Bulgarian to create example sentences for all non-projective structures with 4 nodes (every larger non-projective structure contains such a structure) and found that about 85% of these structures could be instantiated with a grammatical sentence.³ Just as our experiments, this result indicates that projectivity cannot be used as a language-theoretic universal. Nevertheless, Pericliev and Ilarionov concede that most non-projective analyses in Bulgarian correspond to word orders that are stylistically marked.

³ Pericliev and Ilarionov are misled in assuming that ‘the total number of non-projective situations [in dependency structures with 4 nodes] is 32’ (p. 57): since the number of unrestricted dependency structures with 4 nodes is 64, and the corresponding number of projective structures is 30, there are in fact 34 non-projective dependency structures with four nodes.

CHAPTER 4 | Dependency structures of bounded degree

As we have seen in the previous chapter, the phenomenon of non-projectivity cannot be ignored in practical applications. At the same time, the step from projectivity to unrestricted non-projectivity is quite a dramatic one. In this chapter, we study non-projective dependency structures under a gradual relaxation of projectivity, the *block-degree restriction*.

Structure of the chapter. Following our program from the previous chapter, we first characterize the class of dependency structures with restricted block-degree in terms of a structural constraint (Section 4.1), then build an algebraic framework for this class (Section 4.2), next present an efficient algorithm that encodes dependency structures into terms (Section 4.3), and finally evaluate the practical relevance of the block-degree restriction on treebank data (Section 4.4).

4.1 | The block-degree measure

In projective dependency structures, each yield forms a set that is convex with respect to the precedence relation. In non-projective structures, yields may be discontinuous. In this section, we develop a formal measure that allows us to classify dependency structures based on their degree of non-projectivity: the minimal number of convex sets needed to cover all nodes of a yield.

4.1.1 | Blocks and block-degree

The formal cornerstone of our measure is the notion of a *congruence relation* on a chain. In general, congruence relations (or simply: congruences) are equivalence relations that are compatible with certain properties of the underlying mathematical structure. For chains, a natural notion of congruence is obtained by requiring each equivalence class to form a convex set.

Definition Let $\mathcal{C} = (A; \preceq)$ be a chain, and let $S \subseteq A$ be a set. An equivalence relation on S is called a *congruence* on S , if each of its classes is convex with respect to \mathcal{C} . \dashv

401

Example Let \mathcal{C}_6 be the set $[6]$, equipped with the standard order on natural numbers, and consider the set $S_1 := \{1, 2, 3\}$. There are four possible congruence relations on S_1 . Using shades to mark the elements of S_1 , and boxes to mark equivalence classes, these relations can be visualized as in Figure 4.1a. Similarly, there are two possible congruence relations on the set $S_2 := \{2, 3, 6\}$ (depicted in Figure 4.1b), and one congruence on the set $S_3 := \{1, 3, 6\}$ (Figure 4.1c). \dashv

402

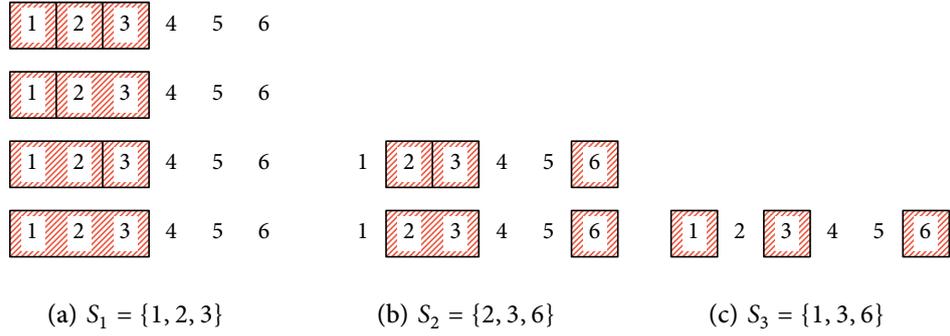


Figure 4.1: Examples for congruence relations

The quotient of a set S by a congruence relation forms a partition of S in which every class is convex; we call such partitions *convex partitions*. Congruences on the same chain can be compared with respect to the *coarseness* of their quotients: given a set S and two partitions Π_1, Π_2 of S , we say that Π_1 is *coarser* than Π_2 (and that Π_2 is *finer* than Π_1), if for every class $C_2 \in \Pi_2$, there is a class $C_1 \in \Pi_1$ such that $C_2 \subseteq C_1$. The set of convex partitions of a given set together with the ‘coarser-than’ relation forms a complete lattice. As a consequence, there is a *coarsest congruence* on a given set.

403 **Example (continued)** The lowermost congruence relation in Figure 4.1a is the coarsest congruence on the set S_1 in \mathfrak{C}_6 (all other congruence relations on S_1 have more equivalence classes), the topmost relation is the finest congruence on S_1 . ←

The coarsest congruence relation on a set can also be characterized directly:

404 **Lemma** Let $\mathfrak{C} = (A; \preceq)$ be a chain, and let $S \subseteq A$ be a set. Define a binary relation on S by putting $a \equiv_S b$ if and only if $\forall c \in [a, b]. c \in S$. Then \equiv_S is the coarsest congruence relation on S . ←

The cardinality of the quotient of a set S modulo the coarsest congruence relation on S provides us with a way to measure the ‘non-convexity’ of S : the more convex sets we need to cover all the elements of S , the less convex it is.

405 **Definition** Let $\mathfrak{C} = (A; \preceq)$ be a chain, and let $S \subseteq A$ be a set. A *block* of S with respect to \mathfrak{C} is an element of the quotient S/\equiv_S . The *block-degree* of S with respect to \mathfrak{C} is the cardinality of the set S/\equiv_S , that is, the number of different blocks of S . ←

406 **Example (continued)** In the pictures in Figure 4.1, the blocks of a set are visualized as contiguous shaded regions, and the block-degree corresponds to the number of these regions. The block-degree of the set S_1 in \mathfrak{C}_6 is 1: the coarsest congruence relation on S_1 has only one block. More generally, every convex set has block-degree 1. The block-degree of the set S_2 is 2: there is no way to cover S_2 with less than 2 blocks. Finally, the block-degree of the set S_3 is 3. ←

Rather than counting the number of blocks of a set, we can also count the number of discontinuities or *gaps* between the blocks. Formally, these concepts can be defined on the complement of a set relative to its *convex hull*.

Definition Let $\mathcal{C} = (A; \preceq)$ be a chain, and let $S \subseteq A$ be a set. The *convex hull* of S , $\mathbf{H}(S)$, is the smallest convex superset of S . The elements of the set $\bar{S} := \mathbf{H}(S) - S$ are called *holes* in S . 407

Applying the definition of blocks and block-degree to sets of holes, we say that a *gap* in S is a class in the quotient $\bar{S}/\equiv_{\bar{S}}$, and the *gap-degree* of S is the cardinality of the quotient $\bar{S}/\equiv_{\bar{S}}$. The gap-degree of a set is obtained as its block-degree, minus 1.

Example (continued) We return to Figure 4.1, where gaps are visualized as contiguous non-shaded regions between blocks. The convex hull of the set S_1 is the set S_1 itself; thus, the set \bar{S}_1 is empty, and there are no gaps in S . The convex hull of the set S_2 is $\mathbf{H}(S_2) = \{2, 3, 4, 5, 6\}$, and the set $\{4, 5\}$ is the set of holes in S_2 . This set also forms a gap in S_2 , so the gap-degree of S_2 is 1. Finally, for the set S_3 we have $\mathbf{H}(S_3) = [6]$ and $\bar{S}_3 = \{2, 4, 5\}$; the gap-degree of S_3 is 2. 408

In the following, we usually talk about blocks and block-degree, but all our results could also be expressed in terms of gaps and the gap-degree measure.

4.1.2 | A hierarchy of non-projective dependency structures

We now apply the block-degree measure to dependency structures. With the definition of projectivity in mind, the interesting congruences on dependency structures are the coarsest congruences on their yields: two nodes v_1, v_2 belong to the same block of a yield $\lfloor u \rfloor$, if all nodes between v_1 and v_2 belong to $\lfloor u \rfloor$ as well. The maximal number of blocks per yield is a measure for the non-projectivity of a dependency structure.

Definition Let D be a dependency structure, and let u be a node of D . The set of *blocks* of u is the set $\lfloor u \rfloor / \equiv_{\lfloor u \rfloor}$, where the congruence relation $\equiv_{\lfloor u \rfloor}$ is defined relative to the precedence relation underlying D . 409

Definition Let D be a dependency structure. The *block-degree* of a node u of D is the number of blocks of u . The block-degree of D is the maximum among the block-degrees of its nodes. 410

Example Figure 4.2 shows two examples of non-projective dependency structures. For both structures, consider the yield of the node 2. In structure D_1 , the yield $\lfloor 2 \rfloor$ falls into two blocks, $\{2, 3\}$ and $\{6\}$. Since this is also the maximal number of blocks per yield, the block-degree of D_1 is 2. In structure D_2 , the yield $\lfloor 2 \rfloor$ consists of three blocks, $\{1\}$, $\{3\}$, and $\{6\}$; the block-degree of D_2 is 3. 411

4 DEPENDENCY STRUCTURES OF BOUNDED DEGREE

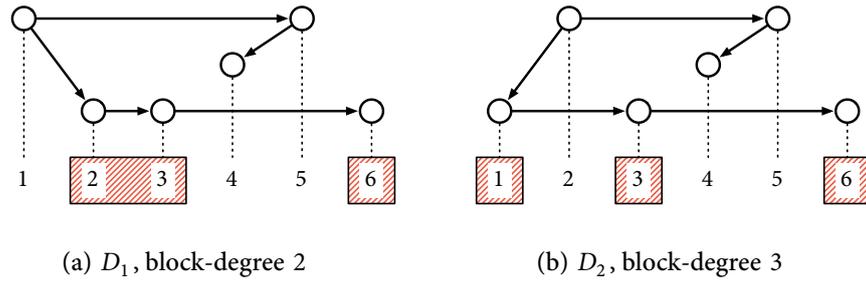


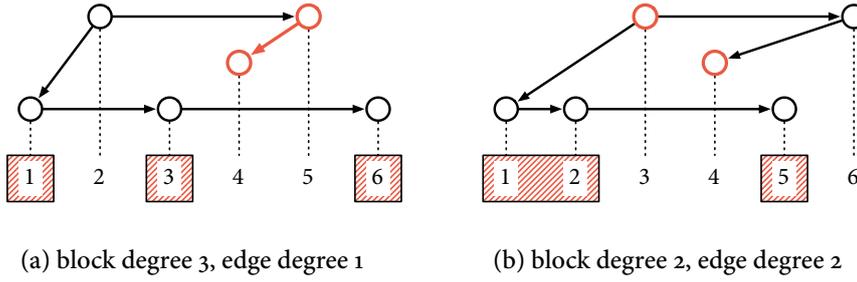
Figure 4.2: Two non-projective dependency structures

Let us say that a dependency structure is *block k* , if its block-degree is at most k . We write \mathcal{D}_k for the class of all dependency structures that are block k . It is immediate from this definition that the class \mathcal{D}_k is a proper subclass of the class \mathcal{D}_{k+1} , for all $k \in \mathbb{N}$. It is also immediate that a dependency structure is projective if and only if it belongs to the class \mathcal{D}_1 . Thus, the block-degree measure induces an infinite hierarchy of ever more non-projective dependency structures, with the class of projective structures at the lowest level of this hierarchy. This is interesting because it allows us to scale the complexity of our formal models with the complexity of the data: the transition from projectivity to full non-projectivity becomes gradual. A crucial question is, of course, whether block-degree is a useful measure in practice. To answer this question, we evaluate the practical relevance of the block-degree measure in Section 4.4.

4.1.3 | Related work

✧ The gap-degree measure (and hence, the block-degree measure) is intimately related to the notion of *node-gaps complexity*, due to Holan et al. (1998). Node-gaps complexity was originally introduced as a complexity measure for derivations in a dependency grammar formalism. Later, it was also applied to the empirically more transparent *results* of these derivations, objects essentially the same as our dependency structures. In this latter application, node-gaps complexity and gap-degree are identical. Note however that some authors (Hajičová et al., 2004; Zeman, 2004) use the term ‘gap’ to refer to a node—rather than a set of nodes—between two blocks of a yield. This is what we have called a ‘hole’ in Definition 407. Havelka (2007a) defines the ‘gap’ of a dependency edge $u \rightarrow v$ as the set of holes in the set $[u] \cap (u, v)$.

✧ We can view the block-degree of a set as a descriptive complexity measure, similar to Kolmogorov complexity in algorithmic information theory: Once a chain $(A; \leq)$ is given, a convex subset of A can be represented by a pair of elements from A , namely the minimal and the maximal element of the set. In this way, even very large sets can be represented with little information. However, the higher the block-degree of a set, the more elements of A we need to represent it, and the less benefit an interval representation has over an explicit representation.



(a) block degree 3, edge degree 1

(b) block degree 2, edge degree 2

Figure 4.3: Block degree and edge degree are incomparable

§ The block-degree measure quantifies the non-projectivity of a dependency structure by counting the number of contiguous blocks in the yields of the structure. A similar measure, based on edges, was introduced by Nivre (2006a). For an edge $e = u \rightarrow v$ in a dependency structure D , let us write \mathfrak{F}_e for the forest that results from restricting the governance relation in D to the nodes in the open interval (u, v) . The *degree* of the edge e (in the sense of Nivre) is the number of those components in \mathfrak{F}_e that are not governed by u in D ; the degree of D is the maximum among the degrees of its edges. This degree measure is incomparable to our block-degree measure. To see this, consider the two dependency structures depicted in Figure 4.3. The left structure (Figure 4.3a) has block-degree 3 and edge-degree 1, as the open interval $(3, 6)$ that corresponds to the edge $3 \rightarrow 6$ contains one component not governed by 3, and this is the maximal number of components per edge. On the other hand, the right structure (Figure 4.3b) has block-degree 2 and edge-degree 2, as the edge interval $(2, 5)$ contains two distinct components not governed by the node 2.

4.2 | Algebraic framework

In this section, we generalize the algebraic framework developed in Section 3.2 to dependency structures with restricted block-degree.

4.2.1 | Traversal of block-ordered trees

One of the essential properties of our procedure for the traversal of treelet-ordered trees is that for each node u of the input tree T , the call `TREELET-ORDER-COLLECT(u)` returns a linearization of the nodes in the subtree of T that is rooted at u . This property ensures that we can interpret the result of a call as a dependency structure, but at the same time constrains this structure to be projective. We now develop a procedure `BLOCK-ORDER-COLLECT` that returns a linearization not of a complete yield, but only of some given block of that yield. To do so, we allow the procedure to be called on a node more than once: the i th call on u produces a linearization of the i th block of u , where blocks are assumed to be numbered in the order of their precedence.

```

BLOCK-ORDER-COLLECT( $u$ )
1   $L \leftarrow \text{NIL}; \text{calls}[u] \leftarrow \text{calls}[u] + 1$ 
2  foreach  $v$  in  $\text{order}[u][\text{calls}[u]]$ 
3      do if  $v = u$ 
4          then  $L \leftarrow L \cdot [u]$ 
5          else  $L \leftarrow L \cdot \text{BLOCK-ORDER-COLLECT}(v)$ 
6  return  $L$ 
    
```

Table 4.1: Traversal of a block-ordered tree

Figure 4.1 shows pseudo-code for an implementation of BLOCK-ORDER-COLLECT. The implementation assumes the existence of a global array *calls* that records for each node *u* the number of times that the procedure has been called on *u*. It further assumes that each node *u* is annotated with lists *order*[*u*][*i*] of nodes in the treelet rooted at *u*.

412 **Example** Figure 4.4 shows an order-annotated tree and the dependency structure induced by its traversal according to the procedure in Table 4.1 when called on the root node of that tree. We assume that the array *calls* is initialized with all zeros. The tuples at the ends of the dotted lines represent the annotated orders. Specifically, the list *order*[*u*][*i*] can be found as the *i*th component of the tuple annotated at the node *u*. To give an example, the order annotations for the node 2 are *order*[2][1] = 23 and *order*[2][2] = 3.

Our first aim in this section is to show that suitably annotated trees together with the procedure BLOCK-ORDER-COLLECT are expressive enough to fully characterize the class of dependency structures with finite block-degree—just as treelet-ordered trees and our procedure for traversing them are expressive enough to fully characterize the class of projective structures. More specifically, we want to show that trees in which no node is annotated with more than *k* lists are expressive enough to characterize the class of structures with block-degree at most *k*.

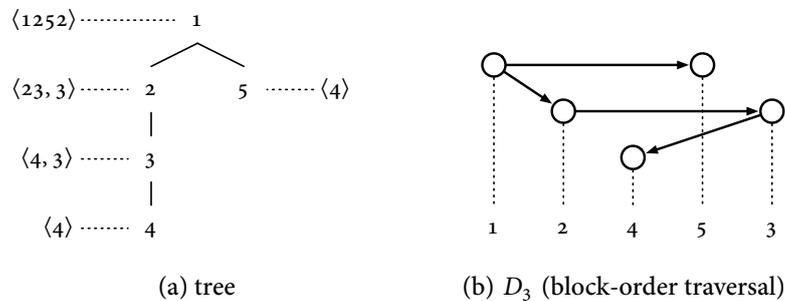


Figure 4.4: A block-ordered tree and its corresponding dependency structure

For our proofs to go through, we need to be more specific about the exact form of the order annotations in the inputs to `BLOCK-ORDER-COLLECT`. Without further constraints, the procedure may fail to induce dependency structures:

- * Assume that two distinct calls `BLOCK-ORDER-COLLECT(u)` return lists that contain the node u , or that a single call returns a list that contains the node u more than once. In both of these cases, the result of the traversal does not qualify as a linearization of the input tree. A similar situation arises if none of the calls `BLOCK-ORDER-COLLECT(u)` returns a list that contains the node u .
- * Assume that `BLOCK-ORDER-COLLECT(u)` is called more often than the number of lists annotated at u . In this case, the results of some of the calls are undefined. Similarly, if `BLOCK-ORDER-COLLECT(u)` is called less often than there are lists annotated at u , then the linearization may be incomplete.

To prevent these problems, we require the inputs to `BLOCK-ORDER-COLLECT` to be well-typed, in the following sense: For each node u , let $k(u)$ be the number of lists annotated at u . We require that, in union, these lists contain exactly one occurrence of the node u , and exactly $k(v)$ occurrences of the node v , for all children v of u . These restrictions ensure that the mapping from trees to dependency structures is well-defined. It is not necessarily injective:

- * Assume that some list $order[u][i]$ is empty. Then we can modify the order annotations without altering the induced dependency structure as follows: delete the list $order[u][i]$, and reindex the remaining annotations at u accordingly.
- * Assume that some list $order[u][i]$ contains two adjacent occurrences of some child v of u . Then we can modify the order annotations without altering the induced dependency structure as follows: delete the second occurrence, append the corresponding order annotation to the annotation corresponding to the first occurrence, and reindex the remaining order annotations at v accordingly.

To prevent these ambiguities, we require that no list $order[u][i]$ is empty, and that no list $order[u][i]$ contains two or more adjacent occurrences of the same node v . We call trees that satisfy all of these requirements *block-ordered trees*.

4.2.2 | Segmented dependency structures

There is one more thing that we need to take care of. Consider a block-ordered tree in which the root node is annotated with more than one list of nodes. When we call `BLOCK-ORDER-COLLECT` on the root node of this tree, all but the first of these lists are ignored, and hence, the linearization of the tree is incomplete, and fails to induce a dependency structure. One way to remedy this problem is to require the root nodes of block-ordered trees to be annotated with exactly one list;

but this would break inductive arguments like the one that we used in the proof of Lemma 308. We therefore opt for another solution, motivated by the following observation: Let T be a block-ordered tree with root node r , and let u be a non-root node of T . For notational convenience, put $k := k(u)$. The well-typedness conditions ensure that the calls to `BLOCK-ORDER-COLLECT` on u can be understood as a tuple $\langle \vec{v}_i \mid i \in [k] \rangle$ of lists of nodes, where \vec{v}_i is the result of the i th call to `BLOCK-ORDER-COLLECT`(u), for $i \in [k]$, and $\vec{v}_1 \cdots \vec{v}_k$ forms a linearization of the subtree rooted at u . Only for the root node r , the procedure is called only once, independently of the number of annotated lists. In order to do away with this asymmetry, we stipulate that the call of `BLOCK-ORDER-COLLECT` on r should return the $k(r)$ -tuple $\langle \text{BLOCK-ORDER-COLLECT}(r)_i \mid i \in [k(r)] \rangle$, where `BLOCK-ORDER-COLLECT`(r) $_i$ stands for the i th call to the node r . Of course, in order for this to make sense, we need to say what such an output should mean in the context of dependency structures. This gives rise to the notion of *segmented dependency structures*, which essentially are dependency structures where even the root nodes can have block-degrees greater than one.

413 **Definition** Let $D = (V; \preceq, \leq)$ be a dependency structure, and let \equiv be a congruence relation on D . The *segmentation* of D by \equiv is the structure $D' := (V; \preceq, \leq, R)$, where R is a new ternary relation on V defined as follows:

$$(u, v_1, v_2) \in R \iff v_1 \equiv v_2 \wedge \forall w \in [v_1, v_2]. w \in [u].$$

The elements of the set V/\equiv are called the *segments* of D' . ←

We write $v_1 \equiv_u v_2$ instead of $(u, v_1, v_2) \in R$.

414 **Example** Figure 4.5 shows how we visualize segmented dependency structures: we use boxes to group nodes that belong to the same segment; all other congruences \equiv_u are uniquely determined by this choice. As an example, $4 \equiv_1 3$ holds in D_4 because both 4 and 3 lie in the same segment of D_4 , and all nodes between them are governed by the node 1. At the same time, $4 \not\equiv_2 3$ holds in D_4 : while both 4 and 3 belong to the same segment of D_4 , the node 5, which is situated between 4 and 3, is not governed by 2. The non-congruence $4 \not\equiv_2 3$ also holds in the substructure $D_4/2$, this time because 4 and 3 do not even lie in the same segment. ←

For each node u , the relation \equiv_u is the coarsest congruence on $[u]$ that is finer than \equiv . Based on this observation, we adapt our definition of blocks (Definition 409):

415 **Definition** Let D be a segmented dependency structure, and let u be a node of D . The set of *blocks* of u is the set $[u]/\equiv_u$. ←

In segmented dependency structures with just one segment, this definition coincides with the old one. In other structures, the new definition ensures that elements from different segments belong to different blocks of all nodes of the structure.

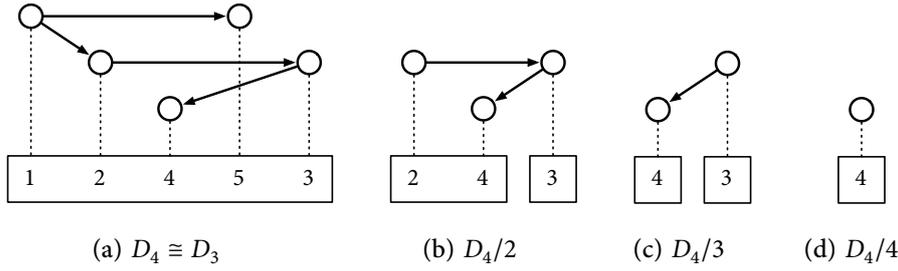


Figure 4.5: Segmented dependency structures

We call the number of segments of a segmented dependency structure the *sort* of that structure, and write \mathcal{D}_k^{\equiv} for the class of all segmented dependency structures of sort k . By our definition of block-degree, the block-degree of a segmented dependency structure is at least as high as its sort. The class \mathcal{D}_1^{\equiv} is essentially the same as the class \mathcal{D} of all dependency structures, and it will be convenient not to distinguish them too carefully. We now connect segmented dependency structures to our modified tree traversal.

Definition Let T be a tree, and let $k \in \mathbb{N}$. A *linearization* of T with k components is a k -tuple $L = \langle \vec{u}_i \mid i \in [k] \rangle$ such that $\vec{u} := \vec{u}_1 \cdots \vec{u}_k$ is a list of the nodes of T in which each node occurs exactly once. The segmented dependency structure *induced* by a linearization L of T is the structure in which the governance relation is isomorphic to T , the precedence relation is isomorphic to \vec{u} , and the segments are isomorphic to the tuple components of L . 416

We can now show the correspondents of the Lemmata 308, 310, and 312 for projective structures. To do so, we regard the functions `lin` and `dep` as *sorted* functions: given a block-ordered tree T in which the root node is annotated with k lists, the function `lin` maps T to the linearization of T with k components that is computed by the traversal of the input tree according to `BLOCK-ORDER-COLLECT`, and the function `dep` maps T to the segmented dependency structure of sort k that is induced by `lin`(T). Apart from this change, the proofs carry over without larger modifications. We therefore get the following Lemma:

Lemma For every segmented dependency structure D , there exists exactly one block-ordered tree T such that `dep`(T) = D . Furthermore, if T is a block-ordered tree in which each node is annotated with at most k lists, for some $k \in \mathbb{N}$, then `dep`(T) is a segmented dependency structure with block-degree at most k . 417

4.2.3 | Order annotations

Next we take the step from the algorithmic to the algebraic and encode block-ordered trees as terms over an extended set of order annotations. This encoding is a relatively

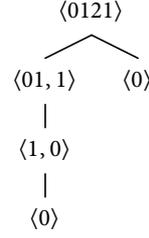


Figure 4.6: The term for the block-ordered tree from Figure 4.4

straightforward generalization of the procedure that we presented in Section 3.2.3. The major novelty comes from the requirement that we need to ensure that all decodings of terms satisfy the well-typedness conditions. To do so, we now understand the set Ω of order annotations as a *sorted* set; as sorts, we use the natural numbers.

Let T be a block-ordered tree with root node u , and put $k := k(u)$. We need the following auxiliary sequences: first, the vector $\vec{v} = v_1 \cdots v_m$ obtained from the concatenation $order[u][1] \cdots order[u][k]$ of the list-based order annotations by removing the node u and all but the first occurrence of each other node; second, for each $j \in [k]$, the string \vec{i}_j obtained from the list $order[u][j]$ by replacing every child of u by its position in \vec{v} , and u itself by the symbol 0. For each $j \in [m]$, put $k_j := k(v_j)$. The term $t(T)$ corresponding to T is then defined recursively as

$$t(T) := \langle \vec{i}_1, \dots, \vec{i}_k \rangle (t(T/v_1), \dots, t(T/v_m)).$$

In this definition, the string $\langle \vec{i}_1, \dots, \vec{i}_k \rangle$ is understood as a term constructor of type $k_1 \times \cdots \times k_m \rightarrow k$. From now on, we use Ω to denote the set of all such constructors. We define the *degree* of a symbol $\omega \in \Omega$, $\deg(\omega)$, as the maximum over its input and output sorts, and put $\Omega(k) := \{ \omega \in \Omega \mid \deg(\omega) \leq k \}$. Note that, by this definition, the set $\Omega(k)$ is exactly what we need in order to encode the set of all block-ordered trees with up to k lists per node, and therefore, the set of all segmented dependency structures with block-degree at most k . Specifically, the set $\Omega(1)$ is essentially identical to our previous definition of Ω for projective dependency structures.

418

Example For the dependency structure D_3 shown in Figure 4.4, the generic procedure to extract the order annotations from a dependency structure described in the proof of Lemma 310 (page 27) yields the following list-based order annotations:

$$\begin{array}{lll}
 \text{for node 1:} & L = 12453, & L' = 12252, & order[1] = 1252; \\
 \text{for node 2:} & L = 12453, & L' = \square 23 \square 3, & order[2] = 23 \square 3; \\
 \text{for node 3:} & L = 12453, & L' = \square \square 4 \square 3, & order[3] = 4 \square 3; \\
 \text{for node 4:} & L = 12453, & L' = \square \square 4 \square \square, & order[4] = 4; \\
 \text{for node 5:} & L = 12453, & L' = \square \square \square 5 \square, & order[5] = 5.
 \end{array}$$

If we read the values of the strings $order[u]$ as tuples, where the symbol \square separates tuple components, we obtain the annotations shown in Figure 4.4. From these, by the construction above, we construct the following order annotations:

$$\langle 0121 \rangle : 2 \times 1 \rightarrow 1, \quad \langle 01, 1 \rangle : 2 \rightarrow 2, \quad \langle 1, 0 \rangle : 1 \rightarrow 1, \quad \langle 0 \rangle : 1.$$

Figure 4.6 shows a term built over these constructors; this term encodes the block-ordered tree from Figure 4.4. \dashv

4.2.4 | Dependency structure algebras

We have now reached a situation very similar to the situation at the beginning of Section 3.2.4: we have identified a sorted set Ω and a sorted bijection $dep: T_\Omega \rightarrow \mathcal{D}^\equiv$ between terms over Ω and segmented dependency structures. This means that we can give the set \mathcal{D}^\equiv an algebraic structure.

Definition Let $\omega: k_1 \times \dots \times k_m \rightarrow k$ be an order annotation. The *composition operation* corresponding to ω is the function $f_\omega: \mathcal{D}_{k_1}^\equiv \times \dots \times \mathcal{D}_{k_m}^\equiv \rightarrow \mathcal{D}_k^\equiv$ defined as 419

$$f_\omega(D_1, \dots, D_m) := dep(\omega(\text{term}(D_1), \dots, \text{term}(D_m))). \quad \dashv$$

Each composition operation f_ω simulates a single step of the block-order traversal: given a sequence of argument structures, it returns the segmented dependency structure that can be decomposed into the given argument structures in the way that is specified by the order annotation ω .

Definition Let $\Sigma \subseteq \Omega$ be a finite set of order annotations. The *dependency algebra* over Σ is the Σ -algebra that has the \mathcal{S}_Σ -indexed set $\langle dep(\Sigma_i) \mid i \in [k] \rangle$ as its carrier, and interprets each $\omega \in \Sigma$ by the composition operation corresponding to ω . \dashv 420

Theorem Let $\Sigma \subseteq \Omega$ be a finite set of order annotations. Then the dependency structure algebra over Σ is isomorphic to the (many-sorted) term algebra over Σ . \dashv 421

4.3 | Algorithmic problems

In this section, paralleling Section 3.3, we address three of the algorithmic problems related to our algebraic view on non-projective dependency structures: decoding, encoding, and computing the block-degree for a given dependency structure. The decoding problem is essentially solved by the procedure `BLOCK-ORDER-COLLECT` that we gave in Table 4.1. The main contribution of this section is an efficient algorithm to encode a non-projective dependency structure into a term. The algorithm that computes the block-degree is a straightforward extension of the encoding algorithm.

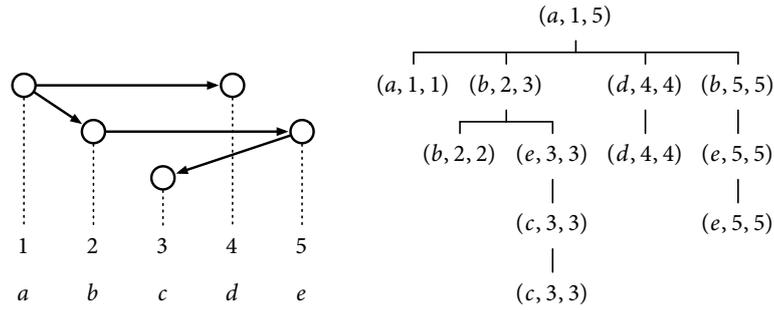


Figure 4.7: A non-projective dependency structure and the corresponding span tree

4.3.1 | Encoding

On page 27, we described a generic procedure to extract the order annotations from a dependency structure. A naïve implementation of this procedure takes time cubic in the number of nodes of the input structure. The algorithm that we present in this section may perform significantly better; it runs in time linear in the number of blocks in the input structure, of which there are *at most* quadratically many, but often less.

The crucial component of our encoding algorithm is a procedure that transforms the input structure into a certain tree representation called *span tree*, from which all order annotations can be easily read off. The span tree T for a dependency structure D is a labelled, ordered tree in which each node has one of two types: it can be a *block node* or an *anchor*. The block nodes of T stand in one-to-one correspondence with the blocks of D ; the anchors stand in one-to-one correspondence with the singletons. For a node u of T , we write $S(u)$ for the set of nodes in D that corresponds to u . The dominance relation of T represents an inclusion relation in D : u strictly dominates v if and only if $S(u) \supseteq S(v)$ and either u is a block node and v is an anchor, or both u and v are block nodes. By this relation, all block nodes are inner nodes, and all anchors are leaf nodes in T . The precedence relation of T represents a precedence relation in D : u strictly precedes v if and only if all nodes in $S(u)$ precede all nodes in $S(v)$.

422 **Example** Figure 4.7 shows the dependency structure from Figure 4.4 and the span tree for this structure. Consider the root node of the structure, the node a . The yield of a consists of a single block, which contains all the nodes of the structure, positions 1 to 5 in the precedence order. This information is represented in the span tree in that the root node of this tree is labelled with the triple $(a, 1, 5)$. The block of a decomposes into four components; read in their order of precedence, these are: the singleton $\{a\}$, the first block of the node b , the block of the node d , and the second block of b . Note that, since b contributes two blocks to the block of a , we find two nodes of the form (b, i, j) as children of $(a, 1, 5)$ in the span tree. Note furthermore that the precedence order on the components of the block of a is reflected by the sibling order on their corresponding nodes in the span tree. –1

```

STEP( $u$ )
1   $T \leftarrow \bigcup \{ trees[v] \mid v \in children[u] \} \cup \{(u, pos[u], pos[u], NIL)\}$ 
2  foreach ( $w, i, j, S$ ) in  $T$ 
3      do  $right[i] \leftarrow j$ ;  $sub[i] \leftarrow$  if  $i = pos[u]$  then  $[(u, pos[u], pos[u])]$  else  $S$ 
4  foreach ( $w, i, j, S$ ) in  $T$ 
5      do  $k \leftarrow j + 1$ 
6          while  $right[k] \neq \perp$ 
7              do  $right[i] \leftarrow right[k]$ ;  $sub[i] \leftarrow sub[i] \cdot sub[k]$ 
8                   $right[k] \leftarrow \perp$ ;  $sub[k] \leftarrow \perp$ ;  $k \leftarrow right[i] + 1$ 
9  foreach ( $w, i, j, S$ ) in  $T$ 
10     do if  $right[i] \neq \perp$ 
11         then  $trees[u] \leftarrow trees[u] \cup \{(u, i, right[i], sub[i])\}$ 
12          $right[i] \leftarrow \perp$ ;  $sub[i] \leftarrow \perp$ 

```

Table 4.2: Constructing the span tree for a dependency structure

We now describe the general structure of an algorithm that transforms a dependency structure into its span tree. We assume that the input structure is given to us as a collection of nodes, where each node u is equipped with a set $children[u]$ of its children and an integer $pos[u]$ that represents the position of u with respect to the precedence relation. Our algorithm can be separated into two phases:

- * In the first phase, we allocate two global data structures: an array $right$ that will map the left endpoints of blocks to their right endpoints, and an array sub that will record the component structures of blocks. Each element of these arrays is initialized to the void value, which we write as \perp .
- * In the second phase of the algorithm, we perform a post-order traversal of the input structure. For each node u , we compute a set $trees[u]$ that contains the span trees for the blocks of the yield of u . An element of $trees[u]$ is a four-tuple of the form (w, i, j, S) , where w is a node in the treelet rooted at u , the position i is the left endpoint of the span represented by the tree, the position j is the right endpoint, and the list S contains the trees for the components in their left-to-right order. If u is a leaf node, then $trees[u]$ consists of the single tree $(u, pos[u], pos[u], [(u, pos[u], pos[u], NIL)])$. If u is an inner node, then the set $trees[u]$ is obtained by exhaustive *merging* of this trivial tree and the span trees that were constructed for the children of u . Two trees can be merged if they correspond to adjacent spans of positions.

In the following, we concentrate on the procedure $STEP(u)$ that processes a single node u during the post-order traversal; pseudo-code for this procedure is given in

Table 4.2. In line 1, we collect the trees for the subblocks of u . In lines 2–3, we register these trees in the global data structures—in particular, for each tree (w, i, j, S) , we register the position j as the right endpoint of a span that starts at position i . In lines 4–8, we merge the spans of adjacent trees into larger spans: we try each tree as a trigger to merge all right-adjacent spans into a new span, removing all traces of the old spans (line 8). The result of a merger spans the positions from the left endpoint of the trigger to the right endpoint of the span that was right-adjacent to it. In lines 9–12, we construct the set $trees[u]$ from the spans that remain after merging, and remove all traces from the global data structures.

We now look at the asymptotic complexity of the algorithm. The following invariant is essential for the analysis (and for the correctness proof, which we will omit):

Every element $right[i]$ that receives a non-void value during a call to the procedure STEP is void again when this call finishes.

This can be seen as follows. The only places in a call to STEP where non-void values are assigned to $right$ are in lines 3 and 7. The assignment in line 3 is witnessed by a tree from the set T , which is not altered after line 1; the assignment in line 7 merely overwrites a previous assignment. For all trees in T , it is checked in line 10 whether the element $right[i]$ is assigned, and if so, the element is made void. Therefore, every element of $right$ is void when a call to STEP finishes. Since every element is void before the first call to STEP, and is not altered in between calls to STEP, it is void both before and after any call to STEP. A similar argument holds for the array sub .

423

Lemma Let D be a dependency structure with n nodes and g gaps. Then the span tree corresponding to D can be constructed in time $O(n + g)$. –

Proof. Let us assume that we can use the algorithm that we have outlined above to construct the span tree for D . The algorithm breaks down into two phases: the initialization of the global data structures, and the tree traversal with calls to STEP. The first phase takes time $O(n)$; the second phase can be implemented to take time linear in the sum over the times taken by the calls to STEP. We show that each such call can be implemented to take time linear in the number of subblocks of the visited node. The total number of subblocks is asymptotically identical to the number of blocks in the input structure, which is $n + g$: each node has one more block than gaps.

Fix some node u of D , and consider the call $STEP(u)$. Let m be the number of subblocks of u ; this number is obtained as the number of blocks of the children of u , plus 1 for the trivial block containing the singleton u . Note that $m = |T|$. We now check the runtime of the constructs in Table 4.2: The assignment in line 1 can be implemented to run in time $O(m)$; all other elementary operations can be implemented to run in time $O(1)$. Each of the **foreach** loops is executed $O(m)$ times. Finally, we analyze the runtime of the **while** loop. During each iteration of that loop, one of the

elements of *right* is made void. Because of the invariant given above, all elements of *right* are void before the call $\text{STEP}(u)$, and since only m fields of *right* are initialized in line 3, the **while** loop cannot make more than m iterations *during a single call* to STEP . Putting everything together, we see that every call to STEP takes time $O(m)$. \dashv

This finishes our discussion of the transformation of the input dependency structure into a span tree. We now explain how to read off the order annotations from this tree. This can be done in a single pre-order traversal.¹ We first initialize an array $\text{order}[u]$ that maps nodes to their order annotations. For each non-leaf node (u, i, j, S) in the span tree, we construct a list L that contains the first components of the children of that node in their left-to-right order, and add this list to the array $\text{order}[u]$. At the end of the tree traversal, each list $\text{order}[u]$ contains the full order annotation for the node u .

Example (continued) Consider the root node of the span tree depicted in Figure 4.7, the node $(a, 1, 5)$. The children of this node are $(a, 1, 1)$, $(b, 2, 3)$, $(d, 4, 4)$, and $(b, 5, 5)$. The corresponding order annotation is $\text{order}[a] = \langle abdb \rangle$. Similarly, for node b , we get $\text{order}[b] = \langle be, e \rangle$. This yields the order annotations $\langle 0121 \rangle$ and $\langle 01, 1 \rangle$, which can also be found at the corresponding nodes of the term for the encoded dependency structure (Figure 4.6). \dashv

424

In conclusion, we get the following result:

Lemma Let D be a dependency structure with n nodes and g gaps. Then the term $\text{term}(D)$ can be computed in time $O(n + g)$. \dashv

425

Note that, in a dependency structure with block-degree k , the number $n + g$ of blocks is bounded by the number $k \cdot n$. Therefore, a coarser bound on the complexity of our encoding algorithm is $O(k \cdot n)$.

4.3.2 | Computing the block-degree of a dependency structure

It is very easy to extend the encoding algorithm into an algorithm that computes the block-degree of a dependency structure. Since there is a direct correspondence between the block-degree of a node u in a dependency structure D and the degree of the order annotation at u in the term $\text{term}(D)$ that encodes D (see our discussion on page 44), it suffices to compute the encoding and count the relevant numbers. This can be done in time $O(n + g)$.

Lemma Let D be a dependency structure with n nodes and g gaps. Then the block-degree of D can be computed in time $O(n + g)$. \dashv

426

¹ Note that, given that the order in the span tree reflects the precedence order of the blocks in the dependency structure, during a pre-order traversal we visit the blocks from left-to-right.

4 DEPENDENCY STRUCTURES OF BOUNDED DEGREE

block-degree	DDT		PDT 1.0		PDT 2.0	
1 (projective)	3 730	84.95%	56 168	76.85%	52 805	77.02%
2	654	14.89%	16 608	22.72%	15 467	22.56%
3	7	0.16%	307	0.42%	288	0.42%
4	–	–	4	0.01%	1	< 0.01%
5	–	–	1	< 0.01%	1	< 0.01%
TOTAL	4 391	100.00%	73 088	100.00%	68 562	100.00%

Table 4.3: Dependency structures of various block-degrees in three treebanks

It is also possible to parameterize the algorithm that constructs the span tree for the input structure by a constant k such that it terminates as soon as it discovers that the tree to be constructed contains at least one node with a block-degree that exceeds k . In this way, a test whether a given dependency structure has block-degree k can be carried out in time $O(k \cdot n)$.

4.4 | Empirical evaluation

Using the algorithms developed in the previous section, we now evaluate the coverage of different block-degrees on treebank data. Specifically, we check how many and how large a percentage of the structures in the three treebanks that we used in the experiments reported in Section 3.4 have a block-degree of exactly k , for increasing values of k . Table 4.3 shows the results of the evaluation. The numbers and percentages for block-degree 1 reiterate the results for projective dependency structures from Table 3.4; the counts for structures with block-degree greater than 1 partition the figures for non-projective structures in that table.

The general impression that we get from the experiments is that even a small step beyond projectivity suffices to cover virtually all of the data in the three treebanks. Specifically, it is sufficient to go up to block-degree 2: the structures with block-degree greater than 2 account for less than half a percent of the data in any treebank. These findings confirm similar results for other measures of non-projectivity, such as Nivre’s *degree* (Nivre, 2006a) and Havelka’s *level types* (Havelka, 2007a). Together, they clearly indicate that to contrast only projective and non-projective structures may be too coarse a distinction, and that it may be worthwhile to study classes of dependency structures with intermediate degrees of non-projectivity. The class of dependency structures with block-degree at most 2 appears to be a promising starting point.

CHAPTER 5 | Dependency structures without crossings

The block-degree of a dependency structure is a quantitative property—it measures the independence of governance and precedence along an infinite scale of possible values. In this chapter, we study *well-nestedness*, a property related not to the degree, but to the form of non-projectivity in dependency structures.

Structure of the chapter. To motivate the well-nestedness restriction, we first look at another structural constraint, *weak non-projectivity*, and investigate its entanglement with the block-degree measure (Section 5.1). From there we are led to the relational and the algebraic characterization of well-nestedness, and to an efficient test for this property (Section 5.2). At the end of the chapter, we evaluate and compare the empirical adequacy of weak non-projectivity and well-nestedness (Section 5.3).

5.1 | Weakly non-projective dependency structures

Let us go back to our motivation of projectivity in Chapter 3. Recall that we had to refine the slogan that projectivity should ‘disallow dependency analyses with pictures that contain crossing edges’ because some pictures of dependency structures *with* crossing edges can be ‘fixed’ by changing the vertical positions of some nodes—to witness, consider the two pictures of the non-projective structure D_1 that are shown in Figures 5.1a and 5.1b. In this section, we return to our original slogan, and take a closer look at dependency structures that can be redrawn without crossing edges. We call such structures *weakly non-projective*, a term that we borrow from the Russian tradition of dependency grammar (Dikovskiy and Modina, 2000).

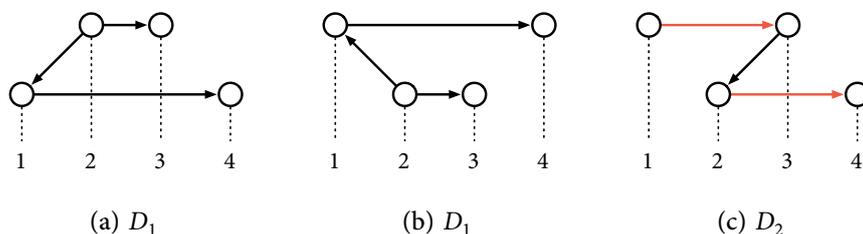


Figure 5.1: Three pictures of non-projective dependency structures. The structure shown in pictures (a) and (b) is weakly non-projective; the structure shown in picture (c) contains overlapping edges.

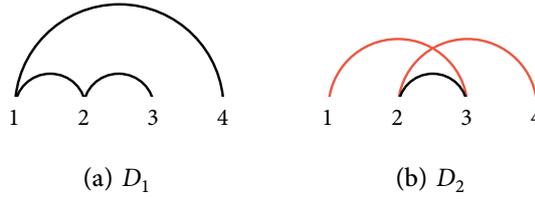


Figure 5.2: Alternative pictures for the dependency structures from Figure 5.1

5.1.1 | Definition of weak non-projectivity

In Chapter 3, we used the term ‘crossing edge’ to refer to a dependency edge that crosses a projection line. There is another natural meaning: that a ‘crossing edge’ is a dependency edge that crosses another edge. We formalize this situation as follows.

501 **Definition** Let $\mathcal{C} = (A; \preceq)$ be a chain. Two \mathcal{C} -intervals B and C overlap, if

$$\min B < \min C < \max B < \max C \quad \text{or} \quad \min C < \min B < \max C < \max B .$$

We write $B \checkmark C$ to assert that B and C overlap.

502 **Definition** Let D be a dependency structure. Two edges $v_1 \rightarrow v_2, w_1 \rightarrow w_2$ in D overlap, if $[v_1, v_2]$ and $[w_1, w_2]$ overlap as intervals. ←

503 **Definition** A dependency structure D is called *weakly non-projective*, if it does not contain overlapping edges. ←

504 **Example** The structure D_2 depicted in Figure 5.1c is not weakly non-projective: the edges $1 \rightarrow 3$ and $2 \rightarrow 4$ overlap. On the other hand, the structure D_1 depicted in Figures 5.1a and 5.1b does not contain overlapping edges; it is weakly non-projective. The difference between the two structures can be seen more clearly when drawing them with undirected edges as in Figure 5.2. ←

A dependency structure that contains overlapping edges cannot be drawn without edges that cross a projection line. Conversely, the only non-projective edges $u \rightarrow v$ that a weakly non-projective dependency structure can contain are such that all ancestors of the node u and all nodes governed by these ancestors, except for the nodes also governed by u , fall into the interval $[u, v]$. Such structures can be redrawn by moving the material in the gap of the edge to vertical positions below the node u .

505 **Example (continued)** The picture in Figure 5.1a contains one edge that crosses a projection line: the edge $1 \rightarrow 4$. When we take the set of all ancestors of 1 (the set $\{2\}$) and all nodes governed by these ancestors ($\{1, 2, 3, 4\}$), minus the nodes also governed by 1 ($\{1, 4\}$), we obtain the set $\{2, 3\}$. By moving these nodes to vertical positions below the node 1, we get to the picture in Figure 5.1b. ←

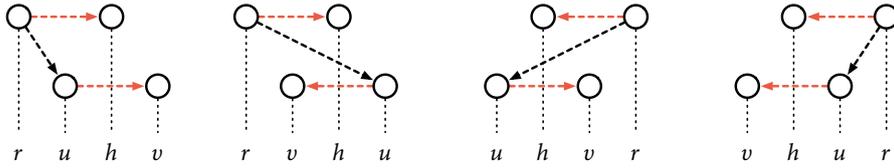


Figure 5.3: The four cases in the proof of Lemma 506

5.1.2 | Relation to the block-degree measure

We now show that, from the perspective of the block-degree measure, weak non-projectivity is a weak extension of projectivity indeed: it does not even take us up one step in the block-degree hierarchy. We first prove an auxiliary lemma.

Lemma Let D be a weakly non-projective dependency structure. Then every gap in D contains the root node of D . 506

Proof. We show the contrapositive of the statement: If at least one gap in D does not contain the root node, then D is not weakly non-projective. Let D be a dependency structure with a gap that does not contain the root node. We can then choose four pairwise distinct nodes r, u, h, v as follows: Choose r to be the root node of D . Choose u to be a node such that G is a gap of u and $r \notin G$; since the root node does not have a gap, it is certain that $r \neq u$. Choose $h \in G$; we then have $h \notin [u]$ and $h \neq r$. Choose v to be a node governed by u such that $h \in (u, v)$; since the root node is governed only by itself, we have $v \neq r$. Based on the relative precedences of the nodes, we now distinguish four cases, shown schematically in Figure 5.3. In all four cases, some edges on the paths from r to h and from u to v overlap. ◻

Note that the converse of Lemma 506 does *not* hold: there are dependency structures that are not weakly non-projective, but in which all gaps contain the root node. As an example, consider the structure D_2 depicted in Figure 5.1c.

Lemma Every weakly non-projective dependency structure has a block-degree of at most 2. Furthermore, there is at least one dependency structure with block-degree 2 that is not weakly non-projective. 507

Proof. To see the inclusion, let D be weakly non-projective. By Lemma 506, we know that every gap in D contains the root node. Since distinct gaps of one and the same node are set-wise disjoint, this implies that no node in D can have more than one gap. Therefore, the structure D has a block-degree of at most 2. To see that the inclusion is proper, consider again structure D_2 in Figure 5.1c: this structure has block-degree 2, but contains overlapping edges. ◻

5 DEPENDENCY STRUCTURES WITHOUT CROSSINGS

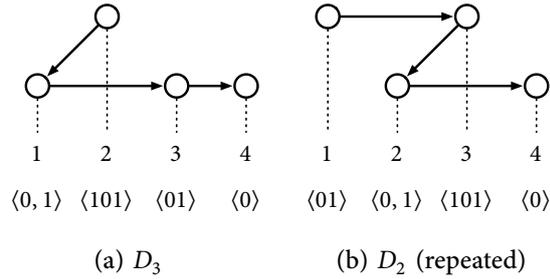


Figure 5.4: Weak non-projectivity is algebraically opaque

Weak non-projectivity appears like a very interesting constraint at first sight, as it expresses the intuition that ‘crossing edges are bad’, but still allows a limited form of non-projectivity. On the other hand, the property stated in Lemma 506 seems rather peculiar. To get a better understanding of the explanatory force of weak non-projectivity, we evaluate its empirical relevance in Section 5.3.

5.1.3 | Algebraic opaqueness

From the perspective of our algebraic setting, there is a fundamental difference between weak non-projectivity and the block-degree restriction. Recall from Section 4.3.2, that in order to check whether a dependency structure D has block-degree at most k , it suffices to check whether the corresponding term $\text{term}(D)$ only contains symbols from the sub-signature $\Omega(k)$ of order annotations with degree at most k . In this sense, the block-degree measure is *transparent*: it is directly related to the set of composition operations used to build a structure. The class of weakly non-projective dependency structures cannot be characterized in the same way. To see this, consider Figure 5.4, which shows a weakly non-projective structure (D_3) and a structure that is not weakly non-projective (D_2). Both of these structures are composed using the same set of algebraic operations.

5.1.4 | Related work

✦ In the Western literature, weak non-projectivity is more widely known as *planarity* (Sleator and Temperley, 1993). Unfortunately, the latter term clashes with the concept of planarity known from graph theory, for at least two reasons: First, while a planar graph is a graph that can be drawn into the plane such that no edges intersect, a ‘planar’ dependency structure is a graph that is drawn into the *half plane* above the words of the sentence. Second, to show that a graph is planar in the graph-theoretic sense, its nodes may be rearranged on the plane in any arbitrary way; in the context of dependency structures, the order of the nodes is fixed. Due to these incompatibilities, it seems wise to avoid the term ‘planarity’, and use a less biased name instead.

§ Projectivity and weak non-projectivity are closely related. Some authors in fact *define* projectivity by requiring the weak non-projectivity of the structure that is obtained when the dependency structure proper is extended by an artificial root node, preceding all other nodes (see e.g. McDonald et al., 2005). In Link Grammar for example, the artificial root node is called *the wall*, ‘an invisible word which the [parser] program inserts at the beginning of every sentence’ (Sleator and Temperley, 1991). As a corollary of Lemma 506, every weakly non-projective dependency structure in which the root node occupies the leftmost position in the precedence order is projective. Another way to enforce the projectivity of weakly non-projective dependency structures is to require that no edge covers the root node (Hudson, 1984; Mel’čuk, 1988).

§ Yli-Jyrä (2003) proposes a generalization of weak non-projectivity, and evaluates its empirical adequacy using data from the Danish Dependency Treebank. He calls a dependency structure *m-planar*, if its governance relation can be partitioned into *m* sets, called *planes*, such that each of the substructures induced by such a plane is weakly non-projective. Since every dependency structure is *m-planar* for some sufficiently large *m* (put each edge onto a separate plane), an interesting question in the context of multiplanarity is about the *minimal* values for *m* that occur in real-world data. To answer this question, one not only needs to show that a dependency structure *can* be decomposed into *m* weakly non-projective structures, but also, that this decomposition is the one with the smallest possible number of planes. Up to now, no tractable algorithm to find the minimal decomposition has been given, so it is not clear how to evaluate the significance of the concept as such. The evaluation presented by Yli-Jyrä (2003) makes use of additional constraints that are sufficient to make the decomposition unique.

§ In combinatorics, weakly non-projective dependency structures are known as *non-crossing rooted trees*. Their number is given by sequence A001764 in Sloane (2007).¹ Using Lemma 506, we see that every projective dependency structure can be decomposed into two halves—one with the root node at the right border, one with the root node at the left—such that each half is a non-crossing rooted tree. One can then obtain the number of projective dependency structures as the convolution of sequence A001764 with itself; this is sequence A006013 in Sloane (2007).

5.2 | Well-nested dependency structures

In this section, we develop the notion of *well-nestedness*. Well-nestedness is similar to weak non-projectivity in that it disallows certain ‘crossings’, but different in that it is transparent at the level of our algebraic signatures. We first introduced well-nestedness in Bodirsky et al. (2005). It was subsequently studied in detail by Möhl (2006).

¹ Sequence A001764 actually gives the number of non-crossing *unrooted* trees. In order to get the right numbers for weakly non-projective structures, one has to read the sequence with an offset of 1.

5 DEPENDENCY STRUCTURES WITHOUT CROSSINGS

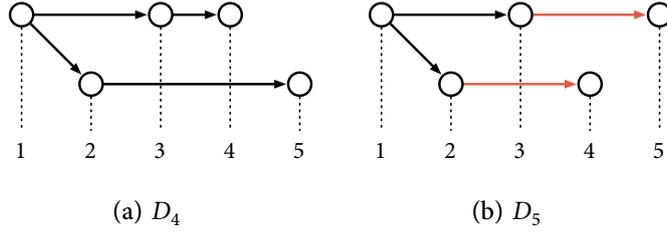


Figure 5.5: Two dependency structures: one well-nested, the other one ill-nested

5.2.1 | Definition of well-nestedness

The definition of well-nestedness specializes the definition of weak non-projectivity in that it bans overlapping edges only if they belong to disjoint subtrees. Overlapping configurations in which one of the edges governs the other are allowed.

508 **Definition** A dependency structure D is called *well-nested*, if the following implication holds for all edges $v_1 \rightarrow v_2, w_1 \rightarrow w_2$ in D :

$$[v_1, v_2] \not\ll [w_1, w_2] \implies v_1 \trianglelefteq w_1 \vee w_1 \trianglelefteq v_1.$$

Dependency structures that are not well-nested are called *ill-nested*. ⊣

We write \mathcal{D}_{wn} for the class of all well-nested dependency structures. From the definition, it is straightforward that every weakly non-projective dependency structure is also well-nested. As the following example shows, the converse does not hold.

509 **Example** Figure 5.5 shows pictures of two non-projective dependency structures. Structure D_4 is not weakly non-projective, as the edges $1 \rightarrow 3$ and $2 \rightarrow 5$ overlap; however, it is well-nested, as $1 \trianglelefteq 2$. Structure D_5 is not even well-nested: the spans $2 \rightarrow 4$ and $3 \rightarrow 5$ overlap, but 2 and 3 belong to disjoint subtrees. ⊣

In contrast to weak non-projectivity, well-nestedness is independent of the block-degree measure: it is not hard to see that for every block-degree $k > 1$, there are both well-nested and ill-nested dependency structures of degree k . Projective structures are both weakly non-projective and well-nested. In summary, we obtain the following hierarchy of classes of dependency structures:

$$\text{projective} \not\subseteq \text{weakly non-projective} \not\subseteq \text{well-nested} \not\subseteq \text{unrestricted}.$$

5.2.2 | Non-crossing partitions

Our next aim is to show that well-nestedness is algebraically transparent. To do so, we develop an alternative relational characterization of well-nestedness based on the notion of *non-crossing partitions*.

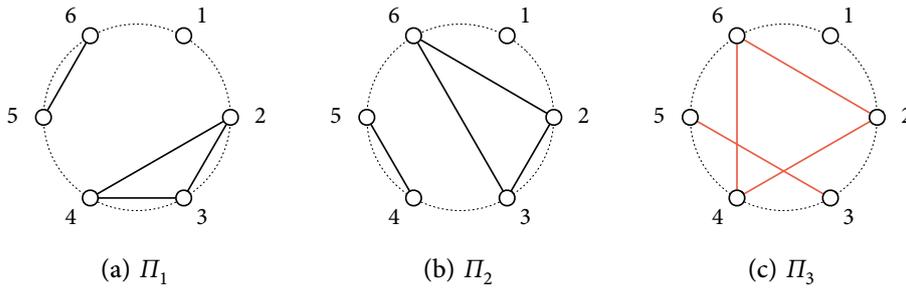


Figure 5.6: Two non-crossing and one crossing partition

Definition Let $\mathfrak{C} = (A; \leq)$ be a chain. A partition Π of A is called *non-crossing*, if whenever there exist four elements $a_1 < b_1 < a_2 < b_2$ in A such that a_1 and a_2 belong to the same class of Π , and b_1 and b_2 belong to the same class of Π , then these two classes coincide. A partition that is not non-crossing is called *crossing*. 510 \dashv

Non-crossing partitions enjoy a number of interesting formal properties. In particular, the number of non-crossing partitions of a chain with n elements is the Catalan number, $C_n = \frac{1}{n+1} \binom{2n}{n}$, and by this property, non-crossing partitions are connected to a large family of mathematical structures—such as binary trees, Catalan paths in the plane, pattern-avoiding permutations, and (most important in the context of this study) well-bracketed strings and children-ordered trees. Consequently, non-crossing partitions appear in a large number of mathematical applications. Simion (2000) provides a comprehensive overview.

Example Consider the following partitions on the chain $([6]; \leq)$: 511

$$\begin{aligned} \Pi_1 &= \{\{1\}, \{2, 3, 4\}, \{5, 6\}\}, & \Pi_2 &= \{\{1\}, \{2, 3, 6\}, \{4, 5\}\}, \\ \Pi_3 &= \{\{1\}, \{2, 4, 6\}, \{3, 5\}\}. \end{aligned}$$

Both Π_1 and Π_2 are non-crossing. Partition Π_3 is crossing, as witnessed by the sequence $2 < 3 < 4 < 5$: the elements 2 and 4 and the elements 3 and 5 belong to the same class of Π_3 , but these two classes do not coincide. \dashv

Example Let $n \in \mathbb{N}$. A neat way to visualize a partition Π on the canonical chain $([n]; \leq)$ goes as follows: Consider a regular n -gon inscribed into a circle, and assume that the points where the n -gon touches the circle are numbered clockwise from 1 to n . Now, for every class of Π of size k , connect the corresponding points on the circle with straight lines to form a convex k -gon. The partition Π is non-crossing if and only if no of these k -gons intersect. Figure 5.6 shows such pictures for the partitions from example 511. 512 \dashv

We now use non-crossing partitions to characterize well-nestedness.

513 **Lemma** A dependency structure D is well-nested if and only if for every node u of D , the set $\mathcal{C}(u)$ of constituents of u (see Definition 309) is non-crossing with respect to the chain $\mathcal{C} := ([u]; \leq|_{[u]})$. \dashv

Proof. We prove the contrapositive of the claim: a dependency structure D is ill-nested if and only if there exists a node $u \in \text{dom}(D)$ such that the partition $\mathcal{C}(u)$ forms a crossing partition with respect to \mathcal{C} . The proof falls into two parts.

\Rightarrow Assume that D is ill-nested. In this case there exist overlapping edges $v_1 \rightarrow v_2$ and $w_1 \rightarrow w_2$ such that $v_1 \perp w_1$. Let u be the greatest (farthest from the root node) common ancestor of v_1 and w_1 . The node sets $\{v_1, v_2\}$ and $\{w_1, w_2\}$ belong to the yields of distinct children of u , and hence, to distinct constituents of u . Furthermore, the intervals $[v_1, v_2]$ and $[w_1, w_2]$ overlap with respect to \mathcal{C} . Thus we deduce that $\mathcal{C}(u)$ is crossing with respect to \mathcal{C} .

\Leftarrow Let $u \in \text{dom}(D)$ be a node, and assume that the partition $\mathcal{C}(u)$ is crossing with respect to \mathcal{C} . In this case, there exist two distinct constituents C_v and C_w in $\mathcal{C}(u)$ and elements $v_1, v_2 \in C_v$, $w_1, w_2 \in C_w$ such that $[v_1, v_2] \not\bowtie [w_1, w_2]$. By the definition of $\not\bowtie$, both C_v and C_w have a cardinality of at least 2; therefore, they correspond to the yields of distinct and hence disjoint children of the node u , say $C_v = [v]$ and $C_w = [w]$. For every arrangement of the nodes v_1, v_2 and w_1, w_2 , we can choose edges $v'_1 \rightarrow v'_2$ in $[v]$ and $w'_1 \rightarrow w'_2$ in $[w]$ such that these edges overlap. Furthermore, by construction we have $v_1 \perp w_2$, and hence, $v'_1 \perp w'_1$. Thus we deduce that D is ill-nested. \dashv

514 **Example (continued)** Consider the constituents of the root nodes in Figure 5.5:

$$D_4 : \mathcal{C}(1) = \{\{1\}, \{2, 5\}, \{3, 4\}\} \quad D_5 : \mathcal{C}(1) = \{\{1\}, \{2, 4\}, \{3, 5\}\}$$

The first of these partitions is crossing, the second non-crossing. \dashv

Lemma 513 shows that by restricting ourselves to composition operations that arrange their arguments into non-crossing partitions, we produce exactly the well-nested dependency structures. In this sense, well-nestedness is a transparent property.

5.2.3 | Algebraic characterization

We now give an explicit characterization of the composition operations that generate the well-nested dependency structures. More specifically, we state a syntactic restriction on order annotations that identifies a sub-signature Ω_{wn} of Ω such that the dependency structures that are obtained as the values of the terms over Ω_{wn} are exactly the well-nested dependency structures. The syntactic restriction ensures that all constituents form non-crossing partitions.

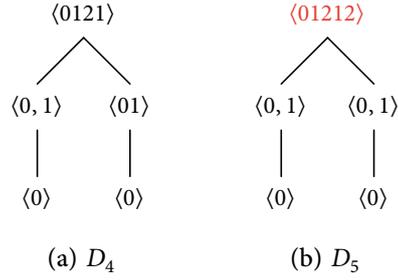


Figure 5.7: Terms for the dependency structures in Figure 5.5

Let $\vec{x} \in A^*$ be a string. We say that \vec{x} contains the string \vec{y} as a *scattered substring*, if, for some $k \in \mathbb{N}$, \vec{x} can be written as

$$\vec{x} = \vec{z}_0 \cdot \vec{y}_1 \cdot \vec{z}_1 \cdots \vec{z}_{k-1} \cdot \vec{y}_k \cdot \vec{z}_k,$$

where $\vec{z}_0 \cdots \vec{z}_k \in A^*$, and $\vec{y}_1 \cdots \vec{y}_k = \vec{y}$.

Definition An order annotation $\omega \in \Omega$ is called *well-nested*, if it does not contain a string of the form $ijij$ as a scattered substring, for $i \neq j \in \mathbb{N}$. 515

We write Ω_{wn} for the set of all well-nested order annotations.

Example The order annotation $\langle 0121 \rangle$ is well-nested, the annotation $\langle 01212 \rangle$ is not: it contains the string 1212 as a scattered substring. Figure 5.7 shows terms that make use of the two order annotations; these terms evaluate to the dependency structures shown in Figure 5.5. 516

We now present the main result of this section:

Theorem A dependency structure D is well-nested if and only if $\text{term}(D) \in T_{\Omega_{wn}}$. 517

Proof. This is a corollary of Lemma 513. The presence of the scattered substring $ijij$ implies that for some node u , there exist distinct constituents $C_1, C_2 \in \mathcal{C}(u)$ and nodes $v_1, v_2 \in C_1, w_1, w_2 \in C_2$ such that $v_1 < w_1 < v_2 < w_2$; then, $\mathcal{C}(u)$ would be crossing. Conversely, the encoding algorithm translates every constituent set that is crossing into an order annotation that contains the forbidden pattern. 517

5.2.4 | Testing whether a dependency structure is well-nested

Given that the class of well-nested dependency structures forms a subclass of the class of all dependency structures, the algorithmic problems of encoding and decoding can be solved using the algorithms that we have presented in Section 4.3. Here we address the problem of testing whether a given dependency structure D is well-nested.

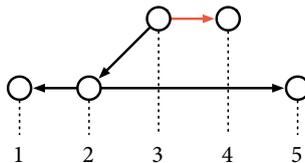


Figure 5.8: Pseudo-projectivity in the sense of Nasr (1995)

518 **Lemma** Let D be a dependency structure with n nodes and g gaps. The question whether D is well-nested can be decided in time $O(n + g)$. \dashv

Proof. To check whether D is well-nested, we first encode D into a term using the algorithm presented in Section 4.3; this takes time $O(n + g)$. In a traversal over this term, we then check whether any of the order annotations contains the forbidden scattered substring (see Theorem 517); using a stack data structure, this can be done in time linear in the accumulated size of the order annotations, which is again $O(n + g)$. By Theorem 517, the structure D is well-nested if and only if we do not find the forbidden substring. \dashv

Similar to the situation in Lemma 425, we can also bound the complexity of the algorithm as $O(k \cdot n)$, where k is the block-degree of D .

5.2.5 | Related work

✎ Nasr (1995) proposes a restriction on non-projective dependency structures that he calls the *pseudo-projectivity principle*.² Formally, ‘a dependency [edge] is pseudo-projective, if its dependent D is not situated, in the linear sequence, between two dependents of a node that is not an ancestor of D .’ Pseudo-projectivity is incomparable with both weak non-projectivity and well-nestedness. To see this, consider the dependency structure in Figure 5.8. The edge $3 \rightarrow 4$ in this structure is not pseudo-projective, as the node 4 is situated between two dependents (1 and 5) of a node that is not an ancestor of 4 (the node 2). On the other hand, the structure is weakly non-projective, and therefore well-nested. The structure D_5 (Figure 5.5b) is not well-nested, but pseudo-projective.

✎ The first algorithms for a well-nestedness test were presented by Möhl (2006). His first algorithm is based on a characterization of well-nestedness in terms of *interleaving yields*. The algorithm performs a tree traversal of the input structure, in which it first computes the yield of each node, and then checks for each pair of sibling nodes whether their yields interleave. In doing so, it uses $O(m^2 \cdot n)$ operations on bit vectors, where m

² This principle should not be confused with the notion of pseudo-projectivity introduced by Kahane et al. (1998), and subsequently used in (Gerdes and Kahane, 2001).

ALL DEPENDENCY STRUCTURES						
	DDT		PDT 1.0		PDT 2.0	
projective	3 730	84.95%	56 168	76.85%	52 805	77.02%
weakly non-proj.	3 794	86.40%	60 048	82.16%	56 367	82.21%
well-nested	4 386	99.89%	73 010	99.89%	68 481	99.88%
TOTAL	4 391	100.00%	73 088	100.00%	68 562	100.00%

NON-PROJECTIVE DEPENDENCY STRUCTURES ONLY						
	DDT		PDT 1.0		PDT 2.0	
weakly non-proj.	64	9.68%	3 880	22.93%	3 562	22.61%
well-nested	597	90.32%	16 842	99.54%	15 676	99.49%
TOTAL	661	100.00%	16 920	100.00%	15 757	100.00%

Table 5.1: The number of weakly non-projective and well-nested dependency structures in three treebanks

is the out-degree of D . Möhl’s second algorithm is build on the notion of the *gap graph*. The gap graph for a dependency structure is an extension of the governance relation by information about interleaving yields. Möhl shows that a dependency structure D is well-nested if and only if the gap graph for D contains a cycle. The size of the gap graph for a dependency structure D is tightly bounded by the square of the size of D , and the existence of a cycle in a graph can be checked in time linear in the size of that graph. Consequently, the run-time of Möhl’s second algorithm is $O(n^2)$.

✂ Havelka (2007b) studies the relationship between well-nestedness and the *level types* of non-projective edges (Havelka, 2005) and presents an algorithm that tests for well-nestedness in time $O(n^2)$.

5.3 | Empirical evaluation

To conclude this chapter, we now evaluate and compare the empirical adequacy of weak non-projectivity and well-nestedness on the treebank data. The corresponding counts and percentages are given in Table 5.1.

The experimental results for weak non-projectivity mirror its formal restrictiveness: enforcing weak non-projectivity excludes more than 75% of the non-projective data in both versions of the PDT, and 90% of the data in the DDT. Given these figures, weak non-projectivity appears to be of little use as a generalization of projectivity. The relatively large difference in coverage between the two treebanks may at least partially

be explained with their different annotation schemes for sentence-final punctuation: In the DDT, sentence-final punctuation marks are annotated as dependents of the main verb of a dependency subtree. This places severe restrictions on permitted forms of non-projectivity in the remaining sentence, as every discontinuity that includes the main verb must also include the dependent punctuation marks (see the discussion in Section 5.1). On the other hand, in the PDT, a sentence-final punctuation mark is annotated as a separate root node with no dependents. (Analyses in the PDT may be forests.) This scheme does not restrict the remaining discontinuities at all.

In contrast to weak non-projectivity, the well-nestedness constraint appears to constitute a very attractive extension of projectivity. For one thing, the almost perfect coverage of well-nestedness on both DDT and PDT (around 99.89%) could by no means be expected on purely combinatorial grounds: only 7% of all possible dependency structures for sentences of length 17 (the average sentence length in the PDT), and only slightly more than 5% of all possible dependency structures for sentences of length 18 (the average sentence length in the DDT) are well-nested.³ Similar results have been reported on other data sets (Havelka, 2007a). Moreover, a cursory inspection of the few problematic cases at least in the DDT indicates that violations of the well-nestedness constraint may, at least in part, be due to properties of the annotation scheme, such as the analysis of punctuation in quotations. However, a more detailed analysis of the data from both treebanks is needed before any stronger conclusions can be drawn concerning well-nestedness.

³ The number of unrestricted dependency structures on n nodes is given by sequence A000169, the number of well-nested dependency structures is given by sequence A113882 in Sloane (2007). The latter sequence was discovered by the author and Manuel Bodirsky. It can be calculated using a recursive formula derivable from the correspondence indicated in Example 512.

CHAPTER 6 | Structures and grammars

In the last three chapters, we have developed an algebraic framework for dependency structures. We now put this framework to use and classify several lexicalized grammar formalisms with respect to the classes of dependency structures that are induced by derivations in these formalisms.

Structure of the chapter. Each section of this chapter associates a grammar formalism with a class of dependency structures:

SECTION	FORMALISM	CLASS
Section 6.1	Context-Free Grammar (CFG)	\mathcal{D}_1
Section 6.2	Linear Context-Free Rewriting Systems (LCFRS(k))	\mathcal{D}_k
Section 6.3	Coupled Context-Free Grammar (CCFG(k))	$\mathcal{D}_k \cap \mathcal{D}_{wn}$
Section 6.4	Tree Adjoining Grammar (TAG)	$\mathcal{D}_2 \cap \mathcal{D}_{wn}$

6.1 | Context-free grammars

Let us go back to the notion of *induction* that we sketched in Chapter 1. Consider a derivation d of a terminal string by means of a context-free grammar. A *derivation tree* for d is a tree in which the nodes are labelled with (occurrences of) the productions used in the derivation, and the edges indicate how these productions were combined. If the underlying grammar is lexicalized, then there is a one-to-one correspondence between the nodes in the derivation tree and the positions in the derived string: every production that participates in the derivation contributes exactly one terminal symbol to this string. If we now order the nodes of the derivation tree according to the string positions of their corresponding anchors, then we get a dependency structure. We say that this dependency structure is *induced* by the derivation d . Induction identifies the governance relation of the induced dependency structure with the derivation relation, and the precedence relation with the left-to-right order in the derived string: the dependency structure contains an edge $u \rightarrow v$ if and only if the production that corresponds to the node v was used to rewrite some non-terminal in the production that corresponds to the node u ; the node u precedes the node v if and only if the anchor contributed by the production that corresponds to u precedes the anchor contributed by the production that corresponds to v . In this section, we formalize the correspondence between derivations and induced dependency structures and show that the class of dependency structures that can be induced by context-free grammars is exactly the class of projective dependency structures.

6.1.1 | Definition

We start with the familiar definition of a context-free grammar.

601 **Definition** A *context-free grammar* is a construct $G = (N, T, S, P)$, where N and T are alphabets of *non-terminal* and *terminal symbols*, respectively, $S \in N$ is a distinguished *start symbol*, and $P \subseteq N \times (N \cup T)^*$ is a finite set of productions. \dashv

We use indexed symbols (N_G, T_G, S_G, P_G) to refer to the components of a specific context-free grammar G .

602 **Example** To illustrate the ideas and constructions of this section, we use the following grammar $G = (N, T, S, P)$ as a running example:

$$N = \{S, B\}, \quad T = \{a, b\}, \quad P = \{S \rightarrow aSB, S \rightarrow aB, B \rightarrow b\}.$$

This grammar generates the string language $\{a^n b^n \mid n \in \mathbb{N}\}$. \dashv

Following the approach of Goguen et al. (1977), we treat context-free grammars as many-sorted algebras, in the following way. Let $G = (N, T, S, P)$ be a context-free grammar. For every string $\vec{x} \in (N \cup T)^*$, we define $\text{res}_N(\vec{x})$ to be the restriction of \vec{x} to letters in N . Specifically, res_N is the homomorphism from $(N \cup T)^*$ to N^* that is defined by $\text{res}_N(A) := A$ for $A \in N$, and $\text{res}_N(a) := \varepsilon$ for $a \in T$. We now turn the set P of productions of G into an N -sorted set $\Sigma(G)$ by defining

$$\text{type}_{\Sigma(G)}(A \rightarrow \vec{x}) := \text{res}_N(\vec{x}) \cdot A,$$

for every production $A \rightarrow \vec{x}$ in P . The set $T_{\Sigma(G)}$ of terms over the sorted set $\Sigma(G)$ forms an N -sorted algebra. This algebra represents the set of all derivations of G : the sortedness enforces the i th child of a node labelled with a production p to be labelled with a production that can be used to rewrite the i th non-terminal in p . More formally, there is a one-to-one correspondence between $T_{\Sigma(G)}$ and the set of all leftmost derivations in G . The set $T_{\Sigma(G), S}$ of all terms of sort S (the start symbol of the grammar) then corresponds to the set of all complete derivations in G . We call $T_{\Sigma(G)}$ the *derivation algebra* for G , and the terms $T_{\Sigma(G)}$ the *derivation trees* of G .

603 **Example (continued)** For our example grammar, we get the following sorted set:

$$[S \rightarrow aSB] : S \times B \rightarrow S, \quad [S \rightarrow aB] : B \rightarrow S, \quad [B \rightarrow b] : B.$$

(We enclose productions in square brackets to avoid ambiguities.) A derivation tree of the grammar is shown in the left half of Figure 6.1. \dashv

604 **Definition** A context-free grammar is called *lexicalized*, if each of its productions contains exactly one terminal symbol, called the *anchor* of that production. \dashv

We only consider lexicalized grammars. Each production in such a grammar has the form $A \rightarrow A_1 \cdots A_{k-1} \cdot a \cdot A_k \cdots A_m$, for some $m \in \mathbb{N}_0$ and $k \in [m + 1]$. Note that our example grammar is lexicalized.

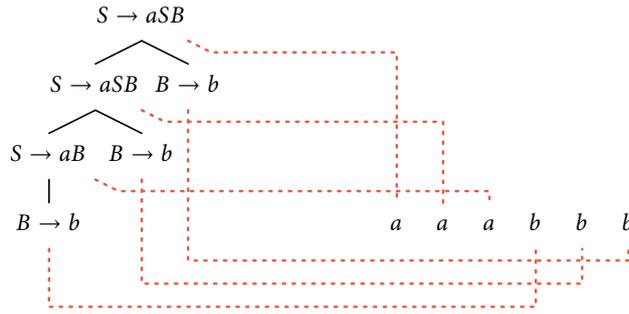


Figure 6.1: A derivation in a lexicalized context-free grammar

6.1.2 | String semantics

An immediate benefit of our algebraic take on context-free grammars is that we can use every $\Sigma(G)$ -algebra \mathfrak{A} as a semantic domain for the derivations of G : since $T_{\Sigma(G)}$ is a term algebra, it gives us the unique homomorphism $\llbracket \cdot \rrbracket_{\mathfrak{A}} : T_{\Sigma(G)} \rightarrow \mathfrak{A}$ that evaluates the derivation trees of G in \mathfrak{A} . In this way it is straightforward to derive the usual notion of the string language generated by a grammar G :

Definition Let G be a context-free grammar. The *string algebra* for G is the $\Sigma(G)$ -algebra \mathfrak{A} in which $\text{dom}(\mathfrak{A})_A = T_G^+$, for all $A \in N_G$, and

605

$$f_p(\vec{a}_1, \dots, \vec{a}_m) = \vec{a}_1 \cdots \vec{a}_{k-1} \cdot a \cdot \vec{a}_k \cdots \vec{a}_m,$$

for each production $p = A \rightarrow A_1 \cdots A_{k-1} \cdot a \cdot A_k \cdots A_m$ in $\Sigma(G)$. The *string language* generated by G is the set $L(G) := \llbracket T_{\Sigma(G), S_G} \rrbracket_{\mathfrak{A}}$. \dashv

Each composition operation f_p of the string algebra for a grammar G concatenates the anchor of p and the strings obtained from the subderivations in the order specified by the production p . This implements the usual rewriting semantics for context-free grammars. In the following, given a derivation tree $t \in T_{\Sigma(G)}$ of some CFG G , we write $\llbracket t \rrbracket_{\mathfrak{S}}$ for the evaluation of that tree in the string algebra for G . We also extend this notation to sets of derivation trees in the obvious way.

Example (continued) The right half of Figure 6.1 shows the string corresponding to the evaluation of the derivation that is shown in the left half in the string algebra for our example grammar. \dashv

606

6.1.3 | Linearization semantics

In a lexicalized context-free grammar, there is a one-to-one correspondence between the nodes in a derivation t and the positions of the string $\llbracket t \rrbracket_{\mathfrak{S}}$: every production participating in a derivation contributes exactly one terminal to the derived string.

607 **Example (continued)** In Figure 6.1, the one-to-one correspondence between the nodes of the derivation tree and the positions of the string is indicated by dashed lines. \dashv

We now show how to compute the mapping between nodes in the derivation and positions in the derived string that forms the basis of our notion of induction. To do so, we evaluate derivations t not in the string algebra, but in an algebra of *term linearizations*. A term linearization is a list of the nodes of a term in which each node occurs exactly once. In the following, we write $V := \mathbb{N}^*$ for the set of all nodes in terms; V^* then stands for the set of all strings over nodes. To avoid ambiguity, we use the symbol \circ for the concatenation operation on \mathbb{N} (which builds nodes), and \cdot for the concatenation operation on V (which builds strings of nodes). For every $i \in \mathbb{N}$, let pfx_i be the function that prefixes every node u in a given string by the number i . More formally, pfx_i is the string homomorphism from V to V that is defined by $\text{pfx}_i(u) = i \circ u$.

608 **Definition** Let G be a context-free grammar. The *linearization algebra* for G is the $\Sigma(G)$ -algebra \mathfrak{A} in which $\text{dom}(\mathfrak{A})_A = V^+$, for all $A \in N_G$, and

$$f_p(\vec{u}_1, \dots, \vec{u}_m) = \text{pfx}_1(\vec{u}_1) \cdot \dots \cdot \text{pfx}_{k-1}(\vec{u}_{k-1}) \cdot \varepsilon \cdot \text{pfx}_k(\vec{u}_k) \cdot \dots \cdot \text{pfx}_m(\vec{u}_m),$$

for each production $p = A \rightarrow A_1 \cdots A_{k-1} \cdot a \cdot A_k \cdots A_m$ in $\Sigma(G)$. The *linearization language* generated by G is the set $\Lambda(G) := \llbracket T_{\Sigma(G), S_G} \rrbracket_{\mathfrak{A}}$. \dashv

Each composition operation f_p of a linearization algebra concatenates a root node (representing the anchor of p) and the appropriately prefixed linearizations for the subderivations in the same order as they would be concatenated in the string algebra. Since the grammar G is lexicalized, the result of the evaluation in the linearization algebra defines a bijection between the set $\text{nod}(t)$ of nodes in the derivation tree t , and the set $\text{pos}(\llbracket t \rrbracket_{\mathfrak{S}})$ of positions in the derived string. Similar to the case of string algebras, for a derivation tree $t \in T_{\Sigma(G)}$ of some CFG G , we write $\llbracket t \rrbracket_{\mathfrak{L}}$ for the evaluation of the tree t in the linearization algebra for G .

609 **Example (continued)** For the derivation tree t shown in Figure 6.1, we get

$$\llbracket t \rrbracket_{\mathfrak{L}} = \varepsilon \cdot 1 \cdot 11 \cdot 111 \cdot 12 \cdot 2.$$

This linearization defines a mapping from the set $\text{nod}(t)$ of nodes in t to the set $\text{pos}(\llbracket t \rrbracket_{\mathfrak{S}})$ of positions in the derived string and back in the obvious way. Notice that, if we read the anchors of the productions in the derivation tree in the order specified by this linearization, then we obtain the string $\llbracket t \rrbracket_{\mathfrak{S}}$. \dashv

6.1.4 | Dependency semantics

With the mapping between nodes in the derivation tree and positions in the derived string at hand, we can now formalize the notion of induction:

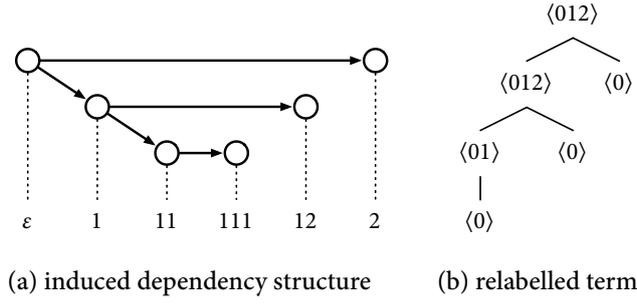


Figure 6.2: The dependency structure induced by the derivation in Figure 6.1

Definition Let G be a context-free grammar, and let $t \in T_{\Sigma(G)}$ be a derivation tree. 610
 The dependency structure *induced* by t is the structure $D := (\text{nod}(t); \preceq, \preceq)$ where

$$\begin{aligned}
 u \preceq v & \quad \text{if and only if} & \quad u \text{ dominates } v \text{ in } t, \text{ and} \\
 u \preceq v & \quad \text{if and only if} & \quad u \text{ precedes } v \text{ in } \llbracket t \rrbracket_{\mathbf{L}}.
 \end{aligned}
 \quad \dashv$$

Example Figure 6.2a shows the dependency structure induced by the derivation given in Figure 6.1. To illustrate the correspondence with the linearization, we have labelled the nodes with their addresses in the derivation tree. 611
 †

It is straightforward that the linearization semantics of derivations in CFGs directly mirrors our procedure for the traversal of treelet-ordered trees from Section 3.2.2. More specifically, we can understand the productions of a CFG as order annotations, and each derivation tree as a treelet-ordered tree. The behaviour of TREELET-ORDER-COLLECT is reflected in the evaluation of derivation trees in their corresponding linearization algebras. Taken together, this allows us to show that the class of dependency structures that is induced by lexicalized context-free grammars is exactly the class of projective dependency structures. In the following, we write $\mathcal{D}(\text{CFG})$ for the class of all dependency structures that can be induced by a context-free grammar.

Theorem $\mathcal{D}(\text{CFG}) = \mathcal{D}_1$ 612
 †

Proof. The proof of this statement is essentially identical to the proofs of Lemma 308 and Lemma 310. †

On a formal level, the translation between derivation trees of a CFG and projective dependency structures (represented by their corresponding terms) can be understood as a simple form of *relabelling*. Specifically, we can replace each production by an order annotation as follows, while maintaining the typing information:

$$A \rightarrow A_1 \cdots A_{k-1} \cdot a \cdot A_k \cdots A_m \quad \longleftrightarrow \quad \langle 1 \cdots (k-1) \cdot 0 \cdot k \cdots m \rangle.$$

The two terms are equivalent with respect to their linearization semantics. The only information that we lose is the label of the anchor, but that information is irrelevant with respect to induction anyway. (In Chapter 8, we will consider labelled dependency structures, where this information can be preserved.) As the essence of our discussion, we can define a dependency semantics for context-free grammars as follows. For each production p of a given context-free grammar G , let us write $\text{relab}(p)$ for the relabelling defined above.

613 **Definition** Let G be a context-free grammar. The *dependency algebra* for G is the $\Sigma(G)$ -algebra \mathfrak{D} in which $\text{dom}(\mathfrak{D})_A = \mathfrak{D}_1$, for all $A \in N_G$, and

$$f_p(D_1, \dots, D_m) = \text{dep}((\text{relab}(p))(\text{term}(D_1), \dots, \text{term}(D_m))),$$

for each production $p = A \rightarrow A_1 \cdots A_{k-1} \cdot a \cdot A_k \cdots A_m$. The *dependency language* generated by G is the set $\mathfrak{D}(G) := \llbracket T_{\Sigma(G), S_G} \rrbracket_{\mathfrak{D}}$. \dashv

6.2 | Linear Context-Free Rewriting Systems

We now extend our results from context-free grammars to the class of Linear Context-Free Rewriting Systems, LCFRS (Vijay-Shanker et al., 1987; Weir, 1988). This class was proposed as a generalization of a broad range of mildly context-sensitive formalisms. In this section, we show that the dependency structures induced by LCFRS are exactly the dependency structures of bounded degree. More specifically, we see that the block-degree measure for dependency structures is the structural correspondent of the *fan-out* measure that is used to identify sub-classes of LCFRS.

6.2.1 | Definition

Linear Context-Free Rewriting Systems can be understood as generalizations of context-free grammars in which derivations evaluate to tuples of strings. Our formal definition of LCFRS is essentially the same as the definitions proposed by Vijay-Shanker et al. (1987) and Satta (1992). In contrast to these, we make use of an explicit typing regime; this will simplify both the presentation and our formal arguments.

614 **Definition** Let A be an alphabet, and let $m \in \mathbb{N}_0$, and $\langle k_i \mid i \in [m] \rangle \in \mathbb{N}^m$, and $k \in \mathbb{N}$. A *generalized concatenation function* over A of type $k_1 \times \cdots \times k_m \rightarrow k$ is a function

$$f: (A^*)^{k_1} \times \cdots \times (A^*)^{k_m} \rightarrow (A^*)^k$$

that can be defined by an equation of the form

$$f(\langle x_{1,1}, \dots, x_{1,k_1} \rangle, \dots, \langle x_{m,1}, \dots, x_{m,k_m} \rangle) = \langle \vec{y}_1, \dots, \vec{y}_k \rangle,$$

where $\vec{y}_1 \cdots \vec{y}_k$ is a string over the variables on the left-hand side of the equation and the alphabet A in which each variable $x_{i,j}$, $i \in [m]$, $j \in [k_i]$, appears exactly once. \dashv

The semantics of a generalized concatenation function of type $k_1 \times \cdots \times k_m \rightarrow k$ is that it takes m tuples of strings and arranges the components of these tuples and a constant number of symbols from the alphabet A into a new k -tuple. The arity of the i th argument tuple is specified by the sort k_i . We regard generalized concatenation functions as syntactic objects and identify them with their defining equations. We call the right-hand sides of these equations the *bodies* of the (defining equations of) the corresponding generalized concatenation functions.

Definition A *linear context-free rewriting system* is a construct $G = (N, T, S, P)$, where N is an alphabet of *non-terminal symbols*, each of which is associated with a number $\varphi(A) \in \mathbb{N}$ called the *fan-out* of A ; T is an alphabet of *terminal symbols*; $S \in N$ is a distinguished *start symbol* with $\varphi(S) = 1$; and P is a finite set of productions of the form $A \rightarrow f(A_1, \dots, A_m)$, $m \in \mathbb{N}_0$, where $A, A_i \in N$, $i \in [m]$, and f is a generalized concatenation function over T of type $\varphi(A_1) \times \cdots \times \varphi(A_m) \rightarrow \varphi(A)$. \dashv

615

Example The following productions define an LCFRS. We use this LCFRS as our running example in this section. \dashv

616

$$\begin{array}{ll}
 S \rightarrow f_1(A) & f_1(\langle x_{1,1} \rangle) := \langle x_{1,1} b \rangle \\
 S \rightarrow f_2(A, B) & f_2(\langle x_{1,1} \rangle, \langle x_{2,1}, x_{2,2} \rangle) := \langle x_{1,1} x_{2,1} b x_{2,2} \rangle \\
 A \rightarrow f_3 & f_3 := \langle a \rangle \\
 B \rightarrow f_4(A, B) & f_4(\langle x_{1,1} \rangle, \langle x_{2,1}, x_{2,2} \rangle) := \langle x_{1,1} x_{2,1}, b x_{2,2} \rangle \\
 B \rightarrow f_5(A) & f_5(\langle x_{1,1} \rangle) := \langle x_{1,1}, b \rangle
 \end{array}$$

This LCFRS generates the string language $\{ a^n b^n \mid n \in \mathbb{N} \}$. \dashv

Our algebraic view on grammars generalizes to LCFRS without greater problems. Let $G = (N, T, S, P)$ be an LCFRS. We turn the set P of productions of G into a sorted set $\Sigma(G)$ by defining type $\Sigma(G)(A \rightarrow f(A_1, \dots, A_m)) := A_1 \cdots A_m \cdot A$, for every production $A \rightarrow f(A_1, \dots, A_m)$ in P . Just as in the case of context-free grammars, the set of all terms over $\Sigma(G)$ forms an N -sorted algebra that represents the set of all derivation trees of G .

Example (continued) For our example grammar, we get the following set $\Sigma(G)$: \dashv

617

$$\begin{array}{l}
 [S \rightarrow f_1(A)] : A \rightarrow S, \quad [S \rightarrow f_2(A, B)] : A \times B \rightarrow S, \quad [A \rightarrow f_3] : A, \\
 [B \rightarrow f_4(A, B)] : A \times B \rightarrow B, \quad [B \rightarrow f_5(A)] : A \rightarrow B.
 \end{array}$$

A corresponding derivation tree is shown in Figure 6.3a. \dashv

The concept of lexicalization is as for context-free grammars:

Definition An LCFRS is called *lexicalized*, if each of its productions contains exactly one terminal symbol. \dashv

618

6 STRUCTURES AND GRAMMARS

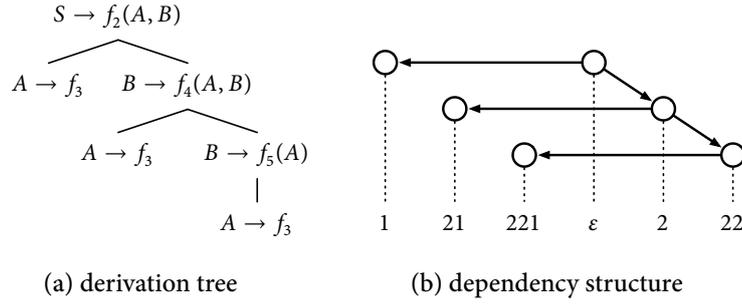


Figure 6.3: A derivation tree of an LCFRS, and the induced dependency structure

6.2.2 | String semantics

We now give LCFRS their usual string semantics. The string language generated by an LCFRS can be defined in terms of the evaluation of its derivation trees in an algebra over tuples of strings over the terminal alphabet. The arity of these tuples is specified by the fan-out of the non-terminal symbols that derive them. In the following, we use the Greek letters α and γ to refer to tuple components.

619 **Definition** Let $G = (N, T, S, P)$ be an LCFRS. The *string algebra* for G is the $\Sigma(G)$ -algebra \mathfrak{A} in which $\text{dom}(\mathfrak{A})_A = (T^*)^k$, for all $A \in N$ and $k = \varphi(A)$, and

$$f_p(\vec{\alpha}_1, \dots, \vec{\alpha}_m) = \vec{\gamma}[x_{i,j} \leftarrow \alpha_{i,j} \mid i \in [m] \wedge j \in [k_i]],$$

for each production $p = A \rightarrow f(A_1, \dots, A_m)$ with $f : k_1 \times \dots \times k_m \rightarrow k$ and body $\vec{\gamma}$. The *string language* generated by G is the set

$$L(G) := \{ \vec{a} \mid \exists t \in T_{\Sigma(G), S} \cdot \langle \vec{a} \rangle \in \llbracket t \rrbracket_{\mathfrak{A}} \}. \quad \dashv$$

Each composition operation f_p of the string algebra uses the body of the production p to construct a new tuple of strings. This tuple is obtained by replacing, for every $i \in [m]$ and $j \in [k_i]$, the variable $x_{i,j}$ with the j th component of the i th argument tuple. The string language generated by the LCFRS is obtained by extracting the strings from the (necessarily unitary) tuples derived from the start symbol of the grammar. We use the notation $\llbracket t \rrbracket_{\mathfrak{S}}$ as for context-free grammars.

6.2.3 | Non-essential concatenation functions

Unfortunately, the construction of the linearization semantics of LCFRS does not go through as smoothly as in the case of CFG. The fundamental problem is the fact that the mapping from derivation trees to strings is not injective in the case of LCFRS—to phrase it as a slogan, in LCFRS there are ‘more derivations than denotations’. At the root of this problem we find two types of ambiguity in LCFRS:

- (1) The first type of ambiguity concerns concatenation functions like

$$f = \langle a, \varepsilon \rangle \quad \text{and} \quad f(\langle x, y \rangle) = \langle axy \rangle.$$

Such function definitions cannot be translated into order annotations. In fact, we specifically excluded them in our definition of block-ordered trees (page 41) because they are essentially superfluous from the perspective of the linearization semantics: in the first function, the second component of the result does not contain any string material and could have been omitted to begin with; in the second definition, the two argument components are concatenated and would not have needed to be distributed over two components from the start.

- (2) The second ambiguity is inherent in the syntactic definition of LCFRS. In each derivation tree of a CFG, the left-to-right order on the subterms corresponds the left-to-right order on the material derived in these subterms. In LCFRS, this is not necessarily the case. Consider the following productions:

$$\begin{aligned} A \rightarrow f_1(B, C) & & f_1(\langle x_{1,1} \rangle, \langle x_{2,1} \rangle) & := \langle ax_{1,1}x_{2,1} \rangle \\ A \rightarrow f_2(C, B) & & f_2(\langle x_{1,1} \rangle, \langle x_{2,1} \rangle) & := \langle ax_{2,1}x_{1,2} \rangle \end{aligned}$$

For each pair of strings \vec{a}_1, \vec{a}_2 , both of the following two derivations evaluate to the same string $\vec{a}_1 \cdot \vec{a}_2$:

$$[A \rightarrow f_1(B, C)](\langle \vec{a}_1 \rangle, \langle \vec{a}_2 \rangle), \quad [A \rightarrow f_2(C, B)](\langle \vec{a}_2 \rangle, \langle \vec{a}_1 \rangle).$$

Let us call the body of a concatenation function $f: k_1 \times \dots \times k_m \rightarrow k$ *monotone*¹, if (i) the variables of the form $x_{i,1}$, $i \in [m]$, are numbered in increasing sequence from left-to-right, such that the i th occurrence of such a variable has the name $x_{i,1}$; (ii) for each $i \in [m]$, the variables of the form $x_{i,j}$, $j \in [k_i]$, are numbered in increasing sequence from left-to-right, such that the j th occurrence of such a variable has the name $x_{i,j}$. By requiring every definition of a concatenation function to be monotone, we can avoid the order ambiguity of our example: we simply disallow the definition of the function f_2 .

Let us call concatenation functions that fall into one of the two classes described above *non-essential*. We will now state a lemma (and give a rather technical and arduous proof) that will allow us to assume without loss of generality that an LCFRS is free of non-essential functions. For those functions that remain, the definitions of a linearization semantics go through without further problems.

To prove the lemma, we make use of more complicated versions of the relabelling function that we used for context-free grammars. Such a general relabelling is a function from terms over some input alphabet Σ into terms over an output alphabet Δ ; it works

¹ A slightly weaker condition of the same name is discussed by Kracht, 2003, p. 408.

by replacing, at each node u of the term, the label from the input alphabet by some label from the output alphabet. Both labels must have the same type. The choice of the output label is conditioned on the input label, and on some finite-state information coming from the children (bottom-up) or the parent (top-down) of u . Our formal definition in addition allows the relabelling to perform a permutation of subterms.

620 **Definition** A *bottom-up relabelling* is a construct $M = (Q, \Sigma, \Delta, R)$, where Q is a finite set of *states*, Σ and Δ are sorted alphabets of *input* and *output symbols*, respectively, and R is a finite set of rules such that for every symbol $\sigma :_{\Sigma} s_1 \times \cdots \times s_m \rightarrow s$ and all states $q_1, \dots, q_m \in Q$, R contains exactly one rule of the form

$$\sigma(\langle q_1, x_1 \rangle, \dots, \langle q_m, x_m \rangle) \rightarrow \langle q, \delta(x_{\pi(1)}, \dots, x_{\pi(m)}) \rangle,$$

where $q \in Q$, $\delta :_{\Delta} s_1 \times \cdots \times s_m \rightarrow s$, and π is some permutation $[m] \rightarrow [m]$. Dually, a *top-down relabelling* has rules of the form

$$\langle q, \sigma(x_1, \dots, x_m) \rangle \rightarrow \delta(\langle q_{\pi(1)}, x_{\pi(1)} \rangle, \dots, \langle q_{\pi(m)}, x_{\pi(m)} \rangle). \quad \dashv$$

The *derivation relation* induced by a bottom-up relabelling is the binary relation on the set $T_{\langle Q, T_{\Delta} \rangle \cup \Sigma}$ that is obtained by interpreting the rules of the relabelling as rewriting rules in the obvious way (for details, see Engelfriet, 1975; Engelfriet and Maneth, 2002). Similarly, a top-down relabelling gives rise to a relation on the set $T_{\langle Q, T_{\Sigma} \rangle \cup \Delta}$. In the proof of the next lemma, we will give (informal) descriptions of relabellings by specifying their state sets and explaining the behaviour of their translation rules.

621 **Lemma** Let $k \in \mathbb{N}$. For each lexicalized LCFRS $G \in \text{LCFRS}(k)$, there exists a lexicalized LCFRS $G' \in \text{LCFRS}(k)$ such that the derivation trees of G and G' are isomorphic modulo relabelling, $\llbracket G \rrbracket_{\mathcal{S}} = \llbracket G' \rrbracket_{\mathcal{S}}$, and G' only contains essential concatenation functions. \dashv

Proof. Let $k \in \mathbb{N}$, and let $G \in \text{LCFRS}(k)$ be a lexicalized LCFRS. Furthermore, assume that the rank of G is bounded by m , for some $m \in \mathbb{N}_0$. We define three relabellings on the set of derivation trees of G that transform each derivation tree t into a term over an alphabet Δ that does not contain the non-essential functions.

- (1) The first transformation is a top-down relabelling. As the set of states, use the set of all permutations $\pi: [k'] \rightarrow [k']$, for $k' \in [k]$, and start with the identity permutation on the set $[1]$. Given a state q and a production p with body $\vec{\alpha}$, replace $\vec{\alpha}$ by the tuple $\vec{\alpha}'$ that is obtained from $\vec{\alpha}$ by permuting the components of $\vec{\alpha}$ according to the permutation q , and re-index the variables in p from left to right. Let $\pi: [m] \rightarrow [m]$ be the permutation that maps each $i \in [m]$ to the position of the variable $x_{i,1}$ from $\vec{\alpha}$ in the total order on all variables of the form $x_{j,1}$ in $\vec{\alpha}'$, $j \in [m]$. For each $i \in [m]$, let $q_i: [k_i] \rightarrow [k_i]$ be the permutation that maps each $j \in [k_i]$ to the position of the variable $x_{i,j}$ from $\vec{\alpha}$ in the total order on all variables of the form $x_{i,j'}$ in $\vec{\alpha}'$, $j' \in [k_i]$.

- (2) The second transformation is a top-down relabelling. As the set of states, use the set of all subsets of $[k]$. Start with the empty set. Given a state q and a production p with body $\vec{\alpha}$, replace $\vec{\alpha}$ by the tuple $\vec{\alpha}'$ that is obtained from $\vec{\alpha}$ by (i) merging, for every $i \in q$, the i th component of $\vec{\alpha}$ with the $(i + 1)$ st component, and (ii) deleting, for every $i \in [m]$ and $j \in [k_i]$, each maximal substring of variables of the form $x_{i,j'}$, $j < j' \leq k_i$, that is adjacent to $x_{i,j}$; then re-index the variables of p from left to right. For each $i \in [m]$ let q_i be the set of all indices $j \in [k_i]$ such that the variable $x_{i,j}$ was removed in step (ii), and let π be the identity permutation.
- (3) The third transformation is a bottom-up relabelling. As the set of states, use the set of all subsets of $[m'] \times [k']$, for $m' \in [m]$, $k' \in [k]$. Start with the empty set. Given a production p with body $\vec{\alpha}$ and states $\langle q_i \mid i \in [m] \rangle$ for the subterms, replace $\vec{\alpha}$ by the tuple $\vec{\alpha}'$ that is obtained from $\vec{\alpha}$ by (i) deleting all variables $x_{i,j}$, where $i \in [m]$ and $j \in q_i$, and (ii) deleting all empty components; then, re-index the variables of p from left-to-right. Let $\pi: [m] \rightarrow [m]$ be the permutation that maps each $i \in [m]$ to the position of the variable $x_{i,1}$ from $\vec{\alpha}$ in the total order on all variables of the form $x_{j,1}$ in $\vec{\alpha}'$, $j \in [m]$. Let q be the set of indices of all components deleted in step (ii).

None of these transformations alters the term structure of the original derivation tree (apart from the permutation of subterms, which is inessential with respect to tree-isomorphism), or the string derived from the tree: by induction on the derivation tree t , we can verify that its string semantics remain invariant under the relabelling. \dashv

6.2.4 | Linearization semantics

With all non-essential concatenation functions out of our way, we are now ready to define the linearization semantics for LCFRS.

Definition Let G be an LCFRS. The *linearization algebra* for G is the $\Sigma(G)$ -algebra \mathfrak{A} in which $\text{dom}(\mathfrak{A})_A = (V^+)^k$, for all $A \in N_G$ and $k = \varphi(A)$, and

622

$$f_p(\vec{\alpha}_1, \dots, \vec{\alpha}_m) = \vec{\gamma}[a \leftarrow \varepsilon][x_{i,j} \leftarrow \text{pfx}_i(\alpha_{i,j}) \mid i \in [m] \wedge j \in [k_i]],$$

for each production $p = A \rightarrow f(A_1, \dots, A_m)$ with $f: k_1 \times \dots \times k_m \rightarrow k$, anchor $a \in T$ and body $\vec{\gamma}$. The *linearization language* generated by G is the set

$$\Lambda(G) := \{ \vec{u} \mid t \in T_{\Sigma(G), S_G} \wedge \langle \vec{u} \rangle \in \llbracket t \rrbracket_{\mathfrak{A}} \}. \quad \dashv$$

Each composition operation f_p of a linearization algebra uses the body of the defining equation of p to concatenate a root node (representing the anchor of p) and the appropriately prefixed linearizations for the subderivations in the same order as they would be concatenated in the string algebra.

623 **Example (continued)** For the derivation tree t shown in Figure 6.3a, we get the linearization $\llbracket t \rrbracket_{\mathbf{L}} = 1 \cdot 21 \cdot 221 \cdot \varepsilon \cdot 2 \cdot 22$. \dashv

6.2.5 | Dependency semantics

We now define a dependency semantics for LCFRS.

624 **Definition** Let G be an LCFRS, and let $t \in T_{\Sigma(G)}$ be a derivation tree. The dependency structure *induced* by t is the segmented structure $D := (\text{nod}(t); \preceq, \leq, \equiv)$ where

$$\begin{aligned} u \preceq v & \quad \text{if and only if} & \quad u \text{ dominates } v \text{ in } t, \\ u \leq v & \quad \text{if and only if} & \quad u \text{ precedes } v \text{ in } \llbracket t \rrbracket_{\mathbf{L}}, \text{ and} \\ u \equiv v & \quad \text{if and only if} & \quad u \text{ and } v \text{ appear in the same component of } \llbracket t \rrbracket_{\mathbf{L}}. \end{aligned} \quad \dashv$$

625 **Example (continued)** The derivation tree t shown in Figure 6.3a induces the dependency structure shown in Figure 6.3b. Note that, while the string language generated by the LCFRS is the same as the string language generated by the CFG in the previous section, the dependency structure is fundamentally different. Just as in the case of context-free grammar, however, the generated string $\llbracket t \rrbracket_{\mathbf{S}}$ can be recovered from the linearization by reading the anchors of t in the order specified by $\llbracket t \rrbracket_{\mathbf{L}}$. \dashv

Inspecting the linearization semantics, we see that there is a obvious similarity between the bodies of the generalized concatenation functions used in an LCFRS and the order annotations that we used for block-ordered trees: the j th occurrence of the symbol i in our order annotations has the same semantics as the variable $x_{i,j}$ in the body of a generalized concatenation function. Under this view, the linearization semantics of LCFRS mirrors the behaviour of the procedure BLOCK-ORDER-COLLECT that we gave in Section 4.2.1. Just as in the case of projective dependency structures, the following theorem can be shown by replicating the proofs of Lemma 417. In the following, we write $\text{LCFRS}(k)$ for the class of all LCFRS in which the fan-out of the non-terminal symbols is bounded by k .

626 **Theorem** $\forall k \in \mathbb{N}. \mathcal{D}(\text{LCFRS}(k)) = \mathcal{D}_k$ \dashv

We can also define a relabelling function that translates between productions of an LCFRS and order annotations:

$$\text{relab}(p) := \vec{\alpha}[a \leftarrow 0][x_{i,j} \leftarrow i \mid i \in [m] \wedge j \in [\varphi(A_i)]],$$

for every production $p = A \rightarrow f(A_1, \dots, A_m)$ with anchor a and body $\vec{\alpha}$. Based on this function, the concept of a dependency algebra for an LCFRS can be defined analogously to the corresponding definition for context-free grammars (Definition 613).

6.2.6 | Related work

The string languages generated by LCFRS have many characterizations; among other things, they are generated by Multiple Context-Free Grammars (Seki et al., 1991), and they are the images of regular tree languages under deterministic tree-walking transducers (Weir, 1992) and under finite-copying top-down tree transducers (van Vugt, 1996). They can also be characterized as the yield languages of rational tree relations (Raoult, 1997).

6.3 | Coupled Context-Free Grammars

In this section, we look at the dependency structures that are induced by derivations in Coupled Context-Free Grammars (CCFGs) (Guan, 1992; Hotz and Pitsch, 1996). This formalism can be understood as a variant of LCFRS where rewriting rules are restricted to words over a Dyck language, that is, a language that consists of balanced strings of parentheses.² We show that this syntactic restriction enforces the dependency structures induced by CCFG derivations to be well-nested.

6.3.1 | Definition

We start with a formal definition of CCFGs. Our definition deviates from the one given by Hotz and Pitsch (1996) in that we treat ‘parentheses’ as symbols from a ranked set.

Definition Let Π be a ranked alphabet. The alphabet of *components* of Π is the set $\text{comp}(\Pi) := \{ \langle \pi, i \rangle \in \Pi \times \mathbb{N} \mid 1 \leq i \leq \text{rank}_{\Pi}(\pi) \}$. 627 \dashv

In the following, if no confusion can arise, we write π_i instead of $\langle \pi, i \rangle$.

Definition Let Π be a ranked alphabet, and let A be an alphabet. The *extended semi-Dyck set* over Π and A is the smallest set $\text{ED}(\Pi, A) \subseteq (\text{comp}(\Pi) \cup A)^*$ that is closed under the inference rules given in Table 6.1. 628 \dashv

Example Consider the ranked alphabet $\Pi := \{ \circ \}$ where $\text{rank}_{\Pi}(\circ) = 2$. The components of Π can be understood as opening and closing brackets, respectively. Specifically, let us write $[$ instead of $\langle \circ, 1 \rangle$, and $]$ instead of $\langle \circ, 2 \rangle$. Then the set $\text{ED}(\Pi, \{a, b\})$ consists of the set of all well-bracketed words over $\{a, b\}^*$. For example, the strings $[a][b]$ and $a[a[ba]][a]$ belong to this set, while the strings $[[a]$ and $[b]a]$ do not. 629 \dashv

One important property of an extended semi-Dyck set $\text{ED}(\Pi, A)$ that we will make use of is that, modulo the associativity of the concatenation rule and the introduction of superfluous empty strings, every string $\vec{x} \in \text{ED}(\Pi, A)$ has a unique decomposition in terms of the rules given in Table 6.1.

² Dyck languages are named after the German mathematician Walther von Dyck (1856–1934), whose surname rhymes with ‘week’ rather than ‘spike’.

$$\begin{array}{c}
 \frac{}{\varepsilon \in \text{ED}(\Pi, A)} \qquad \frac{a \in A}{a \in \text{ED}(\Pi, A)} \qquad \frac{\vec{x} \in \text{ED}(\Pi, A) \quad \vec{y} \in \text{ED}(\Pi, A)}{\vec{x} \cdot \vec{y} \in \text{ED}(\Pi, A)} \\
 \frac{\pi \in \Pi_1}{\langle \pi, 1 \rangle \in \text{ED}(\Pi, A)} \qquad \frac{\vec{x}_1 \in \text{ED}(\Pi, A) \quad \cdots \quad \vec{x}_k \in \text{ED}(\Pi, A) \quad \pi \in \Pi_{k+1}}{\langle \pi, 1 \rangle \cdot \vec{x}_1 \cdot \langle \pi, 2 \rangle \cdots \langle \pi, k \rangle \cdot \vec{x}_k \cdot \langle \pi, k+1 \rangle \in \text{ED}(\Pi, A)}
 \end{array}$$

Table 6.1: Inference rules for the extended semi-Dyck set

630 **Definition** A *coupled context-free grammar* is a construct $G = (\Pi, T, S, P)$, where Π is a ranked alphabet of *non-terminal symbols*, T is an alphabet of *terminal symbols*, $S \in \Pi_1$ is a distinguished *start symbol*, and P is a finite, non-empty set of productions of the form $A \rightarrow \langle \vec{x}_1, \dots, \vec{x}_k \rangle$, where $k \in \mathbb{N}$, $A \in \Pi_k$, and $\vec{x}_1 \cdots \vec{x}_k \in \text{ED}(\Pi, T)$. \dashv

631 **Example** We use the following CCFG G as our running example. The alphabet of non-terminal symbols is $\Pi_G := \{S/1, R/2, B/1, C/1, D/1\}$; the alphabet of terminal symbols is $T_G := \{a, b, c, d\}$. The start symbol of G is $S_G := S$. Finally, the set of productions is defined as follows. (We omit subscripts for non-terminals with rank 1.)

$$\begin{aligned}
 S &\rightarrow \langle aR_1BCR_2D \rangle \mid \langle aBCD \rangle \\
 R &\rightarrow \langle aR_1B, CR_2D \rangle \mid \langle aB, CD \rangle \\
 B &\rightarrow \langle b \rangle, \quad \text{and similarly for } C \text{ and } D. \quad \dashv
 \end{aligned}$$

We now show how to construct the derivation algebra for a coupled context-free grammar $G = (\Pi, T, S, P)$. For every string $\vec{x} \in (\text{comp}(\Pi) \cup T)^*$, we define $\text{res}_\Pi(\vec{x})$ to be the restriction of \vec{x} to the first components of Π . More specifically, res_Π is the homomorphism from $(\text{comp}(\Pi) \cup T)^*$ to Π^* that is defined by $\text{res}_\Pi(\langle \pi, 1 \rangle) = \pi$ for $\pi \in \Pi$, and $\text{res}_\Pi(x) = \varepsilon$ for all other symbols. We turn the set P of productions of G into a Π -sorted set $\Sigma(G)$ by defining

$$\text{type}_{\Sigma(G)}(A \rightarrow \langle \vec{x}_1, \dots, \vec{x}_k \rangle) := \text{res}_\Pi(\vec{x}_1 \cdots \vec{x}_k) \cdot A,$$

for every production $A \rightarrow \langle \vec{x}_1, \dots, \vec{x}_k \rangle$, where $A \in \Pi_k$. The set of *derivation trees* of G is defined as for CFG and LCFRS.

632 **Example** Figure 6.4 shows a derivation tree for our example grammar. \dashv

The notion of lexicalization is defined as usual:

633 **Definition** A CCFG is called *lexicalized*, if each of its productions contains exactly one terminal symbol. \dashv

Notice that the example grammar is lexicalized.

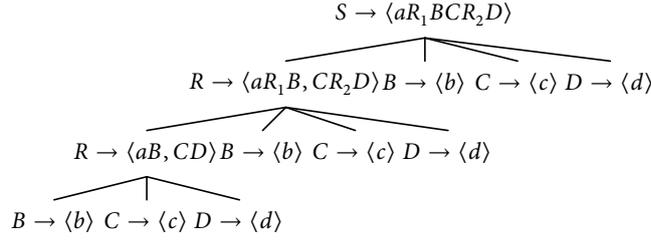


Figure 6.4: A derivation tree for a CCFG

6.3.2 | String semantics

The basic reading of a production p in a CCFG is, that like a production in an LCFRS it describes a generalized concatenation function that arranges material derived in subderivations into a new tuple. However, in order to define the semantics of CCFGs precisely, we need to make explicit which components on the right-hand side $\vec{\alpha}$ of p ‘belong together’ in the sense that they should be replaced with material derived in the same subderivation. In LCFRS, this correspondence is encoded by means of appropriately named variables: the variable $x_{i,j}$ in $\vec{\alpha}$ is a placeholder for the j th component of the i th argument of the generalized concatenation function used in p (see Definition 619). In CCFG, due to the Dyck-ness restriction on the productions, the correspondence can be left more implicit.

Let us say that two occurrences of components in the right-hand side $\vec{\alpha}$ of a production p in a CCFG are *synchronized*, if they were introduced in the same inference step in the derivation of $\vec{\alpha}$ according to the rules given in Table 6.1. Given that every right-hand side $\vec{\alpha}$ has a unique such derivation (modulo inessential ambiguities), the synchronization relation is well-defined; it defines an equivalence relation on the set of all occurrences of components in $\vec{\alpha}$. We can then index the synchronized sets of components according to the position of their leftmost member (a component of the form $\langle \pi, 1 \rangle$), and use the components from the i th group in this sequence as placeholders for the material from the i th subderivation. More formally, we can rewrite $\vec{\alpha}$ into an explicit version $\text{exp}(\vec{\alpha})$ by replacing the j th element of the i th synchronized group of component occurrences by the variable symbol $x_{i,j}$. Based on this explicit version, we can define the string semantics of CCFGs as in the case of LCFRS:

Definition Let $G = (\Pi, T, S, P)$ be a CCFG. The *string algebra* for G is the $\Sigma(G)$ -algebra \mathfrak{A} in which $\text{dom}(\mathfrak{A})_A = (T^*)^k$, for all $A \in \Pi_k$, and

634

$$f_p(\vec{\alpha}_1, \dots, \vec{\alpha}_m) = \text{exp}(\vec{\gamma})[x_{i,j} \leftarrow \alpha_{i,j} \mid i \in [m] \wedge j \in [k_i]],$$

for each production $p = A \rightarrow \vec{\gamma}$. The *string language* generated by G is the set

$$L(G) := \{ \vec{\alpha} \mid \exists t \in T_{\Sigma(G), S} \cdot \langle \vec{\alpha} \rangle \in \llbracket t \rrbracket_{\mathfrak{A}} \}. \quad \dashv$$

635 **Example (continued)** Our example grammar generates the context-sensitive string language $\{ a^n b^n c^n d^n \mid n \in \mathbb{N} \}$. \dashv

It is straightforward that every CCFG can be translated into an LCFRS that generates the same string language. In this translation, both the alphabet of non-terminal symbols, the alphabet of terminal symbols, and the start symbol remain unchanged; the only thing that we need to adapt is the form of the production rules, which can be done in the way that we just explained. In this sense, we can view CCFGs as a syntactically restricted form of LCFRS. We can then define the linearization semantics of CCFGs as for LCFRS.

6.3.3 | Dependency semantics

Given that every CCFG can be seen as a special LCFRS, it is clear that CCFGs cannot induce more dependency structures than LCFRS. In particular, we have the following lemma, which relates the block-degree of the induced dependency structures to the maximal rank of the inducing CCFG. Let us write $\text{CCFG}(k)$ for the class of all CCFGs in which the maximal rank of a non-terminal is k .

636 **Lemma** $\forall k \in \mathbb{N}. \mathcal{D}(\text{CCFG}(k)) \subseteq \mathcal{D}_k$ \dashv

We now show that CCFGs in fact induce a proper subclass of the dependency structures inducible by LCFRS: every dependency structure induced by a CCFG is well-nested.

637 **Lemma** $\mathcal{D}(\text{CCFG}) \subseteq \mathcal{D}_{wn}$ \dashv

Proof. Let $G = (\Pi, T, S, P)$ be a CCFG, and let $t = p(t_1, \dots, t_m)$ be a derivation tree of G , where $p : k_1 \times \dots \times k_m \rightarrow k$ and $t_i \in T_{\Sigma(G), k_i}$, for all $i \in [m]$. The proof proceeds by induction on the depth of t . Consider the linearization $\llbracket t \rrbracket_{\mathbf{L}}$ of t ; this linearization has the form $\llbracket t \rrbracket_{\mathbf{L}} = \langle \vec{u}_1, \dots, \vec{u}_k \rangle$. Now assume that there exist four nodes $v_1 = ui\vec{x}_1$, $v_2 = ui\vec{x}_2$ and $w_1 = uj\vec{y}_1$, $w_2 = uj\vec{y}_2$ in t such that $v_1 \rightarrow v_2$, $w_1 \rightarrow w_2$, and $[v_1, v_2] \not\bowtie [w_1, w_2]$. Furthermore, and without loss of generality, assume that $\min(v_1, v_2) < \min(w_1, w_2)$. Then the string $\vec{u}_1 \cdots \vec{u}_k$ contains a substring of the form

$$ui\vec{x}_1 \cdots uj\vec{y}_1 \cdots ui\vec{x}_2 \cdots uj\vec{y}_2,$$

for some $u \in \text{nod}(t)$, $i, j \in \mathbb{N}$, and $\vec{x}_1, \vec{x}_2, \vec{y}_1, \vec{y}_2 \in V^*$. Distinguish two cases: If $u \neq \varepsilon$, then both v_1, v_2 and w_1, w_2 belong to a proper subterm of t , and by the induction hypothesis, we may assume that $v_1 \trianglelefteq w_1$ or $w_1 \trianglelefteq v_1$. So assume that $u = \varepsilon$. The right-hand side of the production p has the form $\langle \gamma_1, \dots, \gamma_k \rangle$, and the string $\gamma := \gamma_1 \cdots \gamma_k$ is formed according to the inference rules given in Table 6.1. Given the specific form of the string $\vec{u}_1 \cdots \vec{u}_k$, the string γ contains a substring of the form

$$\langle \pi_i, i_1 \rangle \cdots \langle \pi_j, j_1 \rangle \cdots \langle \pi_i, i_2 \rangle \cdots \langle \pi_j, j_2 \rangle,$$

for some $\pi_i, \pi_j \in \Pi$ and $1 \leq i_1 < i_2 \leq \text{rank}_\Pi(\pi_i)$, $1 \leq j_1 < j_2 \leq \text{rank}_\Pi(\pi_j)$, where i is the left-to-right index of the synchronized group of (occurrences of) components to which $\langle \pi_i, i_1 \rangle$ and $\langle \pi_i, i_2 \rangle$ belong, and j is the corresponding index for $\langle \pi_j, j_1 \rangle$ and $\langle \pi_j, j_2 \rangle$. By the inferences rules in Table 6.1, it is then clear that all four occurrences are synchronized; consequently, $i = j$, and either $v_1 \trianglelefteq w_1$ or $w_1 \trianglelefteq v_1$ holds. This shows that the dependency structure induced by t is well-nested. \dashv

As in the case of CFGs and LCFRS, the converses of the preceding lemmata are easy to show. We thus obtain a characterization of CCFG in terms of the dependency structures that it can induce as follows:

Theorem $\forall k \in \mathbb{N}. \mathcal{D}(\text{CCFG}(k)) = \mathcal{D}_k \cap \mathcal{D}_{wn}$ \dashv 638

6.4 | Tree Adjoining Grammar

To conclude this chapter, we now look at the dependency structures that are induced by derivations in Tree Adjoining Grammars (TAG) (Joshi et al., 1975; Joshi and Schabes, 1997). In contrast to LCFRS and CCFGs, TAGs manipulate trees rather than strings.

6.4.1 | Definition

The building blocks of a TAG are called *elementary trees*. These are children-ordered trees in which each node has one of three types: it can be an *anchor* (or *terminal node*), a *non-terminal node*, or a *foot node*. Anchors and foot nodes are required to be leaves; non-terminal nodes may be either leaves or inner nodes. Each elementary tree can have at most one foot node. Elementary trees without a foot node are called *initial trees*; the remaining trees are called *auxiliary trees*. A TAG grammar is *lexicalized*, if each of its elementary trees contains exactly one anchor (Schabes et al., 1988). Trees in TAG can be combined using two operations (see Figure 6.5): *substitution* combines a tree τ with an initial tree τ' by identifying a non-terminal leaf node u of τ with the root node of τ' (Figure 6.5a); *adjunction* identifies an inner node u of a tree τ with the root node of an auxiliary tree τ' ; the subtree of τ that is rooted at u is excised from τ and inserted below the foot node v of τ' (Figure 6.5b). Combination operations are disallowed at root and foot nodes.

Example Figure 6.6 shows an example for how TAGs are specified. The grammar contains 5 elementary trees, named τ_1 to τ_5 . The elementary trees $\tau_1 - \tau_4$ are initial trees. The tree τ_5 is an auxiliary tree; the foot node of this tree is marked with a star. Note that this grammar is lexicalized. By adjoining the tree τ_5 into the tree τ_1 , and then repeatedly into the tree resulting from this first adjunction, we can produce the string language $\{a^n b^n c^n d^n \mid n \in \mathbb{N}\}$. This language is beyond the string-generative capacity of context-free grammars. \dashv 639

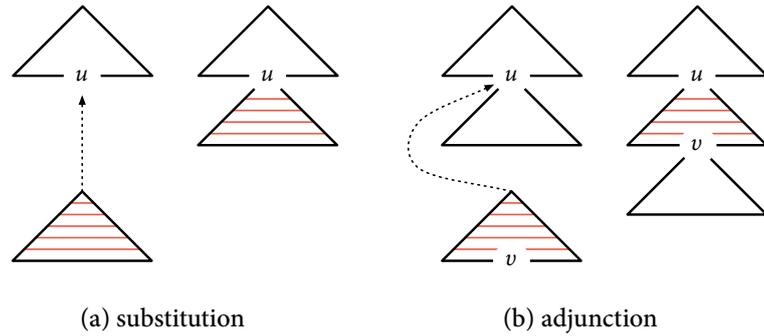


Figure 6.5: Combination operations in TAG

Just as in the case of the other grammar formalisms that we have looked at in this section, TAG *derivation trees* record information about how elementary structures are combined. Formally, derivation trees can be seen as terms over the signature of elementary trees; this set is finite for any given TAG. The root of each derivation tree is an initial tree. By repeated applications of the substitution and adjunction operations, larger and larger trees are built from this tree. TAG *derived trees* represent the results of complete derivations: they are standard children-ordered trees made up from the accumulated material of the elementary trees participating in the derivation. Just as in the case of lexicalized CFG there was a one-to-one correspondence between the nodes of the derivation tree and the positions of the derived string, in lexicalized TAGS there is a one-to-one correspondence between the nodes of the derivation tree and the leaves of the derived tree. Thus, in just the same way as derivations in CFG and LCFRS, derivations in TAG induce dependency structures. The major question that we have to answer in the context of TAGS is how to define the linearization semantics of the derivation trees. Essentially the same question needs to be addressed when trying to relate TAG to other mildly context-sensitive grammar formalisms. (See Boullier, 1999 for a formal version of the construction that we discuss here.)

640 **Example (continued)** Figure 6.7 shows a derivation tree for our example grammar and the dependency structure that is induced by this derivation. ←

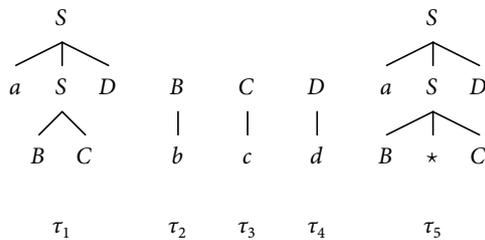


Figure 6.6: A TAG grammar

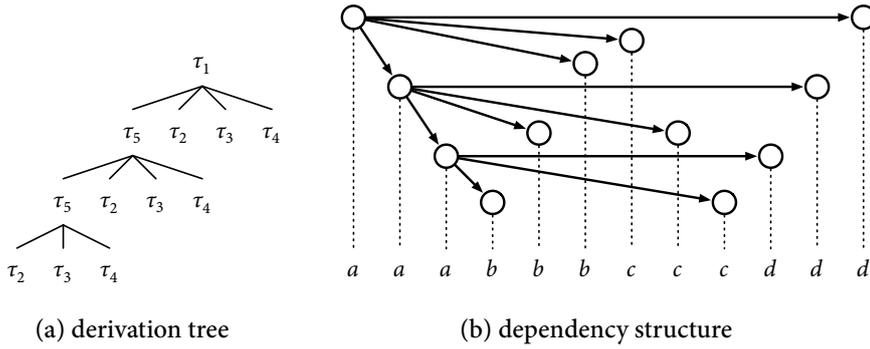


Figure 6.7: A TAG derivation, and the induced dependency structure

6.4.2 | Linearization semantics

To understand the linearization semantics of a TAG elementary tree τ , we must read it with the derived tree in mind that would result from a derivation starting at τ . Let us do so for the elementary tree τ_1 that is shown in Figure 6.6. Since combination operations are disallowed at the root nodes of elementary trees, we see that the leftmost leaf in the derived tree that we produce as the result of the derivation is the anchor of τ_1 , which is labelled with the symbol a . Now assume that an auxiliary tree adjoins at the central node of τ_1 , the node that is labelled with S . Then in the derived tree, all the material in the left half of the adjoined tree precedes the material that is dominated by the adjunction site in τ_1 , while all the material in the right half succeeds it (see again Figure 6.5b). Specifically, let us write S_1 for the material derived from the left half of the auxiliary tree, and S_2 for the material in the right half. Then S_1 precedes the material that gets substituted into the nodes labelled with B and C in τ_1 , and this material in turn precedes S_2 . Finally, at the right edge of the derived tree that we build from τ_1 , we find the material substituted into the node labelled with D . To summarize, we have the following linear sequence of tree material:

$$a S_1 B C S_2 D.$$

Using reasoning similar to this, we see that every elementary tree specifies a generalized concatenation function over tuples of arity at most 2: for sub-derivations that correspond to adjunctions, there is not one single slot per subderivation, but two; the first of these slots is filled by the material in the left half of the adjoined tree, the second slot by the right half of the corresponding tree. A crucial observation now is that the linearization of the elementary trees always produces strings from an extended semi-Dyck set. More specifically, there can never be two distinct adjunction sites A and B such that the linearization of the corresponding elementary tree yields the sequence $A_1 \cdots B_1 \cdots A_2 \cdots B_2$. This is so because all material that is situated between the two slots of a given adjunction corresponds to material that in the elementary tree is dominated by that adjunction site. The forbidden sequence then would mean that

both A dominates B , and B dominates A , which is only possible if $A = B$. Hence, from the linearization point of view, TAG corresponds to the class CCFG(2) of Coupled Context-Free Grammars with rank at most 2.

641 **Theorem** $\mathcal{D}(\text{TAG}) = \mathcal{D}_2 \cap \mathcal{D}_{wn}$ ←

This result was first shown by Bodirsky et al. (2005).

6.4.3 | Related work

✂ The same reasoning that we have used for the linearization semantics is needed when designing left-to-right parsing algorithms for TAGS (Joshi and Schabes, 1997). It was also used by Guan (1992) to link TAGS to CCFGs of rank at most 2

✂ Tree Adjoining Grammars are special forms of Ranked Node Rewriting Grammars (Abe, 1988) and Context-Free Tree Grammars (Mönnich, 1997; Fujiyoshi and Kasai, 2000; Kepser and Mönnich, 2006). Our results carry over to these extended formalisms. In particular, k -adic (non-deleting) Linear Context-Free Tree Grammars correspond to CCFGs with rank at most k , and induce well-nested dependency structures with block-degree at most k .

Summary

In this chapter, we have presented a classification of lexicalized grammar formalisms in terms of the dependency structures that these formalisms can induce. Our classification provides a new measure for the generative capacity of a grammar formalism that is attractive as an alternative to both string-generative capacity and tree-generative capacity: dependency structures are more informative than strings, but less formalism-specific than parse trees (see Kallmeyer, 2006, for a similar argument).

Together with the treebank studies that we presented in the previous three chapters, our classification also provides new insights into the practical relevance of grammar formalisms: If we accept our conclusion that the class of projective dependency structures is insufficient to cover all the data in the three treebanks that we looked at, then by Theorem 612, the same holds for lexicalized context-free grammars. At the same time, our treebank studies revealed that only a small step beyond projectivity is necessary to cover virtually all of the practically relevant data. Together with Theorem 626, we can interpret this result as saying that we only need LCFRS with a very small fan-out, say the class LCFRS(2). Perhaps most interestingly, we find that close to 99.5% of all the dependency analyses in the treebanks are well-nested and have a block-degree of at most 2. Given Theorem 641, this means that one should at least in theory be able to write a TAG that induces all the structures in the three treebanks.

CHAPTER 7 | Regular dependency languages

In the first part of this dissertation, we have looked at formal properties of individual dependency structures. In this chapter, we turn our attention to sets of such structures, or *dependency languages*. Specifically, we investigate the languages that arise when we equip dependency structures with a ‘regular’ means of syntactic composition.

Structure of the chapter. We start by defining regular dependency languages as the recognizable subsets in dependency algebras and provide natural notions of automata and grammars for this class of languages (Section 7.1). We then develop a powerful pumping lemma for regular dependency languages (Section 7.2) and apply it to show that the languages in this class are of constant growth, a property characteristic for mildly context-sensitive languages (Section 7.3).

7.1 | Regular sets of dependency structures

The primary goal of this dissertation is to illuminate the connections between language-theoretic properties such as generative capacity and parsing complexity on the one hand, and graph-theoretic properties such as block-degree and well-nestedness on the other. Specifically, we started with the question,

Which grammars induce which sets of dependency structures?

At this point, we have already come quite close to an answer to this question. Consider the class of lexicalized context-free grammars (LCFG) for example. In the previous chapter, we have seen that LCFG is linked to projectivity in the sense that every LCFG can induce only projective structures, and every such structure can be induced by some LCFG. However, this result does not yet provide a full answer to our question, which asks about classes of *languages*, not classes of *structures*. The step from structures to languages is non-trivial: it is not true that every *set* of projective dependency structures can be induced by an LCFG; in particular, the set of LCFGs is denumerable, the set of all subsets of \mathcal{D}_1 is not. In this chapter, we identify a class of dependency languages that *can* be induced by the grammar formalisms that we have discussed. We call this class the *regular* dependency languages. As we will see, the condition of regularity provides the missing link between grammar formalisms and dependency structures.

7.1.1 | Algebraic recognizability

We define regularity by referring to the concept of *algebraic recognizability*. This notion was introduced by Mezei and Wright (1967), following previous work by Richard

Büchi, John Doner, Calvin Elgot, and James Thatcher. It generalizes the definitions of regular string languages to arbitrary abstract algebras and provides a canonical way to characterize regular sets of objects. Recognizability was originally defined for single-sorted algebras. Here we use a many-sorted version, due to Courcelle (1996).

701 **Definition** Let \mathfrak{A} be a Σ -algebra, and let $s \in \mathcal{S}_\Sigma$ be a sort. A set $L \subseteq \text{dom}(\mathfrak{A})_s$ is called *recognizable*, if there exists a finite Σ -algebra \mathfrak{B} , a homomorphism $h: \mathfrak{A} \rightarrow \mathfrak{B}$, and a set $F \subseteq \text{dom}(\mathfrak{B})_s$ such that $L = h^{-1}(F)$. \dashv

We want to call a set of dependency structures ‘regular’, if it is recognizable in some dependency algebra \mathfrak{D} . In this case, since \mathfrak{D} is initial, we do not have a choice about the homomorphism h in the above definition: it is the uniquely determined evaluation homomorphism $\llbracket \cdot \rrbracket_{\mathfrak{B}}$. This leads to the following definition of regularity:

702 **Definition** Let \mathfrak{D} be a dependency algebra with signature Σ , and let $i \in \mathcal{S}_\Sigma$ be a sort. A set $L \subseteq \text{dom}(\mathfrak{D})_i$ is called *regular*, if there exists a finite Σ -algebra \mathfrak{B} and a set $F \subseteq \text{dom}(\mathfrak{B})_i$ such that $L = \llbracket F \rrbracket_{\mathfrak{B}}^{-1}$. \dashv

The pair $M = (\mathfrak{B}, F)$ is called a (deterministic, complete) *automaton* for L . We can understand the signature of \mathfrak{B} as an *input alphabet*, the domains $\text{dom}(\mathfrak{B})$ as sets of *states*, the (finitely many) possible combinations of input and output values for the composition operations of \mathfrak{B} as a *transition function*, and the set F as a set of *final states* of M . The behaviour of M can be described as follows: A *run* of M is a bottom-up traversal of a dependency structure D during which each node u gets labelled with a state $q \in \text{dom}(\mathfrak{B})$. The label for u is chosen conditional on both the local order at u (represented by an order annotation $\sigma \in \Sigma$), and the state labels at the children of u . More specifically, when M visits a node u that is annotated with a symbol σ , and the children of the node u have previously been labelled with states q_1, \dots, q_m , then the automaton labels u with the state $f_\sigma(q_1, \dots, q_m)$. The automaton M *recognizes* D if, at the end of the run, the root node of D is labelled with a state $q \in F$.

703 **Example** To illustrate the definition of regularity, we show that, for every dependency algebra \mathfrak{D} with signature Σ , every $k \in \mathcal{S}_\Sigma$, and every set $G \subseteq \{f_\sigma \mid \sigma \in \Sigma\}$, the set L of those structures of sort k that are composed using only operations from the set G is regular. We do so by constructing an automaton $M = (\mathfrak{B}, F)$ for L as follows. For each sort $i \in \mathcal{S}_\Sigma$, the state set $\text{dom}(\mathfrak{B})_i$ is the set $\{1, 0\}$. For each order annotation $\sigma: s_1 \times \dots \times s_m \rightarrow s$ in Σ and each tuple $\langle b_1, \dots, b_m \rangle$ of values, we put $f_\sigma(b_1, \dots, b_m) := (\bigwedge_{i=1}^m b_i) \wedge b$, where $b = 1$ if and only if $f_\sigma \in G$. As the set of final states F , we choose $\{1\} \subseteq \text{dom}(\mathfrak{B})_k$. For each dependency structure $D \in \text{dom}(\mathfrak{D})$, the evaluation of D in \mathfrak{B} returns 1 if and only if it was composed using only composition operations from the set G . Thus, the set $\llbracket F \rrbracket_{\mathfrak{B}}^{-1}$ contains exactly those dependency structures with this property that have sort k . \dashv

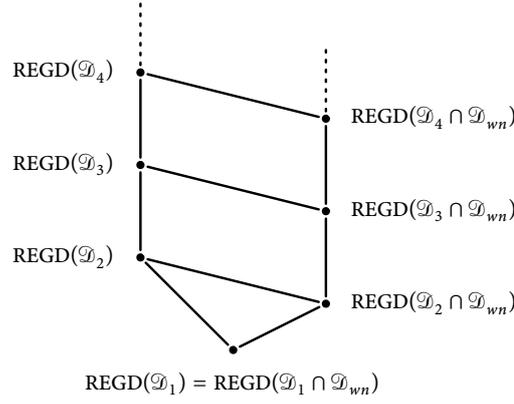


Figure 7.1: The hierarchy of regular dependency languages

We write REGD for the class of all regular dependency languages. For a given class \mathcal{D} of dependency structures, we write $\text{REGD}(\mathcal{D})$ for the class of all regular dependency languages that are subsets of \mathcal{D} .

7.1.2 | Elementary properties

We now review some of the elementary formal properties of regular dependency languages. All of these properties are immediate consequences of our definitions and general results about recognizable subsets.

Lemma The empty set is a regular dependency language. Furthermore, REGD is closed under union, intersection, difference, and inverse homomorphisms. 704

Proof. See e.g. Courcelle (1996), Proposition 4.6. 705

Lemma The following relations hold for all $k \in \mathbb{N}$:

- * $\text{REGD}(\mathcal{D}_k) \subsetneq \text{REGD}(\mathcal{D}_{k+1})$
- * $\text{REGD}(\mathcal{D}_k \cap \mathcal{D}_{wn}) \subsetneq \text{REGD}(\mathcal{D}_{k+1} \cap \mathcal{D}_{wn})$
- * $\text{REGD}(\mathcal{D}_1) = \text{REGD}(\mathcal{D}_1 \cap \mathcal{D}_{wn})$
- * $k \neq 1 \implies \text{REGD}(\mathcal{D}_k \cap \mathcal{D}_{wn}) \subsetneq \text{REGD}(\mathcal{D}_k)$ 705

Proof. Each of the classes of dependency structures mentioned in this lemma coincides with a specific sub-signature of the set Ω of all order annotations. All relations therefore can be reduced to the corresponding relations on the signatures. 705

This lemma is visualized in Figure 7.1. It shows that the structural restrictions imposed by the block-degree measure and the well-nestedness condition induce two infinite, related hierarchies of ever more expressive regular dependency languages. The only point where these two hierarchies coincide is the case $k = 1$.

Finite degree Since each regular dependency language is built using a finite set of composition operations (a finite signature), and since there is a direct correspondence between the type of a composition operation and the measures of out-degree and block-degree, no regular dependency language can be unbounded in either measure.

706 **Definition** Let $L \subseteq \mathcal{D}$ be a dependency language. We say that L is of *finite degree*, if there are numbers $m \in \mathbb{N}_0$, $k \in \mathbb{N}$ such that no structure in L has an out-degree of more than m or a block-degree of more than k . \dashv

707 **Lemma** Every regular dependency language is of finite degree. \dashv

The property that regular dependency languages are of finite out-degree separates them from dependency frameworks that allow an arbitrary number of children per node (see e.g. Duchier and Debusmann, 2001). The restriction to finite block-degree formalizes the rather informal notion of ‘limited cross-serial dependencies’ that is characteristic for mildly context-sensitive language (Joshi, 1985). At the same time, this restriction implies that regular dependency languages are not able to account for linguistic phenomena that require arbitrary degrees of non-projectivity, such as the phenomenon of scrambling in German subordinate clauses (Becker et al., 1992).

Connection with regular term languages Recall from Theorem 421 that every dependency algebra is isomorphic to its corresponding term algebra. Therefore the following lemma applies:

708 **Lemma** Let \mathcal{D} be a dependency algebra with signature Σ , and let $i \in \mathcal{S}_\Sigma$ be a sort. Then a set $L \subseteq \text{dom}(\mathcal{D})_i$ is regular if and only if $\text{term}(L)$ is T_Σ -recognizable. \dashv

The major benefit of this connection is that it allows us to study dependency languages using the tools and results of the well-developed formal theory of recognizable term languages, which is more widely known as the class of *regular term languages*. This is a very natural and robust class with many different characterizations: it is recognized by finite tree automata, generated by regular term grammars, and definable in monadic second-order logic (Gécseg and Steinby, 1997). For our purposes, the characterization in terms of grammars is the most convenient. It allows us to characterize regular dependency languages as the images of the term languages generated by regular term grammars over the signature Ω of order annotations.

7.1.3 | Regular term grammars

A *term grammar* specifies a rewriting system for terms over an alphabet of terminal and non-terminal symbols. In each step of the rewriting process, a non-terminal symbol is replaced by a term; this yields a new term. Regular term grammars are the generative correspondents of the algebraic automata from Definition 702: Essentially, they are

obtained by reading the composition functions of an automaton as a directed rewrite system—whenever an automaton defines $f_\sigma(q_1, \dots, q_m) = q$, the corresponding term grammar contains a rule $q \rightarrow f_\sigma(q_1, \dots, q_m)$; the states of the automaton become the non-terminal symbols of the grammar. Regular term grammars are distinguished from other term grammars by the restriction that non-terminal symbols may occur only at the leaves of a term, which implies that derivations correspond to sequences of simple substitution operations, just as in context-free grammars.¹ To formalize this restriction, we introduce the following notation:

Definition Let \mathcal{S} be a set of sorts. Let Σ be an \mathcal{S} -sorted set, and let A be an \mathcal{S} -indexed family of sets. The set of terms over Σ indexed by A , denoted by $T_\Sigma(A)$, is the \mathcal{S} -indexed set of all terms over $\Sigma \cup A$. 709

Regular term grammars are usually defined for single-sorted algebras (Gécseg and Steinby, 1997; Denecke and Wismath, 2001); here we adapt their definition to the many-sorted case. This extension is straightforward: instead of one non-terminal alphabet and set of productions, we need one such alphabet and set of productions per sort. The set of productions is set up such that a given non-terminal symbol can only be rewritten by a term of the same sort.

Definition Let \mathcal{S} be a finite set of sorts. A *regular term grammar* (over \mathcal{S}) is a construct $G = (N, \Sigma, S, P)$, where N is an \mathcal{S} -indexed family of *non-terminal alphabets*, Σ is an \mathcal{S} -sorted *terminal alphabet*, $S \in N$ is a distinguished *start symbol*, and $P \subseteq N \times T_\Sigma(N)$ is a finite, \mathcal{S} -indexed family of sets of *productions*. 710

We use indexed symbols $(N_G, \Sigma_G, S_G, P_G)$ to refer to the components of specific regular term grammars G . For a production $p = (A, t)$, we call A the *left-hand side* and t the *right-hand side* of p . Just as in conventional string grammars, we usually write $A \rightarrow t$ instead of (A, t) . The *derivation relation* associated to a regular term grammar $G = (N, \Sigma, S, P)$ is the binary relation \Rightarrow_G on $T_\Sigma(N)$ that is defined by the following inference rule:

$$\frac{t \in T_\Sigma(N) \quad t/u = A \quad (A \rightarrow t') \in P}{t \Rightarrow_G t[u \leftarrow t']}$$

Using this relation, the definition of the *term language* generated by G is completely analogous to the definition of the language generated by a string grammar—it is the set of all terms without non-terminals that can eventually be derived from the trivial term formed by the start symbol of the grammar: $L(G) = \{ t \in T_\Sigma \mid S \Rightarrow_G^* t \}$. Two grammars are *equivalent*, if they generate the same language. Notice that all terms in the language generated by a regular term grammar are of one and the same sort; this is the sort of the start symbol S .

¹ In fact, one can show that a language is context-free if and only if it is the frontier of a set of terms generated by a regular term grammar (Gécseg and Steinby, 1997, p. 33).

711 **Definition** A regular term grammar $G = (N, \Sigma, S, P)$ over \mathcal{S} is called *normalized*, if every production has the form $A \rightarrow \sigma(A_1, \dots, A_m)$, where $A \in N_s$, $\sigma: s_1 \times \dots \times s_m \rightarrow s$, and $A_i \in N_{s_i}$, for every $i \in [m]$ and some sort $s \in \mathcal{S}$. \dashv

712 **Lemma** For every regular term grammar, there exists an equivalent regular term grammar that is normalized. \dashv

Proof. A proof of this result can be found in Gécseg and Steinby (1997). The proof is a standard grammar transformation, as it is also known from context-free grammars: we delete rules of the form $A \rightarrow B$ and rules where the right-hand side is a term with a depth greater than 1, and replace them by new rules and non-terminal symbols that jointly simulate the old rules. \dashv

7.1.4 | Regular dependency grammars

We now define *regular dependency grammars* as regular term grammars that generate term languages over the signature Ω of order annotations. This restriction ensures that terms manipulated by regular dependency grammars can be interpreted as (segmented) dependency structures.

713 **Definition** Let $k \in \mathbb{N}$. A *regular dependency grammar* of degree k is a construct $G = (N, S, P)$, where N is a $[k]$ -indexed family of *non-terminal alphabets*, $S \in N$ is a distinguished *start symbol*, and $P \subseteq N \times T_{\Omega(k)}(N)$ is a finite, $[k]$ -indexed family of sets of productions. \dashv

For a given regular dependency grammar G , let Σ be the finite subset of order annotations that occurs in the productions of G ; then the construct (N, Σ, S, P) forms a regular term grammar. Based on this observation, we make free use of all terminology for regular term grammars when talking about regular dependency grammars. We will only work with regular dependency grammars in which $S \in N_1$. This restriction ensures that the languages generated by regular dependency grammars can be interpreted as sets of proper dependency structures.

714 **Definition** Let G be a regular dependency grammar. The *dependency language* generated by G is the set $\mathcal{D}(G) := \text{dep}(L(G))$. \dashv

715 **Example** To illustrate the definitions, we give two examples of regular dependency grammars. The dependency languages generated by these grammars mimic the verb-argument relations found in German and Dutch subordinate clauses, respectively (Huybregts, 1984; Shieber, 1985; Rambow, 1994): grammar G_1 generates structures with nested dependencies, grammar G_2 generates structures with crossing dependencies. Figure 7.2 shows terms generated by these grammars, and the corresponding structures.

7.1 REGULAR SETS OF DEPENDENCY STRUCTURES

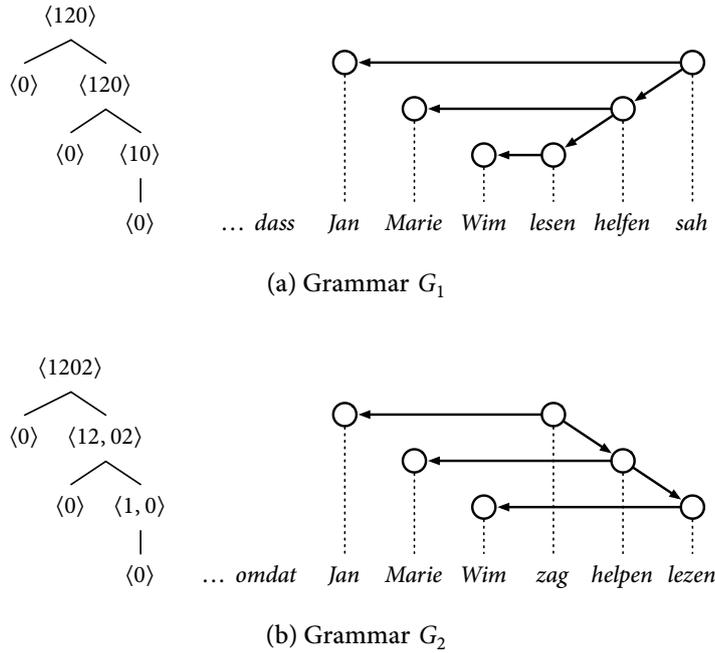


Figure 7.2: Terms and structures generated by two regular dependency grammars

GRAMMAR $G_1 := (N_1, S, P_1)$ (degree 1)

$$N_1 := \{1 \mapsto \{S, N, V\}\}$$

$$P_1 := \{S \rightarrow \langle 120 \rangle(N, V), V \rightarrow \langle 120 \rangle(N, V), V \rightarrow \langle 10 \rangle(N), N \rightarrow \langle 0 \rangle\}$$

GRAMMAR $G_2 := (N_2, S, P_2)$ (degree 2)

$$N_2 := \{1 \mapsto \{S, N\}, 2 \mapsto \{V\}\}$$

$$P_2 := \{S \rightarrow \langle 1202 \rangle(N, V), V \rightarrow \langle 12, 02 \rangle(N, V), V \rightarrow \langle 1, 0 \rangle(N), N \rightarrow \langle 0 \rangle\} \dashv$$

We now state the main result of this section:

Theorem A dependency language is regular if and only if it is generated by a regular dependency grammar. 716 \dashv

Proof. This is a direct consequence of the isomorphism between regular dependency languages and regular term languages (Lemma 708) and the standard result that a term language is regular if and only if it is generated by a regular term grammar. A proof of this result can be found in Denecke and Wismath (2001); it proceeds by translating every regular term grammar into an equivalent algebraic automaton and vice versa. The major difference between grammars and automata is that automata are complete

and deterministic (exactly one value per function-argument pairing), while grammars may be incomplete and indeterministic (more than one rule per non-terminal). These differences can be removed by grammar normalizations on the one hand, and a standard subset construction for automata on the other. \dashv

7.1.5 | Dependency languages and lexicalized grammars

With the concept of regularity at hand, we can now lift our results from the previous chapter to the level of languages. Given a class \mathcal{G} of grammars, let us write $\mathcal{D}(\mathcal{G})$ for the class of all dependency languages induced by grammars in \mathcal{G} .

717 **Theorem** The following statements hold for all $k \in \mathbb{N}$:

- (1) $\text{REGD}(\mathcal{D}_k) = \mathcal{D}(\text{LCFRS}(k))$
- (2) $\text{REGD}(\mathcal{D}_k \cap \mathcal{D}_{wn}) = \mathcal{D}(\text{CCFG}(k))$ \dashv

To put it into words: The dependency languages induced by the class of LCFRS with fan-out at most k are exactly the regular dependency languages over dependency structures with block-degree at most k . Similarly, the dependency languages induced by the class of CCFGs with rank at most k are exactly the regular dependency languages over well-nested dependency structures with block-degree at most k .

Proof. The proof falls into two parts:

- \supseteq Let G be a lexicalized LCFRS. The set of derivation trees of G forms a regular term language over some finite signature of concatenation functions. By Lemma 621, we can transform this language into an equivalent (modulo relabelling) term language L that only uses essential concatenation functions. Crucially, the elimination of non-essential functions uses bottom-up and top-down relabellings, which preserve regularity (Engelfriet, 1975); therefore, the transformed language L still is a regular term language, say $L \subseteq T_\Sigma$. We have furthermore seen (on page 74) how to define a bijective function $\text{relab}: \Sigma \rightarrow \Omega$ from the set of essential concatenation functions to the set of order annotations such that a derivation $t \in T_\Sigma$ induces the dependency structure $\text{dep}(\text{relab}(t))$, for all $t \in L$. Since the mapping relab is injective, we can translate L into a regular term language L' over some finite signature $\Delta \subseteq \Omega$, and hence, modulo decoding, into a regular dependency language.
- \subseteq Let G be a regular dependency grammar. We can construct an LCFRS G' such that the derivations of G' induce the dependency language generated by G by reversing the above relabelling on a per-rule basis.

In both directions, the relabelling maintains the signature restrictions: essential concatenation functions of fan-out k are translated into order annotations of degree k , and vice versa. The relabelling also maintains the well-nestedness restriction. \dashv

7.2 | Pumping lemmata

Since even infinite regular term languages (and hence: regular dependency languages) can be represented by finite grammars, these languages, very much like regular or context-free string languages, must have a periodic structure. In this section, we prove a number of novel pumping lemmata for regular term languages that make this observation precise. These lemmata provide the keys to our results about the growth of regular dependency languages (Section 7.3) and their string-generative capacity (Section 8.2).

7.2.1 | The pumping lemma for regular term languages

Recall the standard pumping lemma for context-free string languages:

For every context-free language $L \subseteq A^$, there is a number $p \in \mathbb{N}$ such that any string $z \in L$ of length at least p can be written as $z = uvwxy$ such that $1 \leq |vx|$, $|vwx| \leq p$, and $uv^nwx^n y \in L$, for every $n \in \mathbb{N}_0$.*

This result is usually proved using a combinatorial argument about the derivations of a grammar that generates the language L . Essentially the same argument can be used to show a pumping lemma for regular term grammars (see e.g. Proposition 5.2 in Gécseg and Steinby, 1997):

Lemma For every regular term language $L \subseteq T_\Sigma$, there is a number $p \in \mathbb{N}$ such that any term $t \in L$ of size at least p can be written as $t = c' \cdot c \cdot t'$ such that $1 \leq |c| \leq p$, $|c \cdot t'| \leq p$, and $c' \cdot c^n \cdot t' \in L$, for every $n \in \mathbb{N}_0$. 718

The number p in this lemma is called *pumping number*. The phrase ‘the term t can be written as $t = c' \cdot c \cdot t'$ ’ is an abbreviation for the formal assertion that ‘there exist contexts $c' \in C_\Sigma$ and $c \in C_\Sigma$ and a term $t' \in T_\Sigma$ such that $t = c' \cdot c \cdot t'$ ’. The contrapositive of the pumping lemma specifies a strategy for proofs that a certain term language is *not* regular:

To show that a term language $L \subseteq T_\Sigma$ is non-regular, show that, for all $p \in \mathbb{N}$, there exists a term $t \in L$ of size at least p such that for all partitions $t = c' \cdot c \cdot t'$ in which $1 \leq |c| \leq p$ and $|c \cdot t'| \leq p$, there is an $n \in \mathbb{N}_0$ such that $c' \cdot c^n \cdot t' \notin L$.

We can think of such a proof as a game against an imagined adversary. Our objective in this game is to prove that a given term language L is non-regular; our adversary’s objective is to foil us. The game consists of four alternating turns: In the first turn, ADVERSARY must choose a number $p \in \mathbb{N}$. In the second turn, we must respond to this choice by providing a term $t \in L$ of size at least p . In the third turn, ADVERSARY must choose a decomposition of t into fragments $c' \cdot c \cdot t'$ such that $1 \leq |c| \leq p$ and $|c \cdot t'| \leq p$. In the fourth turn, we must provide a number $n \in \mathbb{N}_0$ such that

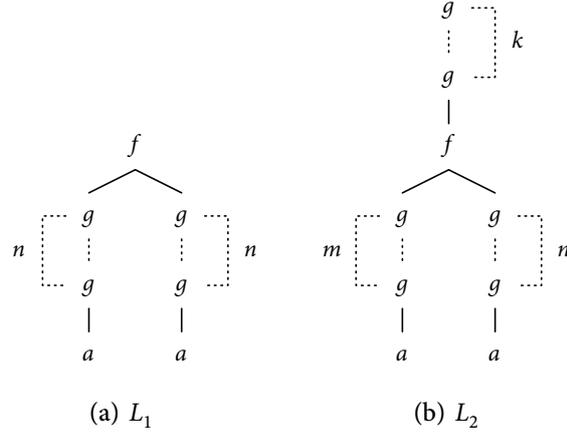


Figure 7.3: Two term languages that are not regular

$c' \cdot c^n \cdot t' \notin L$. If our last turn succeeds, we win the game; otherwise, ADVERSARY wins. The language L is non-regular, if we have a winning strategy for this game.

719 **Example** Consider the following term language (see Figure 7.3a for a schema):

$$L_1 := \{ f(g^n \cdot a, g^n \cdot a) \mid n \in \mathbb{N}_0 \}$$

We show that L_1 is non-regular by stating a winning strategy for the game associated with L_1 . Assume that ADVERSARY has chosen the number $p \in \mathbb{N}$. Then we can always win the game by responding with the term $t = f(g^p \cdot a, g^p \cdot a)$. It is clear that t is a valid term in L_1 , and that $|t| \geq p$. In whatever way ADVERSARY decomposes t into segments $c' \cdot c \cdot t'$, the term $s := c' \cdot c^2 \cdot t'$ does not belong to L_1 . In particular, if c lies below the ‘branching point’ of t (the node labelled with the symbol f), then the lengths of the two branches in s are out of sync. Thus we deduce that L_1 is non-regular. \dashv

Unfortunately, Lemma 718 sometimes is too blunt a tool to show the non-regularity of a term language, as the following example demonstrates.

720 **Example** Consider the following term language (see Figure 7.3b for a schema):

$$L_2 := \{ g^k \cdot f(g^m \cdot a, g^n \cdot a) \mid k, m, n \in \mathbb{N}_0 \wedge (k \neq 0 \implies m = n) \}$$

The terms in L_2 can be classified into two categories: A term in which $k \neq 0$ can be regarded as a term from L_1 that is prefixed with a non-empty chain of nodes labelled with the symbol g . In a term in which $k = 0$ however, the lengths of the two branches are not required to be synchronized as in L_1 .

The language L_2 is non-regular (as we will show at the end of this section), but it is impossible to prove this using Lemma 718. To see this, assume that ADVERSARY has chosen the number $p \in \mathbb{N}$ in the game associated to L_2 . Then whatever term $t \in L_2$

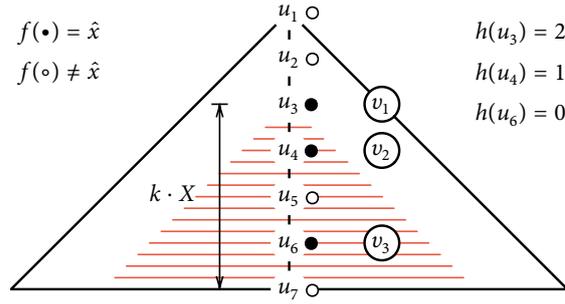


Figure 7.4: An example for the situation in the proof of Lemma 721

we provide, ADVERSARY can win the game: if $k = 0$, he can decompose t as $c' \cdot c \cdot t'$ such that c lies somewhere below the branching point of t ; if $k \neq 0$, he can choose c such that it lies on the prefix above the branching point. In both cases, both deleting and repeating c yield only valid terms in L_2 . \dashv

7.2.2 | Ogden's lemma for regular term languages

Our goal in this section is to develop a pumping lemma for regular term languages that is more powerful than Lemma 718 and allows us to prove (among other things) that the language L_2 in the example above is non-regular. This lemma (Lemma 723) can be seen as the natural correspondent of Ogden's lemma for context-free languages (Ogden, 1968). While Lemma 718 merely asserts that a pumpable context does exist *somewhere* in the term, our lemma gives us a way to delimit the exact region. This added flexibility proves to be very useful in the application of the lemma.

Lemma Let Σ be a finite signature. For every term language $L \subseteq T_\Sigma$, every non-empty, finite set X , and every $k \in \mathbb{N}$, there is a number $p(X, k) \in \mathbb{N}$ such that for any term $t \in L$ with size $|t| \geq p(X, k)$ and any function $f: \text{nod}(t) \rightarrow X$, there exist a strictly increasing sequence \vec{v} of $k + 1$ nodes of t and an object $\hat{x} \in X$ such that

721

- (1) $\forall i \in [k + 1]. f(v_i) = \hat{x}$;
- (2) $\forall i \in [k + 1]. |t/v_i| \leq p(X, k)$; and
- (3) $\forall i \in [k]. 1 \leq |t/v_i| - |t/v_{i+1}| \leq p(X, k)$. \dashv

Proof. The proof is based on the elementary observation that the size of a term with rank at most m and depth at most d is bounded by the value $g(m, d) := \sum_{i=0}^d m^i$. Put $m_\Sigma := \max_{\sigma \in \Sigma} \text{rank}_\Sigma(\sigma)$; this value is an upper bound for the rank of the terms in L . We will show that we can choose $p(X, k) := g(m_\Sigma, k \cdot |X|)$.

- * Let $t \in L$ be a term with $|t| \geq p(X, k)$; in this case, t has a depth of at least $k \cdot |X|$. Let \vec{u} be a maximal path in t , and for every node u on \vec{u} , let $h(u)$ denote the number of those nodes u' on \vec{u} that succeed u and for which $f(u') = f(u)$ holds

7 REGULAR DEPENDENCY LANGUAGES

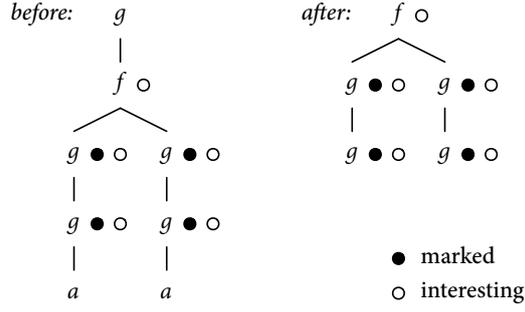


Figure 7.5: A term before and after the transformation from the proof of Lemma 722

(see Figure 7.4). Since $|\vec{u}| > k \cdot |X|$, there is at least one element $x \in X$ and at least $k + 1$ nodes u on \vec{u} for which $f(u) = x$, and hence, there is at least one node u on \vec{u} for which $h(u) = k$. Without loss of generality, assume that u is the greatest node with this property, and let \vec{v} be the sub-sequence of \vec{u} that is formed by u and the k nodes u' on \vec{u} that succeed u and for which $f(u') = f(u)$ holds.

- * We now check that the sequence \vec{v} satisfies the required properties. Condition (1) is satisfied with $\hat{x} := f(u)$, and the first half of condition (3) is satisfied by construction. For the remaining conditions, consider the node v_1 . This node is succeeded on \vec{u} by exactly k nodes v' with $f(v') = f(v_1)$, and at most $k \cdot (|X| - 1)$ nodes v' with $f(v') \neq f(v_1)$. Since the path \vec{u} is maximal in t , the depths of the nodes on \vec{u} coincide with their depths in t . We therefore have $d(v_1) \leq k + k \cdot (|X| - 1) = k \cdot |X|$. Using the relation between depth and size, we then see that $|t/v_1| \leq p(X, k)$. Since every node on \vec{v} is at most as deep as v_1 , the claim is proved. \dashv

722 Lemma For every normalized regular term grammar G and every $k \in \mathbb{N}$, there is a number $p(k) \in \mathbb{N}$ such that for any term $t \in L(G)$ in which at least $p(k)$ nodes are marked as distinguished, there exists a non-terminal $A \in N_G$ such that every derivation of t can be written as

$$S \Rightarrow_G^* c' \cdot A \Rightarrow_G^* c' \cdot c_1 \cdot A \Rightarrow_G^* \cdots \Rightarrow_G^* c' \cdot c_1 \cdots c_k \cdot A \Rightarrow_G^* c' \cdot c_1 \cdots c_k \cdot t' = t$$

such that at most $p(k)$ nodes in t' and at least one and at most $p(k)$ nodes in c_i are marked, for every index $i \in [k]$.

Proof. Let G be a normalized regular term grammar, and let $k \in \mathbb{N}$. We will show that we can choose $p(k) := p(|N_G|, k)$, where the latter value is the number asserted by Lemma 721. We start with some terminology. Let $t \in L(G)$ be a term. We call a node $u \in \text{nod}(t)$ *interesting*, if either it is marked, or at least two of its children dominate a marked node. We call t *interesting*, if at least one of its nodes is interesting.

- * We define a partial function τ on the term language $L(G)$ as follows. For an interesting term $t \in L(G)$, construct the term $\tau(t)$ by exhaustive application of the following transformations to nodes u (see Figure 7.5): if u is unmarked and does not have a child that dominates a marked node, delete the subterm rooted at u ; if u is unmarked and has exactly one child v that dominates a marked node, replace u with the subterm rooted at v . We can verify the following properties:
 - (a) The dominance relation of $\tau(t)$ is isomorphic to the restriction of the dominance relation of t to the interesting nodes of t .
 - (b) The rank of the term $\tau(t)$ is bounded by the rank of t . This ensures that the image of $L(G)$ under τ is a term language over a finite signature.
- * Let $t \in L(G)$ be a term in which at least $p(k)$ nodes are marked as distinguished, and let d be a derivation of t . Define a function $f: \text{nod}(t) \rightarrow N_G$ such that, for every node $u \in \text{nod}(t)$, the symbol $f(u)$ is the left-hand side of the production that was applied to introduce u into t . This function is well-defined because G is normalized, and there thus is a one-to-one correspondence between the nodes in t and the rewriting steps in d . The restriction of f to the interesting nodes in t defines a function $g: \text{nod}(\tau(t)) \rightarrow N_G$. Then by Lemma 7.21, there exists a strictly increasing sequence \vec{v} of $k+1$ nodes in $\tau(t)$ and a non-terminal $A \in N_G$ such that
 - * $\forall i \in [k+1]. g(v_i) = A$;
 - * $\forall i \in [k+1]. |t/v_i| \leq p(k)$; and
 - * $\forall i \in [k]. 1 \leq |t/v_i| - |t/v_{i+1}| \leq p(k)$.

Because of the isomorphism property (a) above, this implies that there exists a strictly increasing sequence \vec{w} of $k+1$ nodes in t such that $f(w_i) = A$ holds for all $i \in [k+1]$. We use this sequence to identify contexts and subterms in t as follows: Let c' be the context identified by the root node of t and the node w_i . For every $i \in [k]$, let c_i be the context identified by the nodes w_i and w_{i+1} ; then by the isomorphism property (a), this context contains at least one and at most $p(k)$ interesting nodes. Finally, put $t' := t/v_{k+1}$; then again by isomorphism, this term contains at most $p(k)$ interesting nodes.

- * We now check the required properties. To see that both every context c_i ($i \in [k]$) and the term t' contains at most $p(k)$ marked nodes, it suffices to notice that the number of interesting nodes is an upper bound for the number of marked nodes. To see that every context c_i ($i \in [k]$) contains at least one marked node, let u be one of the interesting nodes in c_i , and assume that u is not itself marked. In this case, at least two of its children dominate a marked node. At most one of these marked nodes is dominated by the hole of c_i ; the other marked node is a node in c_i . Finally, since $f(w) = A$ holds for all nodes w in the sequence \vec{w} , and because of the one-to-one correspondence between the nodes in t and the rewriting steps in d , the derivation d can indeed be linearized into the required form. –1

7 REGULAR DEPENDENCY LANGUAGES

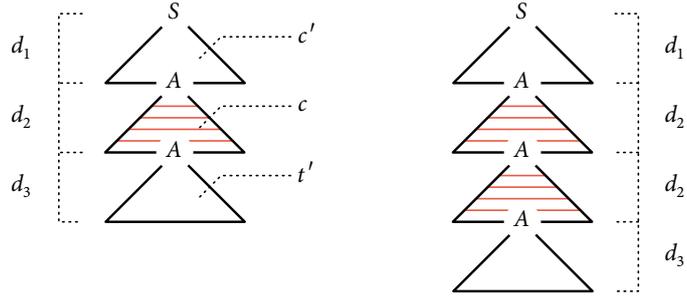


Figure 7.6: Surgery on derivation trees

Instantiating the previous lemma, we obtain our main result:

723 **Lemma** For every regular term language $L \subseteq T_\Sigma$, there is a number $p \in \mathbb{N}$ such that any term $t \in L$ in which at least p nodes are marked as distinguished can be written as $t = c' \cdot c \cdot t'$ such that at least one and at most p nodes in c are marked, at most p nodes in $c \cdot t'$ are marked, and $c' \cdot c^n \cdot t' \in L$, for all $n \in \mathbb{N}_0$. \dashv

Note that, in the special case where all nodes of a term are marked as distinguished, this lemma is equivalent to Lemma 718.

Proof. Let $L \subseteq T_\Sigma$ be a regular term language, and let G be a normalized regular tree grammar with $L(G) = L$. We will show that we can choose $p := p(1)$, where the latter value is the number asserted by Lemma 722. Let $t \in L(G)$ be a term in which at least p nodes are marked as distinguished. Then by Lemma 722, every derivation of t can be written as $S \Rightarrow_G^* c' \cdot A \Rightarrow_G^* c' \cdot c \cdot A \Rightarrow_G^* c' \cdot c \cdot t' = t$ such that the context c contains at least one and at most p marked nodes and the term t' contains at most p marked nodes. This means that the derivation can be partitioned into three sub-derivations

$$S \Rightarrow_G^* c' \cdot A \ (d_1), \quad A \Rightarrow_G^* c \cdot A \ (d_2), \quad \text{and} \quad A \Rightarrow_G^* t' \ (d_3)$$

(see the schema in the left half of Figure 7.6). Iterating the second of these sub-derivations n times (where n may be zero), we obtain a new valid derivation (see the schema in the right half of Figure 7.6). \dashv

The contrapositive of Lemma 723 corresponds to the following modified game for term languages L : In the first turn, ADVERSARY must choose a number $p \in \mathbb{N}$. In the second turn, we choose a term $t \in L$ and mark at least p nodes in t . In the third turn, ADVERSARY chooses a decomposition $t = c' \cdot c \cdot t'$ such that at least one and at most p nodes in c and at most p nodes in $c \cdot t'$ are marked. In the fourth and final turn, we must choose a number $n \in \mathbb{N}_0$ such that $c' \cdot c^n \cdot t' \in L$.

724 **Example (continued)** Using Lemma 723, we are now ready to show that the lan-

guage L_2 from the example above is non-regular: we show that we have a winning strategy for the game associated with L_2 . Assume that ADVERSARY has chosen the number $p \in \mathbb{N}$. Then we can always win the game by responding with the term $t = g^1(f(g^p \cdot a, g^p \cdot a))$ and marking all nodes in the subterm rooted at the branching point as distinguished. In whatever way ADVERSARY decomposes t into segments $c' \cdot c \cdot t'$ such that c contains at least one marked node, the context c contains at least one node below the branching point. Then, the term $s := c' \cdot c^2 \cdot t'$ does not belong to L_2 . \dashv

7.3 | Constant growth

We have claimed (in our discussion of Lemma 707) that the block-degree restriction inherent to regular dependency languages formalizes the notion of ‘limited cross-serial dependencies’ that is characteristic for the class of mildly context-sensitive languages. In this section, we show that regular dependency languages also have another characteristic property of this class, *constant growth*. This property is usually attributed to string languages (Joshi, 1985); here, we define it for term languages.

7.3.1 | Constant growth and semilinearity

Informally, a language has the constant growth property, if there are no arbitrarily long leaps in its size progression. More formally, let L be a term language, and let \vec{n} be the sequence of distinct sizes of terms in L , sorted in ascending order. If L is of constant growth, then adjacent elements of \vec{n} differ by at most a constant (see Kracht, 2003, Definition 5.1).

Definition A term language $L \subseteq T_\Sigma$ is of *constant growth*, if either L is finite, or there is a number $c \in \mathbb{N}$ such that for each term $t \in L$, there exists a term $t' \in L$ such that $|t| < |t'| \leq |t| + c$. \dashv 725

Example We look at an example for a term language that does *not* have the constant growth property. Let Σ be a signature, and let L be the set of all complete binary terms over Σ . Given a term $t_i \in L$ with size $|t_i|$, a ‘next larger’ term $t_{i+1} \in L$ is obtained by adding two children to every leaf in t_i . This procedure results in a linear size progression: we see that $|t_{i+1}| = 2 \cdot |t_i| + 1$. In particular, there is no number c such that $|t_{i+1}| \leq |t_i| + c$ holds for all indices $i \in \mathbb{N}$. \dashv 726

Constant growth is closely related to a property known as *semilinearity* (Parikh, 1966). This property is concerned with the interpretation of the elements of a language as multisets of labels. For the following definition, given a term $t \in T_\Sigma$ and a symbol $\sigma \in \Sigma$, we write $\#_\sigma(t)$ for the number of occurrences of σ in t .

727 **Definition** Let Σ be a signature, and put $n := |\Sigma|$. We fix an (arbitrary) order on Σ and write σ_i for the i th symbol with respect to this order, for $i \in [n]$. The *Parikh mapping* for terms over Σ (with respect to this order) is the function $\psi_\Sigma: T_\Sigma \rightarrow \mathbb{N}_0^n$ defined as

$$\psi_\Sigma(t) := \langle \#_{\sigma_1}(t), \dots, \#_{\sigma_n}(t) \rangle.$$

We extend ψ_Σ to languages $L \subseteq T_\Sigma$ by putting $\psi_\Sigma(L) := \{ \psi_\Sigma(t) \mid t \in L \}$. \dashv

The order on Σ with respect to which a Parikh mapping is defined is convenient for our formal treatment of semilinearity, but irrelevant for our results. Therefore, we refer to *the* Parikh mapping for terms over Σ .

The Parikh mapping reduces a term to the multiset of its labels. A measure that is preserved under this reduction is the size of a term. More specifically, define the following norm on \mathbb{N}_0^n : $\|\vec{x}\| := \sum_{i=1}^n x_i$. Then for all terms $t \in T_\Sigma$, it holds that $\|\psi_\Sigma(t)\| = |t|$. The relevance of this observation is that it allows us to recast constant growth as a property of the image of a term language under its Parikh mapping.

728 **Lemma** A term language $L \subseteq T_\Sigma$ is of constant growth if and only if either $\psi_\Sigma(L)$ is finite, or there is a number $c \in \mathbb{N}$ such that for each term $t \in L$, there exists a term $t' \in L$ such that $\|\psi_\Sigma(t)\| < \|\psi_\Sigma(t')\| \leq \|\psi_\Sigma(t)\| + c$. \dashv

We now give a formal definition of semilinearity. To do so, we equip each set \mathbb{N}_0^n with two operations: component-wise addition of two vectors ($\vec{x} + \vec{y}$), and multiplication of a vector by a scalar $a \in \mathbb{N}_0$ ($a \cdot \vec{x}$).

729 **Definition** Let $n \in \mathbb{N}$. A set $S \subseteq \mathbb{N}_0^n$ is called *linear*, if there exists a vector $\vec{x}_0 \in \mathbb{N}_0^n$, a number $k \in \mathbb{N}_0$, and an indexed set $\langle \vec{x}_i \in \mathbb{N}_0^n \mid i \in [k] \rangle$ of vectors such that

$$S = \left\{ \vec{x}_0 + \sum_{i=1}^k c_i \cdot \vec{x}_i \mid c_i \in \mathbb{N}_0 \right\}.$$

A set is called *semilinear*, if it is a finite union of linear sets. A language $L \subseteq T_\Sigma$ is called *linear (semilinear)*, if $\psi_\Sigma(L)$ is a linear (semilinear) set of vectors. \dashv

Each element of a semilinear language is the outcome of one of a finite number of generative processes. Such a process is specified by a single ‘base structure’ and a finite set of ‘additives’. Its outcome is the set of all structures that can be obtained by combining the base structure with any number (including zero) of one or more additives. In this way, semilinearity is closely related to pumpability.

730 **Lemma** Each semilinear term language has the constant growth property. \dashv

Proof. Let $L \subseteq T_\Sigma$ be a term language. Put $n := |\Sigma|$, and $P := \psi_\Sigma(L)$. We show that, if P is linear or semilinear, then it satisfies the conditions of Lemma 728.

- * Assume that P is linear. In this case, there exists a vector $\vec{x}_0 \in \mathbb{N}_0^n$, a number $k \in \mathbb{N}_0$, and an indexed set $\langle \vec{x}_i \in \mathbb{N}_0^n \mid i \in [k] \rangle$ of base vectors such that

$$P = \left\{ \vec{x}_0 + \sum_{i=1}^k a_i \cdot \vec{x}_i \mid a_i \in \mathbb{N}_0 \right\}.$$

A vector $\vec{x} \in \mathbb{N}_0^n$ is called *null*, if $\|\vec{x}\| = 0$. Distinguish two cases: If all base vectors are null, then P is finite. Otherwise, let \vec{x} be a base vector that is not null and for which $\|\vec{x}\|$ is minimal among all non-null base vectors. Put $c := \|\vec{x}\|$, and let $\vec{y} \in P$. Since P is linear, the vector $\vec{z} := \vec{y} + \vec{x}$ is an element of P . Since \vec{x} is not null, we have $\|\vec{y}\| < \|\vec{z}\| \leq \|\vec{y}\| + c$.

- * Now, assume that P is semilinear. In this case, there exists a number $m \in \mathbb{N}$ and an indexed family $\langle P_i \mid i \in [m] \rangle$ of linear sets such that $P = \bigcup_{i \in [m]} P_i$. Assume that P is non-finite. For every $i \in [m]$ for which the set P_i is non-finite, let c_i be the number constructed in the previous item. Put $c := \max_{i \in [m]} c_i$, and let $\vec{y} \in P$. Then there exists an index $i \in [m]$ such that $\vec{y} \in P_i$, and by the previous item, there exists a vector $\vec{z} \in P_i \subseteq P$ such that $\|\vec{y}\| < \|\vec{z}\| \leq \|\vec{y}\| + c_i \leq \|\vec{y}\| + c$.

Using Lemma 728, we conclude that L has the constant-growth property. ←

7.3.2 | Regular term languages are semilinear

We now show that regular term languages are semilinear. Semilinearity of a language is routinely proven by providing an encoding of that language into a context-free language with the same Parikh image, and referring to Parikh's theorem (Parikh, 1966):

Proposition Every context-free language is semilinear. ← 731

Unfortunately, the standard proof of this theorem is rather opaque. In particular, it does not elucidate the close connection between semilinearity and pumpability. Therefore, we give a direct proof of the semilinearity of regular term languages, following a similar proof for context-free languages (Goldstine, 1977).

Theorem Every regular term language is semilinear. ← 732

Proof. Let L be a regular term language, and let G be a normalized regular term grammar with $L(G) = L$. For each set $M \subseteq N_G$ of non-terminals that contains S_G , let $L_M \subseteq L(G)$ be the subset of $L(G)$ that consists of all terms $t \in L(G)$ for which there is at least one derivation $S_G \Rightarrow_G^* t$ that uses exactly the non-terminals in M . Since there are only finitely many such sets L_M , and since their union is $L(G)$, it suffices to show that every set L_M is semilinear. Therefore, let us fix a set $M \subseteq N_G$, and put $m := |M|$ and $p := p(m)$, where the latter value is the constant from Lemma 722. We write \Rightarrow for the restriction of the derivation relation \Rightarrow_G to derivations that use only rules of

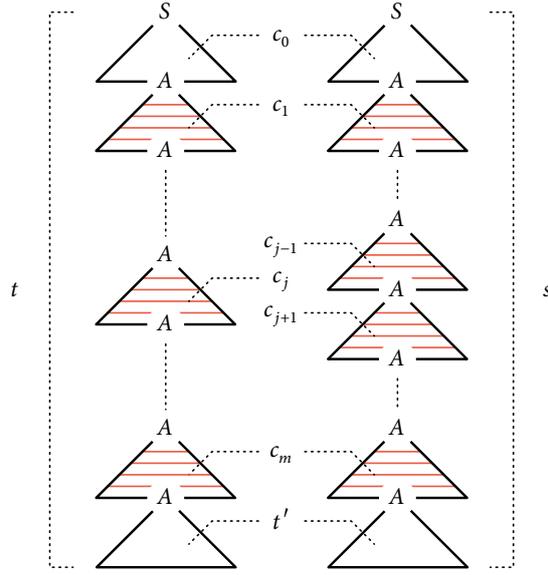


Figure 7.7: Semilinearity

the form $A \rightarrow t$, where $A \in M$ and $t \in T_\Sigma(M)$. By the definition of L_M , it then holds that $t \in L_M$ if and only if $S \Rightarrow^* t$. Put

$$\begin{aligned} T &:= \{t \in L_M \mid |t| < p\}, \\ C &:= \{c \in C_\Sigma \mid 1 \leq |c| \leq p \wedge \exists A \in M. A \Rightarrow^* c \cdot A\}, \quad \text{and} \\ X &:= \left\{ \psi_\Sigma(t) + \sum_{c \in C} a_c \cdot \psi_\Sigma(c) \mid t \in T \wedge a_c \in \mathbb{N}_0 \right\}. \end{aligned}$$

We start by noticing that the set X is semilinear: it is a finite union of linear sets, one for each term $t \in T$. To prove that the set L_M is semilinear, we show that $\psi_\Sigma(L_M) = X$.

⊆ Let $t \in L_M$ be a term. We show that $\psi_\Sigma(t) \in X$ by induction on the size $|t|$ of t . First assume that $|t| < p$. In this case, we see that $t \in T$, and $\psi_\Sigma(t) \in X$ by the definition of X . Now assume that $|t| \geq p$. In this case, if we mark all nodes in t as distinguished, then by Lemma 722, each derivation $S \Rightarrow^* t$ can be written as

$$S \Rightarrow^* c_0 \cdot A \Rightarrow^* c_0 \cdot c_1 \cdot A \Rightarrow^* \cdots \Rightarrow^* c_0 \cdot c_1 \cdots c_m \cdot A \Rightarrow^* c_0 \cdot c_1 \cdots c_m \cdot t',$$

where $A \in M$, $c_0 \in C_\Sigma$, $c_i \in C$ for all $i \in [m]$, and $t' \in T_\Sigma$. (See the left half of Figure 7.7.) Let us write d_i for the sub-derivation $A \Rightarrow^* c_i \cdot A$, and let $M' \subseteq M$ be the set of those non-terminals in $M - \{A\}$ that are used in some sub-derivation d_i , $i \in [m]$. For each $B \in M'$, choose some $i \in [m]$ such that B is used in d_i . Then, since $|M'| < m$, some $j \in [m]$ is not chosen at all. Therefore, if the corresponding sub-derivation d_j is deleted, every non-terminal in M (including A) is still present in the resulting derivation. In this way, we obtain a new valid derivation for a term

$s \in L_M$ with $|s| < |t|$. (See the right half of Figure 7.7.) By the induction hypothesis, we may assume that $\psi_\Sigma(s) \in X$. We see that $\psi_\Sigma(t) = \psi_\Sigma(s) + \psi_\Sigma(c_j)$, and so, $\psi_\Sigma(t) \in L_M$. Thus, in all cases, we have shown that $\psi_\Sigma(L_M) \subseteq X$.

- ⊇ Let $\vec{x} \in X$ be a vector. By the definition of X , there exists a term $t \in T$ and an indexed set $\langle a_c \in \mathbb{N}_0 \mid c \in C \rangle$ such that $\vec{x} = \psi_\Sigma(t) + \sum_{c \in C} a_c \cdot \psi_\Sigma(c)$. We show that there exists a term $s \in L_M$ with $\psi_\Sigma(s) = \vec{x}$ by induction on $n := \sum_{c \in C} a_c$. First, assume that $n = 0$. In this case, we have $\vec{x} = \psi_\Sigma(t)$, and since $t \in L_M$, we deduce that $\psi_\Sigma(t) \in X$. Now, assume that $n > 0$. In this case, there exists a context $c \in C$ and a vector $\vec{x}' \in X$ such that $\vec{x} = \vec{x}' + \psi_\Sigma(c)$, and by the induction hypothesis, we may assume that there exists a term $t' \in L_M$ with $\vec{x}' = \psi_\Sigma(t')$. From the definition of C , we see that there is a non-terminal $A \in M$ and a derivation $A \Rightarrow^* c \cdot A$. Since the derivation $S \Rightarrow^* t'$ uses every non-terminal $B \in M$ (including A), it can be written as $S \Rightarrow^* c' \cdot A \Rightarrow^* c' \cdot t'' = t'$, for some context $c' \in C_\Sigma$ and term $t'' \in T_\Sigma(M)$. In particular, we have $A \Rightarrow^* t''$. Plugging everything together, we can construct a valid derivation for a new term $s \in L_M$:

$$S \Rightarrow^* c' \cdot A \Rightarrow^* c' \cdot c \cdot A \Rightarrow^* c' \cdot c \cdot t'' = s.$$

We see that $\vec{x} = \vec{x}' + \psi_\Sigma(c) = \psi_\Sigma(t') + \psi_\Sigma(c) = \psi_\Sigma(s)$. Thus, in all cases, we have shown that $X \subseteq \psi_\Sigma(L_M)$. \dashv

Corollary Every regular term language has the constant-growth property. \dashv 733

As an immediate consequence of this corollary and the isomorphism between regular dependency languages and regular term languages (Lemma 708), we obtain the main result of this section:

Theorem Every regular dependency language is of constant growth. \dashv 734

7.3.3 | Related work

⚡ To claim that every formal language that adequately models natural language should have the constant growth property is not claiming very much: Kracht (2003) remarks that ‘it seems that for every natural language $[L]$ there is a number d_L such that for every $n \geq d_L$ there is a string of length n in L ’. Semilinearity, on the other hand, may be too strong a restriction to impose on mathematical models of natural language: Michaelis and Kracht (1997) show that an infinite progression of ‘case stacking’ in Old Georgian² would mean that this language is not semilinear. However, since there are no speakers of Old Georgian, there is no possibility to test whether this theoretical progression is actually possible.

² Old Georgian is an extinct Caucasian language that was spoken roughly between the 4th and 11th century AD. It has a rich literary tradition.

7 REGULAR DEPENDENCY LANGUAGES

✧ The class of semilinear subsets of \mathbb{N}_0^n is interesting in its own right. Among other things, it is closed under union, intersection, and complement. More generally, Ginsburg and Spanier (1966) show that a subset of \mathbb{N}_0^n is semilinear if and only if it is definable in Presburger arithmetic³. The class of languages with semilinear Parikh images forms an abstract family of languages, except that it is not closed under intersection with regular string languages (Kracht, 2003, Theorem 2.93).

³ Presburger arithmetic is the first-order theory of the natural numbers with addition. It was named in honour of Mojżesz Presburger (1904–1943), who proved its decidability in 1929.

CHAPTER 8 | Generative capacity and parsing complexity

In this chapter, we complete our study of regular dependency languages by investigating their string-generative capacity and parsing complexity. Specifically, we study the connection between these two measures and the structural constraints discussed in the first part of this dissertation.

Structure of the chapter. We start by explaining how regular dependency grammars can be extended to generators of sets of strings (Section 8.1). We then show that, for the string languages generated by these extended grammars, the block-degree measure induces an infinite hierarchy of expressivity, and that the well-nestedness restriction properly decreases expressivity on nearly all levels of this hierarchy (Section 8.2). Finally, we discuss the complexity of the parsing problem of the string languages generated by regular dependency grammars. In particular, we show that the well-nestedness condition can make the change between tractable and intractable parsing (Section 8.3).

8.1 | Projection of string languages

Up to this point, dependency structures were defined solely in terms of their governance and precedence relations. However, for many practical applications we are interested in *labelled* structures, where apart from the nodes and the edges, we also have ways to encode non-structural information such as word forms, grammatical functions, or edge probabilities. In this section, we extend our notion of dependency structures and dependency languages to accommodate such information. In particular, we show how dependency grammars can be understood as generators of string languages.

8.1.1 | Labelled dependency structures

The extension to labelled structures is straightforward:

Definition Let A be some alphabet. An A -labelled dependency structure is a pair (D, lab) , where D is a dependency structure, and $\text{lab}: \text{dom}(D) \rightarrow A$ is a total function on the nodes of D , called the *labelling function*. 801

Just as unlabelled dependency structures can be represented as terms over the alphabet Ω of order-annotations, A -labelled dependency structures can be represented as terms over the product alphabet $\langle \Omega, A \rangle$ in which each constructor $\langle \omega, a \rangle$ inherits the type of ω . For terms over this alphabet, we can extend the function dep in the natural way: the first component of a term constructor $\langle \omega, a \rangle$ carries the information

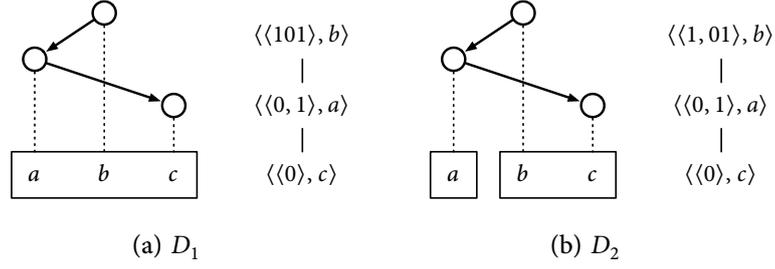


Figure 8.1: Two labelled dependency structures and their terms

about the dependency structure as such, the second component determines the label for the root node of the structure. In this way, each term over the signature $\langle\Omega, A\rangle$ denotes an A -labelled dependency structure. It is straightforward to extend our notion of dependency algebra accordingly. We can also define a string semantics for labelled dependency structures as follows. Recall that we use the notation $i\#j$ to refer to the j th occurrence of a symbol i in an order annotation.

802 **Definition** Let $\Sigma \subseteq \Omega$ be a finite set of order annotations, and let A be an alphabet. The *string algebra* over Σ and A is the $\langle\Sigma, A\rangle$ -algebra in which $\text{dom}(\mathfrak{A})_i = (A^+)^i$, for every $1 \leq i \leq \text{deg}(\Sigma)$, and

$$f_{\langle\omega, a\rangle}(\vec{\alpha}_1, \dots, \vec{\alpha}_m) = \omega[0 \leftarrow a][i\#j \leftarrow \alpha_{i,j} \mid i \in [m] \wedge j \in [k_i]],$$

for each constructor $\langle\omega, a\rangle : k_1 \times \dots \times k_m \rightarrow k$ in $\langle\Omega, A\rangle$. –

Let $\langle\Sigma, A\rangle$ be some finite signature, where $\Sigma \subseteq \Omega$. Given a term d over this signature, we write $\llbracket d \rrbracket_{\mathfrak{S}}$ for the evaluation of d in the string algebra over Σ and A and say that the labelled dependency structure that is denoted by d *projects* $\llbracket d \rrbracket_{\mathfrak{S}}$. Notice that, if d denotes a dependency structure of sort k , then the projection of d is a k -tuple of (non-empty) strings over the alphabet A . For the case $k = 1$, we identify the set of one-component tuples of strings with the set of strings.

803 **Example** Figure 8.1 shows two examples for labelled (segmented) dependency structures and their corresponding terms. Note that, in pictures of labelled structures, we annotate labels at the end of the corresponding projection lines. –

8.1.2 | String-generating regular dependency grammars

With our algebraic framework in mind, it is straightforward to extend regular dependency grammars into generators of string languages (via projection). The only thing that we need to add to the existing signature is the alphabet of labels, now called (*surface*) *terminal symbols*.

Definition Let $k \in \mathbb{N}$. A *string-generating regular dependency grammar* of degree k is a construct $G = (N, T, S, P)$, where N is a $[k]$ -indexed family of *non-terminal alphabets*, T is an alphabet of *terminal symbols*, $S \in N$ is a distinguished *start symbol*, and $P \subseteq N \times T_{\langle \Omega(k), A \rangle}(N)$ is a k -indexed family of finite sets of productions. \dashv

804

The derivation relation and the notion of the dependency language generated by a grammar are defined as usual, except that we are now dealing with labelled structures. The string algebra corresponding to a string-generating regular dependency grammar G is the string algebra over Σ_G and T_G , where Σ_G is the collection of those order annotations $\omega \in \Omega(k)$ that occur in the productions of G .

Definition Let G be a string-generating regular dependency grammar. The *string language* projected by G is defined as $\llbracket G \rrbracket_{\mathfrak{S}} := \{ \llbracket d \rrbracket_{\mathfrak{S}} \mid d \in L(G) \}$. \dashv

805

Example We give examples for two regular dependency grammars that generate the string language $\{ a^n b^n \mid n \in \mathbb{N} \}$. The dependency structures generated by the first grammar are projective (‘nested dependencies’), the structures generated by the second grammar may have block-degree 2 (‘cross-serial dependencies’). Both grammars use the same terminal-alphabet $\{a, b\}$. We only state the productions of the grammars; sample terms and generated dependency structures are shown in Figure 8.2.

806

G_1 (projective dependency structures):

$$S \rightarrow \langle \langle 012 \rangle, a \rangle (S, \langle \langle 0 \rangle, b \rangle) \qquad S \rightarrow \langle \langle 01 \rangle, a \rangle (\langle \langle 0 \rangle, b \rangle)$$

G_2 (dependency structures with block-degree 2):

$$\begin{array}{ll} S \rightarrow \langle \langle 0121 \rangle, a \rangle (R, \langle \langle 0 \rangle, b \rangle) & S \rightarrow \langle \langle 01 \rangle, a \rangle (\langle \langle 0 \rangle, b \rangle) \\ R \rightarrow \langle \langle 01, 21 \rangle, a \rangle (R, \langle \langle 0 \rangle, b \rangle) & R \rightarrow \langle \langle 0, 1 \rangle, a \rangle (\langle \langle 0 \rangle, b \rangle) \end{array}$$

8.1.3 | String-generative capacity

It is apparent that our results on the equivalences between the dependency languages induced by various lexicalized grammar formalisms on the one hand and classes of regular dependency languages over mildly non-projective dependency structures on the other hand can be transferred to string languages without any problems: the linearization semantics that we used for the unlabelled structures is fully compatible with the string semantics that we now use for labelled structures. However, one thing to note is, that all our results crucially depend on the grammars being lexicalized, meaning that each production in these grammars contributes an overt lexical item to the derived string—without this restriction, the notion of ‘induced dependency structure’ as we have used it here is ill-defined. Nevertheless, for some of the formalisms that we have studied, lexicalization is not really a restriction after all:

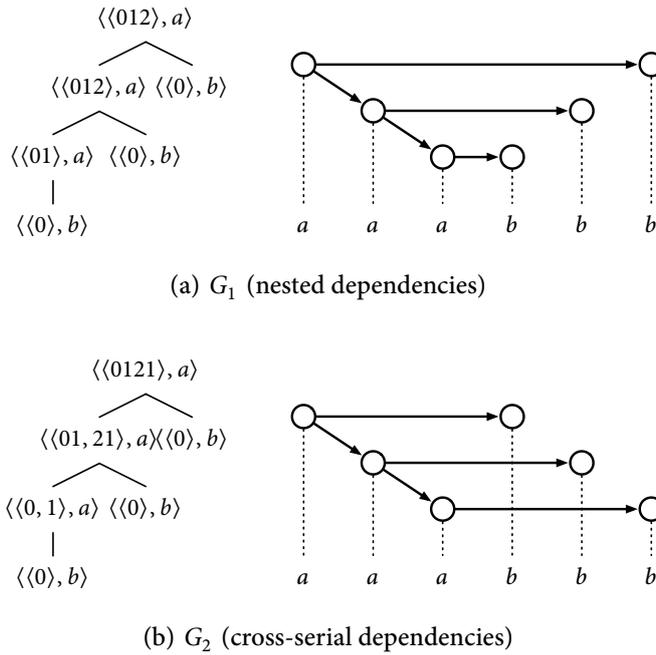


Figure 8.2: Derivations in two string-generating grammars

807 **Lemma** The string languages projected by $\text{REGD}(\mathcal{D}_1)$ are exactly the context-free languages.¹ ←

Proof. This follows from our previous results in combination with the result that every context-free grammar can be put into a lexicalized normal form, such as Greibach normal form (Greibach, 1965) or Rosenkrantz normal form (Rosenkrantz, 1967). One caveat is that this transformation changes the structure of the derivation trees, and thus the dependency structures that we get out from these lexicalized grammars do not necessarily encode the same syntactic dependencies as the original grammar. ←

808 **Lemma** The string languages projected by $\text{REGD}(\mathcal{D}_2 \cap \mathcal{D}_{wn})$ are exactly the string languages generated by TAGs. ←

Proof. This follows from our previous results in combination with the results that every TAG can be put into a lexicalized normal form (Schabes, 1990). ←

For LCFRS and CCFG, the problems whether every grammar can be put into some lexicalized normal form are open. These problems make an interesting topic for research for themselves, but are beyond the scope of this dissertation.

¹ There is one minor difference: The context-free language that contains only the empty word cannot be projected by any regular dependency language.

8.2 | String languages and structural properties

In this section, we study the impact of structural constraints on the string-generative capacity of regular dependency languages. We present two results: first, that the string-language hierarchy known for LCFRS can be recovered in our framework by controlling the block-degree parameter; second, that additionally requiring well-nestedness leads to a proper decrease in generative capacity on nearly all levels of this hierarchy.

String-language hierarchies are usually proven using formalism-specific pumping lemmata. For more powerful formalisms, pumping arguments tend to become rather difficult and technical (see Seki et al., 1991 or Guan, 1992) because they need to reason about the combinatorial structure of the derivation and the order of the derived material at the same time. Our hierarchy proofs are novel in that they clearly separate these two issues: for the combinatorial aspect of the argument, we use only one powerful pumping lemma (Lemma 723); to reason about the order of the derived material, we use our knowledge about structural properties. With this proof technique, we can show that certain string languages ‘enforce’ certain structural properties in regular dependency languages that project them. The usefulness of this approach is witnessed by our result about the language hierarchy for well-nested languages, which solves an open problem concerning the relation between LCFRS and CCFG.

8.2.1 | Masked strings

To prepare our proofs, we first show two elementary results about congruence relations on strings. Recall (from Definition 401) that a congruence relation on a chain \mathcal{C} is an equivalence relation in which all blocks are convex with respect to \mathcal{C} . Congruences on strings can be represented as lists of pairwise disjoint intervals of positions.

Definition Let $s \in A^*$ be a string, and let $n \in \mathbb{N}_0$. A *mask* for s of length n is a list $M = [i_1, j_1] \cdots [i_n, j_n]$ of pairwise disjoint intervals of positions in s such that $j_k < i_{k+1}$, for all $k \in [n-1]$. It is understood that $i_k \leq j_k$, for all $k \in [n]$. \dashv

809

We call the intervals $[i, j]$ the *blocks* of the mask M , and write $|M|$ to denote their number. In slight abuse of notation, we write $B \in M$, if B is a block of M . Given a string s and a mask M for s , the set of *positions* corresponding to M is defined as

$$\text{pos}([i_1, j_1] \cdots [i_n, j_n]) := \{i \in \text{pos}(s) \mid \exists k \in [n]. i \in [i_k, j_k]\}.$$

For a set P of positions in a given string s , we put $\bar{P} := \text{pos}(s) - P$, and write $[P]$ for the smallest mask for s such that $\text{pos}(M) = P$. We say that P *contributes* to a block B of some mask, if $P \cap B \neq \emptyset$. Finally, for masks M with an even number of blocks, we define the *fusion* of M as

$$F([i_1, j_1][i'_1, j'_1] \cdots [i_n, j_n][i'_n, j'_n]) := [i_1, j'_1] \cdots [i_n, j'_n].$$

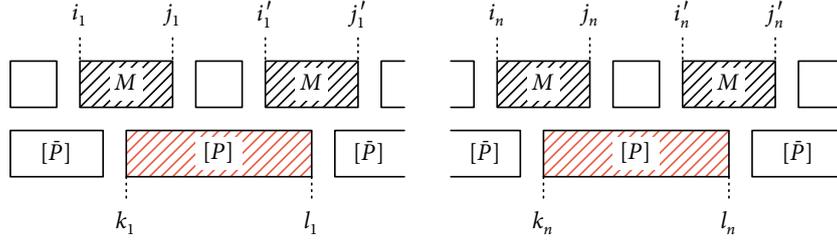


Figure 8.3: The situation in the proof of Lemma 8.10

8.10 **Lemma** Let $s \in A^*$ be a string, let M be a mask for s with an even number of blocks, and let P be a set of positions in s such that both P and \bar{P} contribute to every block of M . Then $|[P]| \geq |M|/2$. Furthermore, if $|[P]| \leq |M|/2$, then $P \subseteq \text{pos}(F(M))$. \dashv

Proof. For every block $B \in [P]$, let $n(B)$ be the number of blocks in M that B contributes to. We make two observations: First, since P contributes to each block of M , $|M| \leq \sum_{B \in [P]} n(B)$. Second, since \bar{P} contributes to each block of M , no block $B \in [P]$ can fully contain a block of M ; therefore, $n(B) \leq 2$ holds for all blocks $B \in [P]$. Putting these two observations together, we deduce that

$$|M| \leq \sum_{B \in [P]} n(B) \leq \sum_{B \in [P]} 2 = 2 \cdot |[P]|.$$

For the second part of the lemma, let

$$M = [i_1, j_1][i'_1, j'_1] \cdots [i_n, j_n][i'_n, j'_n] \quad \text{and} \quad [P] = [k_1, l_1] \cdots [k_n, l_n].$$

Then, each block of $[P]$ contributes to exactly two blocks of M . More precisely, for each $h \in [n]$, the block $[k_h, l_h]$ of $[P]$ contributes to the blocks $[i_h, j_h]$ and $[i'_h, j'_h]$ of M . (This situation is depicted in Figure 8.3.) Because \bar{P} also contributes to $[i_h, j_h]$ and $[i'_h, j'_h]$, the interval $[k_h, l_h]$ is a proper subset of $[i_h, j_h]$, which is a block of the fusion $F(M)$. Hence, $P \subseteq \text{pos}(F(M))$. \dashv

8.2.2 | Enforcing a given block-degree

We now show our first result: for every natural number $k \in \mathbb{N}$, there exists a string language $L(k)$ that forces every regular dependency language that projects $L(k)$ to contain structures of block-degree k . For our proof, we use the string languages from the infinite family

$$\text{COUNT}(k) := \{a_1^n b_1^n \cdots a_k^n b_k^n \mid n \in \mathbb{N}\}.$$

We note that the language $\text{COUNT}(1)$ is homomorphic to the context-free language $\{a^n b^n \mid n \in \mathbb{N}\}$ for which we have seen regular dependency grammars in Example 8.06, and that for every $k > 1$, the language $\text{COUNT}(k)$ is not context-free; this can be easily proved using the standard pumping lemma for context-free languages.

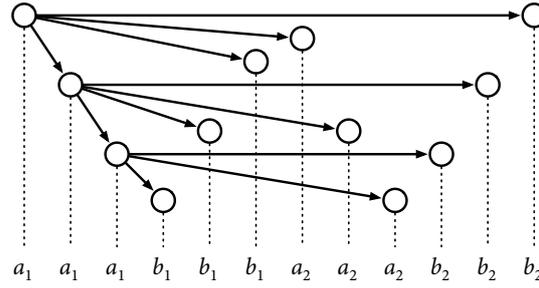


Figure 8.4: A dependency structure for the language $\text{COUNT}(2)$

Example The following grammar generates a dependency language that projects the string language $\text{COUNT}(2)$; it is not hard to see how to modify the grammar so that it generates languages $\text{COUNT}(k)$, for $k > 2$. The grammar is essentially identical to the TAG grammar that we gave in Figure 6.6. We only list the productions. 811

$$\begin{aligned} S &\rightarrow \langle \langle 012314 \rangle, a_1 \rangle (R, \langle \langle 0 \rangle, b_1 \rangle, \langle \langle 0 \rangle, a_2 \rangle, \langle \langle 0 \rangle, b_2 \rangle) \\ S &\rightarrow \langle \langle 0123 \rangle, a_1 \rangle (\langle \langle 0 \rangle, b_1 \rangle, \langle \langle 0 \rangle, a_2 \rangle, \langle \langle 0 \rangle, b_2 \rangle) \\ R &\rightarrow \langle \langle 012, 314 \rangle, a_1 \rangle (R, \langle \langle 0 \rangle, b_1 \rangle, \langle \langle 0 \rangle, a_2 \rangle, \langle \langle 0 \rangle, b_2 \rangle) \\ R &\rightarrow \langle \langle 01, 23 \rangle, a_1 \rangle (\langle \langle 0 \rangle, b_1 \rangle, \langle \langle 0 \rangle, a_2 \rangle, \langle \langle 0 \rangle, b_2 \rangle) \end{aligned}$$

Figure 8.4 shows a dependency structure generated by this grammar. We note that the structure is well-nested. ⊣

In the following proofs, we freely identify (segmented) labelled dependency structures with their corresponding terms. Given a term $d \in T_{\langle \Omega, A \rangle}$, we use the notation $\text{alph}(d)$ to refer to the set of all labels from the alphabet A in d .

Lemma Let $k \in \mathbb{N}$. Every regular dependency language that projects $\text{COUNT}(k)$ contains structures with a block-degree of at least k . 812

Proof. Let $L \in \text{REGD}$ be a regular dependency language that projects $\text{COUNT}(k)$. For notational convenience, put $X := \{x_i \mid x \in \{a, b\} \wedge i \in [k]\}$.

- * We start with a simple auxiliary observation: Let s_1 and s_2 be two strings in $\llbracket L \rrbracket_{\mathcal{S}}$. If $|s_1| < |s_2|$, then every symbol from X occurs more often in s_2 than in s_1 .
- * Let p be the pumping number from Lemma 718, and let $d_1 \in L$ be a dependency structure with $\llbracket d_1 \rrbracket_{\mathcal{S}} = a_1^n b_1^n \cdots a_k^n b_k^n$, where $n = \lceil p/2k \rceil$. Due to the isomorphism between $\llbracket d_1 \rrbracket_{\mathcal{S}}$ and the precedence relation of d_1 , we have $|d_1| = 2k \cdot n \geq p$. In this case, Lemma 718 asserts that d_1 can be written as $d_1 = c' \cdot c \cdot t'$ such that c contains at least one node, and $d_2 := c' \cdot c \cdot c \cdot t'$ belongs to L (see the upper part of Figure 8.5). Now, let u be the uniquely determined node in d_2 for which

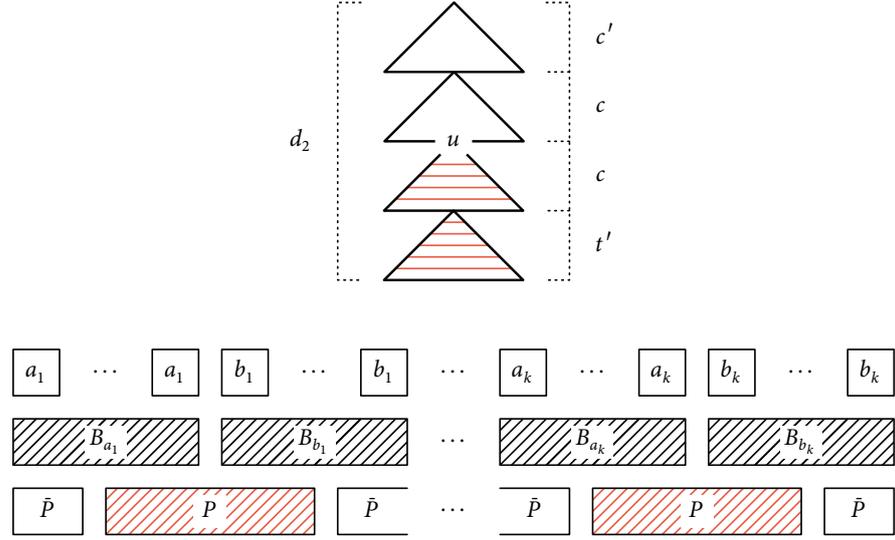


Figure 8.5: The situation in the proof of Lemma 8.12

$d_2/u = c \cdot t'$ holds. As a consequence of the first item and the construction of d_2 , we deduce that every symbol from X occurs in c . Hence, $X \subseteq \text{alph}(c) \subseteq \text{alph}(d_2/u)$.

- * We now show that u has block-degree k . Let $M = B_{a_1} B_{b_1} \cdots B_{a_k} B_{b_k}$ be the uniquely determined mask for $\llbracket d_2 \rrbracket_S$ in which each block B_{x_i} contains exactly those positions that correspond to occurrences of the symbol x_i (see the lower part of Figure 8.5), and let P be the set of those positions that correspond to the yield $\llbracket u \rrbracket$. Since every symbol from X occurs in both P and its complement, both sets contribute to every block of M . With the first part of Lemma 8.10, we deduce that $|\llbracket P \rrbracket| \geq k$. Due to the isomorphism between $\llbracket d_2 \rrbracket_S$ and the precedence relation of d_2 , this means that the yield $\llbracket u \rrbracket$ is distributed over at least k blocks in d_2 . \dashv

8.2.3 | Enforcing ill-nestedness

We now show that even the well-nestedness constraint has an impact on the string-generative capacity of regular dependency languages. More specifically, for every natural number $k \in \mathbb{N}$, there exists a string language $L(k)$ that forces every regular dependency language over structures with a block-degree of at most k that projects $L(k)$ to contain ill-nested structures. For our proof, we use the languages from the family

$$\text{RESP}(k) := \{ a_1^m b_1^m c_1^n d_1^n \cdots a_k^m b_k^m c_k^n d_k^n \mid m, n \in \mathbb{N} \}.$$

Similar to $\text{COUNT}(k)$, the language $\text{RESP}(k)$ is projected by a regular dependency language over the class \mathcal{D}_k of dependency structures with block-degree at most k .

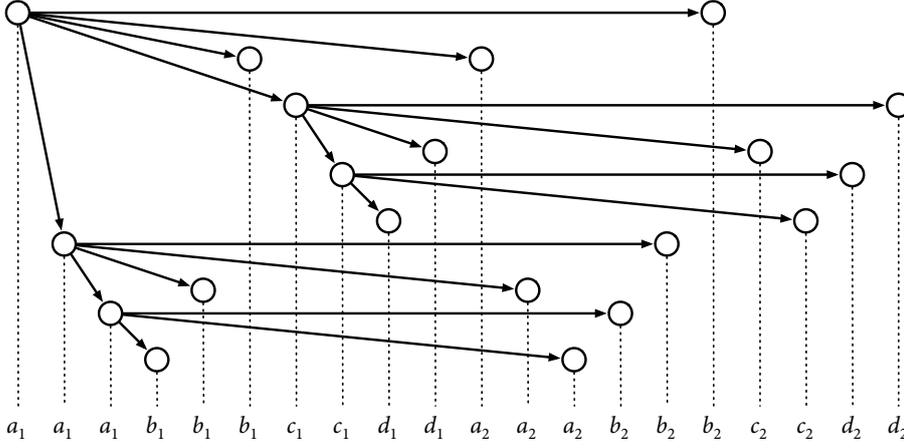


Figure 8.6: A dependency structure for the language $\text{RESP}(2)$

Example The following grammar generates a dependency language that projects the string language $\text{RESP}(2)$. 813

$$\begin{aligned}
 S &\rightarrow \langle \langle 01234153 \rangle, a_1 \rangle (R_1, \langle \langle 0 \rangle, b_1 \rangle, R_2, \langle \langle 0 \rangle, a_2 \rangle, \langle \langle 0 \rangle, b_2 \rangle) \\
 S &\rightarrow \langle \langle 012342 \rangle, a_1 \rangle (\langle \langle 0 \rangle, b_1 \rangle, R_2, \langle \langle 0 \rangle, a_2 \rangle, \langle \langle 0 \rangle, b_2 \rangle) \\
 R_1 &\rightarrow \langle \langle 012, 314 \rangle, a_1 \rangle (R_1, \langle \langle 0 \rangle, b_1 \rangle, \langle \langle 0 \rangle, a_2 \rangle, \langle \langle 0 \rangle, b_2 \rangle) \\
 R_1 &\rightarrow \langle \langle 01, 23 \rangle, a_1 \rangle (\langle \langle 0 \rangle, b_1 \rangle, \langle \langle 0 \rangle, a_2 \rangle, \langle \langle 0 \rangle, b_2 \rangle) \\
 R_2 &\rightarrow \langle \langle 012, 314 \rangle, c_1 \rangle (R_2, \langle \langle 0 \rangle, d_1 \rangle, \langle \langle 0 \rangle, c_2 \rangle, \langle \langle 0 \rangle, d_2 \rangle) \\
 R_2 &\rightarrow \langle \langle 01, 23 \rangle, c_1 \rangle (\langle \langle 0 \rangle, d_1 \rangle, \langle \langle 0 \rangle, c_2 \rangle, \langle \langle 0 \rangle, d_2 \rangle)
 \end{aligned}$$

Figure 8.6 shows a dependency structure generated by this grammar. Note that this structure is ill-nested. This is mirrored in the grammar by the fact that the first order annotation contains the forbidden substring 1313. ⊣

Lemma Let $k > 1$. Every regular dependency language $L \in \text{REGD}(\mathcal{D}_k)$ that projects $\text{RESP}(k)$ contains ill-nested structures. 814 ⊣

Proof. Let $L \in \text{REGD}(\mathcal{D}_k)$ be a regular dependency language that projects $\text{RESP}(k)$. Define the following two sets of symbols:

$$X := \{x_i \mid x \in \{a, b\} \wedge i \in [k]\}, \quad \text{and} \quad Y := \{y_i \mid y \in \{c, d\} \wedge i \in [k]\}.$$

* We start with a simple observation: Let d_1 and d_2 be dependency structures contained in L . If at least one symbol from X occurs more often in d_2 than in d_1 , then every symbol from X does. A symmetric argument holds for Y .

8 GENERATIVE CAPACITY AND PARSING COMPLEXITY

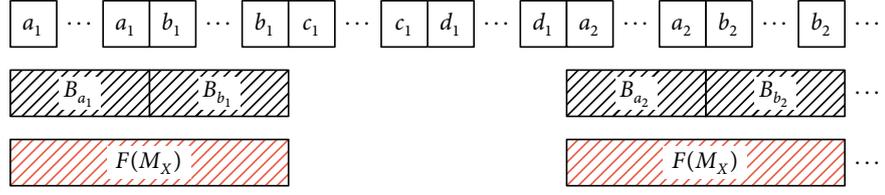


Figure 8.7: Enforcing ill-nestedness

- * Let p be the pumping number from Lemma 723, and let $d \in L$ be a dependency structure with $\llbracket d \rrbracket_{\mathcal{S}} = a_1^m b_1^m c_1^n d_1^n \cdots a_k^m b_k^m c_k^n d_k^n$, where $m = n = \lceil p/2k \rceil$. By this choice, the structure d contains $2k \cdot m \geq p$ occurrences of symbols from X , and equally many occurrences of symbols from Y .
- * Lemma 723 asserts that the structure d can be written as $d = c' \cdot c \cdot t'$ such that the context c contains at least one occurrence of a symbol from X , and the ‘pumped’ structure $d_X := c' \cdot c \cdot c \cdot t'$ is contained in L . Let u_X be the uniquely determined node in d for which $d/u_X = c \cdot t'$. We want to show that $\text{alph}(d/u_X) = X$.
 - ⊇ Since c contains at least one occurrence of a symbol from X , at least one symbol from X occurs more often in d_X than in d . Then, by our observation above, every symbol from X occurs more often in d_X than in d . By the construction of d_X , this implies that $X \subseteq \text{alph}(c) \subseteq \text{alph}(d/u_X)$.
 - ⊆ Let $M_X = B_{a_1} B_{b_1} \cdots B_{a_k} B_{b_k}$ be the uniquely determined mask for $\llbracket d_X \rrbracket_{\mathcal{S}}$ in which each block B_{x_i} contains exactly those positions that are labelled with the symbol x_i . Furthermore, let u be the uniquely determined node in d_X for which $d_X/u = c \cdot t'$, and let P be the set of those positions in $\llbracket d_X \rrbracket_{\mathcal{S}}$ that correspond to the yield $\lfloor u \rfloor$. We now apply Lemma 810: given that $X \subseteq \text{alph}(d_X/u)$, both the set P and its complement contribute to every block of M ; given that $d_2 \in \mathcal{D}_k$, we have $|[P]| \leq k$. From this, we deduce that $P \subseteq \text{pos}(F(M_X))$. Since every position in the set $\text{pos}(F(M_X))$ is labelled with a symbol from X (see Figure 8.7), we conclude that $\text{alph}(d_X/u) = \text{alph}(d/u_X) \subseteq X$.
- * Put $Y := \{x_i \mid x \in \{c, d\} \wedge i \in [k]\}$. In symmetry to the argument above, we can show the existence of a node u_Y in d_1 for which $\text{alph}(d_1/u_Y) = Y$.
- * Due to the isomorphism between $\llbracket d_1 \rrbracket_{\mathcal{S}}$ and the precedence relation of d_1 , the yields $\lfloor u_X \rfloor$ and $\lfloor u_Y \rfloor$ interleave. Since the sets X and Y are disjoint, neither $u_X \preceq u_Y$ nor $u_Y \preceq u_X$ holds. We conclude that d_1 is ill-nested. \dashv

8.2.4 | Hierarchies of string languages

We are now ready to present the main result of this section: the hierarchy on regular dependency languages from Lemma 705 carries over to string languages.

Theorem The following relations hold for all $k \in \mathbb{N}$:

- * $\llbracket \text{REGD}(\mathcal{D}_k) \rrbracket_{\mathcal{S}} \subsetneq \llbracket \text{REGD}(\mathcal{D}_{k+1}) \rrbracket_{\mathcal{S}}$
- * $\llbracket \text{REGD}(\mathcal{D}_k \cap \mathcal{D}_{wn}) \rrbracket_{\mathcal{S}} \subsetneq \llbracket \text{REGD}(\mathcal{D}_{k+1} \cap \mathcal{D}_{wn}) \rrbracket_{\mathcal{S}}$
- * $\llbracket \text{REGD}(\mathcal{D}_1) \rrbracket_{\mathcal{S}} = \llbracket \text{REGD}(\mathcal{D}_1 \cap \mathcal{D}_{wn}) \rrbracket_{\mathcal{S}}$
- * $k \neq 1 \implies \llbracket \text{REGD}(\mathcal{D}_k \cap \mathcal{D}_{wn}) \rrbracket_{\mathcal{S}} \subsetneq \llbracket \text{REGD}(\mathcal{D}_k) \rrbracket_{\mathcal{S}}$

Proof. The inclusions in the first two items as well as the third item are immediate consequences of Lemma 705. The properness of the inclusions and the last item follow from Lemmata 812 and 814 and the facts that $\text{COUNT}(k) \in \llbracket \text{REGD}(\mathcal{D}_k \cap \mathcal{D}_{wn}) \rrbracket_{\mathcal{S}}$ and $\text{RESP}(k) \in \llbracket \text{REGD}(\mathcal{D}_k) \rrbracket_{\mathcal{S}}$, as witnessed by the grammars we gave above. \dashv

The hierarchy established by the first item corresponds to the string-language hierarchy known for LCFRS (Weir, 1988; Groenink, 1997) and other formalisms that generate the same string languages (see e.g. Engelfriet et al., 1980; Seki et al., 1991; van Vugt, 1996; Rambow and Satta, 2004). The result established by the fourth item is original to this dissertation: given the weak equivalence between $\text{REGD}(\mathcal{D}_k)$ and $\text{LCFRS}(k)$ on the one hand and $\text{REGD}(\mathcal{D}_k \cap \mathcal{D}_{wn})$ and $\text{CCFG}(k)$ on the other, it shows that the generative power of CCFG is strictly below that of LCFRS.

8.2.5 | Related work

- ✧ The language $\text{RESP}(2)$ was first considered by Weir (1988), who speculated that it separates the string-languages generated by LCFRS with fan-out 2 from the languages generated by TAG. This was subsequently proved by Seki et al. (1991).
- ✧ Gramatovici and Plátek (2006) study a string-language hierarchy on a dependency formalism in which derivations can be controlled by the node-gaps complexity parameter that we discussed in Section 4.1.3.

8.3 | Parsing complexity

The parsing problem of regular dependency languages is the problem to find, given a grammar and a string of terminal symbols, (a compact description of) the set of all dependency structures generated by the grammar that project the string. In this section, we show that regular dependency languages can be parsed in time polynomial in the length of the input string, but that the parsing problem in which the grammar is part of the input is NP-complete even for a fixed block-degree. However, we also show that the same problem becomes polynomial when grammars are restricted to well-nested order annotations, and hence, to well-nested dependency languages. Together with the treebank evaluation that we presented in Chapter 5, this result provides strong evidence that our interest in the well-nestedness condition is justified.

8.3.1 | Membership problems

Instead of looking at the parsing problem of regular dependency languages directly, we restrict ourselves to a slightly simpler problem: the problem to decide, given a grammar and a string, whether the grammar generates any dependency structure at all that projects the string. This problem is the *membership problem* of the projected string language. For the vast majority of the algorithms that solve membership problems for generative grammars, including the ones that we discuss here, there are standard ways to extend them into full parsers, so the restriction to the membership problem is minor. The membership problem comes in two variants, depending on whether we consider the grammar to be part of the input to the problem or not:

816

Definition The (*standard*) *membership problem* for a regular dependency grammar G is the following decision problem: given a string \vec{a} , is $\vec{a} \in \llbracket L(G) \rrbracket_{\mathcal{S}}$? The *uniform membership problem* for a class \mathcal{G} of regular dependency grammars is the following decision problem: given a grammar $G \in \mathcal{G}$ and a string \vec{a} , is $\vec{a} \in \llbracket L(G) \rrbracket_{\mathcal{S}}$? \dashv

The uniform membership problem is at least as difficult as the standard membership problem, but it may be more difficult. In particular, every polynomial-time algorithm that solves the uniform membership problem also solves the standard membership problem in polynomial time. On the other hand, an algorithm for the standard membership problem may take an amount of time that is exponential in size factors that depend on the grammar. In this case, it does not yield a polynomial-time algorithm for the universal membership problem.

In the computer science literature, the run-time of parsing algorithms is usually given as a function of the length of the input string, which is informative only for the standard membership problem. One of the reasons for the disinterest in the size of the grammar may be that, in many applications, grammars are small, and the candidate string is long—consider the grammar of a programming language for example, which usually only fills a few pages, but may be used in compilers that process ten thousands lines of code. This situation does not apply to computational linguistics, where rather the opposite is true: sentences are short, not more than a hundred words, while grammars are huge, with several hundreds of thousands of entries. Thus, for the parsing of natural language, the important measure in the analysis of parsing algorithms is not the length of the input string, but the size of the grammar (cf. Lee, 2002). This holds true in particular when we consider lexicalized grammars, where all productions are specialized for individual words. At the same time, these grammars have the advantage that parsing needs to consider only those productions that are associated with the words in the input string (Schabes et al., 1988). While this strategy reduces the parsing time in all practical cases, it also introduces an additional factor into the complexity analysis of parsing algorithms that depends on the length of the input string (cf. Eisner and Satta, 1999).

$$\begin{aligned}
\mathcal{A} &:= \{ [a, [i, i]] \in T \times B \mid a_i = a \} \\
\mathcal{I} &:= \{ [A, M] \in N \times B^* \mid |M| = \deg(A) \} \\
\mathcal{G} &:= \{ [S, [1, n]] \}
\end{aligned}$$

Figure 8.8: Axioms, items and goals for the grammatical deduction system

8.3.2 | The standard membership problem

As our first technical result of this section, we now show that the standard membership problem for regular dependency grammars is polynomial. To prove this, we construct a generic recognition algorithm for a regular dependency grammar G in the framework of deductive parsing (Shieber et al., 1995). Let us write k_G for the degree of G (which corresponds to the maximal block-degree among the structures in the language generated by G), and m_G for the maximal rank of G (which corresponds to the maximal out-degree of the language generated by G).

Lemma The membership problem of string languages that are projected by regular dependency languages is in time $O(|P| \cdot n^e)$, where $e = k_G \cdot (m_G + 1)$. 817

Proof. Let L be a regular dependency language, and let \vec{a} a string over some alphabet. Furthermore, let $G = (N, T, S, P)$ be a normalized regular dependency grammar that generates L . To decide whether $\vec{a} \in \llbracket L \rrbracket_{\mathcal{S}}$, we construct a *grammatical deduction system* for G , and use a generic implementation of this system in the framework of deductive parsing (Shieber et al., 1995).

* Put $n := |\vec{a}|$, and let $B \subseteq [n] \times [n]$ be the set of all intervals of positions in the string \vec{a} . A grammatical deduction system consists of four components: a set \mathcal{A} of *axioms*, a set \mathcal{I} of *items*, a set $\mathcal{G} \subseteq \mathcal{I}$ of *goal items*, and a finite collection of *inference rules* over \mathcal{A} and \mathcal{I} . The sets of axioms, items and goal items of our deduction system are defined in Figure 8.8.

The axioms represent the information about which position in \vec{a} is labelled by which terminal symbol. An item $[A, [i_1, j_1] \cdots [i_k, j_k]]$ asserts that there is a dependency structure $d \in L(G)$ such that $\llbracket d \rrbracket_{\mathcal{S}} = \langle a_{i_1} \cdots a_{j_1}, \dots, a_{i_k} \cdots a_{j_k} \rangle$; in particular, the goal item asserts that $\vec{a} \in \llbracket L \rrbracket_{\mathcal{S}}$. The set of inference rules is constructed as follows. For each production $A \rightarrow \langle \omega, a \rangle (A_1, \dots, A_m)$ with $\omega : k_1 \times \cdots \times k_m \rightarrow k$, we use an inference rule of the form

$$\frac{[a, b_{0,1}] \quad [A_1, b_{1,1} \cdots b_{1,k_1}] \quad \cdots \quad [A_m, b_{m,1} \cdots b_{m,k_m}]}{[A, b_1 \cdots b_k]}$$

This rule is subject to the following side conditions, which reflect the semantics of the order annotation ω . Assume that $\omega = \langle \vec{i}_1, \dots, \vec{i}_k \rangle$. We write ℓ_x for the left endpoint of the interval b_x , and r_x for the corresponding right endpoint.

$$r_{0,1} = \ell_{0,1} \quad \Leftarrow \quad (8.1)$$

$$\ell_{i_2, j_2} = r_{i_1, j_1} + 1 \quad \Leftarrow \quad \exists h \in [k]. \vec{i}_h = \vec{x} \cdot i_1 \# j_1 \cdot i_2 \# j_2 \cdot \vec{y} \quad (8.2)$$

$$\ell_h = \ell_{i,j} \quad \Leftarrow \quad \vec{i}_h = i \# j \cdot \vec{x} \quad (8.3)$$

$$r_h = r_{i,j} \quad \Leftarrow \quad \vec{i}_h = \vec{x} \cdot i \# j \quad (8.4)$$

The first condition reflects the semantics of the axioms. The second condition ensures that blocks that are adjacent in ω correspond to intervals of positions that are adjacent in \vec{a} . The third and fourth condition ensure that blocks that are extremal in ω correspond to extremal intervals in \vec{a} . Taken together, the conditions ensure that each inference rule is sound with respect to the intended semantics. Their completeness is obvious. Thus, we have $\vec{a} \in \llbracket L \rrbracket_{\mathcal{S}}$ if and only if starting from the axioms, we can deduce the goal item.

- * The asymptotic runtime of the generic, chart-based implementation of the grammatical deduction system for G is $O(|P| \cdot n^e)$, where e is the maximal number of free variables per inference rule that range over the domain $[n]$ (see McAllester, 2002). To determine e , we inspect the schema for the inference rules above. The total number of variables over $[n]$ in this schema is $2 + 2k + \sum_{i=1}^m 2k_i$. Each non-free variable is determined by exactly one of the side conditions. Therefore, to determine the number of free variables in the rule schema, it suffices to count the instantiations of the schemata for the side conditions, and to subtract this number from the total number of variables. Schema 8.1 has 1 instantiation. Schemata 8.3 and 8.4 each have k instantiations; this is the number of leftmost and rightmost positions in the blocks of ω , respectively. Finally, schema 8.2 has $1 - k + \sum_{i=1}^m k_i$ instantiations: the string \vec{i} has $1 + \sum_{i=1}^m k_i$ positions; k of these mark the end of a block and thus do not have a neighbouring symbol. Then the number of free variables is

$$\left(2 + 2k + \sum_{i=1}^m 2k_i \right) - \left(1 + 2k + 1 - k + \sum_{i=1}^m k_i \right) = k + \sum_{i=1}^m k_i \leq k_G \cdot (m_G + 1).$$

Thus, $e \leq k_G \cdot (m_G + 1)$. ←

818 **Theorem** The membership problem of $\llbracket \text{REGD} \rrbracket_{\mathcal{S}}$ is in PTIME. ←

This result, together with our previous result about the constant growth of the languages in REGD (Theorem 734), allows us to call regular dependency languages mildly context-sensitive, according to Joshi's (1985) characterization.

8.3.3 | The uniform membership problem

The complexity of the generic parsing algorithm for regular dependency languages that we gave in the previous section is exponential both in the degree and in the rank of the grammar that is being processed. This means that we are punished both for languages with a high degree of non-projectivity, and for languages with a high number of dependents per node. A natural question to ask is, whether we can get substantially better than this. Unfortunately, at least in the general case, the answer to this question is probably negative: in this section, we show that the uniform membership problem for the class of regular dependency grammars is NP-complete. Given the close connection between regular dependency grammars and LCFRS, this result does not come entirely unexpected: Satta (1992) showed that the uniform membership problem of both LCFRS with restricted fan-out (our block-degree) and restricted rank is NP-hard. Unfortunately, we cannot directly apply his reduction (of the 3SAT problem) to the membership problem of regular dependency languages, as this reduction makes essential use of concatenation functions with empty components, which we have excluded (see Section 6.2.3).

Instead, we provide a polynomial reduction of the EXACT COVER problem to the uniform membership problem of regular dependency grammars, UNIFORM-REGD. An instance of EXACT COVER is given by a finite set U and a finite collection $\mathcal{F} \subseteq \mathfrak{P}(U)$ of subsets of U . The decision to make is, whether there is a subset $\mathcal{C} \subseteq \mathcal{F}$ such that the sets in \mathcal{C} are disjoint, and their union is U .

Lemma EXACT COVER \leq_p UNIFORM-REGD †

819

Proof. Let $I = (U, \mathcal{F})$ be any instance of the EXACT COVER problem. Put $n := |U|$, and $m := |\mathcal{F}|$, and assume that the elements of U and \mathcal{F} are numbered from 1 to n and 1 to m , respectively. We write u_i to refer to the i th element of U , and S_i to refer to the i th element of \mathcal{F} with respect to these numberings. The main idea behind the following reduction is to construct a regular dependency grammar $G = (N, T, S, P)$ and a string \vec{a} such that each dependency structure that is generated by G and projects \vec{a} represents a partition $\mathcal{C} \subseteq \mathcal{F}$ of U . The string \vec{a} has the form $\$ \cdot \vec{x}_1 \cdots \vec{x}_m \cdot \vec{x}$, where the substring \vec{x} is a representation of the set U , and each substring \vec{x}_i , $i \in [m]$, controls whether the set S_i is included in \mathcal{C} . The grammar G is designed such that each substring \vec{x}_i can be derived in only two possible ways and only as the projection of the first block of a dependency structure with block-degree 2; the second block of this structure projects material in the string \vec{x} . In this way, each derivation corresponds to a guess which sets of \mathcal{F} to include into \mathcal{C} . The string \vec{x} is set up to ensure that this guess is consistent.

We first describe the construction of the string \vec{a} . Each string \vec{x}_i , $i \in [m]$, has the form $\$ y_1 \# \cdots \# y_n \$$, where for all $j \in [n]$, $y_j = u_j$ if $u_j \in S_i$, and $y_j = \bar{u}_j$ otherwise. The string \vec{a} then has the form $\$ \cdot \vec{x}_1 \cdots \vec{x}_m \cdot \bar{u}_1^m u_1 \bar{u}_1^m \cdots \bar{u}_n^m u_n \bar{u}_n^m$.

Next, we describe the construction of the grammar. The non-terminal and terminal alphabets are defined as follows:

$$\begin{aligned} N &:= \{1 \mapsto \{S\}, 2 \mapsto \{[S_i, u_j] \mid i \in [m] \wedge j \in [n]\}\} \\ T &:= \{\$, \#\} \cup \{u_i \mid i \in [n]\} \cup \{\bar{u}_i \mid i \in [n]\} \end{aligned}$$

The start symbol is S . Before we give the set of production, we introduce the following abbreviating notation: for every terminal symbol $a \in T$, put

$$\begin{aligned} \langle a\bar{a}, \bar{a}' \rangle &:= \langle \langle 01, 1 \rangle, a \rangle \langle \langle \bar{a}, \bar{a}' \rangle \rangle & \langle a, \bar{a} \rangle &:= \langle \langle 0, 1 \rangle, a \rangle \langle \langle \bar{a} \rangle \rangle \\ \langle a\bar{a} \rangle &:= \langle \langle 01 \rangle, a \rangle \langle \langle \bar{a} \rangle \rangle & \langle a \rangle &:= \langle \langle 0 \rangle, a \rangle \end{aligned}$$

Now for each set $S \in \mathcal{F}$, each element $u \in U$, and all $i, j \in [m]$, we introduce the following productions:

$$\begin{aligned} [S, u] &\rightarrow \langle \$u, \bar{u}^i u \bar{u}^j \rangle & [S, u] &\rightarrow \langle \$\bar{u}, \bar{u} \rangle & (\text{first selected/not selected}) \\ [S, u] &\rightarrow \langle \#u, \bar{u}^i u \bar{u}^j \rangle & [S, u] &\rightarrow \langle \#\bar{u}, \bar{u} \rangle & (\text{selected/not selected}) \\ [S, u] &\rightarrow \langle \#u\$, \bar{u}^i u \bar{u}^j \rangle & [S, u] &\rightarrow \langle \#\bar{u}\$, \bar{u} \rangle & (\text{last selected/not selected}) \end{aligned}$$

We also need the production $S \rightarrow \langle \langle 0\bar{x}\bar{y} \rangle, \$ \rangle ([S_1, u_1], \dots, [S_m, u_m])$, where \bar{x} is the row-wise reading of the $n \times m$ -matrix in which each cell (i, j) , $i \in [n]$, $j \in [m]$, contains the value $i + n \cdot (j - 1)$, and \bar{y} is the column-wise reading of this matrix.

We now claim that each substring \bar{x}_i , $i \in [m]$, can be derived in only two possible ways: either by rules from the group ‘selected’, or by rules from the group ‘not selected’. Within such a group, each terminal can only be generated by exactly one rule, depending on the position of the terminal in the sub-string (first, inner, last) and the form of the terminal (u, \bar{u}) . In this way, each derivation of \bar{x}_i corresponds to a choice whether S_k should be part of \mathcal{C} or not. If it is chosen, the second components of the rules consume the single terminal \bar{u} in the right half of the string, along with all ‘garbage’ (in the form of superfluous symbols \bar{u}) adjacent to it. No terminal u in the right half of the string can be consumed twice; this reflects the fact that the S_k must be disjoint. If the derivation is complete, all terminals on the right side have been consumed; this reflects the fact that the union of the S_k makes the complete set. \dashv

Note that the grammar constructed in the proof of this lemma has degree 2, but that its maximal rank grows with the input.

820 Example Figure 8.9 shows an example for the encoding in the proof of Lemma 819 for the instance $U = \{u_1, u_2\}$, $\mathcal{F} = \{\{u_1\}, \{u_2\}\}$. \dashv

821 Theorem The uniform string membership problem for the class of normalized regular dependency grammars is NP-complete. \dashv

linear in the size of the grammar, and without an increase in size: essentially, we replace all productions of the dependency grammar by their equivalents under the relabelling that we presented in Section 6.3. The membership problem of G' can be decided in time $O(|G'|^2 \cdot n^{3k})$ (Hotz and Pitsch, 1996), where $|G'|$ is the size of the grammar. Consequently, the membership problem of G can be decided in time $O(|G|^2 \cdot n^{3k})$. \dashv

Together with the treebank evaluation that we presented in Chapter 5, this result is strong evidence that the well-nestedness condition is a very relevant condition on dependency structures indeed. A promising objective of future research is to understand the fundamental differences between well-nested and ill-nested grammars in more detail. In particular, it would be interesting to see how the generic algorithm for the parsing of regular dependency grammars can be modified to make use of the well-nestedness condition, and to give a more fine-grained complexity analysis that reverences the lexicalized nature of dependency grammars.

8.3.5 | Related work

Besides the result that we present here, there are several other results that show that grammar-driven dependency parsing can be intractable. Neuhaus and Bröker (1997) use a reduction of the VERTEX COVER problem to prove that the uniform membership problem for a class of linguistically adequate, minimal dependency grammars is NP-complete. Koller and Striegnitz (2002) show a corresponding result for the formalism Topological Dependency Grammar (Duchier and Debusmann, 2001) using a reduction of HAMILTON CYCLE.

CHAPTER 9 | Conclusion

The primary goal of this dissertation was to connect two research areas in computational linguistics: the search for grammar formalisms with adequate generative capacity, and the search for dependency representations with well-balanced structural flexibility. In this chapter, we summarize our main contributions to this goal, and identify some avenues for future research.

Summary

In the first part of the dissertation, we studied three classes of dependency structures: projective, block-restricted, and well-nested structures. Each of these classes were originally defined in terms of a structural constraint on dependency graphs. Bridging the gap between dependency representations and generative grammar, we complemented this graph-based perspective with an algebraic framework that encapsulates the structural constraints in the composition operations by which dependency graphs can combine. An important technical result on the way to the algebraic framework was the encoding of dependency structures into terms over a signature of *order annotations* in such a way that the three classes that we considered could be characterized through the subsets of this signature that are identified by their encodings. With the one-to-one correspondence between dependency structures and terms, we were able to define the concept of a *dependency algebra* and show that these algebras are isomorphic to their corresponding term algebras. The relevance of this result is that composition operations on dependency structures can be simulated by corresponding operations on terms, which provide us with a well-studied and convenient data structure.

At the end of the first part of the dissertation, we put our algebraic framework to use and classified several lexicalized mildly context-sensitive grammar formalisms with respect to the dependency structures that they induce. Taking the algebraic approach as we did, we could formalize the notion of induction as the evaluation of derivations trees in *linearization algebras*. We showed the following results: The class of dependency structures that is induced by Context-Free Grammar is the class of projective dependency structures. Linear Context-Free Rewriting Systems (LCFRS) induce the class of block-restricted dependency structures; more specifically, the maximal block-degree of a dependency structure induced by an LCFRS is directly correlated to the measure of ‘fan-out’ that is used to sub-classify these systems. Adding well-nestedness to the block-degree restriction corresponds to the step from LCFRS to Coupled Context-Free Grammar (CCFG). As a special case, the class of well-nested dependency structures

9 CONCLUSION

with a block-degree of at most 2 is exactly the class of dependency structures that is induced by Tree Adjoining Grammar (TAG). With these connections, we have effectively quantified the generative capacity of lexicalized grammar formalisms along an infinite hierarchy of ever more non-projective dependency structures. This measure is attractive as an alternative to string-generative capacity and tree-generative capacity because dependency representations are more informative than strings, but less formalism-specific than parse trees.

The algebraic perspective on dependency structures also led to efficient algorithms for deciding whether a given structure is projective or well-nested, or has a certain block-degree. We used these algorithms to evaluate the empirical relevance of the three structural constraints on data from three widely-used dependency treebanks. The outcome of these experiments was, that while the class of projective dependency structures is clearly insufficient for many practical applications, one only needs to go a small step beyond projectivity in order to cover virtually all of the data. In particular, the class of TAG-inducible dependency structures covers close to 99.5% of all the analyses in both the Danish Dependency Treebank (Kromann, 2003) and two versions of the Prague Dependency Treebank (Hajič et al., 2001, 2006), the largest corpus of dependency analyses currently available.

In the second part of the dissertation, we developed the theory of regular sets of dependency structures, or *regular dependency languages*. Our approach to define the notion of ‘regularity’ for dependency structures was completely canonical: the regular dependency languages are the recognizable subsets of dependency algebras in the sense of Mezei and Wright (1967). By this definition, we obtained natural notions of automata and grammars for dependency structures on the basis of which we were able to study language-theoretic properties. We proved a powerful pumping lemma for regular dependency languages, and used it to show that they are semi-linear, a property also characteristic for mildly context-sensitive languages. As another application of our pumping lemma, we showed that, under the constraint of regularity, there is a direct correspondence between block-degree and well-nestedness on the one hand, and string-generative capacity on the other. More specifically, the block-degree parameter induces an infinite hierarchy of ever more powerful string languages, and on almost every level of this hierarchy, the string languages that correspond to well-nested dependency structures form a strict subclass.

Finally, we investigated the parsing complexity of regular dependency languages. While the restriction to a fixed block-degree is sufficient to render the parsing problem polynomial in the length of the input sentence, we found that it does not suffice for practical applications: the parsing problem where the size of the grammar is taken into account is NP-complete for unrestricted regular dependency grammars. Interestingly, the corresponding problem for well-nested grammars is polynomial. Together with our treebank studies, this results provides further evidence that the well-nestedness constraint has relevance for practical applications.

Outlook

There are several aspects of the work reported in this dissertation, both theoretical and practical, that can be elaborated in future research.

Qualitative evaluation of structural constraints The treebank evaluation reported in Chapters 3–5 was a purely quantitative evaluation: we simply counted the structures that satisfied or did not satisfy a given structural constraint. While this was helpful in getting an estimate of the practical relevance of the structural constraints that we discussed in this dissertation, it would be highly desirable to complement our results with a qualitative analysis of the treebank data. In particular, it would be very interesting to see whether there is any systematic connection between structural constraints and specific linguistic phenomena or typological properties.

MSO-definable sets of dependency structures We have characterized regular dependency languages both in terms of algebraic automata and in terms of regular grammars. For the class of regular term languages, there is another well-known characterization: a term language is regular if and only if it can be defined by a formula of monadic second-order logic (Thatcher and Wright, 1968). It should be possible to obtain such a characterization even for regular dependency languages. Based on our definition of (expanded) dependency structures, a natural choice for the atomic predicates of the logic are $x \triangleleft y$ (governance), $x \preceq y$ (precedence), and $x \equiv_z y$ (congruence).

Unranked dependency languages A major limitation in the practical applicability of regular dependency languages is the fact that every such language has bounded out-degree (Lemma 707). This implies that regular dependency languages cannot account for phenomena in which a given word accepts an arbitrary number of modifiers, such as nouns accept chains of adjectives. It would be interesting to extend the concept of regularity to sets of unranked trees. This would require us to find a different algebraic structure for dependency structures, in which structures can not only combine along a vertical dimension, essentially using substitution operations, but also along a horizontal dimension, essentially using concatenation, like in the work of Carme et al., 2004.

Application to parsing Recent work on data-driven dependency parsing has shown that parsing with unrestricted non-projective dependency graphs is intractable under all but the simplest probability models (McDonald and Pereira, 2006; McDonald and Satta, 2007). On the other hand, projective dependency parsers combine favourably with more complicated models. It would be interesting to see whether the structural constraints that we have discussed in this dissertation

9 CONCLUSION

can be exploited to obtain efficient yet well-informed parsing algorithms even for certain classes of non-projective dependency graphs. An attractive such class to look at is the class of well-nested dependency structures with a block-degree of at most 2. As we mentioned above, this class covers close to 99.5% of the data in the three treebanks that we evaluated it on. At the same time, we showed that this class corresponds to the dependency structures induced by Tree Adjoining Grammar (Section 6.4). Thus, it should be possible to tap the literature on efficient parsing algorithms for lexicalized TAG to develop parsing algorithms for this class of dependency representations.

Bibliography

- Abe, Naoki. Feasible learnability of formal grammars and the theory of natural language acquisition. In *Twelfth International Conference on Computational Linguistics (COLING)*, pages 1–6, Budapest, Hungary, 1988.
DOI: 10.3115/991635.991636. ↗ 82
- Becker, Tilman, Rambow, Owen, and Niv, Michael. The derivational generative power of formal systems, or: Scrambling is beyond LCFRS. IRCS Report 92-38, University of Pennsylvania, Philadelphia, USA, 1992. ↗ 86
- Bodirsky, Manuel, Kuhlmann, Marco, and Möhl, Mathias. Well-nested drawings as models of syntactic structure. In *Tenth Conference on Formal Grammar and Ninth Meeting on Mathematics of Language*, Edinburgh, UK, 2005. ↗ 55 and 82
- Boullier, Pierre. On TAG parsing. In *Traitement Automatique des Langues Naturelles (TALN)*, pages 75–84, Cargèse, France, 1999. ↗ 80
- Candito, Marie-Hélène and Kahane, Sylvain. Can the TAG derivation tree represent a semantic graph? An answer in the light of Meaning-Text Theory. In *Fourth Workshop on Tree-Adjoining Grammars and Related Formalisms (TAG+4)*, pages 21–24, Philadelphia, USA, 1998. ↗ 5
- Carme, Julien, Niehren, Joachim, and Tommasi, Marc. Querying unranked trees with stepwise tree automata. In van Oostrom, Vincent, editor, *Rewriting Techniques and Applications. 15th International Conference, RTA 2004, Aachen, Germany, June 3–5, 2004. Proceedings*, volume 3091 of *Lecture Notes in Computer Science*, pages 105–118. Springer, 2004.
DOI: 10.1007/b98160. ↗ 10 and 123
- Cormen, Thomas H., Leiserson, Charles E., Rivest, Ronald L., and Stein, Clifford. *Introduction to Algorithms*. MIT Press, 2nd edition, 2001. ↗ 23
- Courcelle, Bruno. Basic notions of universal algebra for language theory and graph grammars. *Theoretical Computer Science*, 163(1–2):1–54, 1996.
DOI: 10.1016/0304-3975(95)00145-X. ↗ 84 and 85
- Covington, Michael A. A fundamental algorithm for dependency parsing. In *39th Annual ACM Southeast Conference*, pages 95–102, Athens, GA, USA, 2001. ↗ 34

BIBLIOGRAPHY

- Culotta, Aron and Sorensen, Jeffrey. Dependency tree kernels for relation extraction. In *42nd Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 423–429, Barcelona, Spain, 2004.
DOI: 10.3115/1218955.1219009. § 3
- Davey, Brian A. and Priestley, Hilary A. *Introduction to Lattices and Order*. Cambridge University Press, 2nd edition, 2001. § 13
- Denecke, Klaus and Wismath, Shelly L. *Universal Algebra and Applications in Theoretical Computer Science*. Chapman and Hall/CRC, 2001. § 13, 87, and 89
- Diestel, Reinhard. *Graph Theory*, volume 173 of *Graduate Texts in Mathematics*. Springer, 3rd edition, 2005.
URL: <http://www.math.uni-hamburg.de/home/diestel/books/graph.theory/>. § 13
- Dikovsky, Alexander and Modina, Larissa. Dependencies on the other side of the curtain. *Traitement automatique des langues*, 41(1):67–96, 2000. § 22, 32, and 51
- Duchier, Denys and Debusmann, Ralph. Topological dependency trees: A constraint-based account of linear precedence. In *39th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 180–187, Toulouse, France, 2001.
DOI: 10.3115/1073012.1073036. § 86 and 120
- Eisner, Jason and Satta, Giorgio. Efficient parsing for bilexical context-free grammars and Head Automaton Grammars. In *37th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 457–464, College Park, MD, USA, 1999.
DOI: 10.3115/1034678.1034748. § 4 and 114
- Engelfriet, Joost. Bottom-up and top-down tree transformations – a comparison. *Theory of Computing Systems*, 9(2):198–231, 1975.
DOI: 10.1007/BF01704020. § 72 and 90
- Engelfriet, Joost and Maneth, Sebastian. Output string languages of compositions of deterministic macro tree transducers. *Journal of Computer and System Sciences*, 64(2):350–395, 2002.
DOI: 10.1006/jcss.2001.1816. § 72
- Engelfriet, Joost, Rozenberg, Grzegorz, and Slutzki, Giora. Tree transducers, L systems, and two-way machines. *Journal of Computer and System Sciences*, 20(2):150–202, 1980.
DOI: 10.1016/0022-0000(80)90058-6. § 113
- Frank, Robert. *Phrase Structure Composition and Syntactic Dependencies*. Number 38 in *Current Studies in Linguistics*. MIT Press, 2002. § 5

- Fujiyoshi, Akio and Kasai, Takumi. Spinal-formed context-free tree grammars. *Theory of Computing Systems*, 33(1):59–83, 2000.
DOI: 10.1007/s002249910004.  82
- Gaifman, Haim. Dependency systems and phrase-structure systems. *Information and Control*, 8:304–337, 1965.  6, 9, and 32
- Gécseg, Ferenc and Steinby, Magnus. Tree languages. In Rozenberg, Grzegorz and Salomaa, Arto, editors, *Handbook of Formal Languages*, volume 3, pages 1–68. Springer, 1997.  13, 86, 87, 88, and 91
- Gerdes, Kim and Kahane, Sylvain. Word order in German: A formal dependency grammar using a topological hierarchy. In *39th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 220–227, Toulouse, France, 2001.
DOI: 10.3115/1073012.1073041.  60
- Ginsburg, Seymour and Spanier, Edwin H. Semigroups, Presburger formulas, and languages. *Pacific Journal of Mathematics*, 16(2):285–296, 1966.
URL: <http://projecteuclid.org/euclid.pjm/1102994974>.  102
- Goguen, Joseph A., Thatcher, James W., Wagner, Eric G., and Wright, Jesse B. Initial algebra semantics and continuous algebras. *Journal of the Association for Computing Machinery*, 24(1):68–95, 1977.
DOI: 10.1145/321992.321997.  64
- Goldstine, Jonathan. A simplified proof of Parikh’s theorem. *Discrete Mathematics*, 19(3):235–239, 1977.
DOI: 10.1016/0012-365X(77)90103-0.  99
- Gramatovici, Radu and Plátek, Martin. A natural infinite hierarchy by free-order dependency grammars. In *Information Technologies – Applications and Theory*, pages 51–56, Bystrá dolina, Slovak Republic, 2006.  113
- Greibach, Sheila. A new normal-form theorem for context-free phrase structure grammars. *Journal of the Association for Computing Machinery*, 12(1):42–52, 1965.
DOI: 10.1145/321250.321254.  106
- Groenink, Annius V. *Surface without Structure. Word Order and Tractability Issues in Natural Language Analysis*. PhD thesis, Utrecht University, The Netherlands, 1997.
 113
- Guan, Yonggang. *Klammergrammatiken, Netzgrammatiken und Interpretationen von Netzen*. Doctoral dissertation, Saarland University, Saarbrücken, Germany, 1992.  75, 82, and 107

BIBLIOGRAPHY

- Hajič, Jan, Hladka, Barbora Vidova, Panevová, Jarmila, Hajičová, Eva, Sgall, Petr, and Pajas, Petr. Prague Dependency Treebank 1.0. Linguistic Data Consortium, 2001T10, 2001. ↗ 33 and 122
- Hajič, Jan, Panevová, Jarmila, Hajičová, Eva, Sgall, Petr, Pajas, Petr, Štěpánek, Jan, Havelka, Jiří, and Mikulová, Marie. Prague Dependency Treebank 2.0. Linguistic Data Consortium, 2006T01, 2006. ↗ 33 and 122
- Hajičová, Eva, Havelka, Jiří, Sgall, Petr, Veselá, Kateřina, and Zeman, Daniel. Issues of projectivity in the Prague Dependency Treebank. *Prague Bulletin of Mathematical Linguistics*, 81:5–22, 2004. ↗ 33 and 38
- Havelka, Jiří. Projectivity in totally ordered rooted trees. *Prague Bulletin of Mathematical Linguistics*, 84:13–30, 2005. ↗ 22, 31, and 61
- Havelka, Jiří. Beyond projectivity: Multilingual evaluation of constraints and measures on non-projective structures. In *45th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 608–615, Prague, Czech Republic, 2007a. URL: <http://www.aclweb.org/anthology/P/P07/P07-1077.pdf>. ↗ 4, 34, 38, 50, and 62
- Havelka, Jiří. Relationship between non-projective edges, their level types, and well-nestedness. In *Human Language Technologies: The Conference of the North American Chapter of the Association for Computational Linguistics (HLT-NAACL)*, pages 61–64, Rochester, NY, USA, 2007b. URL: <http://www.aclweb.org/anthology/N07-2016>. ↗ 61
- Hays, David G. Grouping and dependency theory. In *National Symposium on Machine Translation*, pages 258–266, Englewood Cliffs, NY, USA, 1960. ↗ 22
- Hays, David G. Dependency theory: A formalism and some observations. *Language*, 40(4):511–525, 1964. DOI: 10.2307/411934. ↗ 3 and 32
- Holan, Tomáš, Kuboň, Vladislav, Oliva, Karel, and Plátek, Martin. Two useful measures of word order complexity. In *Workshop on Processing of Dependency-Based Grammars*, pages 21–29, Montréal, Canada, 1998. ↗ 4 and 38
- Hotz, Günter and Pitsch, Gisela. On parsing coupled-context-free languages. *Theoretical Computer Science*, 161(1–2):205–233, 1996. DOI: 10.1016/0304-3975(95)00114-X. ↗ 9, 75, and 120
- Hudson, Richard. *Word Grammar*. Basil Blackwell, Oxford, UK, 1984. ↗ 22 and 55
- Huybregts, Riny. The weak inadequacy of context-free phrase structure grammars. In de Haan, Ger, Trommelen, Mieke, and Zonneveld, Wim, editors, *Van periferie naar kern*, pages 81–99. Foris, Dordrecht, The Netherlands, 1984. ↗ 2 and 88

- Joshi, Aravind K. Tree Adjoining Grammars: How much context-sensitivity is required to provide reasonable structural descriptions? In *Natural Language Parsing*, pages 206–250. Cambridge University Press, 1985. [§](#) 5, 86, 97, and 116
- Joshi, Aravind K. and Schabes, Yves. Tree-Adjoining Grammars. In Rozenberg, Grzegorz and Salomaa, Arto, editors, *Handbook of Formal Languages*, volume 3, pages 69–123. Springer, 1997. [§](#) 2, 9, 79, and 82
- Joshi, Aravind K., Levy, Leon S., and Takahashi, Masako. Tree Adjunct Grammars. *Journal of Computer and System Sciences*, 10(2):136–163, 1975. [§](#) 79
- Kahane, Sylvain, Nasr, Alexis, and Rambow, Owen. Pseudo-projectivity: A polynomially parsable non-projective dependency grammar. In *36th Annual Meeting of the Association for Computational Linguistics and 18th International Conference on Computational Linguistics (COLING-ACL)*, pages 646–652, Montréal, Canada, 1998. DOI: 10.3115/980451.980953. [§](#) 60
- Kallmeyer, Laura. Comparing lexicalized grammar formalisms in an empirically adequate way: The notion of generative attachment capacity. In *International Conference on Linguistic Evidence*, pages 154–156, Tübingen, Germany, 2006. [§](#) 6 and 82
- Kepser, Stephan and Mönnich, Uwe. Closure properties of linear context-free tree languages with an application to optimality theory. *Theoretical Computer Science*, 354(1):82–97, 2006. DOI: 10.1016/j.tcs.2005.11.024. [§](#) 82
- Koller, Alexander and Striegnitz, Kristina. Generation as dependency parsing. In *40th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 17–24, Philadelphia, USA, 2002. DOI: 10.3115/1073083.1073088. [§](#) 120
- Kracht, Marcus. *The Mathematics of Language*, volume 63 of *Studies in Generative Grammar*. Mouton de Gruyter, 2003. [§](#) 71, 97, 101, and 102
- Kromann, Matthias Trautner. The Danish Dependency Treebank and the underlying linguistic theory. In *Second Workshop on Treebanks and Linguistic Theories (TLT)*, pages 217–220, Växjö, Sweden, 2003. [§](#) 33 and 122
- Kromann, Matthias Trautner. *Discontinuous Grammar. A Model of Human Parsing and Language Acquisition*. Dr. ling. merc. dissertation, Copenhagen Business School, Copenhagen, Denmark, 2005. URL: <http://www.id.cbs.dk/~mtk/thesis/>. [§](#) 33
- Kruijff, Geert-Jan M. Dependency grammar. In *Encyclopedia of Language and Linguistics*, pages 444–450. Elsevier, 2nd edition, 2005. [§](#) 3

BIBLIOGRAPHY

- Kunze, Jürgen. *Die Auslassbarkeit von Satzteilen bei koordinativen Verbindungen im Deutschen*. Akademie-Verlag, Berlin, Germany, 1972. § 32
- Lee, Lillian. Fast context-free grammar parsing requires fast Boolean matrix multiplication. *Journal of the Association for Computing Machinery*, 49(1):1–15, 2002.
DOI: 10.1145/505241.505242. § 114
- Marcus, Solomon. *Algebraic Linguistics: Analytical Models*, volume 29 of *Mathematics in Science and Engineering*. Academic Press, New York, USA, 1967. § 20, 21, 22, and 32
- McAllester, David. On the complexity analysis of static analyses. *Journal of the Association for Computing Machinery*, 49(4):512–537, 2002.
DOI: 10.1145/581771.581774. § 116
- McDonald, Ryan and Pereira, Fernando. Online learning of approximate dependency parsing algorithms. In *Eleventh Conference of the European Chapter of the Association for Computational Linguistics (EACL)*, pages 81–88, Trento, Italy, 2006. § 123
- McDonald, Ryan and Satta, Giorgio. On the complexity of non-projective data-driven dependency parsing. In *Tenth International Conference on Parsing Technologies (IWPT)*, pages 121–132, Prague, Czech Republic, 2007.
URL: <http://www.aclweb.org/anthology/W/W07/W07-2216>. § 4, 6, and 123
- McDonald, Ryan, Pereira, Fernando, Ribarov, Kiril, and Hajič, Jan. Non-projective dependency parsing using spanning tree algorithms. In *Human Language Technology Conference (HLT) and Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 523–530, Vancouver, Canada, 2005.
DOI: 10.3115/1220575.1220641. § 34 and 55
- Mel'čuk, Igor. *Dependency Syntax: Theory and Practice*. State University of New York Press, Albany, USA, 1988. § 55
- Mezei, Jorge E. and Wright, Jesse B. Algebraic automata and context-free sets. *Information and Control*, 11(1–2):3–29, 1967.
DOI: 10.1016/S0019-9958(67)90353-1. § 10, 83, and 122
- Michaelis, Jens and Kracht, Marcus. Semilinearity as a syntactic invariant. In *Logical Aspects of Computational Linguistics. First International Conference, LACL 1996, Nancy, France, September 23–25, 1996. Selected Papers*, volume 1328 of *Lecture Notes in Computer Science*, pages 329–345. Springer, 1997.
DOI: 10.1007/BFb0052165. § 101
- Möhl, Mathias. Drawings as models of syntactic structure: Theory and algorithms. Diploma thesis, Saarland University, Saarbrücken, Germany, 2006. § 32, 55, 60, and 61

- Mönnich, Uwe. Adjunction as substitution. an algebraic formulation of regular, context-free, and tree-adjointing languages. In *Third Conference on Formal Grammar*, pages 169–178, Aix-en-Provence, France, 1997.  82
- Nasr, Alexis. A formalism and a parser for lexicalised dependency grammars. In *Fourth International Workshop on Parsing Technologies (IWPT)*, pages 186–195, Prague, Czech Republic, 1995.  60
- Neuhaus, Peter and Bröker, Norbert. The complexity of recognition of linguistically adequate dependency grammars. In *35th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 337–343, Madrid, Spain, 1997.
DOI: 10.3115/979617.979660.  4 and 120
- Niehren, Joachim and Podelski, Andreas. Feature automata and recognizable sets of feature trees. In Gaudel, Marie-Claude and Jouannaud, Jean-Pierre, editors, *TAPSOFT 1993: Theory and Practice of Software Development. Fourth International Joint Conference CAAP/FASE, Orsay, France, April 13–17, 1993. Proceedings*, volume 668 of *Lecture Notes in Computer Science*, pages 356–375. Springer, 1993.
DOI: 10.1007/3-540-56610-4_76.  10
- Nivre, Joakim. Constraints on non-projective dependency parsing. In *Eleventh Conference of the European Chapter of the Association for Computational Linguistics (EACL)*, pages 73–80, Trento, Italy, 2006a.  4, 34, 39, and 50
- Nivre, Joakim. *Inductive Dependency Parsing*, volume 34 of *Text, Speech and Language Technology*. Springer, 2006b.  32
- Nivre, Joakim and Nilsson, Jens. Pseudo-projective dependency parsing. In *43rd Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 99–106, Ann Arbor, USA, 2005.
DOI: 10.3115/1219840.1219853.  34
- Nivre, Joakim, Hall, Johan, Kübler, Sandra, McDonald, Ryan, Nilsson, Jens, Riedel, Sebastian, and Yuret, Deniz. The CONLL 2007 shared task on dependency parsing. In *Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CONLL)*, pages 915–932, Prague, Czech Republic, 2007.
URL: <http://www.aclweb.org/anthology/D/Do7/Do7-1096>.  3
- Ogden, William. A helpful result for proving inherent ambiguity. *Mathematical Systems Theory*, 2(3):191–194, 1968.
DOI: 10.1007/BF01694004.  93

BIBLIOGRAPHY

- Parikh, Rohit J. On context-free languages. *Journal of the Association for Computing Machinery*, 13(4):570–581, 1966.
DOI: 10.1145/321356.321364. § 10, 97, and 99
- Pericliev, Vladimir and Ilarionov, Ilarion. Testing the projectivity hypothesis. In *Sixth International Conference on Computational Linguistics (COLING)*, pages 56–58, Bonn, Germany, 1986.
DOI: 10.3115/991365.991380. § 32 and 34
- Pullum, Geoffrey K. and Scholz, Barbara C. Contrasting applications of logic in natural language syntactic description. In *Twelfth International Congress on Logic, Methodology and Philosophy of Science*, pages 481–503, London, UK, 2005. § 1
- Quirk, Chris, Menezes, Arul, and Cherry, Colin. Dependency treelet translation: Syntactically informed phrasal SMT. In *43rd Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 271–279, Ann Arbor, USA, 2005.
DOI: 10.3115/1219840.1219874. § 3
- Rambow, Owen. *Formal and Computational Aspects of Natural Language Syntax*. PhD thesis, University of Pennsylvania, Philadelphia, USA, 1994.
URL: <http://wwwlib.umi.com/dissertations/fullcit/9427601>. § 88
- Rambow, Owen and Satta, Giorgio. A two-dimensional hierarchy for parallel rewriting systems. Technical Report IRCS-94-02, University of Pennsylvania, Philadelphia, USA, 2004. § 113
- Raoult, Jean-Claude. Rational tree relations. *Bulletin of the Belgian Mathematical Society*, 4:149–176, 1997. § 75
- Robinson, Jane J. Dependency structures and transformational rules. *Language*, 46(2): 259–285, 1970.
DOI: 10.2307/412278. § 22
- Rosenkrantz, Daniel J. Matrix equations and normal forms for context-free grammars. *Journal of the Association for Computing Machinery*, 14(3):501–507, 1967.
DOI: 10.1145/321406.321412. § 106
- Satta, Giorgio. Recognition of Linear Context-Free Rewriting Systems. In *30th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 89–95, Newark, DE, USA, 1992.
DOI: 10.3115/981967.981979. § 68, 117, and 119
- Schabes, Yves. *Mathematical and Computational Aspects of Lexicalized Grammars*. PhD thesis, University of Pennsylvania, Philadelphia, USA, 1990.
URL: <http://wwwlib.umi.com/dissertations/fullcit/9101213>. § 106

- Schabes, Yves, Abeillé, Ann, and Joshi, Aravind K. Parsing strategies with ‘lexicalized’ grammars: Application to Tree Adjoining Grammars. In *Twelfth International Conference on Computational Linguistics (COLING)*, pages 578–583, Budapest, Hungary, 1988.
DOI: 10.3115/991719.991757. ↗ 5, 79, and 114
- Seki, Hiroyuki, Matsumura, Takashi, Fujii, Mamoru, and Kasami, Tadao. On Multiple Context-Free Grammars. *Theoretical Computer Science*, 88(2):191–229, 1991.
DOI: 10.1016/0304-3975(91)90374-B. ↗ 75, 107, and 113
- Shieber, Stuart M. Evidence against the context-freeness of natural language. *Linguistics and Philosophy*, 8(3):333–343, 1985.
DOI: 10.1007/BF00630917. ↗ 2 and 88
- Shieber, Stuart M., Schabes, Yves, and Pereira, Fernando. Principles and implementation of deductive parsing. *Journal of Logic Programming*, 24(1–2):3–36, 1995.
DOI: 10.1016/0743-1066(95)00035-I. ↗ 115
- Simion, Rodica. Noncrossing partitions. *Discrete Mathematics*, 217(1–3):367–409, 2000.
DOI: 10.1016/S0012-365X(99)00273-3. ↗ 57
- Sleator, Daniel and Temperley, Davy. Parsing English with a Link Grammar. Technical report, Carnegie Mellon University, Pittsburgh, USA, 1991.
URL: <http://arxiv.org/abs/cmp-lg/9508004v1>. ↗ 55
- Sleator, Daniel and Temperley, Davy. Parsing English with a Link Grammar. In *Third International Workshop on Parsing Technologies (IWPT)*, pages 277–292, Tilburg, The Netherlands, and Durbuy, Belgium, 1993. ↗ 54
- Sloane, Neil J. A. The on-line encyclopedia of integer sequences, 2007.
URL: <http://www.research.att.com/~njas/sequences/>. ↗ 55 and 62
- Tesnière, Lucien. *Éléments de syntaxe structurale*. Klincksieck, Paris, France, 1959. ↗ 3
- Tesnière, Lucien. *Grundzüge der strukturalen Syntax*. Klett-Cotta, Stuttgart, Germany, 1980. ↗ 3
- Thatcher, James W. and Wright, Jesse B. Generalized finite automata theory with an application to a decision problem of second-order logic. *Mathematical Systems Theory*, 2(1):57–81, 1968.
DOI: 10.1007/BF01691346. ↗ 10 and 123
- van Vugt, Nikè. Generalized context-free grammars. Master’s thesis, Universiteit Leiden, The Netherlands, 1996. ↗ 75 and 113

BIBLIOGRAPHY

- Veselá, Kateřina, Havelka, Jiří, and Hajičová, Eva. Condition of projectivity in the underlying dependency structures. In *20th International Conference on Computational Linguistics (COLING)*, pages 289–295, Geneva, Switzerland, 2004.
DOI: 10.3115/1220355.1220397. § 4 and 22
- Vijay-Shanker, K., Weir, David J., and Joshi, Aravind K. Characterizing structural descriptions produced by various grammatical formalisms. In *25th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 104–111, Stanford, CA, USA, 1987.
DOI: 10.3115/981175.981190. § 9 and 68
- Weir, David J. *Characterizing Mildly Context-Sensitive Grammar Formalisms*. PhD thesis, University of Pennsylvania, Philadelphia, USA, 1988.
URL: <http://wwwlib.umi.com/dissertations/fullcit/8908403>. § 9, 68, and 113
- Weir, David J. Linear context-free rewriting systems and deterministic tree-walking transducers. In *30th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 136–143, Newark, DE, USA, 1992.
DOI: 10.3115/981967.981985. § 75
- Yli-Jyrä, Anssi. Multiplanarity – a model for dependency structures in treebanks. In *Second Workshop on Treebanks and Linguistic Theories (TLT)*, pages 189–200, Växjö, Sweden, 2003. § 4 and 55
- Zeman, Daniel. *Parsing with a Statistical Dependency Model*. PhD thesis, Charles University, Prague, Czech Republic, 2004. § 34 and 38