

**ALE**

**The Attribute Logic Engine**  
**User's Guide**  
**with TRALE extensions**

Version 4.0 Beta  
December 2003

©1992–1995, Bob Carpenter and Gerald Penn  
©1998, 1999, 2001 Gerald Penn  
©2003 Gerald Penn and Mohammad Haji-Abdolhosseini  
(except in extensions as indicated)

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>TRALE Introduction</b>	<b>T2-1</b>
<b>3</b>	<b>Prolog Preliminaries</b>	<b>3</b>
3.1	Terms . . . . .	3
3.2	Space and Comments . . . . .	4
3.3	Running Prolog . . . . .	4
3.4	Queries . . . . .	4
3.5	Running ALE . . . . .	4
T3.1	Running TRALE . . . . .	T4-1
3.6	Compiling a Grammar . . . . .	5
3.7	Exiting Prolog and Breaking . . . . .	5
3.8	Saved States . . . . .	5
<b>4</b>	<b>Feature Structures, Types and Descriptions</b>	<b>6</b>
4.1	Inheritance Hierarchies . . . . .	6
4.2	Feature Structures . . . . .	8
4.3	Subsumption and Unification . . . . .	10
4.3.1	Subsumption . . . . .	10
4.3.2	Unification . . . . .	11
4.4	Inequations . . . . .	12
4.5	Type System . . . . .	13
4.6	Interaction with the signature . . . . .	17
T4.1	TRALE Signatures . . . . .	T17-1
T4.1.1	Signature specification files . . . . .	T17-1
T4.1.2	Subtype covering . . . . .	T17-3
4.7	Extensionality . . . . .	18
4.8	a_/1 Atoms . . . . .	21
4.9	Attribute-Value Logic . . . . .	22
4.9.1	Enforcement of Inequations . . . . .	26
4.10	Macros . . . . .	28
T4.2	TRALE Descriptions . . . . .	T31-1
T4.2.1	Logical variable macros . . . . .	T31-1
T4.2.2	Macro hierarchies . . . . .	T31-3
T4.2.3	Automatic generation of macros on different types . . . . .	T31-4
T4.2.4	Example . . . . .	T31-7

4.11	Functional Descriptions . . . . .	32
4.12	Type Constraints . . . . .	33
4.13	Example: The Zebra Puzzle . . . . .	34
T4.3	Complex-antecedent constraints . . . . .	T38-1
<b>5</b>	<b>Definite Clauses</b>	<b>39</b>
5.1	Type Constraints Revisited . . . . .	43
5.2	Co-routining . . . . .	44
5.2.1	Shared variables in conditionals . . . . .	45
<b>6</b>	<b>Phrase Structure Grammars</b>	<b>47</b>
6.1	Lexical Entries . . . . .	47
6.2	Empty Categories . . . . .	49
6.3	ALE Lexical Rules . . . . .	51
6.4	Grammar Rules . . . . .	55
6.4.1	Procedural Attachments . . . . .	57
6.4.2	The <code>cats&gt;</code> Operator . . . . .	59
6.4.3	Parsing . . . . .	60
6.4.4	Generation . . . . .	62
T6.1	Test sequences . . . . .	T64-1
T6.1.1	Test files . . . . .	T64-1
T6.1.2	Test queries . . . . .	T64-1
<b>7</b>	<b>TRALE Lexical Rule Compiler</b>	<b>T64-1</b>
T7.1	Introduction . . . . .	T64-1
T7.2	Using the lexical rule compiler . . . . .	T64-1
T7.2.1	Input syntax . . . . .	T64-2
T7.2.2	Interpretation . . . . .	T64-6
<b>8</b>	<b>Topological Parsing</b>	<b>T64-1</b>
T8.1	Introduction . . . . .	T64-1
T8.2	Running TGC . . . . .	T64-1
T8.2.1	Compiling Grammars . . . . .	T64-1
T8.2.2	Using the Parser . . . . .	T64-2
T8.3	Grammars . . . . .	T64-3
T8.3.1	Phenogrammatical Rules . . . . .	T64-3
T8.3.2	Linkage . . . . .	T64-4
T8.3.3	Matching . . . . .	T64-5
T8.3.4	Covering . . . . .	T64-8
T8.3.5	Compaction . . . . .	T64-9
T8.3.6	matrix . . . . .	T64-9
T8.3.7	Tectogrammatical Rules . . . . .	T64-9
T8.4	The Lexicon . . . . .	T64-13
T8.4.1	Lexical Categories . . . . .	T64-13
T8.5	Edges . . . . .	T64-14
T0.1	Syntax of A Topological Grammar . . . . .	T64-16
T0.2	TGC Error and Warning Messages . . . . .	T64-18

<b>A</b>	<b>Compiling ALE Programs</b>	<b>65</b>
1.1	File Management . . . . .	65
1.2	Compiling Programs . . . . .	66
1.3	Compile-Time Error Messages . . . . .	69
<b>B</b>	<b>Running and Debugging ALE Programs</b>	<b>71</b>
2.1	Testing the Signature . . . . .	71
2.2	Evaluating Descriptions . . . . .	73
2.3	Hiding Types and Features . . . . .	76
2.4	Evaluating Definite Clause Queries . . . . .	77
2.5	Displaying Grammars . . . . .	79
2.6	Executing Grammars: Parsing . . . . .	82
2.7	Executing Grammars: Generation . . . . .	87
2.8	Mini-interpreter (parsing only) . . . . .	90
2.9	Subsumption Checking (parsing only) . . . . .	93
T2.1	The Chart Display . . . . .	T94-1
T2.1.1	Installation and Customization . . . . .	T94-1
T2.1.2	Working with the Chart Display . . . . .	T94-2
2.10	Source-Level Debugger . . . . .	95
2.10.1	Running without XEmacs . . . . .	95
2.10.2	Running with XEmacs . . . . .	95
2.10.3	Debugger Commands . . . . .	96
2.10.4	Debugger Ports and Steps . . . . .	97
2.10.5	Leashing . . . . .	98
2.10.6	Skipping . . . . .	99
2.10.7	Breakpoints . . . . .	100
<b>C</b>	<b>[incr TSDB0]</b>	<b>T101-1</b>
T3.1	Installing [incr TSDB0] . . . . .	T101-1
T3.1.1	Setting up Shell Variables . . . . .	T101-1
T3.1.2	Setting up .tsdbrc . . . . .	T101-2
T3.1.3	Setting up PVM . . . . .	T101-2
T3.1.4	Initializing a CPU . . . . .	T101-3
T3.1.5	Creating a Test Suite . . . . .	T101-4
<b>D</b>	<b>Pretty-printing Hooks</b>	<b>102</b>
4.0.1	Portraying feature structures . . . . .	102
4.0.2	Portraying inequations . . . . .	103
4.0.3	A sample pretty-printing hook . . . . .	104
<b>E</b>	<b>TRALE Output</b>	<b>T106-1</b>
T5.1	Saving of outputs . . . . .	T106-1
T5.2	Grisu interface support . . . . .	T106-2
T5.2.1	Unfilling . . . . .	T106-2
T5.2.2	Feature ordering . . . . .	T106-2
T5.2.3	Diff of feature structures . . . . .	T106-3
<b>F</b>	<b>ALE Keyword Summary</b>	<b>107</b>

<b>G Flag Summary</b>	<b>111</b>
<b>H References</b>	<b>114</b>
<b>A Sample Grammars</b>	<b>118</b>
1.1 English Syllabification Grammar . . . . .	118
1.2 Categorical Grammar with Cooper Storage . . . . .	123
1.3 Simple Generation Grammar . . . . .	128
<b>B Error and Warning Messages</b>	<b>132</b>
2.1 Error Messages . . . . .	132
2.2 Warning Messages . . . . .	136
<b>C BNF for ALE</b>	<b>138</b>
T3.1 BNF for TRALE . . . . .	T141-1
<b>D Reference Card</b>	<b>142</b>

# **Preface – Version 4.0**

TO BE ADDED

# Preface – Version 3.0

ALE 3.0 is completely compatible with ALE 2.0 grammars, and adds the following new features:

- A semantic-head-driven generator, based on the algorithm presented in Shieber et al. (1990). The generator was adapted to the logic of typed feature structures by Octav Popescu in his Carnegie Mellon Master's Thesis, Popescu (1996). Octav also wrote most of the generation code for this release. Grammars can be compiled for parsing only, generation only, or both. Some glue-code is also available from the ALE homepage, to parse and generate with different grammars through a unix pipe.
- A source-level debugger with a graphical XEmacs interface. This debugger works only with SICStus Prolog 3.0.6 and higher. A debugger with reduced functionality will be made available to SWI Prolog users in a later release. This debugger builds on, and incorporates the functionality of the code for the SICStus source-level debugger, written by Per Mildner at Uppsala University.
- Functional descriptions. Instead of binding variables in a description and calling a procedural attachment, e.g., `a cons f:X,g:Y,h:Z goal append(X,Y,Z)`, it is now possible to incorporate certain functional relations into descriptions themselves, e.g., `a cons f:X,g:Y,h:append(X,Y)`.
- `a_/1` atoms. There are now an infinite number of atoms (types with no appropriate features), implicitly declared in every signature. These atoms can be arbitrary Prolog terms, including unbound variables, and can be used wherever normal ALE types can, e.g., `f:(a_ p(3.7))`. `a_/1` atoms are extensional as Prolog terms, i.e., are taken to be identical according to the Prolog predicate, `==/2`. In particular, this means that ground atoms behave exactly as ALE extensional types.
- Optional edge subsumption checking. For completeness of parsing, one only needs to ensure that, for every pair of nodes in the chart, the most general feature structure spanning those nodes is stored in the chart. This can reduce the number of edges in many domains.
- An autonomous `intro/2` operator. Features can now be declared on their own in a separate part of the grammar.
- Default specifications for types. These are NOT default types. If a type appears on the right-hand side of a `sub/2` or `intro/2` specification, but not on

the left-hand side of one, ALE will assume this type is maximal, i.e., assume the specification, *Type* sub []. Similarly, if it occurs on a left-hand side, but not on a right-hand side, ALE will assume the type is immediately subsumed by `bot`, the most general type. In both cases, ALE will announce these assumptions during compilation.

- Several bug corrections and more compile-time warning and error messages.
- An SWI Prolog 2.9.7 port. Version 3.0 will be the last version for which we will port the system to Quintus Prolog. We will now support ALE for SICStus Prolog and SWI Prolog. The SICStus version of the system works for SICStus Prolog 3.0.5 and higher, except for the source-level debugger, which requires version 3.0.6.

We would like to thank all of the users who have supplied us with feedback and suggestions over the past few years for how to improve ALE. In particular, we would like to thank Ion Androutsopoulos, Mike Calcagno, Mats Carlsson, Frederik Fouvry, Gertjan van Noord, Peter van Roy, Margriet Verlinden, and Jan Wielemaker for their patient assistance. This release incorporates many, but unfortunately not all of those changes. Quite a few more will be made in the next several releases, along with many performance improvements.

Bob Carpenter and Gerald Penn  
Murray Hill and Tuebingen, March 1998



# Preface – Version 2.0

ALE 2.0 is a proper extension of version 1.0. Specifically, version 2.0 will run any grammar that will run under version 1.0. But version 2.0 includes many extensions to version 1.0, including the following.

- Inequations
- Extensionality
- General Constraints on Types
- Mini-interpreter
- Error-suppression

Inequations allow inequational constraints to be imposed between two structures. Extensionality allows structures of specified types to be identified if they are structurally identical. Together, these provide the ability to simulate Prolog II programs (Colmerauer 1987). ALE 2.0 also allows general constraints to be placed on types, using arbitrary descriptions from the constraint language, including path equations, inequations and disjunctions, and procedural attachments. It also has a mini-interpreter, which allows the user to traverse and edit an ALE parse tree. Error messages for incompatible descriptions are now automatically disabled during lexicon and empty category compilation.

The second release of ALE, Version 2.0, is based on an extension of the first version of ALE, that was completed for Gerald Penn's (1993) MS Project in the Computational Linguistics Program at Carnegie Mellon University.

There are many people whom we would like to thank for their comments and feedback on version 1.0 and  $\beta$ -versions of 2.0. These people have actually used the system in their research and have thus had the best opportunity to provide us with practical feedback. First, we would like to thank the first group of users, housed at Sharp Laboratories of Europe, located in Oxford, England, including Pete Whitelock, Antonio Sanfillipo, and Osamu Nishida. They not only used the system but provided feedback on the code. Secondly, the group at University of Tübingen, who are developing a competing system, Troll, have rigorously tested existing systems, including ALE, both for their ability to express grammars naturally and for efficiency. Specifically, we would like to thank Detmar Meurers, Dale Gerdemann, Thilo Götz, Paul King, John Griffith, and Erhard Hinrichs. John and Thilo also provided the changes necessary for the system to run directly in Quintus Prolog. This group is undoubtedly the best informed when it comes to implemented grammar formalisms. We would also like to thank the grammar

development group at Stanford University, including Ivan Sag, Chris Manning, Suzanne Riehemann. We would further like to thank Bob Kasper, Carl Pollard, and Andreas Kathol of the Ohio State University, for a great deal of feedback on the design of HPSG grammars in general, and ALE implementations of them in particular. Chris Manning, in addition, found a bug in SICStus Prologs prior to 2.1.8, which prevented cyclic structures from being used in completed chart edges, a bug found by both Steven Bird of Edinburgh and C. J. Rupp of IDSIA. Their feedback on Bob Carpenter's prototype implementation of HPSG for English led to the design of Gerald Penn's much more comprehensive implementation of HPSG and was the primary impetus for the importation of general type constraints into version 2.0. Next, we would like to thank Claire Gardent, who has been using ALE to develop discourse grammars in Amsterdam. We should also thank Carsten Guenther and Markus Walther, of the Universities of Hamburg and Düsseldorf, respectively, who have used the system to develop phonological grammars. Finally, we should thank Michael Mastroianni, who implemented a comprehensive approach to constraint-based phonology in ALE (Mastroianni 1993). He suffered through early, buggy versions of the system, thus sparing the rest of us much of that pain. The feedback we received from these users was invaluable.

We would like to thank EAGLES, the European Advisory Group on Linguistic Engineering Standards, for allowing us to present our system at a meeting in Saarbrücken in March 1993 of the European Expert Group on Linguistic Formalisms devoted to implemented formalisms. We learned a great deal from the other participants in the workshop including especially Jochen Dörre, Michael Dorna, and Martin Emele, of Stuttgart, and Andreas Podelski, then associated with the Digital Equipment Paris Research Lab. We also benefitted from discussions with Hans Uszkoreit, Rolf Backofen, and Uli Krieger, of Saarbrücken, Steve Pulman from SRI in Cambridge, and C. J. Rupp and Graham Russell, of ISSCO in Switzerland.

We had many discussions of the ALE formalism at the HPSG workshop running concurrently with the LSA Linguistic Institute in Columbus. We would especially like to thank Gregor Erbach for comments on our system, including benchmark test results. We would also like to thank Hiroshi Tusda, of the Institute for New Gernateion Computer Technology, for discussion of our systems and comparisons to his system, cu-Prolog. We also discussed ALE heavily during the workshop on implementations of attribute-value logics, during the 1993 Summer School on Logic, Language, and Information in Lisbon, Portugal. We especially benefitted from discussions with Suresh Manandhar, of the University of Edinburgh, and Gerrit Rentier of Tilburg University, and Gert Webelhuth of the University of North Carolina, among those we have not already thanked. We also benefitted from discussions with Ed Stabler and Mark Johnson, and from sitting in on their class on the implementation of constraint-based grammars.

We would also like to thank Ann Copestake and Ted Briscoe, of the Cambridge Computing Laboratory, for feedback on the design of the system.

We would like to thank Richard O'Keefe, who provided some invaluable feedback on coding style. Of course, any glitches or failure to follow his excellent example are our own.

We would also like to thank Elizabeth Hinkelman, who runs the Software Registry, and Mark Kantrowitz, who administers the Prolog Resource Guide and the

Prime Time Freeware for AI CD-ROM. They have helped in publicizing the system description as well as providing access.

The extensions we have not made, though would like to, include the addition of:

- Primitive, Atomic Data Types
- Parametric Types
- Partial Type Inference
- Assert Mode in Compiler
- Peephole Code Optimization
- Subsumption Checking of Chart Edges

The incorporation of Prolog data types such as Real, Integer, Character, and String, is straightforward theoretically, but not so straightforward in terms of ALE. The same goes for parametric polymorphism at the type level. Partial type inference could provide a great deal of optimization in some circumstances. We could not figure out how to incorporate these three changes without drastically modifying the underlying representations and algorithms. The remaining changes are lying dormant because we have other obligations.

The next wave of development of attribute-logic grammars should not be in Prolog, but rather through the use of a direct abstract machine. Bob Carpenter has worked on an abstract machine with Yan Qu, in the context of her MS project in the Carnegie Mellon Computational Linguistics Program, and with Shuly Wintner, of the Technion, in Haifa, Israel, who is writing a PhD dissertation on the topic. Such an undertaking is also underway among the LIFE community, led by Hassan Ait-Kaci, Andreas Podelski, and Peter van Roy.

We would like to thank a number of people for discovering bugs and providing comments on Version 2.0: Ingo Schroeder, Frank Morawietz, Detmar Meurers, Rob Malouf, Frederik Fouvry, Jo Calder, and Suresh Manandhar.

Finally, we would like to thank Jo Calder, Chris Brew, Kevin Humphreys, and Mike Reape, who developed the Pleuk grammar development environment as well as interfacing it to ALE. Details of that system can be found in the appropriate Appendix.

This material is based upon work supported under a National Science Foundation Graduate Research Fellowship (for Gerald Penn). Any opinions, findings, conclusions or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Bob Carpenter and Gerald Penn  
Pittsburgh, August 1994

## Preface – Version Beta

A number of people have asked me to make this system, along with its documentation, available to the public. Now that it's available, I hope that it's useful. But a word of caution is in order. The system is still only a prototype, hence the label "version  $\beta$ ."

Any bug reports would be greatly appreciated. But what I'd really like is comments on the functionality of the system, as well as on the utility of its documentation. I am also interested in hearing of any applications that are made of the system. I would also be glad to answer questions about the system. I have tried to document the strategies used by ALE in this guide. I have also tried to comment the code to the point where it might be adaptable by others. I would, of course, be interested in any kind of improvements or extensions that are discovered or developed, and would like to have the chance to incorporate any such improvements in future versions of this package.

In the implementation, I have endeavored to follow the logic programming methodology laid out by O'Keefe (1990), but there are many spots where I have fallen short. Thus the code is not as fast as it could be, even in Prolog. But I view this system more as a prototype, indicating the utility of a typed logic programming and grammar development system. Borrowing techniques from the WAM directly, implementing an abstract machine C, would lead to roughly a 100-fold speedup, as there is no reason that ALE should be slower than Prolog itself.

I would like to acknowledge the help of Gerald Penn in working through many implementation details of a general constraint resolver, which was the inspiration for this implementation. This version of the system is a great improvement on the last version due to Gerald's work on the system. Secondly, I would like to thank Michael Mastroianni, who has actually used the system to develop grammars for phonology. Finally, I would like to thank Carl Pollard and Bob Kasper for looking over a grammar of HPSG coded in ALE and providing the impetus for the inclusion of empty categories and lexical rules.

The system is available without charge from the author. It is designed to run in either SICStus or Quintus Prologs.

Bob Carpenter  
Pittsburgh, 1992



# Chapter 1

## Introduction

This report serves as an introduction to both the ALE formalism and its Prolog implementation. ALE is an integrated phrase structure parsing and definite clause logic programming system in which the terms are typed feature structures. Typed feature structures combine type inheritance and appropriateness specifications for features and their values. The feature structures used in ALE generalize the common feature structure systems found in the linguistic programming systems PATR-II and FUG, the grammar formalisms HPSG and LFG, as well as the logic programming systems Prolog-II and LOGIN. Programs in any of these languages can be encoded directly in ALE.

Terms in grammars and logic programs are specified in ALE using a typed version of Rounds and Kasper's attribute-value logic with variables. At the term level, we have variables, types, feature value restrictions, path equations, inequations, general constraints, and disjunction. The definite clause programs allow disjunction, negation and cut, specified with Prolog syntax. Phrase structure grammars are specified in a manner similar to DCGs, allowing definite clause procedural attachment. The grammar formalism also fully supports empty categories. Lexical development is supported by a very general form of lexical rule which operates on both categories and surface strings. Macros are available to help organize large descriptions, either in programs or in grammars. Both definite clause programs and grammars are compiled into abstract machine instructions. These instructions are then interpreted by an emulator compiled from the type specifications. Like Prolog compilers, a structure copying strategy is used for matching both definite clauses and grammar rules.

For parsing, ALE compiles from the grammar specification a Prolog-optimized bottom-up, dynamic chart parser. Definite clauses are also compiled into Prolog. As it stands, the current version of ALE, running the feature-structure-based naive reverse of a 30-element list on top of the SICStus 3.7 native code compiler, runs at roughly 152,300 logical inferences per second (LIPS) on a Sun Ultra Enterprise 450. This is slightly faster than the speed of the SICStus 3.7 interpreter on the Prolog naive reverse, 60% as fast as SWI Prolog 3.2.7, and about 9% as fast as the SICStus 3.7 compact-code compiler. The definite clause compiler performs last call optimization, but does not index first arguments or use specialised list cells at the WAM level. Thus it will perform relatively well versus non-optimized interpreters, but lag further behind compiled grammars when programs are written in

a more sophisticated manner than naive reverse.

Full details of the theory behind ALE can be found in Carpenter (1992).

The user who is only interested in definite clause programming can skip the material on phrase structure grammars, while those interested in only grammars without procedural attachments may skip the material in the section on definite clauses.

## Chapter 2

# TRALE Introduction

TRALE is an extension of the ALE system that supports some extra functionality (most importantly, complex-antecedent constraints), and that automatically adds extra constraints to grammars in order to enforce the view of subtyping standardly assumed in Head-driven Phrase Structure Grammar. TRALE is implemented as a pre-processor for ALE that intercepts certain grammar clauses at compile-time to generate extra code for the ALE compiler. This code is used internally — no new ALE grammar is generated. As a result of TRALE's different view of subtyping, many grammars will behave very differently with TRALE than they do with ALE .

The sections of this manual marked either with TRALE-EXTENSION (HTML) or with a T page-number prefix (PS, PDF) describe these extra features. Complex-antecedent constraints are discussed in Section T4.3. Subtype covering constraints, and the HPSG style of subtyping that they support are discussed in Section T4.1.2. They also present some extra syntactic constructs that TRALE can use, including signature declaration files (Section T4.1), logical-variable macros, macro hierarchies and macro generation tools (Section T4.2).

The Grisu graphical interface that is distributed with TRALE is described in Chapter E.

Grisu was developed by Holger Wunsch of Eberhard-Karls-Universität Tübingen and Detmar Meurers of Ohio State University. The signature file parser was developed as part of the ConTroll system by Thilo Götz, Stephan Kepser, Dale Gerdemann, and Detmar Meurers at Eberhard-Karls-Universität Tübingen. The macro hierarchy and macro generation components were developed by Detmar Meurers. The lexical rule compiler described in Chapter 7 was developed by Detmar Meurers, Vanessa Metcalf and Markus Dickinson of Ohio State University. Integration of [incr tsdb()] and the DFKI graphical parsing chart display was performed by Stefan Müller of Universität Bremen, based in part upon earlier work by Frederik Fouvry at Eberhard-Karls-Universität Tübingen.



## Chapter 3

# Prolog Preliminaries

While it is not absolutely necessary, some familiarity with logic programming in general, and Prolog in particular, is helpful in understanding the definite clause portion of ALE. Similarly, experience with unification grammar systems such as PATR-II, DCGs, or FUG is helpful in understanding the phrase structure component of the system. In particular, writing efficient programs and grammars in ALE involves the same kinds of strategies necessary for writing efficient programs in Prolog or PATR-II. For those not familiar with Prolog, the sequence of two books by Sterling and Shapiro (1986) and by O'Keefe (1990) are excellent general introductions to the theory and practice of logic programming. For those not familiar with unification-based grammar formalisms, Shieber (1986), Gazdar and Mellish (1987) and Pereira and Shieber (1987) are useful resources.

For those not familiar with Prolog, we need to point out the salient features of the language which will be assumed throughout this report. This section contains all of the information necessary about Prolog required to run ALE.

### 3.1 Terms

A Prolog *constant* is composed of either a sequence of characters and/or underscores, beginning with a lower case letter, a number, or any sequence of symbols surrounded by apostrophes. So, `abc`, `johnDoe`, `b_17`, `123`, `'JohnDoe'`, `'65$'`, and `'_65a.'` are constants, but `A19`, `JohnDoe`, `B_112`, `_au8`, and `[dd,e]` are not (although see the warning at the end of this section). A variable, on the other hand, is any string of letters, underscores or numbers beginning with a capital letter. Thus `C`, `C_foo`, and `TR5ab` are variables, but `1Xa`, `aXX`, and `_Xy`<sup>1</sup> are not.

In general, it is a bad idea to have constants or variables which are only distinguished by the capitalization of some of their letters. For instance, while `aBa` and `aba` are different constants, they should not both be used in one program. One reason for this in the context of ALE is that the output routines adopt standard capitalization conventions which hide the differences between such constants.

**Warning:** As pointed out to us by Ingo Schroeder, constants or atoms beginning with a capital letter are not treated properly by the compiler. Thus constants

---

<sup>1</sup>Technically, a variable may begin with an underscore, but such variables, said to be *anonymous*, have a very different status than those which begin with a capital letter. The use of anonymous variables is discussed later.

such as ' Foo ' should not be used.

## 3.2 Space and Comments

In your own program and grammar files, extra *whitespace* between symbols beyond that needed to separate constants or variables is ignored. Whitespace consists of either spaces, blank lines or line breaks are ignored. This allows you to format your programs in a manner that is readable. Furthermore, any symbols on a line appearing after a % symbol are treated as *comments* and ignored.

## 3.3 Running Prolog

To fire up Prolog locally, you should contact your systems administrator. You should have either SICStus or SWI Prolog, or a Prolog compiler compatible with one of these. Once Prolog is fired up, you will see a *prompt*. The Prolog prompt should look like:

```
| ?-
```

It is important that Prolog be invoked from a directory for which the user has write permission. ALE, in the process of compiling user programs, writes a number of local files.

## 3.4 Queries

What you type after the prompt is called a *query*. Queries should always end with a period and be followed by a carriage return. In fact, all of the grammar rules, definite clauses, macros and lexical entries in your programs should also end with periods. Most of the interface in ALE is handled directly by top-level Prolog queries. Many of these will return *yes* or *no* after they are called, the significance of which within ALE is explained on a query by query basis.

## 3.5 Running ALE

To run ALE, it is only necessary to type the following query:

```
| ?- compile(ale).
```

assuming that `ale.pl` is the file in which ALE resides. SWI Prolog users must type:

```
| ?- consult(ale).
```

If `ale.pl` is not local to the directory from which Prolog was invoked, a path-name can be added in single-quotes. The syntax of pathnames is operating-system specific.

In SICStus Prolog, when ALE is loaded, it turns on Prolog character escapes. ALE will not be able to generate code properly during compilation without this.

## T3.1 Running TRALE

Similarly, to run TRALE, one types:

```
| ?- compile(trale).
```

assuming that `trale.pl` is the file in which TRALE resides, or, in SWI Prolog:

```
| ?- consult(trale).
```

TRALE needs to load several other files (including `ale.pl`). All of these files are assumed to live in the directory from which Prolog is executed. If this is not true, set the `TRALE_HOME` environment variable to point to the directory in which TRALE is installed *before* running Prolog from that shell.

## 3.6 Compiling a Grammar

Once ALE has been loaded, you can load and compile an ALE grammar as follows:

```
| ?- compile_gram(hpsg).
```

where `hpsg.pl`, for example, is the name of your grammar file (or the name of the file that loads different components of a grammar from other files). This consults the grammar as Prolog source, and compiles it as ALE source.

NOTE: if you only load your grammar with a Prolog command:

```
| ?- consult(hpsg).
```

ALE will not be able to parse or execute other queries with it. You must use `compile_gram/0` (if already consulted) or `compile_gram/2`.

## 3.7 Exiting Prolog and Breaking

To exit from Prolog, you can type `halt` at any prompt (followed by a period, of course).

If you find Prolog hanging at some point, and you are working on a standard Unix implementation, typing `control-c` should produce something like the following message:

```
Prolog interruption (h for help)?
```

You should reply with the character `a`, with or without a following period, followed by a carriage return. If this doesn't work, typing `control-z` should take you out of Prolog altogether.

## 3.8 Saved States

All information concerning an ALE state is encoded in the current Prolog state. Thus, any options presented by your version of Prolog to save states should be able to save ALE states. Saving run-time stacks is only necessary to preserve a currently executing parse or other top-level ALE query, not compiled code.

## Chapter 4

# Feature Structures, Types and Descriptions

This section reviews the basic material from Carpenter (1992), Chapters 1–10, which is necessary to use ALE.

### 4.1 Inheritance Hierarchies

ALE is a language with strong typing. What this means is that every structure it uses comes with a type. These types are arranged in an inheritance hierarchy, whereby type constraints on more general types are inherited by their more specific subtypes, leading to what is known as *inheritance-based polymorphism*. Inheritance-based polymorphism is a cornerstone of object-oriented programming. In this section, we discuss the organization of types into an inheritance hierarchy. Thus many types will have *subtypes*, which are more specific instances of the type. For instance, `person` might have subtypes `male` and `female`.

ALE does much of its processing of types at compile time, as it is reading and processing the grammar file. Thus the user is required to declare all of the types that will be used along with the subtyping relationship between them. An example of a simple ALE type declaration is as follows:

```
bot sub [b,c].           % two basic types -- b and c
  b sub [d,e].
    d sub [g,h].
    e sub [].
  c sub [d,f].           % b and c unify to d
  f sub [].
```

There are quite a few things to note about this declaration. The types declared here are `bot`, `b`, `c`, `d`, `e`, `f` and `g`. Note that each type that is mentioned gets its own specification. Of course, the whitespace is not important, but it is convenient to have each type start its own line. A simple type specification consists of the name of the type, followed by the keyword `sub`, followed by a list of its subtypes (separated by whitespace). In this case, `bot` has two subtypes, `b` and `c`, while `f`, `d` and `e` have no subtypes. The subtypes are specified by a Prolog list. In this

case, a Prolog list consists of a sequence of elements separated by commas and enclosed in square brackets. Note that no whitespace is needed between the list brackets and types, between the types and commas, or between the final bracket and the period. Whitespace is only needed between constants. The extra whitespace on successive lines is conventional, indicating the level in the ordering for the user, but is ignored by the program. Also notice that there are comments on two of the lines; recall that comments begin with a % sign and continue the length of the line. Every type (except `a_/1` atoms, discussed below) must have at most one `sub` declaration, i.e., all of the immediate subtypes must be declared in one declaration.

The subtyping relation is only specified by *immediate* subtyping declarations; but subtyping itself is transitive. Thus, in the example, `d` is a subtype of `c`, and `c` is a subtype of `bot`, so `d` is also a subtype of `bot`. The user only needs to specify the direct subtyping relationship. The transitive closure of this relation is computed by the compiler. While redundant specifications, such as putting `d` directly on the subtype list of `bot`, will not alter the behavior of the compiler, they are confusing to the reader of the program and should be avoided. In addition, the derived transitive subtyping relationship must be anti-symmetric. In particular, this means that there should not be two distinct types each of which is a subtype of the other.

There are two additional restrictions on the inheritance hierarchy beyond the requirement that it form a partial order. First, there is a special type `bot`, which must be declared as the unique most general type. In other words, every type must be a subtype of `bot`. If a type is used on the left-hand side of a `sub` declaration, but never declared as a sub-type of anything else, it is assumed that this type is an immediate subtype of `bot`. Similarly, ALE assumes that all types for which no subtypes are declared are maximal, i.e., have no subtypes.

The second and more subtle restriction on type hierarchies is that they be *bounded complete*. Since type declarations must be finite, this amounts to the restriction that every pair of types which have a common subtype have a unique most general common subtype. In the case at hand, `b` and `c` have three common subtypes, `d`, `g`, and `h`. But these subtypes of `b` and `c` are ordered in such a way that `d` is the most general type in the set, as both `g` and `h` are subtypes of `d`. An example of a type declaration violating this condition is:

```
bot sub [a,b].
  a sub [c,d].
    c sub [].
    d sub [].
  b sub [c,d].
```

The problem here is that while `a` and `b` have two common subtypes, namely `c` and `d`, they do not have a most general common subtype, since `c` is not a subtype of `d`, and `d` is not a subtype of `c`. In general, a violation of the bounded completeness condition such as is found in this example can be patched without destroying the ordering by simply adding additional types. In this case, the following type hierarchy preserves all of the subtyping relations of the one above, but satisfies bounded completeness:

```

bot sub [a,b].
  a sub [e].
    e sub [c,d].
      c sub [].
      d sub [].
    b sub [e].

```

In this case, the new type *e* is the most general subtype of *a* and *b*.

This last example brings up another point about inheritance hierarchies. When a type only has one subtype, the system provides a warning message (as opposed to an error message). This condition will not cause any compile-time or run-time errors, and is perfectly compatible with the logic of the system. It is simply not a very good idea from either a conceptual or implementational point of view. For more on this topic, see Carpenter (1992:Chapter 9).

## 4.2 Feature Structures

The primary representational device in ALE is the *typed feature structure*. In phrase structure grammars, feature structures model categories, while in the definite clause programs, they serve the same role as first-order terms in Prolog, that of a universal data structure. Feature structures are much like the frames of AI systems, the records of imperative programming languages like C or Pascal, and the feature descriptions used in standard linguistic theories of phonology, and more recently, of syntax.

Rather than presenting a formal definition of feature structures, which can be found in Carpenter (1992:Chapter 2), we present an informal description here. In fact, we begin by discussing feature structures which are not necessarily well-typed. In the next section, the type system is presented.

A feature structure consists of two pieces of information. The first is a type. Every feature structure must have a type drawn from the inheritance hierarchy. The other kind of information specified by a feature structure is a finite, possibly empty, collection of feature/value pairs. A feature value pair consists of a feature and a value, where the value is itself a feature structure. The difference between feature structures and the representations used in phonology and in GPSG, for instance, is that it is possible for two different substructures (values of features at some level of nesting) to be token identical in a feature structure. Consider the following feature structure drawn from the lexical entry for *John* in the categorial grammar in the appendix, displayed in the output notation of ALE:

```

cat
QSTORE e_list
SYNSEM basic
  SEM j
  SYN np

```

The type of this feature structure is *cat*, which is interpreted to mean it is a category. It is defined for two features, *QSTORE* and *SYNSEM*. As can be seen from this example, we follow the HPSG notational convention of displaying features in all

caps, while types are displayed in lower case. Also note that features and their values are printed in alphabetic order of the feature names. In this case, the value of the QSTORE feature is the simple feature structure of type `e_list`,<sup>1</sup> which has no feature values. On the other hand, the feature SYNSEM has a complex feature as its value, which is of type `basic`, and has two feature values SEM and SYN, both of which have simple values.

This last feature structure doesn't involve any structure sharing. But consider the lexical entry for *runs*:

```

cat
QSTORE e_list
SYNSEM backward
  ARG basic
    SEM [0] individual
    SYN np
  RES basic
    SEM run
    RUNNER [0]
    SYN s

```

Here there is structure sharing between the path SYNSEM ARG SEM and the path SYNSEM RES SEM RUNNER, where a *path* is simply a sequence of features. This structure sharing is indicated by the *tag* [0]. In this case, the sharing indicates that the semantics of the argument of *runs* fills the runner role in the semantics of the result. Also note that a shared structure is only displayed once; later occurrences simply list the tag. Of course, this example only involves structure sharing of a very simple feature structure, in this case one consisting of only a type with no features. In general, structures of arbitrary complexity may be shared, as we will see in the next example.

ALE, like Prolog II and HPSG, but unlike most other systems, allows cyclic structures to be processed and even printed. For instance, consider the following representation we might use for the liar sentence *This sentence is false*:

```

[0] false
  ARG1 [0]

```

In this case, the empty path and the feature ARG1 share a value. Similarly, the path ARG1 ARG1 ARG1 and the path ARG1 ARG1, both of which are defined, are also identical. But consider a representation for the negation of the liar sentence, *It is false that this sentence is false*:

```

false
ARG1 [0] false
  ARG1 [0]

```

Unlike Prolog II, ALE does not necessarily treat these two feature structures as being identical, as it does not conflate a cyclic structure with its infinite unfolding.

<sup>1</sup>Set values, like those employed in HPSG, are not supported by ALE. In the categorial grammar in the appendix, they are represented by lists and treated by attached procedures for union and selection.



We take up the notion of token identical structures in the section below on extensionality.

It is interesting to note that with typed feature structures, there is a choice between representing information using a type and representing the same information using feature values. This is a familiar situation found in most inheritance-based representation schemes. Thus the relation specified in the value of the path SYNSEM RES SEM is represented using a type, in:

```
SEM run
  RUNNER [0]
```

An alternative encoding, which is not without merit, is:

```
SEM unary_rel
  REL run
  ARG1 [0]
```

In general, type information is processed much more efficiently than feature value information, so as much information as possible should be placed in the types. The drawback is that type information must be computed at compile-time and remain accessible at run-time. More types simply require more memory.<sup>2</sup>

### 4.3 Subsumption and Unification

Feature structures are inherently partial in the information they provide. Based on the type inheritance ordering, we can order feature structures based on how much information they provide. This ordering is referred to as the *subsumption* ordering. The notion of subsumption, or information containment, can be used to define the notion of unification, or information combination. Unification conjoins the information in two feature structures into a single result if they are consistent and detects an inconsistency otherwise.

#### 4.3.1 Subsumption

We define subsumption, saying that  $F$  *subsumes*  $G$ , if and only if:

- the type of  $F$  is more general than the type of  $G$
- if a feature  $f$  is defined in  $F$  then  $f$  is also defined in  $G$  such that the value in  $F$  subsumes the value in  $G$
- if two paths are shared in  $F$  then they are also shared in  $G$

Consider the following examples of subsumption, where we let  $<$  stand for subsumption:

```
agr          <  agr
PERS first   PERS first
              NUM plu
```

---

<sup>2</sup>In general, the amount of memory required to represent  $n$  types is proportional to the number of pairs of consistent types. In the worst case, this is  $O(n^2)$  in the number of types.

```

sign                phrase
SUBJ agr            < SUBJ agr
                    PERS pers          PERS first
                                        NUM plu

sign                sign
SUBJ agr            SUBJ [0] agr
                    PERS first        PERS first
                    NUM plu          <  NUM plu
OBJ agr             OBJ [0]
PERS first
NUM plu

false               false                [1] false
ARG1 false          < ARG1 [0] false < ARG1 [1]
                    ARG1 false          ARG1 [0]

```

Note that the second of these subsumptions holds only if `pers` is a more general type than `first`, and `sign` is a more general type than `phrase`. It is also important to note that the feature structure consisting simply of the type `bot` will subsume every other structure, as the type `bot` is assumed to be more general than every other type.

### 4.3.2 Unification

Unification is an operation defined over pairs of feature structures that combines the information contained in both of them if they are consistent and fails otherwise. In ALE, unification is very efficient.<sup>3</sup> Declaratively, unifying two feature structures computes a result which is the most general feature structure subsumed by both input structures. But the operational definition is more enlightening, and can be given by simple conditions which tell us how to unify two structures. We begin by unifying the types of the structures in the type hierarchy. This is why we required the bounded completeness condition on our inheritance hierarchies; we want unification to produce a unique result. If the types are inconsistent, unification fails. If the types are consistent, the resulting type is the unification of the input types. Next, we recursively unify all of the feature values of the structures being unified which occur in both structures. If a feature only occurs in one structure, we copy it over into the result. This algorithm terminates because we only need to unify structures which are non-distinct and there are a finite number of nodes in any input structure.

Some examples of unification follow, where we use `+` to represent the operation:

<sup>3</sup>Using a typed version of the Martelli and Montanari (1982) algorithm, which was adapted to cyclic structures by Jaffar (1984), unification can be performed in what is known as quasi-linear time in the size of the input structures, where in this case, quasi-linear in  $n$  is defined to be  $\mathcal{O}(n \cdot ack^{-1}(n))$ , where  $ack^{-1}$  is the inverse of Ackermann's function, which will never exceed 4 or 5 for structures that can be represented on existing computers. There is also a factor in the complexity of unification stemming from the type hierarchy and appropriateness conditions, which we discuss below.

```

agr          + agr          = agr
PERS first   NUM plu       PERS first
                                NUM plu

sign                sign                sign
SUBJ agr          + SUBJ [0] bot = SUBJ [0] agr
    PERS 1st      OBJ [0]          PERS first
OBJ agr          NUM plu          NUM plu
    NUM plu      OBJ [0]          OBJ [0]

t              t              t
F [0] t + F t          = F [1] t
G [0]          F [1]    F [1]
                G [1]    G [1]

agr          + agr          = *failure*
PERS first   PERS second

e_list + ne_list = *failure*
        HD a
        TL e_list

```

Note that the second example respects our assumption that the type `bot` is the most general type, and thus more general than `agr`. The second example illustrates what happens in a simple case of structure sharing: information is retrieved from both the `SUBJ` and `OBJ` and shared in the result. The third example shows how two structures without cycles can be unified to produce a structure with a cycle. Just as the feature structure `bot` subsumes every other structure, it is also the identity with respect to unification; unifying the feature structure consisting just of the type `bot` with any feature structure  $F$  results simply in  $F$ . The last two unification attempts fail, assuming that the types `first` and `second` and the types `e_list` and `ne_list` are incompatible.

## 4.4 Inequations

Feature structures may also incorporate inequational constraints following (Carpenter 1992), which is in turn based on the notion of inequation in Prolog II (Colmerauer 1987). For instance, we might have the following representation of the semantics of a sentence:

```

SEM binary_rel
  REL know
  ARG1 [0] referent
        GENDER masc
        PERS third
        NUM sing
  ARG2 [1] referent
        GENDER masc

```

```

        PERS third
        NUM sing
[0] =\= [1]

```

Below the feature information, we have included the constraint that the value of the structure [0] is not identical to that of structure [1]. As a result, we cannot unify this structure with the following one:

```

REL know
ARG1 [2]
ARG2 [2]

```

Any attempt to unify the structures [0] and [1] causes failure.

## 4.5 Type System

As we mentioned in the introduction, what distinguishes ALE from other approaches to feature structures and most other approaches to terms, is that there is a strong type discipline enforced on feature structures. We have already demonstrated how to define a type hierarchy, but that is only half the story with respect to typing. The other component of our type system is a notion of feature *appropriateness*, whereby each type must specify which features it can be defined for, and furthermore, which types of values such features can take. The notion of appropriateness used here is similar to that found in object-oriented approaches to typing. For instance, if a feature is appropriate for a type, it will also be appropriate for all of the subtypes of that type. In other words, appropriateness specifications are inherited by a type from its supertypes. Furthermore, value restrictions on feature values are also inherited. Another important consideration for ALE's type system is the notion of type inference, whereby types for structures which are underspecified can be automatically inferred. This is a property our system shares with the functional language ML, though our notion of typing is only first-order. To further put ALE's type system in perspective, we note that type inheritance must be declared by the user at compile time, rather than being inferred. Furthermore, types in ALE are semantic, in Smolka's (1988b) terms, meaning that types are used at run-time. Even though ALE employs semantic typing, the type system is employed statically (at compile-time) to detect type errors in grammars and programs.

As an example of an appropriateness declaration, consider the simple type specification for lists with a head/tail encoding:

```

bot sub [list,atom].
  list sub [e_list,ne_list].
    e_list sub [].
    ne_list sub []
      intro [hd:bot,
            tl:list].
atom sub [a,b].
  a sub [].
  b sub [].

```

This specification tells us that a list can be either empty (`e_list`) or non-empty (`ne_list`). It implicitly tells us that an empty list cannot have any features defined for it, since none are declared directly or inherited from more general types. The declaration also tells us that a non-empty list has two features, representing the head and the tail of a list, and, furthermore, that the head of a list can be anything (since every structure is of type `bot`), but the tail of the list must itself be a list. Note that features must also be Prolog constants, even though the output routines convert them to all caps. The appropriateness declaration, `intro`, can be specified along with subsumption, as shown above, or separately; but for any given type, all features must be declared at once. If no `intro` declaration is given for a type, it is assumed that that type introduces no appropriate features. If an `intro` declaration is made for a type that does not occur on either side of a `sub` declaration, that type is assumed to be an immediate subtype of `bot` with no subtypes of its own. If a value restrictor (such as `list` above for feature `t1`) does not occur on either side of a `sub` declaration, it too is assumed to be maximal and an immediate subtype of `bot`.

In ALE, every feature structure must respect the appropriateness restrictions in the type declarations. This amounts to two restrictions. First, if a feature is defined for a feature structure of a given type, then that type must be appropriate for the feature. Furthermore, the value of the feature must be of the appropriate type, as declared in the appropriateness conditions. The second condition goes the other way around: if a feature is appropriate for a type, then every feature structure of that type must have a value for the feature. A feature structure respecting these two conditions is said to be *totally well-typed* in the terminology of Carpenter (1992, Chapter 6).<sup>4</sup> For instance, consider the following feature structures:

```
list
HD a
TL bot

ne_list
HD bot
TL ne_list
    HD atom
    TL list

ne_list
HD [0] ne_list
    HD [0]
    TL [0]
TL e_list
```

---

<sup>4</sup>The choice of totally well-typed structures was motivated by the desire to represent feature structures as records at run-time, without listing their features. Internally, a feature structure is represented as a term of the form `Tag-Sort(V1, . . . , VN)` where `Tag` represents the token identity of the structure using a Prolog variable, `Sort` is the type of structure, and `V1` through `VN` are the values of the appropriate features, in alphabetical order of the features' names, which are themselves left implicit. Furthermore, the `Tag` is used for forwarding and dereferencing during unification.

The first structure violates the typing condition because the type `list` is not appropriate for any features, only `ne_list` is. But even if we were to change its type to `ne_list`, it would still violate the type conditions, because `bot` is not an appropriate type for the value of `TL` in a `ne_list`. On the other hand, the second and third structures above are totally well-typed. Note that the second such structure does not specify what kind of list occurs at the path `TL TL`, nor does it specify what the `HD` value is, but it does specify that the second element of the list, the `TL HD` value is an `atom`, but it doesn't specify which one.

To demonstrate how inheritance works in a simple case, consider the specification fragment from the categorial grammar in the appendix:

```
functional sub [forward,backward]
    intro [arg:synsem,
          res:synsem].
forward sub [].
backward sub [].
```

This tells us that `functional` objects have `ARG` and `RES` features. Because `forward` and `backward` are subtypes of `functional`, they will also have `ARG` and `RES` features, with the same restrictions.

There are a couple of important restrictions placed on appropriateness conditions in ALE. The most significant of these is the acyclicity requirement. This condition disallows type specifications which require a type to have a value which is of the same or more specific type. For example, the following specification is *not* allowed:

```
person sub [male,female]
    intro [father:male,
          mother:female].
male sub [].
female sub [].
```

The problem here is the obvious one that there are no most general feature structures that are both of type `person` and totally well-typed.<sup>5</sup> This is because any `person` must have a `father` and `mother` feature, which are `male` and `female` respectively, but since `male` and `female` are subtypes of `person`, they must also have `mother` and `father` values. It is significant to note that the acyclicity condition does not rule out recursive structures, as can be seen with the example of lists. The `list` type specification is acceptable because not every list is required to have a head and tail, only non-empty lists are. The acyclicity restriction can be stated graph theoretically by constructing a directed graph from the type specification. The nodes of the graph are simply the types. There is an edge from every type to all of its supertypes, and an edge from every type to the types in the type restrictions in its features. Type specifications are only acceptable if they produce a graph with no cycles. One cycle in the `person` graph is from `male` to `person` (by the supertype relation) and from `person` to `male` (by the `FATHER` feature). On the other hand, there are no cycles in the specification of `list`.

<sup>5</sup>The only finite feature structures that could meet this type system would have to be cyclic, as noted in Carpenter (1992). The problem is that there is no most general such cyclic structure, so type inference cannot be unique.

The second restriction placed on appropriateness declarations is designed to limit non-determinism in much the same way as the bounded completeness condition on the inheritance hierarchy. This second condition requires every feature to be *introduced* at a unique most general type. In other words, the set of types appropriate for a feature must have a most general element. Thus the following type declaration fragment is *invalid*:

```
a sub [b,c,d].
  b sub []
    intro [f:w,
           g:x].
  c sub []
    intro [f:y,
           h:z].
  d sub [].
```

The problem is that the feature *F* is appropriate for types *b* and *c*, but there is not a unique most general type for which it's appropriate. In general, just like the bounded completeness condition, type specifications which violate the feature introduction condition can be patched, without violating any of their existing structure, by adding additional types. In this case, we add a new type between *a* and the types *b* and *c*, producing the equivalent well-formed specification:

```
a sub [e,d].
  e sub [b,c]
    intro [f:bot].
  b sub []
    intro [f:w,
           g:x].
  c sub []
    intro [f:y,
           h:z].
  d sub [].
```

This example also illustrates how subtypes of a type can place additional restrictions on values on features as well as introducing additional features.

As a further illustration of how feature introduction can be obeyed in general, consider the following specification of a type system for representing first-order terms:

```
sem_obj sub [individual,proposition].
  individual sub [a,b].
    a sub [].
    b sub [].
  proposition sub [atomic_prop,relational].
    atomic_prop sub [].
    relational_prop sub [unary_prop,transitive_prop]
      intro [arg1:individual].
    unary_prop sub [].
    transitive_prop sub [binary_prop,ternary_prop]
```

```
        intro [arg2:individual] .
binary_prop sub [].
ternary_prop sub []
        intro [arg3:individual] .
```

In this case, unary propositions have one argument feature, binary propositions have two argument features, and ternary propositions have three argument features, all of which must be filled by individuals.

## 4.6 Interaction with the signature

- `show_approp(<type>)` . shows the appropriateness conditions of a type
- `show_subtypes(<type>)` . shows a mini typehierarchy with the immediate subtypes of <type>
- `show_all_subtypes(<type>)` . shows the complete typehierarchy below <type>
- `show_supertypes(<type>)` . shows the immediate supertypes of <type>
- `show_all_supertypes(<type>)` . shows the hierarchy below the most general type bot and type <type>, including only those types which are direct or indirect supertypes of <type>



## T4.1 TRALE Signatures

### T4.1.1 Signature specification files

As an alternative to the `sub/2` and `intro/2` declarations that ALE grammars use to specify their signatures, TRALE can instead read a special *signature file*, which follows the format discussed in this section.

By default, this alternative is used when TRALE is loaded. You can switch back to `sub/2` and `intro/2` declarations in the grammar by setting the `subintro` flag to `grammar`:

```
| ?- ale_flag(subintro,_,grammar).
```

To reactivate the use of signature files, reset the `subintro` flag to `file`:

```
| ?- ale_flag(subintro,_,file).
```

The `ale_flag/3` command is discussed in more detail in Chapter G.

TRALE signature files are specified using a separate text file with subtyping indicated by indentation as in the following example signature:

```
type_hierarchy
bot
  a f:bool g:bool
    b f:plus g:minus
      c f:minus g:plus
bool
  plus
  minus
.
```

The example shows a type hierarchy with a most general element `bot`, which immediately subsumes types `a` and `bool`. Type `a` introduces two features `F` and `G` whose values must be of type `bool`. Type `a` also subsumes two other types `b` and `c`. The `F` feature value of the former is always `plus` and its `G` feature value is always `minus`. The feature values of type `c` are the inverse of type `b`. Finally, `bool` has two subtypes `plus` and `minus`.

As shown in the example, appropriate features are written after the types that introduce them. The type of the value of a feature is separated from the feature by a colon as in `f:bool`, which says the feature `F` takes a value of type `bool`. Note that subtypes are written at a consistent level of increased indentation. This means that if a type is introduced at column  $C$ , then all its subtypes must be introduced directly below it at column  $C + n$ . The only requirements on the choice of  $n$  are that it be greater than zero, and that one value is consistently used throughout a signature file.  $n = 2$  is recommended. Inconsistent indentation causes an error.

A grammar can only use one signature file, and there can only be one type hierarchy in a signature file. If more than one type hierarchy is declared in the same signature file, only the first one is used by TRALE. By default, the signature file is assumed to be named `signature`, and is assumed to live in the directory from which Prolog was started. The `signature/1` directive in a grammar file tells

TRALE to look for the signature file in the path given as its argument instead. The syntax of this pathname is operating-system specific.

Types are syntactically Prolog atoms, which means that they have to start with a lower-case letter and only consist of letters and numbers (e.g. `bot`, `bool`, `dtr1`). New types are introduced in separate lines. Each informative line must contain at least a type name. A type may never occur more than once as the subtype of the same supertype. It can, however, occur as a subtype of two or more different supertypes — this is called *multiple inheritance*. In this case, one prefixes the subtype with an ampersand (&) to indicate that the multiple inheritance is intended.

There may be zero or more features introduced for each type. As mentioned before, these have the form `feature:value-restriction`. All feature-value pairs are separated by white space and they should appear in the same line. Recall that, `feature` is the name of a feature and `value` is the value restriction for that feature. As with the types, a feature name must always start with a lower case letter. If a feature is specified for a type, all its subtypes inherit this feature automatically. As in ALE, in cases of multiple inheritance, the value restriction on any feature that has been inherited from the supertypes is the union of those value restrictions. A single period (.) in an otherwise blank line signals the end of a signature declaration.

TRALE signature specifications also allow for `a_/1` atoms as potential value restrictions, as described in Section 4.8.

As another example, let us see how the type hierarchy in Figure 4.1 translates into a TRALE signature. As the figure shows, the type `agr`, which is immediately subsumed by  $\perp$ , introduces three features `person`, `number` and `gender`. These features are of types `per`, `num` and `gen` respectively. The TRALE translation of this hierarchy is shown below. Note that `1st`, `2nd` and `3rd` are respectively shown as `first`, `second`, and `third` because a Prolog atom has to begin with a lower-case letter.

type\_hierarchy

bot

per

first

second

third

num

singular

plural




Figure 4.1: A sample type hierarchy


```

gen
  feminine
  masculine
agr person:per number:num gender:gen
.

```

### T4.1.2 Subtype covering

TRALE assumes that subtypes exhaustively cover their supertypes, i.e., that every object of a non-maximal type,  $\tau$ , is also of one of the maximal types subsumed by  $\tau$ . This is only significant when the appropriateness conditions of  $\tau$ 's maximal subtypes on the features appropriate to  $\tau$  do not cover the same products of values as  $\tau$ 's appropriateness conditions. Let us look at the following type hierarchy for example.



In this hierarchy, there are no  $\tau$  objects with identically typed or structure-shared  $F$  and  $G$  values as in the following feature structures:

$$\begin{bmatrix} F & \mathbf{1} \\ G & \mathbf{1} \end{bmatrix} \begin{bmatrix} F & + \\ G & + \end{bmatrix} \begin{bmatrix} F & - \\ G & - \end{bmatrix}$$

Unlike ALE, TRALE does not recognise the existence of feature structures with such combinations of feature values. However, if TRALE detects that only one subtype's product of values is consistent with a feature structure's product of values, it will promote the product to that consistent subtype's product. Thus, in our example, a feature structure:

$$\begin{bmatrix} t \\ F & + \\ G & bool \end{bmatrix}$$

will be promoted automatically to the following. Note that the type itself will not be promoted to  $\tau_1$ .

$$\begin{bmatrix} t \\ F & + \\ G & - \end{bmatrix}$$

To consider a simple example from HPSG, English verbs are typically assumed to have the following two features in order to distinguish auxiliary verbs from main verbs and also to show whether a subject-auxiliary inversion has taken place.

$$\begin{bmatrix} \text{verb} \\ \text{AUX } \textit{bool} \\ \text{INV } \textit{bool} \end{bmatrix}$$

The values for the `AUX` and `INV` features are taken to be of type `bool`. However, note that there cannot be any verbal type with the following combination of feature values:

$$\begin{bmatrix} \text{AUX } - \\ \text{INV } + \end{bmatrix}$$

That is, there are no verbs in English that can occur before the subject and not be auxiliaries. Thus, using TRALE's interpretation of sub-typing, we can prevent this undesirable combination with the following type hierarchy:

```

type_hierarchy
bot
  bool
    plus
    minus
  verb aux:bool inv:bool
    aux_verb aux:plus inv:bool
    main_verb aux:minus inv:minus
.

```

## 4.7 Extensionality

ALE also respects the distinction between *intensional* and *extensional* types (see Carpenter (1992:Chapter 8)). The concept of extensional typing has its origins in the assumption in standard treatments of feature structures, that there can only be one copy of any *atom* (a feature structure with no appropriate features) in a feature structure. Thus, if path  $\pi_1$  leads to atom  $a$ , and path  $\pi_2$  leads to atom  $a$ , then the values for those two paths are token-identical. Token-identity refers to an identity between two feature structures as objects, as opposed to *structure-identity*, which refers to an identity between two feature structures that contain the same information.

Smolka (1988a) partitioned his atoms according to whether more than one copy could exist or not. In ALE, following Carpenter (1992), this notion of copyability has been extended to arbitrary types – loosely speaking, those types which are copyable we call *intensional*, and those which are not we call *extensional*. Thus, it is possible to have two feature structures of the same intensional type which, although they may be structure-identical, are not token-identical. Formally:

Given the set of types,  $\text{Type}$ , defined by an ALE signature, we designate a subset,  $\text{ExtType} \subseteq \text{Type}$ , as the set of extensional types. With the exception of  $a_{\_}/1$  atoms, discussed below, this set consists only of *maximally specific types*, i.e., for each  $\sigma \in \text{ExtType}$ , there is no type  $\tau$  such that  $\sigma$  subsumes  $\tau$ .

The restriction of  $\text{ExtType}$  to maximally specific types is peculiar to ALE, and is levied in order to reduce the computational complexity of enforcing extensionality.<sup>6</sup>

We need one more definition to formally state the effect which an extensional type has on feature structures in ALE.

Given a set of extensional types,  $\text{ExtType}$ , we define an equivalence relation,  $\asymp$ , the *collapsing relation*, on well-typed feature structures, such that  $F_1 \asymp F_2$  for  $F_1 \neq F_2$  only if:

- $F_1$  has the same type,  $\sigma$ , as  $F_2$ , and  $\sigma \in \text{ExtType}$ , and
- for every feature,  $f$ , appropriate to  $\sigma$ ,  $F_1^f$ , the value of  $f$  in  $F_1$ , and  $F_2^f$ , the value of  $f$  in  $F_2$ , are defined, and  $F_1^f \asymp F_2^f$ .

In ALE, all feature structures behave as if they are what Carpenter (1992) referred to as *collapsed*. That is, the only collapsing relation that exists between any two feature structures is the trivial collapsing relation, namely:

$F_1 \asymp F_2$  if and only if  $F_1$  is token-identical to  $F_2$ .

In the case of acyclic feature structures, this definition is equivalent to saying that two feature structures of the same extensional type are token-identical if and only

---

<sup>6</sup>In theory (Carpenter 1992), this set is only required to be *upward closed*, which means that if  $\sigma \in \text{ExtType}$ , and  $\sigma$  subsumes  $\tau$ , then  $\tau \in \text{ExtType}$ . This relaxation of our requirement that extensional types be maximal would actually not be too difficult to implement.

if, for every feature appropriate to that type, their respective values on that feature are token-identical. For example, supposing that we have a signature representing dates, then the two substructures representing dates in the following structure must be token identical.

```

married_person
  BIRTHDAY [1] date
    DAY 12
    MONTH nov
    YEAR 1971
  SPOUSE BIRTHDAY [1] date
    DAY 12
    MONTH nov
    YEAR 1971

```

In other words, this represents a person born on 12 November 1971, who is married to a person with the same birthdate.

Now consider a slightly more complex example, which employs the following signature.

```

bot sub [a,b,c,g].
  a sub []
    intro [f:b,g:c].
  b sub [].
  c sub [].
  g sub []
    intro [h:a,j:a].

```

If a, b, and c are extensional, then the values of H and J in g are always token-identical, i.e., every feature structure of type g satisfies:

```

g
H [0] a
  F b
  G c
J [0]

```

But if only a, and b are extensional, and c is intensional, then the values of H and J are not necessarily token-identical, although they are always structure-identical:

```

g
H a
  F [1] b
  G c
J a
  F [1]
  G c

```

To cite an earlier example, suppose we were to specify that the type *false*, used in the liar sentence and its negation, were extensional. Now the liar sentence's representation is:

```
[0] false
    ARG1 [0]
```

as before, but the negation of the liar sentence would also be represented by:

```
[0] false
    ARG1 [0]
```

since if were still represented by:

```
[1] false
ARG1 [0] false
    ARG1 [0]
```

then we could cite a non-trivial collapsing relation,  $\asymp$ , in which  $[1] \asymp [0]$ .

As a related example, consider:

```
s
A [0] t
  C [0]
B [1] t
  C [1]
```

Assuming that  $t$  is extensional and only appropriate for the feature  $C$ , then the structures  $[0]$  and  $[1]$  in the above structure would be identified.

Extensionality allows the proper representation of feature structures and terms in both PATR-II, the Rounds-Kasper system, and in Prolog and Prolog II. For PATR-II and the Rounds-Kasper system, all atoms (those types with no appropriate features) are assumed to be extensional. Furthermore, in the Rounds-Kasper and PATR-II systems, which are monotyped, there is only one type that is appropriate for any features, and it must be appropriate for all features in the grammar. In Prolog and Prolog II, the type hierarchy is assumed to be flat, and every type is extensional.

Just as with implementations of Prolog, collapsing is only performed as it is needed. As shown by Carpenter (1992), collapsing can be restricted to cases where inequations are tested between two structures, with exactly the same failure behavior. It turns out to be less efficient to collapse structures before asserting them into ALE's parsing chart, primarily because the time to test arbitrary structures for collapsibility is at least quadratic in the size of the structures being collapsed. See the section below on inequations for further discussion. Currently, extensionality is only enforced before the answer to a user query is given.

Extensional types in ALE are specified all at once in a list:

```
ext([ext1, ..., extn]).
```

in the same file in which the subsumption relation is defined. All types that are not mentioned in the `ext` specification are assumed to be intensional, except ALE's `a_/1` atoms, discussed below, which have the same extensionality as Prolog terms, i.e., if they are ground or have the same variables in the same positions.<sup>7</sup> These do not need to be declared as such. If more than one `ext` specification is given, the first one is used. If no `ext` specification is given, then the specification:

<sup>7</sup>This is given by the `==` operator in Prolog.

```
ext( []).
```

is assumed. If a type occurs in  $ext/1$ , but does not appear on the left or right-hand side of a sub declaration, it assumed to be maximal, and immediately subsumed by  $bot$ .

Of course, collapsing is only enforced between feature structures whose life-spans in ALE<sup>8</sup> overlap. So, for example, if one request is made for the representation of the liar sentence:

```
[0] false
    ARG1 [0]
```

and then another is made for that of its negation, the output is not:

```
[0]
```

(referring to the same token above) but rather:

```
[0] false
    ARG1 [0]
```

Every time a new context arises, numbering of structures begins again from [0].

## 4.8 $a_{/1}$ Atoms

ALE also provides an infinite collection of atoms. These are of the form:

$a_{/1} \textit{Term}$

where *Term* is a prolog term. Two  $a_{/1}$  atoms subsume each other if and only if their terms subsume each other as prolog terms. As a result, no two different, ground atoms subsume each other; and the most general atom of this collection is  $a_{/1} \_$ . They implicitly exist in every type hierarchy, with  $a_{/1} \_$  being immediately subsumed by  $bot$ , and with every ground atom being maximal.  $a_{/1}$  atoms are extensional; and non-ground  $a_{/1}$  atoms are extensionally identical as Prolog terms, i.e., if they have the same variables in the same positions. For example,  $a_{/1} f(X)$  and  $a_{/1} g(X)$  are not taken to be the same atom, nor are  $a_{/1} f(X)$  and  $a_{/1} f(f(X))$ , nor are  $a_{/1} f(X)$  and  $a_{/1} f(Y)$ . But  $a_{/1} f(X)$  and  $a_{/1} f(X)$  are. Their status in the type hierarchy should not be explicitly declared, nor should the fact that they bear no features, nor should their extensionality.

Some care must be exercised when using non-ground atoms in chart edges. ALE's chart parser copies edges, including the Prolog variables inside  $a_{/1}$  atoms. When these variables are copied, identity among variables within a single edge is preserved, but identity among variables between different edges may be lost. Because ALE delays the enforcement of extensional type checking, this could result in ALE losing a path equation between two atoms. The best way to avoid this is always to use ground atoms in chart edges. Otherwise, the user should at least

---

<sup>8</sup>The life-span of a feature structure in ALE is the period from its creation to the point when the user command currently being executed finishes, unless that feature structure is asserted as an edge in ALE's chart parser. In this case, the life of the feature structure ends when the edge is removed. Every new request for a parse to ALE removes all of the current edges.



avoid relying on extensional identity when writing grammars by not using the ALE built-in @= or inequations between non-ground atoms from different edges.

If the user requires intensional atoms, they must be explicitly declared. There must also be no user-defined type, `a_`, in the type hierarchy. Certain arguments to `a_/1` cannot be used, such as `a_` itself, and other prolog reserved words, such as `mod`, unless they are used with the proper operator precedence and proper number of arguments to be parsed by Prolog.

Otherwise, `a_/1` atoms can be used wherever a normal type can. They are particularly useful as members of large domains that are too tedious to define, such as phonology attributes in natural language grammars, or to pass extra-logical information around a parse tree, such as numbers representing probabilities. To declare a feature's value as any `a_/1` atom, use `a_ _`:

```
sign intro [phon:(a_ _)].
```

The parentheses are recommended for readability, but not necessary. Because subsumption among `a_/1` atoms mirrors subsumption as prolog terms, one can also declare features appropriate for only certain kinds of atoms. For example:

```
sign intro [phon:(a_ phon(_))].
```

declares PHON appropriate to any atom whose term's functor is `phon/1`.

Structure-sharing between `a_/1` prolog terms in feature appropriateness declarations is ignored by ALE. For example, the declaration:

```
foo intro [f:(a_ X),g:(a_ X)].
```

is treated as:

```
foo intro [f:(a_ _),g:(a_ _)].
```

ALE does respect structure-sharing between `a_/1` prolog terms in descriptions.

## 4.9 Attribute-Value Logic

Now that we have seen how the type system must be specified, we turn our attention to the specification of feature structures themselves. The most convenient and expressive method of describing feature structures is the logical language developed by Kasper and Rounds (1986), which we modify here in three ways. First, we replace the notion of path sharing with the more compact and expressive notion of variable due to Smolka (1988a). Second, we extend the language to types, following Pollard (in press). Finally, we add inequations.

The collection of *descriptions* used in ALE can be described by the following BNF grammar:

```
<desc> ::= <type>
          | <variable>
          | (<feature>:<desc>)
          | (<desc>,<desc>)
          | (<desc>;<desc>)
          | (= \= <desc>)
```

As we have said before, both types and features are represented by Prolog constants. Variables, on the other hand, are represented by Prolog variables. As indicated by the BNF, no whitespace is needed around the feature selecting colon, conjunction comma and disjunction semi-colon, but any whitespace occurring will be ignored.

These descriptions are used for picking out feature structures that satisfy them. We consider the clauses of the definition in turn. A description consisting of a type picks out all feature structures of that type. A variable can be used to refer to any feature structure, but multiple occurrences of the same variable must refer to the same structure. A description of the form  $\langle \text{feature} \rangle : \langle \text{desc} \rangle$  picks out a feature structure whose value for the feature satisfies the nested description. An inequation  $\neq \langle \text{desc} \rangle$  is satisfied by those feature structures that are not token-identical to the feature structure described by  $\langle \text{desc} \rangle$ . Inequations are discussed in more detail below.

There are two ways of logically combining descriptions: following Prolog, the comma represents conjunction and the semi-colon represents disjunction. A feature structure satisfies a conjunction of descriptions just in case it satisfies both conjuncts, while it satisfies a disjunction of descriptions if it satisfies either of the disjuncts.

We should also add to the above BNF grammar the following line:

$$\langle \text{desc} \rangle ::= (\langle \text{path} \rangle == \langle \text{path} \rangle)$$

This is an equational description, of which inequations are the negation. Equational or inequational descriptions are satisfied by the presence or absence, respectively, of token-identity. In particular, an inequation between two structurally-identical feature structures can be satisfied, while a path equation can *only* be satisfied by two structurally-identical feature structures, but is not necessarily satisfied.

All instances of equational descriptions can be captured by using multiple occurrences of variables. For example, the description:

$$([\text{arg1}] == [\text{arg2}])$$

is equivalent to the description:

$$(\text{arg1} : X, \text{arg2} : X).$$

assuming there are no other occurrences of X.

Standard assumptions about operator precedence and association are followed by ALE, allowing us to omit most of the parentheses in descriptions. In particular, equational descriptions bind the most tightly, followed by feature selecting colon, then by inequations, then conjunction and finally disjunction. Furthermore, conjunction and disjunction are left-associative, while the feature selector is right-associative. For instance, this gives us the following equivalences between descriptions:

$$a, b ; c, d ; e = (a,b);(c,d);e$$

$$a,b,c = a,(b,c)$$

$$f : g : \text{bot}, h : j = (f : (g : \text{bot})), (h : j)$$

$$f : g : \neq k, h : j = (f : (g : \neq (k))), (h : j)$$

$$f : [g] == [h], h : j = (f : ([g] == [h])), (h : j)$$

Note that a space must occur between  $\neq$  and other operators such as  $:$ .

A description may be satisfied by no structure, a finite number of structures or an infinite collection of feature structures. A description is said to be *satisfiable* if it is satisfied by at least one structure. A description  $\phi$  *entails* a description  $\psi$  if every structure satisfying  $\phi$  also satisfies  $\psi$ . Two descriptions are *logically equivalent* if they entail each other, or equivalently, if they are satisfied by exactly the same set of structures.

ALE is only sensitive to the differences between logically equivalent formulas in terms of speed. For instance, the two descriptions  $(tl : list, ne\_list, hd : bot)$  and  $hd : bot$  are satisfied by exactly the same set of totally well-typed structures assuming the type declaration for lists given above, but the smaller description will be processed much more efficiently. There are also efficiency effects stemming from the order in which conjuncts (and disjuncts) are presented. The general rule for speedy processing is to eliminate descriptions from a conjunction if they are entailed by other conjuncts, and to put conjuncts with more type and feature entailments first. Thus with our specification for relations above, the description  $(arg1 : a, binary\_proposition)$  would be slower than  $(binary\_proposition, arg1 : a)$ , since  $binary\_proposition$  entails the existence of the feature  $arg1$ , but not conversely.<sup>9</sup>

At run-time, ALE computes a representation of the most general feature structure that satisfies a description. Thus a description such as  $hd : a$  with respect to the list grammar is satisfied by the structure:

```
ne_list
HD a
TL list
```

Every other structure satisfying the description  $hd : a$  is subsumed by the structure given above. In fact, the above structure is said to be a *vague* representation of all of the structures that satisfy the description. The type conditions in ALE were devised to obey the very important property, first noted by Kasper and Rounds (1986), that every non-disjunctive description is satisfied by a unique most general feature structure. Thus in the case of  $hd : a$ , there is no more general feature structure than the one above which also satisfies  $hd : a$ .

The previous example also illustrates the kind of *type inference* used by ALE. Even though the description  $hd : a$  does not explicitly mention either the feature TL or the type  $ne\_list$ , to find a feature structure satisfying the description, ALE must infer this information. In particular, because  $ne\_list$  is the most general type for which HD is appropriate, we know that the result must be of type  $ne\_list$ .

<sup>9</sup>This is because the depth of dereferencing depends on the history of types a given structure is instantiated to. There is no path compression on-line, but it is carried out before an edge is asserted into the chart.

Furthermore, because `ne_list` is appropriate for both the features `HD` and `TL`, ALE must add an appropriate `TL` value. The value type `list` is also inferred, due to the fact that a `ne_list` must have a `TL` value which is a list. As far as type inference goes, the user does not need to provide anything other than the type specification; the system computes type inference based on the appropriateness specification. In general, type inference is very efficient in terms of time. The biggest concern should be how large the structures become.<sup>10</sup> In contrast to a vague description, a disjunctive description is usually *ambiguous*. Disjunction is where the complexity arises in satisfying descriptions, as it corresponds operationally to non-determinism.<sup>11</sup> For instance, the description `hd: (a;b)` is satisfied by two distinct minimal structures, neither of which subsumes the other:

```
ne_list      ne_list
HD a         HD b
TL list      TL list
```

On the other hand, the description `hd: atom` is satisfied by the structure:

```
ne_list
HD atom
TL list
```

Even though the descriptions `hd: atom` and `hd: (a;b)` are not logically equivalent (though the former entails the latter), they have the interesting property of being unifiable with exactly the same set of structures. In other words, if a feature structure can be unified with the most general satisfier of `hd: atom`, then it can be unified with one of the minimal satisfiers of `hd: (a;b)`.

In terms of efficiency, it is very important to use vagueness wherever possible rather than ambiguity. In fact, it is almost always a good idea to arrange the type specification with just this goal in mind. For instance, consider the difference between the following pair of type specifications, which might be used for English gender:

```
gender sub [masc,fem,neut].      gender sub [animate,neut].
  masc sub [].                    animate sub [masc,fem].
  fem sub [].                      masc sub [].
  neut sub [].                     fem sub [].
                                   neut sub [].
```

Now consider the fact that the relative pronouns *who* and *which* are distinguished on the basis of whether they select animate or inanimate genders. In the flatter hierarchy, the only way to select the animate genders is by the ambiguous description `masc; fem`. The hierarchy with an explicit animate type can capture the same

<sup>10</sup>Finding most general satisfiers for non-disjunctive descriptions, even those involving type inference, is quasi-linear in the size of the description. But it should be kept in mind that there is also a factor of complexity determined by the size of the type specification. In practice, this factor is proportional to how large the inferred structure is. In general, the size of the inferred structure is linear in the size of the description, with a constant for the type specification.

<sup>11</sup>It corresponds so closely with non-determinism that satisfiability of descriptions with disjunctions is NP-complete. Furthermore, the algorithm employed by ALE may produce up to  $2^n$  satisfiers for a description with  $n$  disjunctions.

possibilities with the vague description `animate`. An effective rule of thumb is that ALE does an amount of work at best proportional to the number of most general satisfiers of a description and at worst proportional to  $2^n$ , where  $n$  is the number of disjuncts in the description. Thus the ambiguous description requires roughly twice the time and memory to process than the vague description. Whether the amount of work is closer to the number of satisfiers or exponential in the number of disjuncts depends on how many unsatisfiable disjunctive possibilities drop out early in the computation.

### 4.9.1 Enforcement of Inequations

Inequations are persistent in that once that are created, they remain as long as one of the structures being inequated remains. Thus the following two descriptions are logically equivalent:

$$f : (= \setminus = c), f : c$$

$$f : c, f : (= \setminus = c)$$

Both will cause failure; but they are not operationally equivalent. An inequation is evaluated when it arises, and again after high-level unifications in the system; but inequations are not evaluated every time an inequated structure is modified. In an ideal system, inequations would be attached directly to structures so that they could be evaluated on-line during unification. As things stand, ALE represents a feature structure as a regular feature structure with a separate set of inequations. Also, the complexity is sensitive to the conjunctive normal form of inequations at the time at which it is evaluated, though this form may later be reduced.

These sets of inequations are evaluated at run-time at the point they are encountered, before answers are given to top-level queries, before any edge is added to ALE's parsing chart, after every daughter is matched to a description in an ALE grammar rule<sup>12</sup>, and after the head of an ALE definite clause has been matched to a goal. At compile-time, inequations are checked for every empty category, for every lexical entry, and after every lexical rule application.

Inequations are also symmetric. Thus the following two descriptions are logically equivalent:

$$f : (= \setminus = X), g : X$$

$$f : X, g : (= \setminus = X)$$

Both inequate the values of `F` and `G`. Again, these are not operationally equivalent. Because inequations are evaluated at the time they are encountered, the second ordering will normally detect an immediate failure *sooner* than the first.

An inequation between two feature structures is a requirement for them not to be *token-identical*. Thus, if a type is intensional, it is possible for two feature structures to be of that same type, and still satisfy an inequation between them. Thus, any attempt to inequate two structures that should be identical as a result of extensional typing will also cause failure. For instance, consider the following:

<sup>12</sup>In the case of `cats`, they are enforced after the list description itself is matched, and also after every element of the list is matched.

```

s
F [1] t      s
H [1]      +  F [3]
G [2] t      G [4]      = failure
H [1]      [3] =\= [4]

```

The values of the features F and G cannot be inequated because they are extensionally identical (assuming the type  $t$  is declared to be extensional and is only appropriate for the feature H).

When inequations are evaluated, they are reduced. This reduction consists, in part, of partial evaluation of extensionality checking, which is otherwise delayed in ALE. For instance, consider the following:

```

F [1] s
H bot
J bot
G [2] s
H bot
J bot
[1] =\= [2]

```

If the type  $s$  is extensional and appropriate for the features H and J, then the inequation above is reduced to the following:

```

F [1] s
H [3] bot
J [4] bot
G [2] s
H [5]
J [6]
[3] =\= [5] ; [4] =\= [6]

```

The set of inequations is stored in conjunctive normal form. The cost is some space over the re-evaluation of inequations. Of course, if the types on [3] and [4] were more refined than  $bot$ , then the inequations [3] =\= [5] and [4] =\= [6] would further reduce. In addition, when reducing inequations in this way, we eliminate those that are trivially satisfied. The ones that are printed are only the residue after reduction. For instance, consider the following structure:

```

F [1] s
H [3] a
J bot
G [2] s
H [3]
J bot
[1] =\= [2]

```

Since the H values are token-identical, this structure reduces to the following.

```

F s
H [3] a

```

```

    J [4] bot
  G s
    H [3]
    J [5] bot
  [4] =\= [5]

```

If structures [4] and [5] had been of non-unifiable types, then there would be no residual inequation at all — the original inequation would trivially be satisfied.

An important subcase is that of an inequation between extensional atoms. If an atom is extensional, then there is only one instance of it. Thus an inequation between two identical, extensional atoms always fails. For example, if a type signature includes:

```

  bot sub [..., a, b, ...].
  a sub [].
    intro [f:bot].
  b sub [].
  ...
  ext([..., b, ...]).

```

then the description:

```
(a, f :=\= b)
```

is satisfied just in those cases where the value of  $F$  is not of type  $\mathbf{b}$ . If  $\mathbf{b}$  were intensional, then the inequation in this description would essentially have no effect. In fact, the only productive instances of inequations between two intensionally typed feature structures are those used with multiply occurring variables. In all other instances, there is no way for the inequation to be violated, since there is no way to render a structurally-identical copy of an intensionally typed feature structure token-identical to any other structure. ALE detects these trivially satisfied inequations and disposes of them.

## 4.10 Macros

ALE allows the user to employ a general form of parametric macros in descriptions. Macros allow the user to define a description once and then use a shorthand for it in other descriptions. We first consider a simple example of a macro definition, drawn from the categorial grammar in the appendix. Suppose the user wants to employ a description `qstore:e_list` frequently within a program. The following macro definition can be used in the program file:

```

quantifier_free macro
  qstore:e_list.

```

Then, rather than including the description `qstore:e_list` in another description, `@ quantifier_free` can be used instead. Whenever `@ quantifier_free` is used, `qstore:e_list` is substituted.

In the above case, the `<macro_spec>` was a simple atom, but in general, it can be supplied with arguments. The full BNF for macro definitions is as follows:

```

<macro_def> ::= <macro_head> macro <desc>.

<macro_head> ::= <macro_name>
                | <macro_name>(<seq(<var>)>)

<macro_spec> ::= <macro_name>
                | <macro_name>(<seq(<desc>)>)

<seq(X)> ::= X
           | X, <seq(X)>

```

Note that  $\langle \text{seq}(X) \rangle$  is a parametric category in the BNF which abbreviates non-empty sequences of objects of category  $X$ . The following clause should be added to recursive definition of descriptions:

```

<desc> ::= @ <macro\_spec>

```

A feature structure satisfies a description of the form  $@ \langle \text{macro\_spec} \rangle$  just in case the structure satisfies the body of the definition of the macro.

Again considering the categorial grammar in the appendix, we have the following macros with one and two arguments respectively:

```

np(Ind) macro
  syn:np,
  sem:Ind.

n(Restr,Ind) macro
  syn:n,
  sem:(body:Restr,
        ind:Ind).

```

In general, the arguments in the definition of a macro must be Prolog variables, which can then be used as variables in the body of the macro. With the first macro, the description  $@ \text{np}(j)$  would then be equivalent to the description  $\text{syn:np, sem:j}$ . When evaluating a macro, the argument supplied, in this case  $j$ , is substituted for the variable when expanding the macro. In general, the argument to a macro can itself be an arbitrary description (possibly containing macros). For instance, the description:

```

n((and, conj1:R1, conj2:R2), Ind3)

```

would be equivalent to the description:

```

syn:n,
sem:(body:(and, conj1:R1, conj2:R2),
      ind:Ind3)

```

This example illustrates how other variables and even complex descriptions can be substituted for the arguments of macros. Also note the parentheses around the arguments to the first argument of the macro. Without the parentheses, as in  $n(\text{and, conj1:R1, conj2:R2, Ind3})$ , the macro expansion routine would take this



to be a four argument macro, rather than a two argument macro with a complex first argument. This brings up a related point, which is that different macros can have the same name as long as they have the different numbers of arguments.

Macros can also contain other macros, as illustrated by the macro for proper names in the categorial grammar:

```
pn(Name) macro
  synsem: @ np(Name),
  @ quantifier_free.
```

In this case, the macros are expanded recursively, so that the description `pn(j)` would be equivalent to the description

```
synsem:(syn:np,sem:j),qstore:e_list
```

It is usually a good idea to use macros whenever the same description is going to be re-used frequently. Not only does this make the grammars and programs more readable, it reduces the number of simple typing errors that lead to inconsistencies.

As is to be expected, macros can't be recursive. That is, a macro, when expanded, is not allowed to invoke itself, as in the *ill-formed* example:

```
infinite_list(Elt) macro
  hd:Elt,
  tl:infinite_list(Elt)
```

The reason is simple; it is not possible to expand this macro to a finite description. Thus all recursion must occur in grammars or programs; it can't occur in either the appropriateness conditions or in macros.

The user should note that variables in the scope of a macro are not the same as ALE feature structure variables — they denote where macro-substitutions of parameters are made, *not* instances of re-entrancy in a feature structure. If we employ the following macro:

```
blah(X) macro
  b,
  f: X,
  g: X.
```

with the argument `(c,h:a)` for example we obtain the following feature structure:

```
b
F c
  H a
G c
  H a
```

where the values of `F` and `G` are not shared (unless `c` and `a` are extensional). We can, of course, create a shared structure using `blah`, by including an ALE variable in the actual argument to the macro. Thus `blah((Y,c,h:a))` yields:

```

b
F [0] c
  H a
G [0]

```

Because programming with lists is so common, ALE has a special macro for it, based on the Prolog list notation. A description may also take any of the forms on the left, which will be treated equivalently to the descriptions on the right in the following diagram:

<code>[]</code>	<code>e_list</code>
<code>[H T]</code>	<code>(hd:H, tl:T)</code>
<code>[A1,A2,...,AN]</code>	<code>(hd:A1, tl:(hd:A2, tl: ... tl:(hd:AN, tl:e_list)...))</code>
<code>[A1,...,AN T]</code>	<code>(hd:A1, tl:(hd:A2, tl: ... tl:(hd:AN, tl:T)...))</code>

Note that this built-in macro does not require the macro operator `@`. Thus, for example, the description `[a|T3]` is equivalent to `hd:a,tl:T3`, and the description `[a,b,c]` is equivalent to `hd:a,tl:(hd:b,tl:(hd:c,tl:e_list))`. There are many example of this use of Prolog's list notation in the grammars in the appendix.

## T4.2 TRALE Descriptions

The following discusses some additions to the ALE description language which are included in TRALE.

### T4.2.1 Logical variable macros

TRALE's logical variable macros, unlike ALE macros, use logical variables in their definitions rather than true macro variables. Logical variables entail structure-sharing if used more than once in a predicate. For example, the Prolog expression `foo(X,X)` means that the two arguments of `foo` are structure-shared.<sup>13</sup> True macro variables, on the other hand, simply serve as place holders and their multiple occurrence does not entail structure sharing. This makes a difference when a formal parameter to a macro occurs more than once in a macro definition, e.g.:

```
foo(X,Y) macro f:X, g:X, h:Y.
```

In this ALE macro (which uses true macro variables), F's and G's values will not be shared in the result. That is, `foo(a,b)` for types, a and b, will expand to

```
(f:a, g:a, h:b)
```

in which F and G are not structure-shared unless a is extensional.<sup>14</sup> One way to make the two features' values structure-shared is to substitute a (logical) variable as an actual parameter for X, i.e. `foo(A,b)`. Note that the first argument of the macro `foo` here is a variable rather than an atom. This variable is a "logical" variable because it is used as a first-class citizen of ALE's description logic, rather than at the macro level. In this case the macro expands to the following:

```
(f:A, g:A, h:b)
```

TRALE's logical variable macros, on the other hand, automatically interpret macro parameters such as X and Y as logical variables, and thus implicitly enforce structure-sharing with multiple occurrences. For example:

```
foo(X,Y) := f:X, g:X, h:Y.
```

will automatically structure-share the values of F and G. The infix operator `:=/2` indicates a logical variable macro.

Guard declarations for macros can optionally be applied to these parameters by appending the guard with a hyphen:

```
foo(X-a, Y-b) := f:X, g:X, h:Y.
```

This declaration says that X must be consistent with type a, and Y must be consistent with type b for this macro clause to apply. If it does, F's and G's values are shared. Thus `foo(a,b)` expands to the following:

<sup>13</sup>Recall that Prolog variables start with an upper-case letter.

<sup>14</sup>An extensional type cannot be copied and has to be structure-shared if it occurs more than once in a feature structure. Extensional and intensional types are discussed in ALE User's Guide.

```
(f:(X,a), g:(X,a), h:(Y,b))
```

Note that ALE macros can still be declared (with `macro/2`). As in ALE, `@` is used to call a macro in a description. Let us assume that this signature is defined:

```
type_hierarchy

bot
  person gender:gen nationality:nat name:name
  gen
    male
    female
  nat
    american
    canadian
  name
    john
    mary
.
```

The following macros can then be defined in the theory file:

```
man(X-name,Y-nat) :=
  (person, name:X, gender:male, nationality:Y).

woman(X-name,Y-nat) :=
  (person, name:X, gender:female, nationality:Y).
```

The above macros can now be called in feature descriptions using `@` as in these lexical entries:

```
john ---> @ man(john,american).
mary ---> @ woman(mary,canadian).
```

The integrity of lexical entries can be checked by `lex/1`. Given the above information for example, `lex john` results in the following output in TRALE :

```
| ?- lex john.

WORD: john
ENTRY:
person
GENDER male
NAME john
NATIONALITY american
```

```
ANOTHER? n.
```

```
yes
```

In addition, one may check the integrity of macro definitions by `macro/1`. In this case, `macro woman(X,Y)` produces the following output:

```
| ?- macro woman(X,Y).
```

```
MACRO:
```

```
    woman([0] name,
          [1] nat)
```

```
ABBREVIATES:
```

```
    person
    GENDER female
    NAME [0]
    NATIONALITY [1]
```

```
ANOTHER? n.
```

```
yes
```

### T4.2.2 Macro hierarchies

*W. Detmar Meurers, Ohio State University*<sup>15</sup>

Macros can be hierarchically organized by calling one macro in the definition of another. The macro *X* occurring in the definition of a macro *Y* then can be referred to as a supermacro of *X*. And conversely, *Y* is a submacro of macro *X*.

The notions of sub- and supermacro thus in one sense are parallel to the notion of sub- and supertype. But it is important to keep in mind that ontologically macros and types are very different; in particular an object of a type will also be of one of its subtypes, and of exactly one of its most specific subtypes. There is no equivalent to this with macro hierarchies, which are just subsumption hierarchies of some descriptions that were given names. Different from types, macros have no theoretical status; they just serve to write down a theory more compactly.

In terms of the macro hierarchy commands below, note that (parallel to types) the sub- and supermacro relations only include macros on the same level, not those embedded under features.

This file provides the following top-level predicates, where `<(sub/super)macro>` is the macro name (incl. its argument slots). Many of the predicates exist in two version, one that returns single results and can be backtracked into, and the other (same predicate name, but ending in `s`) which returns the list of all results. Note that the list returned by the second kind of predicates is sorted though. Also, it is worth noting that the predicates returning a list will always succeed (they return a `[]` in the case where `setof` would fail).

- `submacro(<macro>,<submacro>)`.
- `submacros(<macro>,<list(submacros)>)`.
- `supermacro(<macro>,<supermacro>)`.

---

<sup>15</sup>Sections T4.2.2-T4.2.4 ©2003, W. Detmar Meurers

- `supermacros (<macro>, <list(supermacros)>)`.
- `show_submacros (<macro>)`.
- `show_all_submacros (<macro>)`.
- `show_supermacros (<macro>)`.
- `show_all_supermacros (<macro>)`.
- `show_all_macros`. shows the entire macro hierarchy
- `macro (<macro>)` shows the most general satisfier of the description abbreviated by the macro (predicate provided by core ale.pl)
- `is_macro (<macro>)` returns every macro that's defined (if tracked back into)
- `macros (<list(macro)>)` returns list of macros that are defined
- `most_specific_macro (<macro>)`
- `most_specific_macros (<list(macro)>)`
- `most_general_macro (<macro>)`
- `most_general_macros (<list(macro)>)`
- `isolated_macro (<macro>)`
- `isolated_macros (<list(macro)>)` Isolated macros are those that are most general and most specific at the same time, i.e. they neither occur in other macro definitions nor are they defined in terms of other macros.

### **T4.2.3 Automatic generation of macros on different types**

Since HPSG theories usually formulate constraints about different kind of objects, the grammar writer usually has to write a large number of macros to access the same attribute, or to make the same specification, namely one for each type of object which this macro is to apply to. For example, when formulating immediate dominance schemata, one wants to access the VFORM specification of a *sign*. When specifying the valence information one wants to access the VFORM specification of a *synsem* object. And when specifying something about non-local dependencies, one may want to refer to VFORM specifications of *local* objects.

TRALE therefore provides a mechanism which derives definitions of macros describing one type of object on the basis of macros describing another type of object – as long as the linguist tells the system which path of attributes leads from the first type of object to the second.

The path from one type of object to another is specified by including a declarations of the following form in the grammar:

```
access_rule (type1, path, type2) .
```

Such a statement is interpreted as: From `type1` objects you get to `type2` objects by following path `path`.

Since only certain macros are supposed to undergo this extension, they are specified using slightly different operators than standard TRALE macros: access macros are specified using the `':=='` instead of the `':'` operator of ordinary macros. The type of the macro is determined on the basis of the access suffix. The typing of each of the arguments (if any) is added using the `-` operator after each argument.

To distinguish the macro names for different type of objects, a naming convention for macros is needed. All macros on objects of a certain type therefore have the same suffix, e.g., `"_s"` for all macros describing *signs*. The type-suffix pairing chosen by the linguist is specified by including declarations of the following form in the grammar:

```
access_suffix(type,suffix).
```

Such a statement declares that the names of macros describing objects of type `type` end in `suffix`.

So the extend access mechanism reduces the work of the linguist to providing

- macros describing the 'most basic' type of object, and
- `access_suffix` and `access_rule` declarations.

Access macros are not compiled directly. Instead the access macros must be translated to ordinary TRALE macros at some point before compiling a grammar using a call to `extend_access(<FileList>,<OutFileName>)`, where `<FileList>` is a Prolog list of Filenames containing access macro declarations and `<OutFileName>` is a single file containing the ordinary TRALE macros. Note that this resulting file needs to be explicitly loaded as part of the theory file that one compiles (as usual, using `compile_gram/1`) in order for those macros to be compiled as part of the grammar.

**A small example** As an example, say we want to have abbreviations to access the VFORM of a *sign*, a *synsem*, *local*, *cat*, and a *head* object. Then we need to define a macro accessing the most basic object having a VFORM, namely *head*:

```
vform_h(X-vform) ::= vform:X.
```

Second, (once per grammar) `access_suffix` and `access_rule` declarations for the grammar need to be provided. The former define a naming convention for the generated macros by pairing types with macro name suffixes. The latter define the rules to be used by the mechanism by specifying the relevant paths from one type of object to another.

```
access_suffix(head,"_h").
access_suffix(cat,"_c").
access_suffix(loc,"_l").
access_suffix(synsem,"_s").
access_suffix(sign,"what_a_great_suffix").
```

```

access_rule(cat,head,head).
access_rule(loc,cat,cat).
access_rule(synsem,loc,loc).
access_rule(sign,synsem,synsem).

```

This results in the following macros to be generated:

```

vform_h(X)                := vform:X.
vform_c(X)                := head:vform_h(X).
vform_l(X)                := cat:vform_c(X).
vform_s(X)                := loc:vform_l(X).
vformwhat_a_great_suffix(X) := synsem:vform_y(X).

```

If we were only interested in a `vform` macro for certain objects, say those of type *sign* and *synsem*, it would suffice to specify access rules for those types instead of the access rules specified above. The following specifications would do the job:

```

access_suffix(head,"_h").
access_suffix(synsem,"_s").
access_suffix(sign,"what_a_great_suffix").
access_rule(synsem,loc:cat:head,head).
access_rule(sign,synsem:loc:cat:head,head).

```

The result would then be:

```

vform_h(X)                := vform:X.
vform_s(X)                := loc:cat:head:vform_h(X).
vformwhat_a_great_suffix(X) := synsem:loc:cat:head:vform_h(X).

```

## Warnings

Several kinds of warnings can appear on user output. Access expansion continues.

1. A `TRALE` macro already defined always has priority over a derived macro, regardless of whether
  - (a) the derived macro is the direct translation of an access macro defined by the user or
  - (b) the derived macro is the result of `access_rule` applications.
2. If a `TRALE` macro has already been derived by translation of an access macro with or without `access_rule` application, an access macro occurring later in the grammar which would derive the same macro is not translated and no further rules are applied to the later access macro. Currently this is also the case if the two predicates differ in arity.



**Errors**

Several types of errors can be detected and a message is printed on user output. Access expansion then aborts.

1. A type occurring in an `access_rule` is not allowed to have
  - (a) multiple suffixes defined for it
  - (b) no suffixes defined for it
2. Two suffixes defined must not
  - (a) be identical
  - (b) have a common ending

**T4.2.4 Example**

Whenever there is an object of type `main_verb`, its `AUX` and `INV` feature values must be set to `minus`. In the case of auxiliaries, their `AUX` feature has to be `plus` but their `INV` feature could be either `plus` or `minus`.

The following example shows a TRALE logical variable macro. This macro assumes subject verb agreement holds.

```
vp(Ind) :=
  synsem:local:(content:index:Ind,
                cat:subcat:[synsem:local:content:index:Ind]
  ).
```

As mentioned in subsection T4.2.1, TRALE treats variables used in TRALE macro definitions (`:=`) as logical variables and therefore, assumes structure-sharing between multiple occurrences of such variables. Using a TRALE logical variable macro, we ensure that the values of the `INDEX` feature of the verb phrase and of the subject are structure-shared. Therefore,

```
vp((person:third, number:singular))
```

is equivalent to:

```
synsem:local:(content:index:(Ind,
                             person:third,
                             number:singular),
              cat:subcat:[synsem:local:content:index:Ind])
```

In the above feature structure, the values of both `INDEX` features are structure-shared. Had we used a regular ALE macro (using `macro/1`), we would have reached a similar result but the values of the `INDEX` features would simply have been structure-*identical*.

We can also use the type guard declaration of TRALE macros to make sure that `Ind` is consistent with the type `ind`. This can be achieved by adding the guard to the head of the macro definition as follows:

```
vp(Ind-ind):=  
  synsem:local:(content:index:Ind,  
    cat:subcat:[synsem:local:content:index:Ind]).
```

## 4.11 Functional Descriptions

ALE also provides the means to define functions mapping descriptions to descriptions. The syntax is:

```
<func_def> ::= <func_spec> +++> <desc> .

<func_spec> ::= <func_name>
               | <func_name>(<seq(desc)>)
```

Functional descriptions are compiled into code that is evaluated at run-time, first adding to each argument (if any) the description given for that argument, and then evaluating to the resulting description, which can itself include other functional descriptions, including recursive calls. As an example, one may consider the `append/2` function:

```
append([],L) +++> L.
append([X|L1],L2) +++> [X|append(L1,L2)].
```

The lists shown here are instances of ALE's list macro notation. Notice that the only type checking in this definition is performed by appropriateness and the list macro, so the first clause could succeed with any `L`. To add more, one could redefine the first clause as:

```
append([],(list,L)) +++> L.
```

Functional descriptions can be used wherever other descriptions can. The user should be particularly careful with ensuring that the arguments to a functional description call will be sufficiently instantiated for the call to terminate; and the clauses, correctly ordered for the description to terminate correctly. Type checking cannot ensure this, in general. For example, in the definition for `append/2`, with or without explicit type-checking on the second argument, both clauses will match a functional description whose first argument or any TL value in the first argument is strictly of type `list`, i.e., a list that is not known to be `elist` or `nelist`. Such a functional description will, thus, evaluate to an infinite number of results. Clauses in functional descriptions are considered in the order they are given; and the search for solutions *always* backtracks into subsequent clauses.

Any functional description that can be defined so that its argument descriptions are only variables, and its result has no (mutual) recursion should be defined as a macro instead, which is completely expanded at compile-time. Remember, however, that it is not always sufficient to replace the `+++>` with `macro:` variable replacement in macros works by true textual replacement, whereas the variables in functional descriptions are ALE descriptions. For example, the functional description:

```
foo(X) +++> (bar,f:X,g:X).
```

evaluates to a feature structure with a re-entrancy. The macro:

```
foo(X) macro (bar,f:X,g:X).
```

does not necessarily evaluate to a feature structure with a re-entrancy. The functional description, `foo/1` is correctly converted to the macro:

```
foo(X) macro (bar, f:(Y,X), g:(Y,X)).
```

The ALE variable, `Y`, establishes the re-entrancy that the macro variable, `X`, does not.

## 4.12 Type Constraints

Our logical language of descriptions can be used with the type system in order to enforce constraints on feature structures of a particular type. Constraints are attached to types, and may consist of arbitrary descriptions. Their effect is to require every structure of the constrained type to always satisfy the constraint description.

Constraints are enforced using the `cons` operator, e.g.:

```
bot sub [a,b].
  a sub []
    intro [f:b,g:b].
  b sub [].
  a cons (f:X,g:=\= X).
```

The constraint on the type `a` (which must occur within parentheses) requires all feature structures of type `a` to have non-token-identical values for features `f` and `g`. Notice that the type `b` has no constraints expressed. This is equivalent to specifying the constraint:

```
b cons bot.
```

which is satisfied by any feature structure (of type `b`). A type constraint may use any of the operators in the description language, including further type descriptions, which may themselves be constrained. The type, `bot`, may not have type constraints, nor may `a_/1` atoms.

It is crucial that the type descriptions be finitely resolvable. Because simple depth-first search is used to evaluate constraints, infinite resolution paths will cause the system to hang. For example, the following signature should not contain the following constraints:

```
bot sub [a,b].
  a sub [c]
    intro [f:bot].
  c sub [].
  b sub []
    intro [g:bot].
  a cons f:b.
  b cons g:c.
```

This is because `a` subsumes `c`. Notice, however, that type constraints can be used to provide additional information regarding value restrictions on appropriate features. In general, ALE performs more efficiently when restrictions are provided

*Puzzle:* Three different houses each contain a different pet, a different drink, and an inhabitant of a different nationality. The following six facts are true about these houses:

1. The man in the third (middle) house drinks milk.
2. The Spaniard owns the dog.
3. The Ukrainian drinks the tea.
4. The Norwegian lives in the first house.
5. The Norwegian lives next to the tea-drinker.
6. The juice-drinker owns the fox.

*Questions:* Who owns the zebra? Who drinks juice?

Figure 4.2: The Zebra Puzzle.

in the appropriateness conditions, rather than in general constraints; but type constraints can encode a greater variety of restrictions. Specifically, they allow constraints to express path equations and inequations, as well as deeper path restrictions. Constraints may include relational constraints, which are defined using definite clauses, which are discussed below. Type constraints are efficiently compiled in the same way as other descriptions. Also, like appropriateness conditions, they are only enforced once for any given structure.

It is also important to note that because of the delay in ALE's inequational enforcement, type constraints that involve recursion that terminates by an inequation failure may go into infinite loops due to this delay in enforcement. Because extensionality is only enforced before the answer to a top-level query is given, recursive type constraints that rely on the extensional identity of two feature structures to terminate on the basis of their type will not terminate.

### 4.13 Example: The Zebra Puzzle

We now provide a simplified form of the Zebra Puzzle (Figure 4.2), a common puzzle for constraint resolution. This puzzle was solved by Aït-Kaci (1984) using roughly the same methods as we use here. The puzzle illustrates the expressive power of the combination of extensional types, inequations and type constraints. Such puzzles, sometimes known as logic puzzles or constraint puzzles, require one to find a state of affairs within some situation that simultaneously satisfies a set of constraints. The situation itself very often implicitly levies certain constraints.

We can represent the simplified Zebra Puzzle in ALE as:

```
% Subsumption
%=====

bot sub [house,descriptor,background].
```

```

descriptor sub [nat_type,ani_type,bev_type].
  nat_type sub [norwegian,ukranian,spaniard].
    norwegian sub [].
    ukranian sub [].
    spaniard sub [].
  ani_type sub [fox,dog,zebra].
    fox sub [].
    dog sub [].
    zebra sub [].
  bev_type sub [juice,tea,milk].
    juice sub [].
    tea sub [].
    milk sub [].

house sub []
  intro [nationality:nat_type,animal:ani_type,beverage:bev_type].

background sub [clue]
  intro [house1:house,house2:house,house3:house].
  clue sub [maximality].
    maximality sub [].

ext([norwegian,ukranian,spaniard,fox,dog,zebra,juice,tea,milk]).

% Constraints
%=====
background cons
  (house1:nationality:N1,                                     % inequational constraints
   house2:nationality:(N2,(=\ N1)),
   house3:nationality:((=\ N1),(=\ N2)),

   house1:animal:A1,
   house2:animal:(A2,(=\ A1)),
   house3:animal:((=\ A1),(=\ A2)),

   house1:beverage:B1,
   house2:beverage:(B2,(=\ B1)),
   house3:beverage:((=\ B1),(=\ B2))).

clue cons
  (house3:beverage:milk,                                     % clue 1

   (house1:nationality:spaniard,house1:animal:dog           % clue 2
    ;house2:nationality:spaniard,house2:animal:dog
    ;house3:nationality:spaniard,house3:animal:dog),

   (house1:nationality:ukranian,house1:beverage:tea        % clue 3
    ;house2:nationality:ukranian,house2:beverage:tea
  
```

```

;house3:nationality:ukranian,house3:beverage:tea),

house1:nationality:norwegian,                                % clue 4

(house1:nationality:norwegian,house2:beverage:tea % clue 5
;house2:nationality:norwegian,house3:beverage:tea
;house2:nationality:norwegian,house1:beverage:tea
;house3:nationality:norwegian,house2:beverage:tea),

(house1:beverage:juice,house1:animal:fox                    % clue 6
;house2:beverage:juice,house2:animal:fox
;house3:beverage:juice,house3:animal:fox)).

maximality cons
(house1:nationality:(norwegian;ukranian;spaniard), % maximality constraints
house2:nationality:(norwegian;ukranian;spaniard),
house3:nationality:(norwegian;ukranian;spaniard),

house1:animal:(fox;dog;zebra),
house2:animal:(fox;dog;zebra),
house3:animal:(fox;dog;zebra),

house1:beverage:(juice;tea;milk),
house2:beverage:(juice;tea;milk),
house3:beverage:(juice;tea;milk)).

```

The subsumption hierarchy is shown pictorially in Figure 4.3. The type, `background`, with the assistance of the types subsumed by `house` and `descriptor`, represents the situation of three houses (the features of `background`), each of which has three attributes (the features of `house`). The implicit constraints levied by the situation appear as constraints on the type, `background`, namely that no two houses can have the same value for any attribute. These are represented by inequations. But notice that, since we are interested in treating the values of attributes as tokens, we must represent those values by extensional types. If we had not done this, then we could still, for example, have two different houses with the beverage, `juice`, since there could be two different feature structures of type `juice` that were not token-identical. Notice also that all of these types are maximal, and hence satisfy the restriction that ALE places on extensional types.

The explicit constraints provided by the clues to the puzzle are represented as constraints on the type `clue`, a subtype of `background`. We also need a subtype of `clue`, `maximality`, to enforce another constraint implicit to all constraint puzzles, namely the one which requires that we provide only maximally specific answers, rather than vague solutions which say, for example, that the beverage for the third house is a type of beverage (`bev_type`), which may actually still satisfy a puzzle's constraints.

To solve the puzzle, we simply type:

```
| ?- mgsat maximality.
```

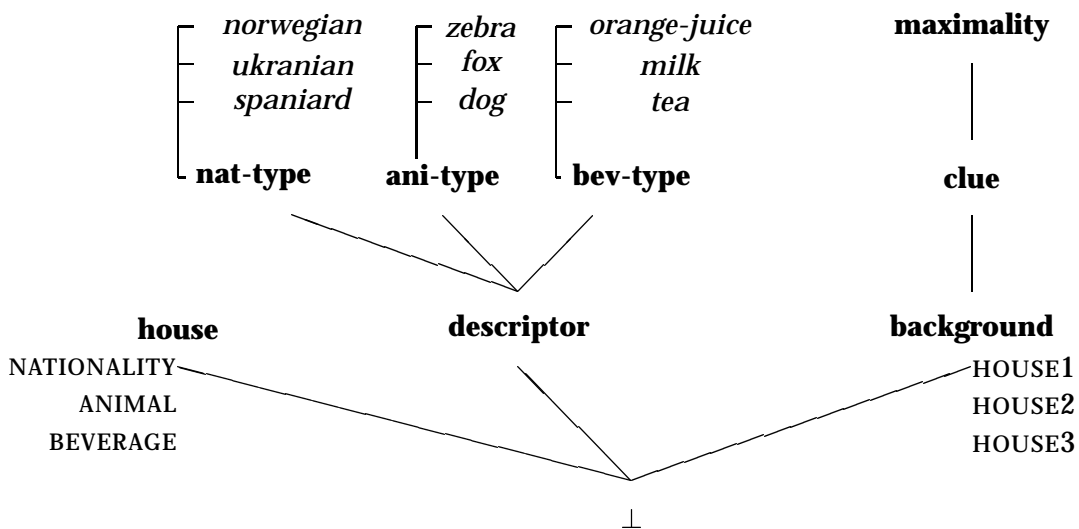


Figure 4.3: Inheritance Network for the Zebra Puzzle.

MOST GENERAL SATISFIER OF: maximality

```

maximality
HOUSE1 house
      ANIMAL fox
      BEVERAGE juice
      NATIONALITY norwegian
HOUSE2 house
      ANIMAL zebra
      BEVERAGE tea
      NATIONALITY ukrainian
HOUSE3 house
      ANIMAL dog
      BEVERAGE milk
      NATIONALITY spaniard
  
```

ANOTHER? y.

```

no
| ?-
  
```

So the Ukrainian owns the zebra, and the Norwegian drinks juice. A most general satisfier of maximality will also satisfy the constraints of its supertypes, background and clue.

Although ALE is capable of representing such puzzles and solving them, it is not actually very good at solving them efficiently. To solve such puzzles efficiently, a system must determine an optimal order in which to satisfy all of the various



constraints. ALE does not do this since it can express definite clauses in its constraints, and the reordering would also be very difficult for the user to keep track of while designing a grammar. A system that does do this is the general constraint resolver proposed by Penn and Carpenter (1993)<sup>16</sup>.

---

<sup>16</sup>This system was actually the precursor to ALE. It implemented a completely reversible constraint-based parser/generator with a weighting on the constraints based on their maximal non-determinism. Re-ordering constraints, however, proved to be insufficient for efficient parsing or generation, compared to a uni-directional system such as ALE.

### T4.3 Complex-antecedent constraints

TRALE theories can use all of the declarations available to ALE. These include relations, lexical entries, lexical rules, extended phrase structure rules, and Prolog-like definite clauses over typed feature structures. They also include one extra, complex-antecedent constraints, which is discussed in the following section.

ALE's constraints are more restricted in the sense that the antecedent (left-hand side) of a constraint can only be a type. ALE constraints have the following forms:

```
t cons Desc.
```

or

```
t cons Desc goal Goal.
```

where the antecedent, *t*, is a type and the consequent (right-hand side), *Desc*, is the description of the constraint. These constraints apply their consequents to all feature structures of type *t*, which is to say that they are universally quantified over all feature structures and are triggered based on subsumption, i.e., *Desc* and *Goal* are not necessarily applied to feature structures consistent with *t*, but rather to those that are subsumed by the most general satisfier of *t*, as given by a constraint-free signature. For example, the constraint:

```
a cons (f:X, g:=\=X)
```

states that the values of the *F* and *G* features of any object of type *a* or a type subsumed by *a* must not be structure-shared. If a goal is specified in a constraint, the constraint is satisfied only if the goal succeeds.

TRALE generalizes this to allow for antecedents with arbitrary function-free, inequation-free antecedents. TRALE constraints are defined using the infix operator, *\*>/2*, e.g.:

```
(f:minus, g:plus) *> X goal foo(X).
```

This constraint executes *foo/1* on feature structures with *F* value *minus* and *G* value *plus*. These are also universally quantified, and they are also triggered based on subsumption by the antecedent, not unification. For example, in the above constraint, if a feature structure does not have an *F* value, *minus*, and a *G* value, *plus*, yet, but is not inconsistent with having them, then TRALE will suspend consideration of the constraint until such a time as it acquires the values or becomes inconsistent with having them. These constraints are thus consistent with a classical view of implication in that the inferences they witness are sound with respect to classical implication, but they are not complete. On the other hand, they operate passively in the sense that the situation can arise in which several constraints may each be waiting for the other for their consequents to apply. This is called *deadlock*. An alternative to this approach would be to use unification to decide when to apply consequents, or some built-in search mechanism that would avoid deadlock but risk non-termination. Of the two, deadlock

is preferable. Additional (constraint) logic programs can be written to search the space of possible solutions in a way suited to the individual case if there are dead-locked or suspended constraints in a solution.

Variables in `*>` constraints are implicitly existentially quantified within the antecedent. Thus:

```
(f:X, g:X) *> Y goal foo(Y).
```

applies the consequent when F's and G's values are structure-shared. In other words, it applies the consequent if there exists an X such that F's and G's values are both X. In addition, the consequent of the above-mentioned constraint implicitly applies only to feature structures for which F and G are both appropriate. Path equations can be used in antecedents to explicitly request structure-sharing as well. The following constraint is equivalent to the one above:

```
([f]==[g]) *> Y goal foo(Y).
```

A singleton variable in the antecedent results in no delaying itself apart from the implicit appropriateness conditions. Variables occurring in the consequent are in fact bound with scope that extends over the consequent and relational attachments. Thus, in the following example:

```
(f:X, g:X) *> W goal foo(X, W).
```

the first argument passed to `foo/2` is the very X that is the value of both F and G. This use of variables on both sides of the implication is a loose interpretation consistent with common practice in linguistics. With a classical interpretation of implication, it is always true that:

$$p \rightarrow q \quad \text{iff} \quad p \rightarrow (p \wedge q)$$

Thus:

$$(\exists x_1, \dots, \exists x_n. p) \rightarrow q \quad \text{iff} \quad (\exists x_1, \dots, \exists x_n. p) \rightarrow ((\exists x_1, \dots, \exists x_n. p) \wedge q)$$

Note that, according to the classical interpretation, *q* is not in the scope of the existential quantifiers. TRALE, however, bends this rule taking the latter to be equivalent to:

$$(\exists x_1, \dots, \exists x_n. p) \rightarrow \exists x_1, \dots, \exists x_n. (p \wedge q)$$

Here, *q* is in the scope of the  $\exists x_1, \dots, \exists x_n$ . Consequently, if the genuine equivalence is desired, one must ensure that  $x_1, \dots, x_n$  do not occur in *q*, in which case  $(\exists x_1, \dots, \exists x_n. p) \wedge q$  and  $\exists x_1, \dots, \exists x_n. (p \wedge q)$  are equivalent. Requiring this extra variable hygiene is the price of permitting this non-equivalent use of implication and quantification as the default, rather than what is logically valid. An example of such a relaxed use of such rules in linguistics is provided below:

```
spec_dtr: (head:noun,
           index:Ind)
*> head_dtr: (head:verb,
             index:Ind).
```

The above constraint is formulated so as to assure subject-verb agreement by enforcing structure-sharing between the INDEX feature of the subject and that of the verb. In the strict interpretation, the second instance of the variable `Ind` would have broad scope (existentially quantified over the entire clause), and thus be possibly different from the `Ind` in the antecedent. TRALE's interpretation of this, however, assumes they are the same. As mentioned above, if the strict interpretation is desired, a different variable name must be used in the consequent of this constraint.

In some languages, the form of the verb depends on the type of eventuality denoted by the sentence. In Czech, for example, it is generally the case that if an event is total (i.e. completed as opposed to simply terminated), the verb that denotes that event surfaces in the perfective form. This rule can be enforced as the following constraint in the grammar. Note that since the constraint applies to a type rather than a particular description, we could use an ALE constraint (`cons/2`), too.

```
(sentence, event:total) *>
    synsem:(local:(cat:(head:(vform:perf))))).
```

This constraint applies its consequent to all feature structures of type `sentence` with the required `EVENT` value.

Another example of a complex-antecedent constraint can be found in HPSG's Binding Theory, which refers to the notions of *o-binding* and *o-commanding*. *O-binding* is defined as follows (see Pollard and Sag 1994, p. 253–54):

“Y (*locally*) *o-binds* Z just in case Y and Z are coindexed and Y (*locally*) *o-commands* Z. If Z is not (*locally*) *o-bound*, then it is said to be (*locally*) *o-free*.”

Pollard and Sag define *o-commanding* as follows:

“Let Y and Z be *synsem* objects with distinct LOCAL values, Y referential. Then Y *locally o-commands* Z just in case Y is less oblique than Z...

“Let Y and Z be *synsem* objects with distinct LOCAL values, Y referential. Then Y *o-commands* Z just in case Y locally *o-commands* X dominating Z.”

Based on these notions, HPSG's Binding Theory is phrased as the following three principles (*ibid*):

HPSG Binding Theory:

Principle A. A locally *o-commanded* anaphor must be locally *o-bound*.

Principle B. A personal pronoun must be locally *o-free*.

Principle C. A nonpronoun must be *o-free*.

A simplified version of Principle A can be written as a constraint over all head-complement structures (`head_comp_struct`) as follows:

```

% BINDING THEORY
% PRINCIPLE A
head_comp_struct *> (spec_dtr:(synsem:content:
                    (X,
                     index:IndX)),
                    comp_dtr:(synsem:content:
                    (Y,ana,
                     index:IndY)))
goal (local_o_command(X,Y) ->
      IndX = IndY; true).

(X = X) if true.

```

The above ALE constraint makes sure that for all `head_comp_struct` type objects, if the complement daughter is anaphoric and locally o-commanded by the specifier, then the two daughters must be coindexed. The definition “(X = X) if true” is provided because ALE does not come with a built-in equality relation defined over descriptions. We leave the definitions of `local_o_command/2` and `o_command/2` to the reader.

An alternative is to formulate the constraint in the following manner:

```

% BINDING THEORY
% PRINCIPLE A (Alternate formulation):
(spec_dtr:(synsem:content:(X,index:IndX)),
 comp_dtr:(synsem:content:(Y,ana,index:IndY)))
*> bot
goal (local_o_command(X,Y) -> IndX = IndY; true).

(X = X) if true.

```

The above constraint applies to any description subsumed by the antecedent. If the first daughter’s index locally o-commands the second’s, then they should be coindexed. `bot` results in no additional description be added. Alternatively, one could use an anonymous variable, “\_”.

Let us now see how Principle B can be written as a TRALE complex-antecedent constraint:

```

% Binding Theory
% PRINCIPLE B:
(spec_dtr:(synsem:content:(X,index:Ind)),
 comp_dtr:(synsem:content:(Y,ppro,index:Ind)))
*> bot
goal (local_o_command(X,Y) -> fail; true).

```

This constraint states that for all descriptions with specifier and complement daughters, if the latter is a personal pronoun and locally o-commanded by the former, then the two daughters must not be coindexed.

Analogously, Principle C can be encoded as follows:

```
% BINDING THEORY
% PRINCIPLE C:
(spec_dtr:(synsem:content:(X,index:Ind)),
 comp_dtr:(synsem:content:(Y,npro,index:Ind)))
*> bot
goal (o_command(X,Y) -> fail; true).
```

This last constraint states that for all descriptions with specifier and complement daughters, if the second one is a nonpronoun (*npro*), then it must not be *o*-commanded by the first and coindexed with it.

## Chapter 5

# Definite Clauses

The next two sections, covering the constraint logic programming and phrase structure components of ALE, simply describe how to write ALE programs and how they will be executed. Discussion of interacting with the system itself follows the description of the programming language ALE provides.

The definite logic programming language built into ALE is a constraint logic programming (CLP) language, where the constraint system is the attribute-value logic described above. Thus, it is very closely related to both Prolog and LOGIN. Like Prolog, definite clauses may be defined with disjunction, negation and cut. The definite clauses of ALE are executed in a depth-first, left to right search, according to the order of clauses in the database. ALE performs last call optimization, but does not perform any clause indexing.<sup>1</sup> Those familiar with Prolog should have no trouble adapting that knowledge to programming with definite clauses in ALE. The only significant difference is that first-order terms are replaced with descriptions of feature structures.

While it is not within the scope of this user's guide to detail the logic programming paradigm, much less CLP, this section will explain all that the user familiar with logic programming needs to know to exploit the special features of ALE. For background, the user is encouraged to consult Sterling and Shapiro (1986) with regard to general logic programming techniques, most of which are applicable in the context of ALE, and Ait-Kaci and Nasr (1986) for more details on programming with sorted feature structures. For more advanced material on programming in Prolog with a compiler, see O'Keefe (1990). The general theory of CLP is developed in a way compatible with ALE in Höhfeld and Smolka (1988). Of course, since ALE is literally an implementation of the theory found in Carpenter (1992), the user is strongly encouraged to consult Chapter 14 of that book for full theoretical details.

The syntax of ALE's logic programming component is broadly similar to that of Prolog, with the only differences being that first-order terms are generally replaced with attribute-value logic descriptions, and a different language of block conditionals for co-routining (see p. 44). The language in which clauses are expressed in ALE is given in BNF as:

```
<clause> ::= <literal> if <goal>.
```

---

<sup>1</sup>Thus, additional cuts might be necessary to ensure determinism, so that last call optimization is effective.

```

<literal> ::= <pred_sym>
           | <pred_sym>(<seq(desc)>)

<goal> ::= true
        | <literal>
        | (<goal>,<goal>)
        | (<goal>;<goal>)
        | (<desc>=@ <desc>)
        | (<cut_free_goal> -> <goal>)
        | (<cut_free_goal> -> <goal> ; <goal>)
        | !
        | (\+ <goal>)
        | prolog(<prolog_goal>)
        | when(<cond>,<goal>)

```

Just as in Prolog, predicate symbols must be Prolog atoms. This is a more restricted situation than the definite clause language discussed in Carpenter (1992), where literals were also represented as feature structures and described using attribute-value logic. Also note that ALE requires every clause to have a body, which might simply be the goal `true`. There must be whitespace around the `if` operator, but none is required around the conjunction comma, the disjunction semicolon, the cut or shallow cut symbols `!`, `->`, or the unprovability symbol `\+`. Parentheses, in general, may be dropped and reconstructed based on operator precedences. The precedence is such that the comma binds more tightly than the semicolon, while the unprovability symbol binds the most tightly. Both the semicolon and comma are right associative.

The operational behavior of ALE is nearly identical to Prolog with respect to goal resolution. That is, it evaluates a sequence of goals depth-first, from the left to right, using the order of clauses established in the program. The only difference arises from the fact that, in Prolog, terms cannot introduce non-determinism. In ALE, due to the fact that disjunctions can be nested inside of descriptions, additional choice points might be created both in matching literals against the heads of clauses and in expanding the literals within the body of a clause. In evaluating these choices, ALE maintains a depth-first left to right strategy.

We begin with a simple example, the `member/2` predicate:<sup>2</sup>

```

member(X,hd:X) if
  true.
member(X,tl:Xs) if
  member(X,Xs).

```

As in Prolog, ALE clauses may be read logically, as implications, from right to left. Thus the first clause above states that `X` is a member of a list if it is the head of a list. The second clause states that `X` is a member of a list if `X` is a member of the tail of the list, `Xs`. Note that variables in ALE clauses are used the same way as in Prolog, due to the notational convention of our description language. Further note that, unlike Prolog, ALE requires a body for every clause. In particular, note that the first

<sup>2</sup>As in Prolog, we refer to predicates by their name and arity.



clause above has the trivial body `true`. The compiler is clever enough to remove such goals at compile time, so they do not incur any run-time overhead.

Given the notational convention for lists built into ALE, the above program could equivalently be written as:

```
member(X, [X|_]) if
  true.
member(X, [_|Xs]) if
  member(X, Xs).
```

But recall that ALE would expand `[X|_]` as `(hd:X, tl:_)`. Not only does ALE not support anonymous variable optimizations, it also creates a conjunction of two descriptions, where `hd:X` would have sufficed. Thus the first method is not only more elegant, but also more efficient.

Due to the fact that lists have little hierarchical structure, list manipulation predicates in ALE look very much like their correlates in Prolog. They will also execute with similar performance. But when the terms in the arguments of literals have more interesting taxonomic structure, ALE actually provides a gain over Prolog's evaluation method, as pointed out by Ait-Kaci and Nasr (1986). Consider the following fragment drawn from the syllabification grammar in the appendix, in which there is a significant interaction between the inheritance hierarchy and the definite clause `less_sonorous/2`:

```
segment sub [consonant,vowel].
  consonant sub [nasal,liquid,glide].
    nasal sub [n,m].
      n sub [].
      m sub [].
    liquid sub [l,r].
      l sub [].
      r sub [].
    glide sub [y,w].
      y sub [].
      w sub [].
  vowel sub [a,e,i].
    a sub [].
    e sub [].
    i sub [].

less_sonorous_basic(nasal,liquid) if true.
less_sonorous_basic(liquid,glide) if true.
less_sonorous_basic(glide,vowel) if true.

less_sonorous(L1,L2) if
  less_sonorous_basic(L1,L2).
less_sonorous(L1,L2) if
  less_sonorous_basic(L1,L3),
  less_sonorous(L3,L2).
```

For instance, the third clause of `less_sonorous_basic/2`, being expressed as a relation between the types `glide` and `vowel`, allows solutions such as `less_sonorous_basic(w,e)`, where `glide` and `vowel` have been instantiated as the particular subtypes `w` and `e`. This fact would not be either as straightforward or as efficient to code in Prolog, where relations between the individual letters would need to be defined. The loss in efficiency stems from the fact that Prolog must either code all 14 pairs represented by the above three clauses and type hierarchy, or perform additional logical inferences to infer that `w` is a glide, and hence less sonorous than the vowel `e`. ALE, on the other hand, performs these operations by unification, which, for types, is a simple table look-up.<sup>3</sup> All in all, the three clauses for `less_sonorous_basic/2` given above represent relations between 14 pairs of letters. Of course, the savings is even greater when considering the transitive closure of `less_sonorous_basic/2`, given above as `less_sonorous/2`, and would be greater still for a type hierarchy involving a greater degree of either depth or branching.

While we do not provide examples here, suffice it to say that cuts, shallow cuts, negation, conjunction and disjunction work exactly the same as they do in Prolog. In particular, cuts conserve stack space representing backtracking points, disjunctions create choice points and negation is evaluated by failure, with the same results on binding as in Prolog.

The definite clause language also allows arbitrary prolog goals, using the predicate, `prolog(<prolog_goal>)`. This is perhaps most useful when used with the Prolog predicates, `assert` and `retract`, which provide ALE users with access to the Prolog database, and with I/O statements, which can be quite useful for debugging definite clauses.

Should `prolog_goal` contain a variable that has been instantiated to an ALE feature structure, this will appear to Prolog as ALE's internal representation of that feature structure. Feature structures can be asserted and retracted, however, without regard to their internal structure. The details of ALE's internal representation of feature structures can be found in Carpenter and Penn (1996), and briefly on p. 85.

Another side effect of not directly attaching inequations to feature structures is that if a feature structure with inequations is asserted and a copy of it is later instantiated from the Prolog database or retracted, the copy will have lost the inequations. In general, passing feature structures with inequations to Prolog hooks should be avoided.

Because the enforcement of extensionality is delayed in ALE, a variable which is instantiated to an extensionally typed feature structure and then passed to a prolog hook may also not reflect token identities as a result of extensionality. Provided that there are no inequations (to which the user does not have direct access), this can be enforced within the hook by calling the ALE internal predicate `extensionalise(FS, [])`.

There is a special literal predicate, `=@`, used with infix notation, which is true when its arguments are token-identical. As with inequations, which forbid token-identity, the `=@` operator is of little use unless multiply occurring variables are

---

<sup>3</sup>Table look-ups involved in unification in ALE rely on double hashing, once for the type of each structure being unified.

used in its arguments' descriptions. Note, however, that while inequations ( $\neq$ ) and path equations ( $\equiv$ ) are part of the description language,  $\neq$  is a definite clause predicate, and cannot be used as a description. It is more important to note that while the negation of the structure-identity operator ( $\neq$ ), namely the inequation ( $\neq$ ), is monotonic when interpreted persistently, the negation of the token-identity operator ( $\neq$ ), achieved by using it inside the argument of the  $\setminus+$  operator, is non-monotonic, and thus its use should be avoided.

It is significant to note that clauses in ALE are truly definite in the sense that only a single literal is allowed as the *head* of a clause, while the *body* can be a general goal. In particular, disjunctions in descriptions of the arguments to the head literal of a clause are interpreted as taking wide scope over the entire clause, thus providing the effect of multiple solutions rather than single disjunctive solutions. The most simple example of this behavior can be found in the following program:

```
foo((b;c)) if true.

bar(b) if true.

baz(X) if foo(X), bar(X).
```

Here the query `foo(X)` will provide two distinct solutions, one where  $X$  is of type `b`, and another where it is of type `c`. Also note that the queries `foo(b)` and `foo(c)` will succeed. Thus the disjunction is equivalent to the two single clauses:

```
foo(b) if true.
foo(c) if true.
```

In particular, note that the query `baz(X)` can be solved, with  $X$  instantiated to an object of type `b`. In general, using embedded disjunctions will usually be more efficient than using multiple clauses in ALE, especially if the disjunctions are deeply nested late in the description. On the other hand, cuts can be inserted for control with multiple clauses, making them more efficient in some cases.

## 5.1 Type Constraints Revisited

The type constraints mentioned in the last chapter can also incorporate relational constraints defined by definite clauses, with the optional operator `goal`. Consider the following example from HPSG:

```
word cons W
    goal (single_rel_constraint(W),
         clausal_rel_prohibition(W)).
```

In this example, the constraint from the description language is simply the variable  $W$ , which is used to match any feature structure of type `word`. That feature structure is then passed as an argument to the two procedural attachments `single_rel_constraint/1` and `clausal_rel_prohibition/1`, which each represent a principle from HPSG which governs words (among other objects). Notice that the goal, when non-literal, must occur within parentheses.

While every type constraint must have a description, procedural attachments are optional. If they do occur, they occur after the description. The syntax is given in BNF as:

```
<cons_spec> ::= <type> cons <desc>
              | <type> cons <desc>
                goal <goal>
```

## 5.2 Co-routining

Co-routining allows programmers to block the execution of a particular thread of execution until some event is witnessed. In logic programming, variables serve as the witnesses to these events — naturally extending their role as representatives of state in the logical setting. Users who are already familiar with SICStus Prolog's `when/2` should exercise some caution with ALE's own `when/2` predicate: as described in further detail below, its conditional expressions generalize SICStus Prolog's conditional expressions to the typed case in such a way that a condition may provably never occur, and its treatment of shared variables in conditional expressions is also necessarily different (see Section 5.2.1). Unlike SICStus Prolog, ALE currently provides no support for residuation or examination of blocked goals.

`when/2` delays execution until some event is witnessed. This is often used to protect very declaratively written goals from non-termination in certain calling contexts, or to make highly disjunctive goals more efficient, by waiting until a condition is observed that reduces the number of satisfying disjuncts. A common example is:

```
foo(X) :-
    ...                % NewVar does not occur here
    when((X=ne_list;X=e_list),append(X,List,NewVar)).

append([],L,L) if true.
append([H|T],L,[H|Res]) if
    append(T,L,Res).
```

The call to `when/2` in `foo/1` waits until `X` is either a non-empty list or an empty list before calling `append/3`. Not to do so with the definition of `append/3` given above, that is, calling `append/3` when `X` is only of type `list` or `bot`, would result in an infinite number of possible answers, because `NewVar` is a new variable. A more compact way of writing the same thing is:

```
foo(X) :-
    ...                % NewVar does not occur here
    when((X=max(list),append(X,List,NewVar)).
```

This waits until `X` has one of the maximally specific types that `list` subsumes. In general, this method is to be preferred because ALE can implement it more efficiently, but there are some disjunctive conditionals that cannot be re-expressed in this fashion without declaring extra types in the signature.

The first argument to `when/2` is called a *conditional expression* or *block condition*; it expresses the event that ALE must wait for before executing the goal in the second argument to `when/2`. Fundamentally, conditional expressions take the form of equations in which the left-hand side is an ALE variable (denoting a typed feature structure), and the right-hand side is a function-free, inequation-free description. Macros can be used, as long as the macros do not expand to a description that contains a function or an inequation. Conditional expressions are closed under conjunction and disjunction, both at the description-level inside the right-hand side of a conditional equation, and outside at the level of the conditional equations themselves. The example above uses an “outside” disjunction, which could have been expressed “inside” as:

```
foo(X) :-
    ...                % NewVar does not occur here
    when(X=(ne_list;e_list), append(X,List,NewVar)).
```

Neither is more efficient than the other, but “outside” disjunctions must be used when the left-hand side variables differ.

When ALE encounters a goal blocked by a conditional equation, it waits until the equation is known to be necessarily true or false. This typically happens as a result of the program or grammar acquiring more information about a variable, such as the structure-sharing of two of its features’ values, or a more specific or join-compatible type. If the conditional equation becomes provably false, i.e., if it becomes known that it will never be true, then the blocked goal is discarded and never executed. If it becomes true, then the blocked goal is executed. If independently blocked goals have the same conditional, or have conditionals that become provably true at the same time, then it is unpredictable in which order they will be executed. If any blocked goal fails, then execution fails back to the point where the goal was *blocked*, not executed. For the most part, this is in keeping with the operational behaviour of `when/2` in SICStus Prolog.

Unlike SICStus Prolog conditionals such as `nonvar/1` or `?=/2`, however, an ALE conditional expression may become provably false. A conditional equation on a type, for example, such as `X=e_list`, can never become true once `X` promotes to an incompatible type, such as `ne_list`. Blocked goals in ALE should be written with this fact in mind. In particular, Prolog goals blocked on `?=/2` in SICStus Prolog must check for equality or inequality because `?=/2` could become true by virtue of either, whereas ALE goals blocked on structure-sharing are executed only as a result of equality. ALE goals cannot be blocked on inequality.

The BNF syntax for conditionals is given in the appendix.

### 5.2.1 Shared variables in conditionals

Particular care must also be taken with shared variables in ALE conditional expressions because there are two scopes within which they can be quantified. The goal:

```
when(X=( [f] == [g] ), bar(X))
```

for example, is not the same as:

```
when(X=(f:Y,g:Y),bar(X))
```

because the former waits until the F and G values of X are equal, whereas the latter waits until the F and G values of X are both equal to the value of some variable Y. Variables such as Y are quantified with scope over the entire clause within which the `when/2` statement occurs, just as any other body variable in ALE or Prolog. Unlike Prolog, however, ALE provides the ability to change this quantification to take narrower scope in block conditions. The goal

```
when(Y^(X=(f:Y,g:Y)),bar(X))
```

is equivalent to a delay on a path equation between F and G, as in the first alternative above. It waits until there exists a Y such that both the F and G values are equal to it. Y is called a *narrowly quantified variable*, and its scope is the `when/2` statement in which it occurs. When `baz/2` is called by this clause, for example:

```
foo(X) :-
  Z = f:Y,
  when(Y^(X=(f:Y,g:Y)),baz(Y,Z)).
```

the first occurrence of Y is not necessarily the same as the F and G values of X (which are both equal to the first argument of `baz/2`). The first occurrence of Z, on the other hand, is the same as the second argument of `baz/2`.

The left-hand side variable of a conditional equation may not be a narrowly quantified variable with scope over the same `when/2` goal. In other words, this goal:

```
when(X^(X=f:t),foo(Y))
```

is not allowed, whereas this goal:

```
when(X^(Y=f:(t,X)),when(X=g:t,foo(Y)))
```

is allowed. X can occur on the left-hand side of the blocked `when/2` goal because it was narrowly quantified in a different `when/2` goal.

## Chapter 6

# Phrase Structure Grammars

The ALE phrase structure processing component is loosely based on a combination of the functionality of the PATR-II system and the DCG system built into Prolog. Roughly speaking, ALE provides a system like that of DCGs, with two primary differences. The first difference stems from the fact that ALE uses attribute-value logic descriptions of typed feature structures for representing categories and their parts, while DCGs use first-order terms (or possibly cyclic variants thereof). The second primary difference is that ALE's parser uses a bottom-up active chart parsing algorithm and a semantic-head-driven generator rather than encoding grammars directly as Prolog clauses and evaluating them top-down and depth-first. In the spirit of DCGs, ALE allows definite clause procedures to be attached and evaluated at arbitrary points in a phrase structure rule, the difference being that these rules are given by definite clauses in ALE's logic programming system, rather than directly in Prolog.

Phrase structure grammars come with two basic components, one for describing lexical entries and empty categories, and one for describing grammar rules. We consider these components in turn.

### 6.1 Lexical Entries

Lexical entries in ALE are specified as rewriting rules, as given by the following BNF syntax:

```
<lex_entry> ::= <word> ---> <desc>.
```

For instance, in the categorial grammar lexicon in the appendix, the following lexical entry is provided, along with the relevant macros:

```
john --->
  @ pn(j).

pn(Name) macro
  synsem: @ np(Name),
  @ quantifier_free.

np(Ind) macro
```

```
syn:np,
sem:Ind.
```

```
quantifier_free macro
qstore:[].
```

Read declaratively, this rule says that the word `john` has as its lexical category the most general satisfier of the description `@ pn(j)`, which is:

```
cat
SYNSEM basic
      SYN np
      SEM j
QSTORE e_list
```

Note that this lexical entry is equivalent to that given without macros by:

```
john --->
  synsem:(syn:np,
          sem:j),
  qstore:e_list.
```

Macros are useful as a method of organizing lexical information to keep it consistent across lexical entries. The lexical entry for the word `runs` is:

```
runs ---> @ iv((run,runner:Ind),Ind).

iv(Sem,Arg) macro
  synsem:(backward,
          arg: @ np(Arg),
          res:(syn:s,
              sem:Sem)),
  @ quantifier_free.
```

This entry uses nested macros along with structure sharing, and expands to the category:

```
cat
SYNSEM backward
      ARG synsem
          SYN np
          SEM [0] sem_obj
      RES SYN s
          SEM run
          RUNNER [0]
QSTORE e_list
```

It also illustrates how macro parameters are in fact treated as variables.

Multiple lexical entries may be provided for each word. Disjunctions may also be used in lexical entries, but are expanded out at compile-time. Thus the first three lexical entries, taken together, compile to the same result as the fourth:



```

bank --->
  syn:noun,
  sem:river_bank.
bank --->
  syn:noun,
  sem:money_bank.
bank --->
  syn:verb,
  sem:roll_plane.

bank --->
  ( syn:noun,
    sem:( river_bank
          ; money_bank
        )
    ; syn:verb,
      sem:roll_plane
  ).

```

Note that this last entry uses the standard Prolog layout conventions of placing each conjunct and disjunct on its own line, with commas at the end of lines, and disjunctions set off with vertically aligned parentheses at the beginning of lines.

The compiler finds all the most general satisfiers for lexical entries at compile time, reporting on those lexical entries that have unsatisfiable descriptions. In the above case of `bank`, the second combined method is marginally faster at compile-time, but their run-time performance is identical. The reason for this is that both entries have the same set of most general satisfiers.

ALE supports the construction of large lexica, as it relies on Prolog's hashing mechanism to actually look up a lexical entry for a word during bottom-up parsing. For generation, ALE indexes lexical entries for faster unification, as described in Penn and Popescu (1997). Constraints on types can also be used to enforce conditions on lexical representations, allowing for further factorization of information.

## 6.2 Empty Categories

ALE allows the user to specify certain categories as occurring without any corresponding surface string. These are usually referred to somewhat misleadingly as *empty categories*, or sometimes as *null productions*. In ALE, they are supported by a special declaration of the form:

```
empty <desc>.
```

Where `<desc>` is a description of the empty category.

For example, a common treatment of bare plurals is to hypothesize an empty determiner. For instance, consider the contrast between the sentences *kids overturned my trash cans* and *a kid overturned my trash cans*. In the former sentence,

which has a plural subject, there is no corresponding determiner. In our categorial grammar, we might assume an empty determiner with the following lexical entry (presented here with the macros expanded):

```
empty @ gdet(some).

gdet(Quant) macro
  synsem: (forward,
    arg: (syn: (n,
      num: plu),
    sem: (body: Restr,
      ind: Ind)),
  res: (syn: (np,
    num: plu),
    sem: Ind),
  qstore: [ (Quant,
    var: Ind,
    restr: Restr) ] .
```

Of course, it should be noted that this entry does not match the type system of the categorial grammar in the appendix, as it assumes a number feature on nouns and noun phrases.

Empty categories are expensive to compute under a bottom-up parsing scheme such as the one used in ALE. The reason for this is that these categories can be used at every position in the chart during parsing (with the same begin and end points). If the empty categories cause local structural ambiguities, parsing will be slowed down accordingly as these structures are calculated and then propagated. Consider the empty determiner given above. It can be used as an inactive edge at every node in the chart, then match the forward application rule scheme and search through every edge to its right looking for a nominal complement. If there are relatively few nouns in a sentence, not many noun phrases will be created by this rule and thus not many structural ambiguities will propagate. But in a sentence such as *the kids like the toys*, there will be an edge spanning *kids like the toys* corresponding to an empty determiner analysis of *kids*. The corresponding noun phrase created spanning *toys* will not propagate any further, as there is no way to combine a noun phrase with the determiner *the*. But now consider the empty slash categories of form  $X/X$  in GPSG. These categories, when coupled with the slash passing rules, would roughly double parsing time, even for sentences that can be analyzed without any such categories. The reason is that these empty categories are highly underspecified and thus have many options for combinations. Thus empty categories should be used sparingly, and preferably in environments where their effects will not propagate.

Another word of caution is in order concerning empty categories: they can occur in constructions with other empty categories. For instance, if we specify categories  $C_1$  and  $C_2$  as empty categories, and have a rule that allows a  $C$  to be constructed from a  $C_1$  and a  $C_2$ , then  $C$  will act as an empty category, as well. These combinations of empty categories are computed at compile-time; but the sheer number of empty categories produced under this closure may be a processing burden if they apply at run-time too productively. Keep in mind that ALE

computes all of the inactive edges that can be produced from a given input string, so there is no way of eliminating the extra work produced by empty categories interacting with other categories, including empty ones.

### 6.3 ALE Lexical Rules

Lexical rules provide a mechanism for expressing redundancies in the lexicon, such as the kinds of inflectional morphology used for word classes, derivational morphology as found with suffixes and prefixes, as well as zero-derivations as found with detransitivization, nominalization of some varieties and so on. The format ALE provides for stating lexical rules is similar to that found in both PATR-II and HPSG.

In order to implement them efficiently, lexical rules, as well as their effects on lexical entries, are compiled in much the same way as grammars. To enhance their power, lexical rules, like grammar rules, allow arbitrary procedural attachment with ALE definite constraints.

The lexical rule system of ALE is productive in that it allows lexical rules to apply sequentially to their own output or the output of other lexical rules. Thus, it is possible to derive the nominal *runner* from the verb *run*, and then derive the plural nominal *runners* from *runner*, and so on. At the same time, the lexical system is leashed to a fixed depth-bound, which may be specified by the user. This bound limits the number of rules that can be applied to any given category. The bound on application of rules is specified by a command such as the following, which should appear somewhere at the beginning of the input file:

```
:-lex_rule_depth(2).
```

Of course, bounds other than 2 can be used. The bound indicates how many applications of lexical rules can be made, and may be 0. If there is more than one such specification in an input file, the last one will be the one that is used. If no specification is given, the default is 2.

The format for lexical rules is as follows:

```
<lex_rule> ::= <lex_rule_name> lex_rule <lex_rewrite>
              morphs <morphs>.

<lex_rewrite> ::= <desc> **> <desc>
                | <desc> **> <desc> if <goal>

<morphs> ::= <morph>
            | <morph>, <morphs>

<morph> ::= (<string_pattern>) becomes (<string_pattern>)
           | (<string_pattern>) becomes (<string_pattern>)
           when <prolog_goal>

<string_pattern> ::= <atomic_string_pattern>
                   | <atomic_string_pattern>, <string_pattern>
```

```

<atomic_string_pattern> ::= <atom>
                        | <var>
                        | <list(<var_char>)>
<var_char> ::= <char>
             | <var>

```

An example of a lexical rule with almost all of the bells and whistles (we put off procedural attachment for now) is:

```

plural_n lex_rule
(n,
 num:sing)
**> (n,
     num:plu)
morphs
goose becomes geese,
[k,e,y] becomes [k,e,y,s],
(X,man) becomes (X,men),
(X,F) becomes (X,F,es) when fricative(F),
(X,ey) becomes (X,[i,e,s]),
X becomes (X,s).

fricative([s]).
fricative([c,h]).
fricative([s,h]).
fricative([x]).

```

We will use this lexical rule to explain the behavior of the lexical rule system. First note that the name of a lexical rule, in this case `plural_n`, must in general be a Prolog atom. Further note that the top-level parentheses around both the descriptions and the patterns are necessary. If the Prolog goal, in this case `fricative(F)`, had been a complex goal, then it would need to be parenthesized as well. The next thing to note about the lexical rule is that there are two descriptions — the first describes the input category to the rule, while the second describes the output category. These are arbitrary descriptions, and may contain disjunctions, macros, etc. We will come back to the clauses for `fricative/1` shortly. Note that the patterns in the morphological component are built out of variables, sequences and lists. Thus a simple rewriting can be specified either using atoms as with `goose` above, with a list, as in `[k,e,y]`, or with a sequence as in `(X,man)`, or with both, as in `(X,[i,e,s])`. The syntax of the morphological operations is such that in sequences, atoms may be used as a shorthand for lists of characters. But lists must consist of variables or single characters only. Thus we could not have used `(X,[F])` in the `fricative` case, as `F` might itself be a complex list such as `[s,h]` or `[x]`. But in general, variables ranging over single characters can show up in lists.

The basic operation of a lexical rule is quite simple. First, every lexical entry, including a word and a category, that is produced during compilation is checked to see if its category satisfies the input description of a lexical rule. If it does, then a new category is generated to satisfy the output description of the lexical

rule, if possible. Note that there might be multiple solutions, and all solutions are considered and generated. Thus multiple solutions to the input or output descriptions lead to multiple lexical entries.

After the input and output categories have been computed, the word of the input lexical entry is fed through the morphological analyzer to produce the corresponding output word. Unlike the categorial component of lexical rules, only one output word will be constructed, based on the first input/output pattern that is matched.<sup>1</sup> The input word is matched against the patterns on the left hand side of the morphological productions. When one is found that the input word matches, any condition imposed by a *when* clause on the production is evaluated. This ordering is imposed so that the Prolog goal will have all of the variables for the input string instantiated. At this point, Prolog is invoked to evaluate the *when* clause. In the most restricted case, as illustrated in the above lexical rule, Prolog is only used to provide abbreviations for classes. Thus the definition for *fricative/1* consists only of unit clauses. For those unfamiliar with Prolog, this strategy can be used in general for simple morphological abbreviations. Evaluating these goals requires the *F* in the input pattern to match one of the strings given. The shorthand of using atoms for the lists of their characters only operates within the morphological sequences. In particular, the Prolog goals do not automatically inherit the ability of the lexical system to use atoms as an abbreviation for lists, so they have to be given in lists. Substituting *fricative(sh)* for *fricative([s,h])* would not yield the intended interpretation. Variables in sequences in morphological productions will always be instantiated to lists, even if they are single characters. For instance, consider the lexical rule above with every atom written out as an explicit list:

```
[g,o,o,s,e] becomes [g,e,e,s,e],
[k,e,y] becomes [k,e,y,s],
(X,[m,a,n]) becomes (X,[m,e,n]),
(X,F) becomes (X,F,[e,s]) when fricative(F),
(X,[e,y]) becomes (X,[i,e,s]),
X becomes (X,[s]).
```

In this example, the *s* in the final production is given as a list, even though it is only a single character.

The morphological productions are considered one at a time until one is matched. This ordering allows a form of suppletion, whereby special forms such as those for the irregular plural of *goose* and *key* to be listed explicitly. It also allows subregularities, such as the rule for fricatives above, to override more general rules. Thus the input word *beach* becomes *beaches* because *beach* matches  $(X,F)$  with  $X = [b,e,a]$  and  $F = [c,h]$ , the goal *fricative([c,h])* succeeds and the word *beaches* matches the output pattern  $(X,F,[e,s])$ , instantiated after the input is matched to  $([b,e,a],[c,h],[e,s])$ . Similarly, words that end in  $[e,y]$  have this sequence replaced by  $[i,e,s]$  in the plural, which is why an irregular form is required for *keys*, which would otherwise match this pattern. Finally, the

<sup>1</sup>Thus ALE's lexical rule system is not capable of handling cases of partial suppletion, where both a regular and irregular morphological form are both allowed. To allow two output forms, one must be coded by hand with its own lexical entry or a separate lexical rule.

last rule matches any input, because it is just a variable, and the output it produces simply suffixes an [s] to the input.

For lexical rules with no morphological effect, the production:

```
X becomes X
```

suffices. To allow lexical operations to be stated wholly within Prolog, a rule may be used such as the following:

```
X becomes Y when morph_plural(X,Y)
```

In this case, when `morph_plural(X,Y)` is called, `X` will be instantiated to the list of the characters in the input, and as a result of the call, `Y` should be instantiated to a ground list of output characters.

We finally turn to the case of lexical rules with procedural attachments, as in the following (simplified) example from HPSG:

```
extraction lex_rule
  local:(cat:(head:H,
            subcat:Xs),
        cont:C),
  nonlocal:(to_bind:Bs,
            inherited:Is)
**> local:(cat:(head:H,
                subcat:Xs2),
      cont:C),
      nonlocal:(to_bind:Bs,
                inherited:[G|Is])
if
  select(G,Xs,Xs2)
morphs
  X becomes X.

select(X,(hd:X),Xs) if true.
select(X,[Y|Xs],[Y|Ys]) if
  select(X,Xs,Ys).
```

This example illustrates an important point other than the use of conditions on categories in lexical rules. The point is that even though only the `LOCAL CAT SUBCAT` and `NONLOCAL INHERITED` paths are affected, information that stays the same must also be mentioned. For instance, if the `cont:C` specification had been left out of either the input or output category description, then the output category of the rule would have a completely unconstrained content value. This differs from the default-based nature of the usual presentation of lexical rules, which assumes all information that hasn't been explicitly specified is shared between the input and the output. As another example, we must also specify that the `HEAD` and `TO_BIND` features are to be copied from the input to the output; otherwise there would be no specification of them in the output of the rule. This fact follows from the description of the application of lexical rules: they match a given category against the input description and produce the most general category(s) matching the output description.

Turning to the use of conditions in the above rule, the `select/3` predicate is defined so that it selects its first argument as a list member of its second argument, returning the third argument as the second argument with the selected element deleted. In effect, the above lexical rule produces a new lexical entry which is like the original entry, except for the fact that one of the elements on the subcat list of the input is removed from the subcat list and added to the inherited value in the output. Nothing else changes.

Procedurally, the definite clause is invoked after the lexical rule has matched the input description against the input category. Like the morphological system, this control decision was made to ensure that the relevant variables are instantiated at the time the condition is resolved. The condition here can be an arbitrary goal, but if it is complex, there should be parentheses around the whole thing. Cuts should not be used in conditions on lexical rules (see the comments on cuts in grammar rules below, which also apply to cuts in lexical rules).

Currently, ALE does not check for redundancies or for entries that subsume each other, either in the base lexicon or after closure under lexical rules. ALE also does not apply lexical rules to empty categories.

## 6.4 Grammar Rules

Grammar rules in ALE are of the phrase structure variety, with annotations for both goals that need to be solved and for attribute-value descriptions of categories. The BNF syntax for rules is as follows:

```

<rule> ::= <rule_name> rule <desc> ==> <rule_body>.

<rule_body> ::= <rule_clause>
              | <rule_clause>, <rule_body>

<rule_clause> ::= cat> <desc>
                | cats> <desc>
                | sem_head> <desc>
                | goal> <goal>
                | sem_goal> <goal>

```

The `<rule_name>` must be a Prolog atom. The description in the rule is taken to be the mother category in the rule, while the rule body specifies the daughters in the rule along with any side conditions on the rule, expressed as ALE goals. A further restriction on rules, which is not expressed in the BNF syntax above, is that there must be at least one category-seeking rule clause in each rule body.<sup>2</sup> Thus empty productions are not allowed and will be flagged as errors at compile time.

A simple example of such a rule, without any goals, is as follows:

```

s_np_vp rule
(syn:s,
 sem:(VPsem,
      agent:NPSem))

```

---

<sup>2</sup>By doubling the size of the BNF for rules, this requirement could be expressed.

```

===>
cat>
  (syn:np,
   agr:Agr,
   sem:NPsem),
cat>
  (syn:vp,
   agr:Agr,
   sem:VPsem).

```

There are a few things to notice about this rule. The first is that the parentheses around the category and mother descriptions are necessary. Looking at what the rule means, it allows the combination of an `np` category with a `vp` type category if they have compatible (unifiable) values for `agr`. It then takes the semantics of the result to be the semantics of the verb phrase, with the additional information that the noun phrase semantics fills the agent role.

Unlike the PATR-II rules, but similar to DCG rules, “unifications” are specified by variable co-occurrence rather than by path equations, while path values are specified using the colon rather than by a second kind of path equation. The rule above is similar to a PATR-II rule which would look roughly as follows:

```

x0 ---> x1, x2 if
  (x0 syn) == s,
  (x1 syn) == np,
  (x2 syn) == vp,
  (x0 sem) == (x2 sem),
  (x0 sem agent) == (x1 sem),
  (x1 agr) == (x2 agr)

```

Unlike lexical entries, rules are not expanded to feature structures at compile-time. Rather, they are compiled down into structure-copying operations involving table look-ups for feature and type symbols, unification operations for variables, sequencing for conjunction, and choice point creation for disjunction.

The descriptions for `cat>` and `cats>` daughters are always evaluated in the order they are specified, from left to right. This is significant when considering goals that might be interleaved with searches in the chart for consistent daughter categories. The order in which the code for the mother’s and semantic head’s descriptions is executed depends on the control strategy used during parsing or generation. These are described in Sections 6.4.3 and 6.4.4, respectively. In theory, the same grammar can be used for both parsing and generation. In practice, a single grammar is rarely efficient in both directions, and can even exhibit termination problems in one, just as a Prolog program may have these problems with queries that have different argument instantiations. So while it is not necessary to fully understand the parsing or generation algorithms used by ALE to exploit its power for developing grammars, practical implementations will order their procedural attachments and distribute their description-level information with these algorithms in mind.

Within a single description, in the case of feature and type symbols, a double-hashing is performed on the type of the structure being added to, as well as either



the feature or the type being added. Additional operations arise from type coercions that adding features or types require. Thus there is nothing like disjunctive normal-form conversion of rules at compile time, as there is for lexical entries. In particular, if there is a local disjunction in a rule, it will be evaluated locally at run time. For instance, consider the following rule, which is the local part of HPSG's Schema 1:

```

schema1 rule
  (cat: (head:Head,
        subcat: []),
   cont:Cont)
==>
cat>
  (Subj,
   cat:head:( subst
              ; spec:HeadLoc,
              )),
cat>
  (HeadLoc,
   cat: (head:Head,
        subcat: [Subj]),
   cont:Cont).

```

Note that there is a disjunction in the `cat:head` value of the first daughter category (the subject in this case). This disjunction represents the fact that the head value is either a substantive category (one of type `subst`), or it has a specifier value which is shared with the entire second daughter. But the choice between the disjuncts in the first daughter of this rule is made locally, when the daughter category is fully known, and thus does not create needless rule instantiations.

ALE's general treatment of disjunction in descriptions, which is an extension of Kasper and Round's (1986) attribute-value logic to phrase structure rules, is a vast improvement over a system such as PATR-II, which would not allow disjunction in a rule, thus forcing the user to write out complete variants of rules that only differ locally. Disjunctions in rules do create local choice points, though, even if the first goal in the disjunction is the one that is solvable.<sup>3</sup> This is because, in general, both parts of a disjunction might be consistent with a given category, and lead to two solutions. Or one disjunct might be discarded as inconsistent only when its variables are further instantiated elsewhere in the rule.

### 6.4.1 Procedural Attachments

A more complicated rule, drawn from the categorial grammar in the appendix is as follows:

```

backward_application rule
  (synsem:Z,
   qstore:Qs)

```

---

<sup>3</sup>In a future release, cuts will be allowed within descriptions, to allow the user to eliminate disjunctive choice points when possible.

```

===>
cat>
  (synsem:Y,
   qstore:Qs1),
cat>
  (synsem:(backward,
           arg:Y,
           res:Z),
   qstore:Qs2),
goal>
  append(Qs1,Qs2,Qs) .

```

Note that the goal in this rule is specified after the two category descriptions. Consequently, it will be evaluated after categories matching the descriptions have already been found, thus ensuring in this case that the variables `Qs1` and `Qs2` are instantiated. The `append(Qs1,Qs2,Qs)` goal is evaluated by ALE's definite clause resolution mechanism. `goal>` attachments are always evaluated in the order they are specified relative to the enforcement of `cat>` and `cats>` daughters, from left to right. All possible solutions to the goal are found with the resulting instantiations carrying over to the rule. These solutions are found using the depth-first search built into ALE's definite constraint resolver. In general, goals may be interleaved with the category specifications, giving the user control over when the goals are fired. Also note that goals may be arbitrary *cut-free* ALE definite clause goals, and thus may include disjunctions, conjunctions, and negations. Cuts may occur, however, within the code for any literal clause specified in a procedural attachment. The attachments themselves must be cut-free to avoid the cut taking precedence over the entire rule after compilation, thus preventing the rule to apply to other edges in the chart or for later rules to apply. Instead, if cuts are desired in rules, they must be encapsulated in an auxiliary predicate, which will restrict the scope of the cut. For instance, in the context of a phrase structure rule, rather than a goal of the form:

```

goal>
  (a, !, b)

```

it is necessary to encode this as follows:

```

goal>
  c

```

where the predicate `c` is defined by:

```

c if
  (a, !, b) .

```

This prevents backtracking through the cut in the goal, but does not block the further application of the rule. A similar strategy should be employed for cuts in lexical rules.

As a programming strategy, rules should be formulated like Prolog clauses, so that they fail as early as possible. Thus the features that discriminate whether

a rule is applicable should occur first in category descriptions. The only work incurred by testing whether a rule is applicable is up to the point where it fails.

Just as with PATR-II, ALE is RE-complete (equivalently, Turing-equivalent), meaning that any computable language can be encoded. Thus it is possible to represent undecidable grammars, even without resorting to the kind of procedural attachment possible with arbitrary definite clause goals. With its mix of depth-first and breadth-first evaluation strategies, ALE is not strictly complete with respect to its intended semantics if an infinite number of edges can be generated with the grammar. This situation is similar to that in Prolog, where a declaratively impeccable program might hang operationally.

### 6.4.2 The `cats>` Operator

The `cats>` operator is used to describe a list of daughters, whose length cannot be determined until run-time. Daughters are not parsed or generated as quickly as part of a `cats>` specification. Note also the interpretation of `cats>` requires that its argument is subsumed by the type `list`, which must be defined, along with `ne_list`, `e_list`, etc., and the features `HD`, and `TL`, which we defined above. This check is not made using unification, so that an underspecified list argument will not work either. If the argument of `cats>` is not subsumed by `list`, then the rule in which that argument occurs will never match any string, and a run-time error message will be given. This operator is useful for so-called “flat” rules, such as HPSG’s Schema 2, part of which is given (in simplified form) below:

```

schema2 rule
  (cat:(head:Head,
        subcat:[Subj]))
===>
cat>
  (cat:(head:Head,
        subcat:[Subj|Comps])),
cats> Comps.

```

Since various lexical items have `SUBCAT` lists of various lengths, e.g., zero for proper nouns, one for intransitive verbs, two for transitive verbs, `cats>` is required in order to match the actual list of complements at run-time.

It is common to require a goal to produce an output for the argument of `cats>`. If this is done, the goal must be placed before the `cats>`. Our use of `cats>` is problematic in that we require the argument of `cats>` to evaluate to a list of fixed length. Thus parsing with the following head-final version of HPSG’s Schema 2 would not work:

```

schema2 rule
  (cat:(head:Head,
        subcat:[SubjSyn]))
===>
cats> Comps,
cat>
  (cat:(head:Head,
        subcat:[Subj|Comps])).

```

One way to work around this is to place some finite upper bound on the size of the `Comps` list by means of a constraint.

```

schema2 rule
(cat:(head:Head,
      subcat:[SubjSyn]))
goal> three_or_less(Comps),
cats> Comps,
cat>
(cat:(head:Head,
      subcat:[Subj|Comps])).

three_or_less([]) if true.
three_or_less([_]) if true.
three_or_less([_,_]) if true.
three_or_less([_,_,_]) if true.

```

The problem with this strategy from an efficiency standpoint is that arbitrary sequences of three categories will be checked at every point in the grammar; in the English case, the search is directed by the types instantiated in `Comps` as well as that list's length. From a theoretical standpoint, it is impossible to get truly unbounded length arguments in this way.

### 6.4.3 Parsing

The ALE system employs a bottom-up active chart parser that has been tailored to the implementation of attribute-value grammars in Prolog. The single most important fact to keep in mind is that rules are evaluated from left to right, with the mother description coming last. Most of the implementational considerations follow from this rule evaluation principle and its specific implementation in Prolog. In parsing, `sem_head>` and `sem_goal>` specifications are treated exactly as `cat>` and `goal>` specifications, respectively.

The chart is filled in using a combination of depth- and breadth-first control. In particular, the edges are filled in from right to left, even though the rules are evaluated from left to right. Furthermore, the parser proceeds breadth-first in the sense that it incrementally moves through the string from right to left, one word at a time, recording all of the inactive edges that can be created beginning from the current left-hand position in the string. For instance, in the string *The kid ran yesterday*, the order of processing is as follows. First, lexical entries for *yesterday* are looked up, and entered into the chart as inactive edges. For each inactive edge that is added to the chart, the rules are also fired according to the bottom-up rule of chart parsing. But no active edges are recorded. Active edges are purely dynamic structures, existing only locally to exploit Prolog's copying and backtracking schemes. The benefit of parsing from right to left is that when an active edge is proposed by a bottom-up rule, every inactive edge it might need to be completed has already been found. This is actually true as long as the direction of traversal through the string is the opposite of the direction of matching daughter categories in a rule; thus the real reason for the right-to-left parsing strategy is to allow the active edges to be represented dynamically, while still evaluating the

rules from left to right. While the overall strategy is bottom-up, and breadth-first insofar as it steps incrementally through the string, filling in every possible inactive edge as it goes, the rest of the processing is done depth-first to keep as many data structures dynamic as possible, to avoid copying other than that done by Prolog's backtracking mechanism. In particular, lexical entries, bottom-up rules, and active edges are all evaluated depth-first, which is perfectly sound, because they all start at the same left point (that before the current word in the right to left pass through the string), and thus do not interact with one another.

ALE computes the closure of its grammar rules under application of the first daughter's description to empty categories at compile-time. This is known as *Empty-First-Daughter closure* or EFD closure. This closure operation has three advantages. First, given ALE's combination of depth-first and breadth-first processing, it is necessary in order to ensure completeness of parsing with empty categories, because any permutation of empty categories can, in principle, be combined to form a new empty category. Second, it works around a problem that many non-ISO-compatible Prologs, including SICStus Prolog, have with asserted predicates that results in empty category leftmost daughters not being able to combine with their own outputs. Third, it allows the run-time parser to establish a precondition that rules only need to be closed with non-empty leftmost daughters at run-time. As a result, when a new mother category is created and closed under rules as the leftmost daughter, it cannot combine with other edges created with the same left node. This allows ALE, at each step in its right-to-left pass through the input string, to copy all of the edges in the internal database back onto the heap before they can be used again, and thus reduces edge copying to a constant two times per edge for non-empty categories. Keeping a copy of the chart on the heap also allows for more sophisticated indexing strategies that would otherwise be overwhelmed by the cost of copying edges with large-sized categories in Prolog before the match. The EFD closure algorithm itself is described in Penn (1999).

EFD closure potentially creates new rules, a prefix of whose daughters have matched empty categories, and new empty categories, formed when every daughter of a rule has matched an empty category. The closure is computed breadth-first.

EFD closure may not terminate. As a result, compilation of some grammars may go into infinite loops. This only occurs, however, with grammars for which every parse would go into an infinite loop at run-time if EFD closure were not applied — specifically, when empty categories alone can produce infinitely many empty categories using the rules of the grammar. Because early versions of ALE did not compute a complete closure of grammar rules over empty categories (even at run-time), some grammars that terminated at run-time under these early versions will not terminate at compile-time under the current version.

Rules can incorporate definite clause goals before, between or after category specifications. These goals are evaluated when they are found. For instance, if a goal occurs between two categories on the right hand side of a rule, the goal is evaluated after the first category is found, but before the second one is. The goals are evaluated by ALE's definite clause resolution mechanism, which operates in a depth-first manner. Thus care should be taken to make sure the required variables in a goal are instantiated before the goal is called. The resolution of all goals

should terminate with a finite (possibly empty) number of solutions, taking into account the variables that are instantiated when they are called.

The parser will terminate after finding all of the inactive edges derivable from the lexical entries and the grammar rules. Of course, if the grammar is such that an infinite number of derivations can be produced, ALE will not terminate. Such an infinite number of derivations can creep in either through recursive unary rules or through the evaluation of goals.

ALE now has an optional mechanism for checking edge subsumption (Section 2.9). This can be used to prevent the propagation of spurious ambiguities through the parse. A category  $C$  spanning a given subsequence is said to be *spurious* if there is another category  $C'$  spanning the same subsequence such that  $C$  is subsumed by  $C'$ . Only the most general category needs to be recorded to ensure soundness. It can also be used to detect two derivations of the same category. Our experience, however, has been that most unification-based grammars do not have any spurious ambiguity. They normally incorporate some notion of thematic or functional structure representing the meaning of a sentence; and in these cases, most structural ambiguities result in semantic ambiguities. For such grammars, subsumption checking is probably not worth the effort, and should be left disabled.

#### 6.4.4 Generation

ALE also contains a generator, based on the *Semantic Head-Driven Generation* algorithm of van Noord (1989), as extended by Shieber et al. (1990), and adapted to the typed feature logic of Carpenter (1992) by Popescu (1996). Its implementation in ALE is described in Penn and Popescu (1997).

Given a description of a feature structure, ALE's generator will non-deterministically generate all the word strings that correspond to its most general satisfier(s). In other words, the generated word strings, when parsed in ALE using the same grammar, will result in feature structures that *unify* with a most general satisfier of the initial description (rather than necessarily be identical). That part of the feature structure which represents semantic information drives the generation process.

##### The `semantics/1` Directive

ALE identifies this part using a user-defined directive, `semantics/1`. This directive distinguishes a binary user-defined definite clause predicate as the predicate to use to find semantic information. The first argument is always the feature structure whose semantics are being identified; and the second argument is always the semantic information. The example below, taken again from the sample generation grammar, simply says that the semantics of a feature structure is the value of its `sem` feature:

```
semantics sem1.
sem1(sem:S,S) if true.
```

In general, the second argument does not need to be a sub-structure of the first — it could have a special type that is used only for the purpose of collect-

ing semantic information, possibly spread over several unrelated sub-structures. The body can be arbitrarily complex; and there can be multiple clauses for the definition of this predicate. The predicate must, however, have the property that it will terminate when only its first argument is instantiated, and when only its second argument is instantiated. ALE will use this predicate in both “directions” — to find semantics information, and in reverse to build templates to find structures that have matching semantic information. There can be only one predicate distinguished by `semantics/1`. If there are multiple directives, ALE will only use the first.

Semantic-head-driven generation makes use of the notion of a *semantic head* of a rule, a daughter whose semantics is shared with the mother. In semantic-head-driven generation, there are two kinds of rules: *chain rules*, which have a semantic head, and *non-chain rules*, which lack such a head. These two subsets are processed differently during the generation process.

Given a feature structure, called the *root goal*, to generate a string for, the generator builds a new feature structure that shares its semantic information (using the user-defined semantics predicate with the second argument instantiated) and finds a *pivot* that unifies with it. The pivot is the lowest node in a derivation tree that has the same semantics as the root. The pivot may be either a *lexical entry* or *empty category* (the base cases), or the *mother category of a non-chain rule*. Once a pivot is identified, one can recursively generate top-down from the pivot using non-chain rules. Since the pivot must be the *lowest*, there can be no lower semantic heads, and thus no lower chain-rule applications. Just as in parsing, the daughters of non-chain rules are processed from left to right.

Once top-down generation from the pivot is complete, the pivot is linked to the root bottom-up by chain rules. At each step, the current chain node (beginning with the pivot) is unified with the semantic head of a chain-rule, its non-head sisters are generated recursively, and the mother becomes the new chain node. The non-head daughters of chain rules are also processed from left to right. The base case is where the current chain node unifies with the root.

To avoid the non-termination problem that arises from infinitely long *semantic head* → *mother* sequences in a grammar, ALE requires the user to specify a bound on the length of chain rule sequences at compile-time. This can be specified with the declaration:

```
:- chain_length(4).
```

Other values than 4 can be used, including 0. The default value is 4. ALE compiles chains of semantic head and mother descriptions of this length to perform the pivot check more efficiently at run-time.

### The `sem_goal`> Operator

For the most part, the generator treats procedural attachments as the parser does — it evaluates them with respect to other daughter specifications in the order given. The one exception to this is `sem_goal`> attachments. These goals are distinguished as attached to the semantic head, and are therefore evaluated either immediately before or immediately after the `sem_head`> description. As a result, `sem_goal`> specification can only occur immediately before or immediately after a

`sem_head`> specification; and thus only in chain rules. `sem_goal`> attachments are not evaluated during the pivot check described above — only the `sem_head`> and mother descriptions. During parsing, `sem_goal`> specifications are treated exactly the same as `goal`> specifications, i.e., evaluated in order.

To summarize, the order of execution for a non-chain rule specification during generation is:

- the mother description, then
- the `cat`>, `cats`>, and `goal`> descriptions, in order, from left to right.

There are no `sem_head`> or `sem_goal`> specifications in a non-chain rule. The order for a chain rule specification during generation is:

- the pre-head `sem_goal`> specification, if it exists,
- the `sem_head`> description,
- the post-head `sem_goal`> specification, if it exists,
- the mother description, then<sup>4</sup>
- the `cat`>, `cats`> and `goal`> specifications, in order, from left to right.

Again, practical grammar implementations will arrange information in rules in such a way as to ensure termination and to force failure as early as possible. For non-chain rules, this means making the mother and early daughters or goals as informative as possible at the description level (that is, up to where type inferencing can take over). For chain rules, the semantic head and its attachments should be maximally informative.

---

<sup>4</sup>The standard head-driven-generation algorithm enforces the mother description after the non-semantic-head-related daughters. We deviate from this order in order to enforce the pivot check, which requires instantiating the mother, more efficiently.



## T6.1 Test sequences

### T6.1.1 Test files

Test items are encoded as t/5 facts:

```
t(Nr, ' 'Test Item' ', Desc, ExpSols, 'Comment' ).
```

- Nr: test item ID number
- Test Item: test string, must be enclosed in double-quotes
- Desc: optional start category description, leave uninstantiated to get all possible parses
- ExpSols: expected number of solutions
- Comment: optional comment, enclosed in single-quotes

### T6.1.2 Test queries

#### Basic test queries

- with pop-up structures
 

```
test(Nr) .
test([From, To]) .
test(all) .
```
- without structures
 

```
testt(Nr) .
testt([From, To]) .
testt(all) .
```

#### Testing with descriptions

- with pop-up structures
 

```
test(Nr, Desc) .
test([From, To], Desc) .
test(all, Desc) .
```
- without structures
 

```
testt(Nr, Desc) .
testt([From, To], Desc) .
testt(all, Desc) .
```

The value of Desc overrides any description given in t/5. If Desc is a variable (or bot or sign) all parses are returned. The expected number of solutions given in t/5 is ignored.

## Chapter 7

# TRALE Lexical Rule Compiler

*Vanessa Metcalf and Detmar Meurers, Ohio State University*<sup>1</sup>

### T7.1 Introduction

In the framework of Head-Driven Phrase Structure Grammar (HPSG, Pollard and Sag 1994) and other current linguistic architectures, linguistic generalizations are often conceptualized at the lexical level. Among the mechanisms for expressing such generalizations, so-called *lexical rules* are used for expressing so-called horizontal generalizations (cf. Meurers 2001, and references cited therein).

Lexical rules have been conceptualized in two distinct ways, either on a meta-level as relating lexical entries (i.e., descriptions) or on the same level as the rest of the grammar, relating lexical objects (i.e., the elements described by the lexical entries). Computationally the meta-level perspective is captured by using the lexical rules to enlarge the set of lexical entries defined in the grammar, whereas under the description-level perspective lexical rules are essentially on a par with unary phrase structure rules and are taken into account at run time. The TRALE system offers both mechanisms for implementing lexical rules. The traditional ALE mechanism for handling lexical rules described elsewhere in this manual corresponds to the meta-level interpretation of lexical entries. It uses the lexical rules at compile-time to generate new lexical entries, using a depth-bound to ensure that this step terminates. The description-level approach is implemented in the lexical rule compiler described in this chapter. It encodes the treatment of lexical rules proposed in Meurers and Minnen (1997). This chapter focuses on how to use the lexical rule compiler. A discussion of this approach to lexical rules can be found in the Reference Manual of the Lexical Rule compiler and the original paper.

### T7.2 Using the lexical rule compiler

The lexical rule compiler is automatically loaded with the main TRALE system, but it's important to understand that this lexical rule compiler is essentially a pre-processor to the normal grammar compilation. Its purpose is to take the lexical

---

<sup>1</sup>©2003, Vanessa Metcalf and Detmar Meurers

rule definitions and the lexical entries and to produce a lexicon definition that includes the lexical rule code that needs to be executed at run time. If the grammar contains lexical rules using the syntax described in the following section, one can call the lexical rule compiler with `compile_lrs(<file(s)>)`, where `<files>` is either a single file or a list of files containing the part of the theory defining the base lexical entries and the lexical rules. The output of the lexical rule compiler is written to the file `lr_compiler_output.pl`. After compilation the user can then access visual representations of both the global finite-state automaton and the word class automata.

The command for viewing the global automaton is `lr_show_global`. The automata for the different classes of lexical entries are shown using the command `lr_show_automata`. Both commands are useful for checking that the expected sequences of lexical rule applications are actually possible for the grammar that was compiled. An example is included at the end of section T7.2.1.

The system can visualize the graphs using either the *graphviz* (<http://www.research.att.com/sw/tools/graphviz/>) or the *vcg* (<http://rw4.cs.uni-sb.de/users/sander/html/gsvcg1.html>) tool, which are freely available and must be installed on your system for the visualization to work. By default, the *graphviz* visualization software is used. One can select *vcg* by calling `set_lr_display_option(vcg)` or by including the statement `lr_display_option(vcg)` in one of the grammar files loaded by the lexical rule compiler.

In order to parse using the output of the lexical rule compiler, one must compile the grammar, without the base lexicon and lexical rules, but including the file generated by the lexical rule compiler. For example, if the grammar without the base lexicon and lexical rules is in the file `theory.pl` and the lexical rule compiler output is in the file `lr_compiler_output.pl` one would call

```
?- compiler_gram([theory,lr_compiler_output]).
```

### T7.2.1 Input syntax

The format of lexical rule specifications for the lexical rule compiler is shown in figure 7.1. Note that this syntax is different from the lexical rule syntax of ALE, which also is provided by the TRALE system. As described in the ALE manual, lexical rules specified using the ALE lexical rule syntax result in expanding out the lexicon at compile time.

```
<lex_rule_name> ===
  <input description>
lex_rule
  <output description>.
```

Figure 7.1: Lexical rule input syntax

A lexical rule consists of a lexical rule name, followed by the infix operator `===`, followed by an input feature description, followed by the infix operator `lex_rule`, followed by an output feature description and ending with a period.

Input and output feature descriptions are ordinary descriptions as defined in the TRALE manual. The lexical compiler currently handles all kinds of descriptions except for path inequalities. Path equalities can be specified within the input or output descriptions, and also between the input and output descriptions.

We illustrate the syntax with the small example grammar from Meurers and Minnen (1997), which is also included with the TRALE system in the subdirectory `lr_compiler/examples`. The signature of this example is shown in figure 7.2; to illustrate this TRALE signature syntax, figure 7.3 shows the type hierarchy in the common graphical notation.

```

type_hierarchy
bot
  t w:bool x:bool y:bool
    t1
    t2 z:list
  word a:val b:bool c:t
  list
  bool
  val
  list
    e_list
    ne_list hd:val tl:list
  bool
    plus
    minus
  val
    a
    b
.

```

Figure 7.2: An example signature




Figure 7.3: A graphical representation of the example type hierarchy

Based on this signature, figure 7.4 shows a set of four lexical rules exemplifying the lexical rule syntax used as input to the lexical rule compiler.

To complete the example grammar, we include three examples for base lexical entries in figure 7.5. These lexical entries can be found in the file `lexicon.pl`.

The user is encouraged to look at this grammar, run the compiler on it, and make sure that the resulting output is consistent with the user's understanding. Visualizing the lexical rule interaction generally is a good way to check whether the intended lexical rule applications do in fact result from by the lexical rules that were specified in the grammar. The visualization obtained by calling `lr_show_global/0` for the example grammar is shown in figure 7.6.

```
lr_one ===
  (b:minus,
   c:y:minus)
lex_rule
  (a:b,
   c:(x:plus,y:plus)).

lr_two ===
  (a:b,
   b:minus,
   c:w:minus)
lex_rule
  (c:w:plus).

lr_three ===
  (c:(t2,
     w:plus,
     x:plus,
     z:t1:One))
lex_rule
  (c:(y:plus,
     z:One)).

lr_four ===
  (b:minus,
   c:(t2,
     w:plus,
     x:plus,
     z:e_list))
lex_rule
  (b:plus,
   c:x:minus).
```

Figure 7.4: An example set of four lexical rule

```

foo ---> (a:b,
          b:minus,
          c:(t2,
             w:minus,
             x:minus,
             y:minus,
             z:(hd:a,
                tl:(hd:b,
                    tl:e_list))))).

bar ---> (a:b,
          b:minus,
          c:(t2,
             w:minus,
             x:minus,
             y:minus,
             z:(hd:a,
                tl:e_list))))).

tup ---> (a:b,
          b:minus,
          c:(t1,
             w:minus,
             x:minus,
             y:minus)).

```

Figure 7.5: An example set of base lexical entries

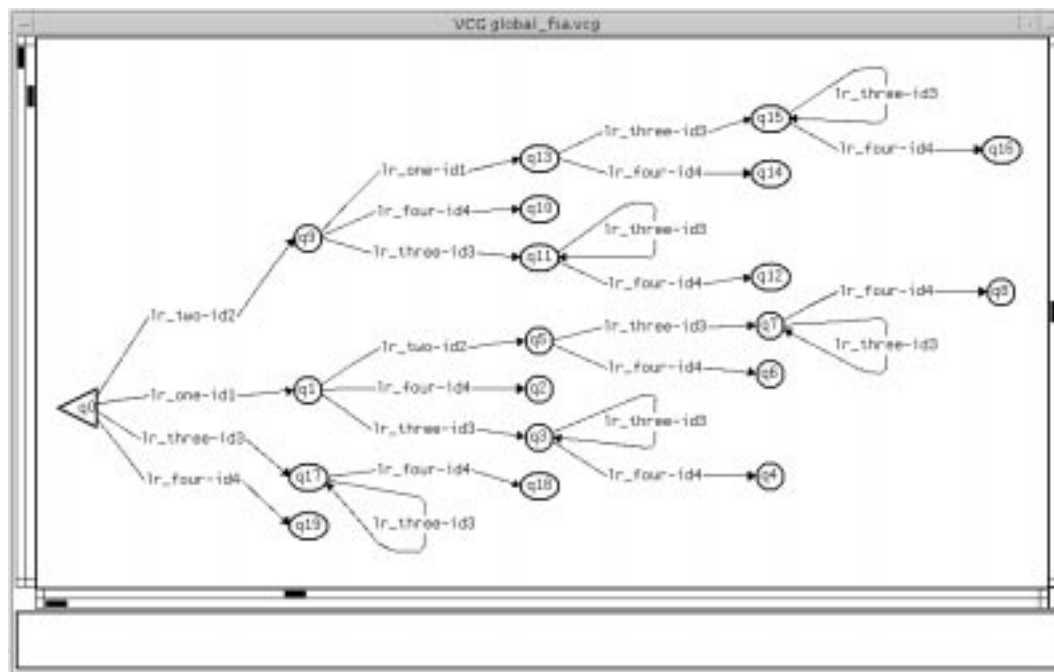


Figure 7.6: Global interaction visualization for the example grammar

The lexical rule interaction which is permitted by a particular lexical class can also be visualized. To view the automaton of an entry with the phonology `Phon` one calls `lr_show_automaton(Phon)`. To view all such automata, the predicate to call is `lr_show_automata/0`. In figure 7.7 we see the visualization obtained for the lexical entry “foo” of our example grammar by calling `show_automaton(foo)`.

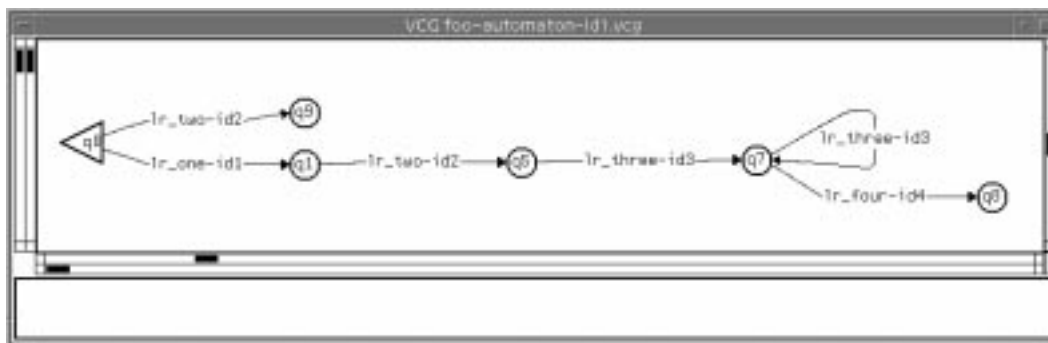


Figure 7.7: Interaction visualization for the entry “foo”

### T7.2.2 Interpretation

While the basic interpretation of lexical rules is straightforward, it turns out to be more difficult to answer the question how exactly the intuition should be spelled out that properties which are not changed by the output of a lexical rule are carried over unchanged, the so-called framing. A detailed discussion of the interpretation of lexical rules and the motivation for this particular interpretation can be found in Meurers (2001); we focus here on the essential ideas needed to sensibly use the lexical rule compiler.

A lexical rule can apply to a variety of lexical entities. While each of these lexical entities must be described by the input of the lexical rule in order for the rule to apply, other properties not specified by the lexical rule can and will vary between lexical entries. Feature structures corresponding to lexical entities undergoing the lexical rule therefore may differ in terms of type value and appropriate features. Frames carrying over properties not changed by the lexical rule need to take into account different feature geometries. Since framing utilizes structure sharing between input and output, we only need to be concerned with the different kinds of objects that can undergo a lexical rule with regard to the paths and subpaths mentioned in the output description. Specifically, when the objects undergoing lexical rule application differ with regard to type value along some path mentioned in the output description, we may need to take into account additional appropriate attributes in framing. Each such possibility will demand its own frame.

The lexical rule compiler provides a truthful procedural realization of the formal interpretation of lexical rules defined in Meurers (2001). Generally speaking, the input description of a lexical rule specifies enough information to capture the class of lexical entries intended by the user to serve as inputs. The output description, on the other hand, specifies what should change in the derivation. All other specifications of the input are supposed to stay the same in the output.

In the spirit of preserving as much information as possible from input to output, we generate frames on the basis of species (= most specific type) pairs; that is, we generate a frame (an IN-OUT pair) on the basis of a maximally specific input type, and a maximally specific output type, subtypes of those specified in, or inferred from, the lexical rule description. In this way we maintain tight control over which derivations we license, and we guarantee that all possible information is transferred, since the appropriate feature list we use is that of a maximally specific type. We create a pair of skeleton feature structures for the species pair, and it is to this pair of feature structures that we add path equalities. We determine the appropriate list of species pairs on the basis of the types of the input and output descriptions.

The first step in this process is determining the types of the input and output of the lexical rule. We then obtain the list of species of the input type, and the list of species of the output type. We refer to these as the input species list, and the output species list, and their members as input and output species. At this point it will be helpful to have an example to work with. Consider the type hierarchy in figure 7.8.




Figure 7.8: An example hierarchy for illustrating the interpretation

We can couch the relationship between the input and output types in terms of type unification, or in terms of species set relations. In terms of unification, there are four possibilities: the result of unification may be the input type, the output type, something else, or unification may fail. In the first case the input type is at least as or more specific, and the input species species will be a subset of the output species. In the second case the output is more specific and the output species will be a subset of the input species. In the third case the input and output types have a common subtype, and the intersection of input and output species is nonempty. In the fourth case the input and output types are incompatible, and the intersection of their species sets is empty.

If a (maximally specific) type value can be maintained in the output, it is. Otherwise, we map that input species to all output species. In terms of set membership, given a set of input species  $X$ , a set of output species  $Y$ , the set of species pairs  $P$  thus can be defined as:

$$P = \{ \langle x, x \rangle \mid x \in X \wedge x \in Y \} \cup \{ \langle x, y \rangle \mid x \in X, x \notin Y \wedge y \in Y \}$$

Given in figure 7.9 are examples of all four cases, using the example signature, showing input and output types, the result of unification, their species lists, and the species-pairs licensed by the algorithm described below. Calling these



	Case 1	Case 2	Case 3	Case 4
Input type	<i>c</i>	<i>a</i>	<i>b</i>	<i>c</i>
Output type	<i>a</i>	<i>c</i>	<i>c</i>	<i>d</i>
Unify to	<i>c</i>	<i>c</i>	<i>f</i>	fail
Input species	<i>f, g</i>	<i>e, f, g, h, i</i>	<i>e, f</i>	<i>f, g</i>
Output species	<i>e, f, g, h, i</i>	<i>f, g</i>	<i>f, g</i>	<i>h, i</i>
Species pairs	<i>f-f, g-g</i>	<i>e-f, e-g, f-f, g-g, h-f, h-g, i-f, i-g</i>	<i>e-f, e-g, f-f</i>	<i>f-h, f-i, g-h, g-i</i>

Figure 7.9: Examples for the four cases of mappings

separate “cases” is misleading, however, since the algorithm for deciding which mappings are licensed is the same in every case.

## Chapter 8

# Topological Parsing

### T8.1 Introduction

Topological Grammar Compiler (TGC) is a tool for compiling topological grammars into Prolog code. The grammar should comply by the specifications formulated in Penn and Haji-Abdolhosseini (2002). The output of the compiler is a parser in Prolog and is based on the grammatical rules and the lexicon provided in the input file. There are however several differences between the specifications in Penn and Haji-Abdolhosseini (2002) and what TGC assumes. These differences are described in pertinent sections throughout this guide.

### T8.2 Running TGC

TGC runs on any computer that has SICStus Prolog version 3.8.7 or higher installed. The compiler also requires the file `header.pl`. The parsed code needs the file `tree.pl` for pretty-printing. These files are included in the TGC package, which also contains two sample grammar files—`german.pl` (for German) and `sc.pl` (for Croatian/Serbian)—as well as two files containing some test sentences for those grammars.

#### T8.2.1 Compiling Grammars

To run TGC, make sure that the working directory of SICStus Prolog is set to where TGC files reside and then type

```
| ?- compile(tgc).
```

at the prompt. The command `topo_compile/2` is used to create a file that contains parser code based on an existing grammar file. For example, assuming that the grammar file `german.pl` exists in the working directory, you may generate the corresponding compiled code for that and keep it in a file called `german.c.pl` by typing the following command:

```
| ?- topo_compile('german.pl','german.c.pl').
```

Note that if you wish to use file extensions, you need to put the file name inside single quotes. This version of the compiler does not add `.pl` to the end of file names.

During compilation several messages are shown on the screen to inform the user of the states that the program enters or completes. During a normal compilation you should see messages similar to the following:

```
| ?- topo_compile('german.pl','german_c.pl').
Analyzing german.pl...
Compiling german.pl. Output to be written in german_c.pl...
Compilation complete.
```

```
yes
| ?-
```

In cases of errors or potential problems, warnings and/or error messages may appear at each state. These will be discussed in the following sections in detail.

### T8.2.2 Using the Parser

You can load the generated parser code into the machine using Prolog's `compile/1` command. Parsing is done with the command `rec/1`, which takes a list of input words in lowercase letters. For example, to parse the sentence 'Der mann lauft', type:

```
| ?- rec([der,mann,lauft]).
```

The output is a tree diagram for every parse that encompasses the whole input string. For the case of the example above, the output looks like this:

```
| ?- rec([der,mann,lauft]).
```

```
s
|_np
| |_det
| | |_der
| | |_nbar
| | |_n
| | |_mann
|_vp
| |_vbar
| | |_iv
| | |_lauft
```

```
yes
| ?-
```

Should the input string yield no such parse, the compiler returns the message: No full parse and fails. Similarly, in case the parser encounters an unrecognised word, it also notifies the user by an error message and fails. Some examples are presented below:

```
| ?- rec([der,mann,will,jochen,sehe]).
```

No full parse.

```
no
| ?- rec([der,mann,will,jochen,sehrn]).
Can't recognise the word sehrn.
```

```
no
| ?-
```

Another command, `time/2` is used to measure the time that it takes to fully process an input string. The number returned shows the processing time in milliseconds. Because this command only measures processing time and does not check for full parses and also because it does not draw any parse trees, you should first make sure that the input string parses using `rec/1`. An example of `time/2` in use is given below:

```
| ?- time([der,mann,sieht,jochen],T).
```

T = 10 ?

```
yes
| ?-
```

## T8.3 Grammars

### T8.3.1 Phenogrammatical Rules

Phenogrammatical rules (pheno-rules) are represented with the infix operator, `topo/2` in a grammar file. The left-hand side of the operator represents a region and the right-hand side a list of the topological field descriptors for the fields inside the region in order from left to right. The syntax of `topo/2` is as follows:

```
<topo-rule> ::= <region> topo <field-descriptor-list>
```

`<region>` is the name of the region and must be a Prolog atom. `<field-descriptor-list>` is a list of field descriptors formulated in Penn and Haji-Abdolhosseini (2002). Field names also have to be Prolog atoms.

The following example shows two rules for `clause` region in German. As in Prolog, these are interpreted disjunctively. Therefore, what the following two rules together mean is that all `clause` regions in German either have the first topology or the second.

```
clause topo [vf, cf, mf*, {vc}, {nf}].
clause topo [ cf, mf+, vc , {nf}].
```

Recall that ‘\*’ describes zero or more fields, ‘+’ at least one field and ‘{ }’ optional fields. A bare field name represents a single obligatory field. According to the first rule above, then, a clause region consists of a single *vf*, one *cf*, zero or more *mf*’s, an optional *vc* and an optional *nf*, and these fields are arranged in the same order as in the list. Note that as in Prolog, every predicate ends in a period.

Two other pheno-rules have been defined in the file `german.pl`:

```
npr topo [sprf, adjf*, nof, postf*].
ppr topo [pf, objf].
```

The first describes the topology of a noun-phrase region and the second that of a prepositional-phrase region.

Note that all pheno-rules must at least be binary branching. If this is not the case an exception is raised at compile time. In addition, the same field name cannot be used ambiguously for more than one region. Ambiguous field names also result in an exception being raised.

A condition implicit in all *topo* rules is that a region must at least span over one word. Thus, in the following rule:

```
r1 topo [{f1},{f2},{f3}].
```

although all the daughter fields of *r1* are optional, we will not have any zero-length *r1*’s in the chart during parsing. An *r1* is introduced in the chart only if it spans at least one word.

### T8.3.2 Linkage

The relationships among various regions are described with linkage constraints. These license links in a pheno-tree with field mothers and region daughters. The syntax for these constraints is as follows:

```
<linking-rule> ::= <region> <--> <field>
                  | <region> <--> (<fields>)
                  | <field> -->> <region>
                  | <field> -->> (<regions>)
```

```
<field> ::= <prolog-atom>
          | matrix
```

```
<region> ::= <prolog-atom>
```

```
<fields> ::= <field>
            | <fields> ; <fields>
```

```
<regions> ::= <region>
             | <regions> ; <regions>
```

These constraints are universally quantified on the left-hand side and existentially quantified on the right-hand side. Therefore, only one such constraint can be declared for each field or region. If more than one linkage constraint is given for a field/region, a warning message is displayed and only the last one is considered.

In a grammar file, we use the two infix operators `<<--/2` and `-->>/2` for linkage rules with regions or fields on the left-hand side respectively. Disjunction on the right-hand side is shown by semicolons (;). Therefore, to show that an `npr` may fit into `vf`, `mf`, `objf` or `nf`, we can use the following declaration. Note that the parentheses around the disjuncts are required. The special field, `matrix` is discussed in section T8.3.6.

```
npr <<-- (vf; mf; objf; nf).
```

If both `<<--/2` and `-->>/2` are needed in a single grammar file, care must be exercised that the two declarations not contradict each other. For example, the following two linkage declarations are inconsistent:

```
r1 <<-- (f1; f2; f3).
```

```
f4 -->> (r1; r2; r3).
```

The first constraint states that all `r1` regions can only be linked to either an `f1`, `f2` or `f3` and nothing else. The second constraint, on the other hand, states that all `f4` fields can only be linked to either `r1`, `r2` or `r3` and nothing else. According to the first constraint, however, `r1` cannot be linked to `f4`. To resolve this inconsistency, we should either add `f4` to the first constraint, or delete `r1` from the second as appropriate. Therefore, both of the following sets of constraints are logically sound:

```
r1 <<-- (f1; f2; f3; f4).
```

```
f4 -->> (r1; r2; r3).
```

or

```
r1 <<-- (f1; f2; f3).
```

```
f4 -->> (r2; r3).
```

### T8.3.3 Matching

```
matches/2
```

The first set of constraints that connect tectogrammar with phenogrammar is the set of `matches` constraints. These constraints say what tectogrammatical categories match in their yields with what phenogrammatical categories. The parser uses these constraints to find what field or region in the phenogrammar each lexical category matches. Therefore, these constraints are crucial to the proper operation of the system. The `matches` constraints have the following syntax:

```

<matches-rule> ::= <phi> matches <field-or-region>
                  | <phi> matches (<fields-or-regions>)
                  | (<phi> matches <field-or-region>) :- <prolog-goals>

<phi> ::= <prolog-atom>
          | <prolog-variable>

<field-or-region> ::= <field>
                    | <region>

<fields-or-regions> ::= <field-or-region>
                       | <fields-or-regions> ; <fields-or-regions>

```

<phi> stands for any tectogrammatical category. In this version, TGC assumes tectogrammatical categories to be either atomic types (e.g. *adv*, *s* or *aux*) or descriptions (e.g. (*n*,*gender*:*G*, *case*:*C*) or (*np*,*person*:*P*, *number*:*N*, *gender*:*G*, *case*:*C*)). The following are, therefore, well-formed *matches* declarations:

```

adv matches mf.
marker matches cf.

```

Assuming that we consider type *n* to have two appropriate features *NUMBER* and *GENDER*, then we may have the following *matches* declaration to mean that a first person, neuter noun matches a ‘noun field’ (*nof*).

```

(n,number:first,gender:neuter) matches nof.

```

As noted above, the right-hand side of these rules may contain a disjunction of fields or regions; therefore, to state that a proper noun (*pn*) matches an *mf*, *vf* or *objf*, the following declaration is used:

```

pn matches (mf; vf; objf).

```

Again, note that the parentheses are required around the disjuncts.

Sometimes, it is necessary to have a condition specified for the *matches* constraints. For example, in the Croatian and Serbian language, a third person singular present tense auxiliary verb (*je*) has its own field, *jef*, and any other auxiliary verb matches an *auxf*. We represent these constraints as follows:

```

((aux,person:third,number:sg,tense:pres) matches jef):- !.
aux matches auxf.

```

What we have done here is place a *matches* constraint inside the head of a Prolog conditional clause. The body of the rule is normal Prolog code. In the above example, the cut (!) prevents *je* from matching *auxf* as well as *jef*.

This strategy of placing a grammatical constraint inside the head of a Prolog conditional only works for *matches/2* and lexical entries (*--->/2* to be discussed in section T8.4 below). The conditional *matches* rules are placed before the unconditional ones in the compiled file. It should be kept in mind that disjunction is not allowed in conditional *matches* constraints.

Universal *matches* rules for non-lexical categories enforce the requirement that all tectogrammatical categories that can unify with <phi> should always match a topologically accessible <field-or-region>.

matched\_by/2

The syntax of `matched_by` constraints is as follows:

```
<matched-by-rule> ::= <field-or-region> matched_by <phi>
                    | <field-or-region> matched_by (<phis>)

<phis> ::= <phi>
          | <phis> ; <phis>
```

These constraints are responsible for predicting tectogrammatical categories (top-down) once the parser has found certain regions or fields. Because they block topological accessibility to higher nodes in the tecto-tree, `matched_by` constraints create inseparable structures that make parsing more efficient. It is, therefore, prudent to try to introduce as many different types of such constraints as possible in addition to clause-level categories. In German, for instance, one can have both of the following constraints at the same time in a single grammar:

```
clause matched_by (s; rp; cp).
ppr matched_by pp.
```

Once a `clause` region is found during pheno-parse, the first constraint results in a top-down tecto-parse in search for an S, RP or CP. The second constraint has the same effect only at the level of PP. The advantage of having the second constraint is that the internal structure of PP's in this case will remain inaccessible to higher nodes and therefore, no active edge higher up in the tree will try to steal the NP or NP's inside PP's, which results in a more efficient parse.

As shown in the previous example, disjunction of tecto-categories on the right-hand side of `matched_by` constraints is also acceptable. As usual, the parentheses around the disjuncts are required.

Because, `matched_by` constraints have universal quantification on their right-hand side, only one such constraint is allowed in each grammar for a field or region. In case more than one `matched_by` constraint is specified for a field/region, a warning message is issued and only the last one gets considered.

<==>/2

A shorthand for

```
phi matches f,
f matched_by phi
```

is `phi <==> f`. Note that no disjunction is allowed on either side of `<==>/2`. The syntax of `<==>/2` is as follows:

```
<bidirectional-matching> ::= <phi> <==> <field-or-region>
```



### T8.3.4 Covering

`covers/2`

Global `covers` declarations have the following syntax:

```
<covers-rule> ::= <phi> covers <field-or-region>
                | <phi> covers (<fields-or-regions>)
```

These constraints state that the yield of all `<phi>` must include all of `<field-or-region>`. As usual, disjunction is represented with semicolons (;) and parentheses are required around disjuncts.

`covered_by/2`

The syntax of `covered_by/2` (presented below) is almost identical to that of `matched_by/2`. It also has similar semantics. The only difference is that `covered_by` constraints require all `<field-or-region>` (declared on the left-hand side) be consumed by some `<phi>` (declared on the right-hand side), but `<phi>` may be larger in its yield than `<field-or-region>`.

```
<covered-by-rule> ::= <field-or-region> covered_by <phi>
                    | <field-or-region> covered_by (<phis>)
```

```
<field-or-region> ::= <field>
                    | <region>
```

```
<phis> ::= <phi> ; <phis>
```

It should be noted that `covered_by` constraints are the least efficient of constraints and grammar writers should try to avoid them by using other more efficient means such as `covers` constraints, structure-sharing, precedence or immediate precedence constraints etc.

Just as `matched_by`, `covered_by` constraints are also universally quantified on the left-hand side. Therefore, if more than one such constraint is declared for a given field or region, a warning message is issued and the last one gets considered.

`<-->/2`

The operator `<-->` in covering is the counterpart of `<==>` in matching. Therefore,

```
phi <--> f.
```

is equivalent to

```
phi covers f.
f covered_by phi.
```

Note that no disjunction is allowed in either side of `<-->/2`. The syntax of `<-->/2` is provided below:

```
<bidirectional-covering> ::= <phi> <--> <field-or-region>
```

### T8.3.5 Compaction

Compaction applies to tectogrammatical categories and states that the elements in the list that is its argument be contiguous strings. As a global constraint, compaction applies this condition to all categories that are consistent with any of the elements in the argument list of the constraint.

```
<compaction-rule> ::= compacts(<phi-list>)
```

In a grammar file, the predicate `compacts/1` is used to show compaction. The argument of `compacts/1` is a list of tectogrammatical categories. For example, in order to show that all NP's in a language are contiguous, we can use the following global constraint:

```
compacts([np]).
```

There can only be one global `compacts/1` declaration in the grammar. Should there be more than one, a warning message is issued and the last one gets considered. Also note that the argument of `compacts/1` must be a list. A non-list argument generates an error.

### T8.3.6 matrix

A special type of field that is recognised by TGC is `matrix`, which must always be maximal in a grammar (i.e. not linked to any higher fields or regions) and it must span over the whole input string. By using `matrix`, a grammar writer may delay prediction of sentences until such time that the whole input sentence has been pheno-parsed. An example of using `matrix` follows:

```
clause topo [vf, cf, mf*, {vc}, {nf}].
clause topo [ cf, mf+, vc , {nf}].

clause <<-- (nf;matrix).

matrix matched_by us.
```

According to the mini-phenogrammar above, the `clause` region is linked to `nf` and `matrix`. The fact that only `matrix` is in charge of predicting an unembedded sentence (US) delays tecto-parsing until all the input sentence has been pheno-parsed and this is because no `matrix` is introduced to the chart until the `clause` that links to it consumes the whole input string.

### T8.3.7 Tectogrammatical Rules

The tectogrammatical rules that TGC uses resemble phrase-structure rules in earlier generative grammars except that the rules do not assume order or contiguity in any way. The operator for tecto-rules is `*-->/2`. The syntax of tecto-rules is defined below:

```

<tecto-rule> ::= <phi> *--> [<phi-conj>]
                | <phi> *--> [<phi-conj>, {<conditions>}]

<phi> ::= <description>
          | <prolog-variable>

<phi-conj> ::= <phi>
               | <phi-conj>, <phi-conj>

<conditions> ::= <condition>
                 | (<conditions>, <conditions>)
                 | (<conditions>; <conditions>)

<condition> ::= <local-covers>
                | <local-matches>
                | <precedence>
                | <immediate-precedence>
                | <local-compaction>
                | <inequation>

```

The category on the left-hand side of a tecto-rule represents the mother node, and the daughters are provided (in no particular order) on the right-hand side. The daughters should be separated by commas. It should be noted that TGC assumes no empty categories.

As is suggested in the BNF above, one may provide some conditions for the applicability of a tecto-rule much in the same manner as DCG's. The conditions are separated from the rule by {}, which appear at the very end of the rule. The acceptable conditions are local covering, local matching, precedence, immediate precedence, local compaction and inequations or other Prolog goals. The following subsections discuss various types of conditions that are allowed in tecto-rules.

### Local Covering

As well as global covering constraints, one can state *covers/2* constraints locally inside tecto-rules. The scope of local constraints does not surpass the tecto-rule in which they are used. The syntax of local covering is as follows:

```

<local-covers> ::= <non-zero-index> covers <field-or-region>

```

Disjunction is not allowed on the right-hand-side of local *covers* constraints. The *<non-zero-index>* stands for a non-zero integer that ranges from 1 to the number of daughters in the applicable tecto-rule. Therefore, the following hypothetical rule is syntactically acceptable:

```

pp *--> [(p, case:C), (nbar, case:C),
          {2 covers npr}].

```

What the above rule says is that a PP consists of a P and an NBar that shares its CASE feature with P and that the second daughter in the rule, namely, NBar covers

an npr. In this context the field or region specified will be one that is topologically accessible to the sponsor of the mother node. Therefore, the npr in the above example has to be topologically accessible to the pheno-edge that resulted in the prediction of that PP.

### Local Matching

Local matching is a special case of local covering. The only difference is that in matching the daughter covering the region cannot be larger than it. The syntax of local matching is provided below:

```
<local-matches> ::= <non-zero-index> matches <field-or-region>
```

An example of a local matches constraint in German grammar is:

```
(nbar, person:P, number:N, gender:G, case:C) *-->
  [(nbar, person:P, number:N, gender:G, case:C), rp,
   { 2 matches nf
     ; 2 matches postf
   }].
```

This rule states that an NBar consists of another NBar with the same PERSON, NUMBER, GENDER and CASE features as well as an RP, which matches either an *nf* or *postf* that is accessible to the sponsor of the mother NBar.

### Precedence

Precedence constraints have this syntax:

```
<precedence> ::= <non-zero-index> < <non-zero-index>
```

What they mean is that the daughter whose index is mentioned on the left-hand side must be entirely located before the daughter whose index is on the right-hand side. Note that this constraint does not make any assumptions about the contiguity of the daughters. As an example, one can provide the following hypothetical rule that states an NP consists of a determiner that precedes but might not be adjacent to an NBar with the same agreement features.

```
(np, person:P, number:N, gender:G, case:C) *-->
  [(det, number:N, gender:G, caes:C),
   (nbar, person:P, number:N, gender:G, case:C),
   {1 < 2}].
```

### Immediate Precedence

A special kind of precedence is immediate precedence, which means that the two daughters mentioned in the constraint have to be adjacent to one another. This only applies to the rightmost word of the first daughter specified in the constraint and the leftmost word of the second one. This constraint does not assume contiguity of daughters either. The syntax of an immediate precedence constraint is as follows:

```
<immediate-precedence> ::= <non-zero-index> << <non-zero-index>
```

The following serves as an example for this and means that an NP consists of a determiner which is immediately followed by an NBar.

```
(np, person:P, number:N, gender:G, case:C) *-->
  [(det, number:N, gender:G, case:C),
   (nbar, person:P, number:N, gender:G, case:C),
   {1 << 2}].
```

### Local Compaction

As is clear from the name of the constraint, a local compaction constraint states that a daughter or the mother node in a rule is a contiguous string. Obviously, you do not need to specify that for lexical categories that are by definition one word long and therefore contiguous. The syntax of local compaction constraints is given below:

```
<local-compaction> ::= compacts(<index>)
```

The variable <index> stands for any integer from 0 to the number of the daughters in the tecto-rule in which the constraint appears. If <index> is set to zero, it means that the mother node should be a contiguous string. For example:

```
(np, person:P, number:N, gender:G, case:C) *-->
  [(det, number:N, gender:G, case:C),
   (nbar, person:P, number:N, gender:G, case:C),
   {compacts(0)}].
```

means that an NP consists of a determiner and an NBar and that the mother NP is contiguous.

Note that the argument of a local compaction constraint is not a list but a single index.

### Inequations and Other Prolog Goals

Sometimes it is necessary to make sure that certain features do not have certain values and sometimes one may need to manipulate symbols or make sure that some other constraints apply. In such cases, it is possible to embed some Prolog code inside the tecto-rules as in DCG's. For example, if you want to make sure that a tenseless verb is not combined with an NP to make a sentence, you can use the following tecto-rule:

```
s *--> [(np, person:P, number:N, case:nom),
        (vp, person:P, number:N, form:F),
        { get_type(F,FType),
          FType\==inf,
          FType\==pastp,
          FType\==presp
        }].
```

This rule states that a sentence contains an NP and a VP who share the same person and number features and that the form of the verb (F) is not infinitival, past-participial or present-participial.

## T8.4 The Lexicon

Lexical entries are introduced using the infix operator `--->/2`. The syntax of a lexical entry is as follows:

```
<lexical-entry> ::= <entry>
                  | <morph-rule>

<entry> ::= <lexical-item> ---> <tecto-category>

<morph-rule> ::= (<entry>) :- <prolog-goals>

<lexical-item> ::= <prolog-atom>
                  | <prolog-variable>
```

A `<lexical-item>` should be either a Prolog atom or variable. A `<tecto-category>` is either a description or variable. An `<entry>` can be placed in the head of a Prolog conditional clause to define morphological processes. The following examples of lexical entries are from the file `german.pl`:

```
will ---> (aux,form:inf).

(sieht ---> (tv,person:third,number:sg,form:pres)):- !.
```

For more examples, refer to the files `german.pl` and `sc.pl`.

### T8.4.1 Lexical Categories

Because the lexicon as defined here is not static and may encode morphological processes, it is an unreliable source of information for inducing lexical categories. TGC, therefore, uses the tecto-rules for this purpose. Any category is considered lexical if no category corresponding to that category appears in the left-hand side of the tecto-rules. In the following example, Aux, Pron and TV are thus considered lexical:

```
s *--> [(np, person:P, number:N, case:nom),
        (vp, person:P, number:N)].

s *--> [(np, person:P, number:N, case:nom),
        (aux, person:P, number:N, form:F),
        (vp, person:no,number:no,form:F)].

(np, person:P, number:N, gender:G, case:C) *-->
  [(pron, person:P, number:N, gender:G, case:C)].
```

```
(vp, person:P, number:N, form:F) *-->
  [(vbar, person:P, number:N, form:F)].
```

```
(vbar, person:P, number:N, form:F) *-->
  [(tv, person:P, number:N, form:F),
   (np, case:acc)].
```

## T8.5 Edges

A TGC-generated parser uses four types of edges: pheno-edge, tecto-aedge, tecto-pedge and tecto-fedge. A pheno-edge is a passive edge corresponding to a found phenogrammatical category. Pheno-edges look like the following in the parser:

```
pheno_edge(L,R,FR,Id)
```

L and R are for the left and right boundaries of the edge. FR records the field or region found, and Id is the unique ID that each pheno-edge has.

Tecto-aedges are active edges that are predicted either from phenogrammar as a result of a `matched_by` constraint or predicted in tectogrammar according to some tecto-rule. They hold the following information:

```
tecto_aedge(Cat, Sponsor, Keys, CanBV, OptBV)
```

Cat is the category predicted. Sponsor is the ID of the pheno-edge that resulted in the prediction of the active edge. Keys is a list of ID's that belong to inaccessible passive or frozen edges. Taking a key by an active edge means that the active edge has promised to consume all of the daughters of the edge it has taken its keys/ID's from. This is done by adjusting the CanBV and OptBV of the active edge. These respectively stand for a bit vector representing the words that an edge is allowed to take and the words that it may optionally take.

A passive tecto edge (tecto-pedge) records a category that has been found and it contains the following:

```
tecto_pedge(Cat, Tree, Keys, DtrIds, BV)
```

Cat is the category that has been found. Tree represents the internal structure of the category. DtrIds is a list of the ID's of the lexical items that a passive edge has consumed. And finally, BV represents the positions of the words consumed by the passive edge as a bit vector.

Frozen edges (tecto-fedge) are introduced by `covered_by` constraints if any. These are basically active edges that are waiting for a CanBV. The format of a tecto-fedge is as follows:

```
tecto_fedge(Cat, Id, ReqBV)
```

Cat is the category predicted. Id is the ID of the pheno-edge that resulted in the prediction of the frozen edge. ReqBV is the bit vector that is required to be consumed for Cat to be made. A frozen edge is unfrozen by an active edge that has access to it.

For more information on the formalism and the working mechanism of a TGC-generated topological parser, refer to Penn and Haji-Abdolhosseini (2002).



# Appendix

## T0.1 Syntax of A Topological Grammar

```

<bidirectional-covering> ::= <phi> <--> <field-or-region>

<bidirectional-matching> ::= <phi> <==> <field-or-region>

<compaction-rule> ::= compacts(<phi-list>)

<condition> ::= <local-covers>
               | <local-matches>
               | <precedence>
               | <immediate-precedence>
               | <local-compaction>
               | <inequation>

<conditions> ::= <condition>
                | (<conditions>, <conditions>)
                | (<conditions>; <conditions>)

<covered-by-rule> ::= <field-or-region> covered_by <phi>
                    | <field-or-region> covered_by (<phis>)

<covers-rule> ::= <phi> covers <field-or-region>
                | <phi> covers (<fields-or-regions>)

<entry> ::= <lexical-item> ---> <tecto-category>

<field> ::= <prolog-atom>
           | matrix

<field-or-region> ::= <field>
                    | <region>

<fields-or-regions> ::= <field-or-region>
                      | <fields-or-regions> ; <fields-or-regions>

<fields> ::= <field>
            | <fields> ; <fields>

<immediate-precedence> ::= <non-zero-index> << <non-zero-index>

<lexical-entry> ::= <entry>
                  | <morph-rule>

<lexical-item> ::= <prolog-atom>

```

```

        | <prolog-variable>

<linking-rule> ::= <region> <<-- <field>
                  | <region> <<-- (<fields>)
                  | <field> -->> <region>
                  | <field> -->> (<regions>)

<local-compactation> ::= compact(<index>)

<local-covers> ::= <non-zero-index> covers <field-or-region>

<local-matches> ::= <non-zero-index> matches <field-or-region>

<matched-by-rule> ::= <field-or-region> matched_by <phi>
                    | <field-or-region> matched_by (<phis>)

<matches-rule> ::= <phi> matches <field-or-region>
                  | <phi> matches (<fields-or-regions>)
                  | (<phi> matches <field-or-region>) :- <prolog-goals>

<morph-rule> ::= (<entry>) :- <prolog-goals>

<phi> ::= <prolog-atom>
         | <prolog-variable>

<phi-conj> ::= <phi>
             | <phi-conj>, <phi-conj>

<phis> ::= <phi>
          | <phis> ; <phis>

<precedence> ::= <non-zero-index> < <non-zero-index>

<region> ::= <prolog-atom>

<regions> ::= <region>
            | <regions> ; <regions>

<tecto-rule> ::= <phi> *--> [<phi-conj>]
              | <phi> *--> [<phi-conj>, {<conditions>}]

<topo-rule> ::= <region> topo <field-descriptor-list>

```

## T0.2 TGC Error and Warning Messages

Ambiguous field name(s) FieldNames in R1 and R2

**R1 and R2 share the field(s) in the list FieldNames**

Argument of compacts/2 not a list in compacts(Phi)

**The argument a global compacts/1 constraint, Phi is not a list.**

covered\_by illigally used in rule Rule

**A covered\_by/2 constraint has been used locally in rule Rule.**

Disjoined field names in conditional matches/2 for Phi matches Field

**Field in a conditional matches/2 constraint is a disjunction of field names.**

Inconsistent linking rules for F and R

**Linking rules -->>/2 and <<--/2 for field F and region R violate the universal quantification condition of the rules.**

Index not in range in Constraint constraint in Rule

**The index used in Constraint in tecto-rule Rule is not in the valid range.**

Index not integer in Constraint constraint in Rule

**The index used in Constraint in tecto-rule Rule is not an integer.**

Matched\_by illigally used in rule Rule

**A matched\_by/2 constraint has been used locally in rule Rule.**

Matrix not maximal in F -->> R

**The linkage rule F -->> R renders matrix non-maximal.**

Matrix not maximal in R <<-- F

**The linkage rule R <<-- F renders matrix non-maximal.**

Multiple Constraint declarations for X

**More than one global constraint Constraint has been found for X. Only the last instance is considered.**

No fields specified in topo rule: R topo []

The right-hand side of a topo rule for R is an empty list.

No linking rules in pheno-grammar

No linking rules (<<--/2 or -->>/2) were found in the grammar.

No matched\_by rules. The parser will not tecto-parse.

The matched\_by rules are responsible for predicting tectogrammatical categories. Without them tecto-parsing is not possible.

Unary-branching topo rule: TopoRule

All topo rules must be at least binary branching. Use linkage rules instead.

## Appendix A

# Compiling ALE Programs

This section is devoted to showing how ALE programs can actually be compiled. ALE was developed to be run with a Prolog compiler, such as SICStus Prolog's. An SWI port of ALE is available, which also has a less extensive compilation phase. We strongly recommend SICStus Prolog. SWI Prolog does not scale up well to large-sized grammars. The local systems administrator should be able to provide help on running Prolog. This documentation only assumes the user has figured out how to run Prolog as well as write and edit files. It is otherwise self-contained.

### 1.1 File Management

After starting up Prolog, the following command should be used to load the ALE system:

```
| ?- compile(AleFile).
```

where *AleFile* is an atom specifying the file name in which ALE resides. For instance, in Unix, you might need to use something like: `compile('/users/carp/Prolog/ALE/ale.pl')`, or a local abbreviation for it like `compile(ale)`. if the system is in a file named `ale.pl` in the local directory (SICStus and SWI can fill in the ".pl" suffix). With SWI Prolog, the command:

```
| ?- consult(AleFile).
```

must be used instead. Note that the argument to `compile` must be an atom, which means it should be single-quoted if it is not otherwise an atom. After the system has compiled, you should see another Prolog prompt. It is necessary to have write permission in the directory from which Prolog is invoked, because ALE creates files during compilation. But note that neither the grammar nor ALE need to be locally defined; it is only necessary to have local write permission.

ALE source code, being a kind of Prolog code, must be organized so that predicate definitions are not spread across files, unless the appropriate `multifile` declarations are made. For instance, the `sub/intro` clauses specifying the type hierarchy must all be in one file. Similarly, the definite clauses must all be in one file, as must the grammar rules and macros.

## 1.2 Compiling Programs

ALE can compile a program incrementally to some extent. In particular, the compiler is broken down into six primary components for compiling the type hierarchy, functional descriptions, type constraints, the attribute-value logic, the definite clauses and the grammar. Compiling the type hierarchy consists of compiling type subsumption, type unification, appropriateness specifications, and extensionality information. The logic compiler compiles predicates which know how to add a type to a feature structure, how to find a feature value in a type and how to perform feature structure unification, as well as the most general satisfiers of every type, with code attached to enforce `cons/2` constraints. Compiling the grammar consists of compiling the lexicon, empty categories, rules and lexical rules, and if compilation for generation is enabled, the `semantics/1` directive. Macros are not compiled, but are rather interpreted during compilation.

There is one predicate `compile_gram/1` that can be used to compile a whole ALE grammar from one file, as follows:

```
| ?- compile_gram(GramFile).
```

where *GramFile* is the name of the file in which the grammar resides. The compiler will display error messages to the screen when it is compiling. But since ALE uses the Prolog compiler to read the files, Prolog might also complain about syntax errors in specifying the files. In either case, there should be some indication of what the error is and which clause of the file contained it.

ALE's compiler creates code for parsing, generation, or both. As of the present version of ALE, only one grammar can be used, even if code for both modes is to be created. Two files, `ale_parse.pl` and `ale_gen.pl`, are included with the distribution, which provide some example glue code to link together two ALE processes running under SICStus Prolog 3.0 or higher in order to parse and generate with two different grammars.

At startup, ALE produces code only for parsing. To produce code for generation only, use the command:

```
| ?- generate.
```

```
compiler will produce code for generation only
```

```
yes
| ?-
```

To produce code for both parsing and generation, use `parse_and_gen` instead. To switch back to producing code for parsing only, use `parse`. Note that these commands modify the behaviour of the compiler, not the compiled code, so grammars may need to be recompiled after these directives are issued.

The following predicates are available to compile grammars and their component parts. They are listed hierarchically, with each command calling all those listed under it. Each higher-level command is nothing more than the combination of those commands below it.

Command	Requires	File	Mode	Clause
<code>compile_gram</code>	nothing	*	both	
<code>compile_sig</code>	nothing	*	both	
<code>compile_sub_type</code>	nothing	*		sub
<code>compile_unify_type</code>	<code>compile_sub_type</code>			
<code>compile_approp</code>	<code>compile_unify_type</code>	*		intro
<code>compile_extensional</code>	<code>compile_approp</code>	*		ext
<code>compile_fun</code>	<code>compile_sig</code>	*	both	+++>
<code>compile_cons</code>	<code>compile_fun</code>	*	both	cons
<code>compile_logic</code>	<code>compile_sig</code>		both	
<code>compile_mgsat</code>	<code>compile_sig</code>			
<code>compile_add_to_type</code>	<code>compile_sig</code>			
<code>compile_featval</code>	<code>compile_add_to_type</code>			
<code>compile_u</code>	<code>compile_sig</code>			
<code>compile_subsume</code>	<code>compile_sig</code>		parse/subtest	
<code>compile_dcs</code>	<code>compile_logic</code>	*	both	if
<code>compile_grammar</code>	<code>compile_logic</code>	*		
<code>compile_lex_rules</code>	<code>compile_logic</code>	*	parse	**>
<code>compile_lex</code>	<code>compile_logic</code>	*	parse	--->
<code>compile_rules</code>	<code>compile_logic</code>	*	parse	===>,empty
	<code>compile_logic</code>	*	gen	===>,empty
				--->,**>
<code>compile_generate</code>	<code>compile_rules</code>	*	gen	semantics

The table above lists which compilations must have already been compiled before the next stage of compilation can begin. Thus before `compile_grammar` can be called, `compile_logic` must be called (or equivalently, the sequence of `compile_mgsat`, `compile_add_to_type`, `compile_featval`, and `compile_u`). Each command with an asterisk in its clauses column in the above table may be given an optional file argument. The file argument should be an atom which specifies the file in which the relevant clauses can be found. The clauses needed before each stage of compilation can begin are listed to the right of the asterisks. For instance, the `if` clauses must be loaded before `compile_dcs` is called. But note that `compile_unify_type` does not require any clauses to be loaded, as it uses the compiled definition of `sub_type` rather than the user specification in its operation. Thus changes to the signature in the source file, even if the source file is recompiled, will not be reflected in `compile_unify_type` if they have not been recompiled by `compile_sub_type` first. If an attempt is made to compile a part of a program where the relevant clauses have not been asserted, an error will result.

Note that `compile_subsume` only compiles code if subsumption checking (p. 93) and parsing have been enabled.

Each of the lowest level commands generates intermediate Prolog source code for that function, which is then compiled further by a Prolog compiler. ALE uses a term-expansion-based compiler in both SICStus and SWI Prologs that avoids the necessity for creating intermediate files. It also improves the speed of intermediate code compilation. Because both SICStus and SWI Prologs require the user to read a file on disk in order to use their compilers, ALE must create a zero-byte file

called `.alec_throw` to throw control to its intermediate code compiler. For that reason, the Prolog process must have write permission in the local directory to create this file, if it does not already exist.

After a grammar is compiled, the system plus grammar code can be saved with the command:

```
| ?- save_program(File).
```

This will save the state of the Prolog database in *File*. SICStus users should normally use this rather than `save/1`, which creates a larger file by saving other information like the state of Prolog's internal stacks. The SWI Prolog command is `qsave_program(File)`. With either Prolog, the state can be reloaded, by executing the saved file directly.

In general, whenever the ALE source program is changed, it should be recompiled from the point of change. For instance, if the definite clauses are the only thing that have changed since the last compilation, then only `compile_dcs(FileSpec)` needs to be run. But if in changing the definite clauses, the type hierarchy had to be changed, then everything must be recompiled.

ALE treats lexicon compilation differently than the other stages. Two commands, `lex_compile/0` and `lex_consult/0`, control whether the intermediate code for the lexicon and empty categories is compiled or consulted (a lesser degree of compilation). Lexicon compilation is usually the most time-consuming stage of grammar compilation in ALE, and consulting the code for this stage can result in a substantial compile-time speed-up. The decrease in run-time performance is only significant in grammars with a high degree of lexical ambiguity, i.e., where one string has a very large number of entries in the lexicon. By default, ALE consults the code for the lexicon. In SWI Prolog, only lexicon consulting is available.

When consulting is chosen, the lexicon and empty categories are also consulted *dynamically*. This means that individual entries can be retracted and added without recompiling the entire lexicon. To retract a lexical entry, use the command:

```
| ?- retract_lex(LexSpec).
```

*LexSpec* can either be a word, or a list of words. The words given to `retract_lex/1` are not closed under morph rules — derived entries with different forms must be retracted explicitly. `retract_lex/1` iterates through all of the entries that match the given word(s), asking for each one whether it should be retracted. `retractall_lex/1` will remove all of them without asking.

To add lexical entries, use the command:

```
| ?- update_lex(UpdateFile).
```

*UpdateFile* is a file containing new lexical entries and empty categories. New lexical entries are closed under lexical rules, as usual.<sup>1</sup>

---

<sup>1</sup>Earlier versions of ALE also permitted incremental updating and retraction of empty categories. Because empty categories are now closed under phrase structure rules at compile-time, this is no longer possible.



## 1.3 Compile-Time Error Messages

There are three sources of compile-time messages generated by ALE: Prolog messages, ALE errors, and ALE warnings.

ALE uses Prolog term input and output, thus requiring the input to be specified as a valid Prolog program. Of course, any ALE program meeting the ALE syntax specification will not cause Prolog errors. If there is a Prolog error generated, there is a corresponding bug in the grammar file(s). Prolog error messages usually generate a message indicating what kind of error it found, and just as importantly, which line(s) of the input the error was found in. The most common Prolog error messages concern missing periods or operators which cannot be parsed. Such errors are usually caused by bad punctuation such as missing periods, misplaced commas, commas before semicolons in disjunctions, etc. These errors are usually easy to track down.

Prolog also generates warnings in some circumstances. In particular, if you only use a variable once in a definition, it will report a singleton variable warning. The reason for this is that variables that only occur once are useless in that they do not enforce any structure sharing. There is little use for singleton variables in ALE outside of the Prolog goals in morphological rules and some macro parameters. Usually a singleton variable indicates a typing error, such as typing `AgNum` in one location and `Agnum` in another. It is standard Prolog practice to replace all singleton variables with *anonymous* variables. An anonymous variable is a variable which begins with the underscore character. For instance, a singleton variable such as `Head` can be replaced with the anonymous variable `_Head`, or even just `_`, to suppress such singleton variable warnings. Two occurrences of the simple anonymous variable `_` are not taken to be co-referential, but two occurrences of something like `_Head` are taken to be co-referential. In particular, the two descriptions, `(foo:X, bar:X)` and `(foo:_X, bar:_X)` are equivalent to each other, but distinct from `(foo:_, bar:_)` in that the latter description does not indicate any structure sharing. The second description above is considered bad style, though, as it uses the anonymous variable `_X` co-referentially.

Besides Prolog syntax errors, there are many errors that ALE is able to detect at compile time. These errors will be flagged during compilation. Most errors give some indication of the program clause in which they are found. Some errors may be serious enough to halt compilation before it is finished. In general, it is a good idea to fix all of the errors before trying to run a program, as the error messages only report serious bugs in the code, such as type mismatches, unspecified types, ill-formed rules, etc.

In certain cases, it is preferable to disable those error messages concerned with ALE's inability to add incompatible descriptions to a feature structure. This is especially true during lexicon and empty category compilation, when, due to the interaction of disjunctions and type constraints, the number of such errors can be overwhelming. In the current version of ALE, these errors are automatically disabled during lexicon and empty category compilation, and enabled otherwise. Commands will be added to future versions so that the user may control when these errors should be displayed.

Less serious problems are flagged with warning messages. Warning messages do not indicate an error, but may indicate an omission or less than optimal ALE

programming style.

The ALE error and warning messages are listed in an appendix at the end of this report, along with an explanation. The manual for the Prolog in which ALE is being run in will probably list the kinds of errors generated by the Prolog compiler.

## Appendix B

# Running and Debugging ALE Programs

After the ALE program compiles without any error messages, it is possible to test the program to make sure it does what it is supposed to. We consider the problem from the bottom-up, as this is the best way to proceed in testing grammars. ALE does not have a sophisticated input/output package, and thus all ALE procedures must be accessed through Prolog queries.

### 2.1 Testing the Signature

Once the signature is compiled, it is possible to test the results of the compilation. To test whether or not a type exists, use the following query:

```
| ?- type(Type).  
  
Type = bot ?;  
  
Type = cat ?;  
  
Type = synsem ?  
  
yes
```

Note that the prompt `| ?-` is provided by Prolog, while the query consists of the string `type(Type).`, including the period and a return after the period. Prolog then responds with instantiations of any variables in the query if the query is successful. Thus the first solution for `Type` that is found above is `Type = bot`. After providing an instantiation representing the solution to the query, Prolog then provides another prompt, this time in the form of a single question mark. After the first prompt above, the user typed a semicolon and return, indicating that another solution is desired. The second solution Prolog found was `Type = cat`. After this prompt, the user requested a third solution. After the third solution, `Type = synsem`, the user simply input a return, indicating that no more solutions were desired. These two options, semicolon followed by return, and a simple return, are the only ones relevant for ALE. If the anonymous variable `_` is used in a

query, no substitutions are given for it in the solution. If there are no solutions to a query, Prolog returns `no` as an answer. Consider the following two queries:

```
| ?- type(bot).

yes

| ?- type(foobar).

no
```

In both cases, no variables are given in the input, so a simple `yes/no` answer, followed by another prompt, is all that is returned.

The second useful probe on the signature indicates type subsumptions and type unifications. To test type subsumption, use the following form of query:

```
| ?- sub_type(X,Y).

X = and,
Y = and ?;

X = backward,
Y = backward ?

yes
```

Note that with two variables, substitutions for both are given, allowing the possibility of iterating through the cases. In general, wherever a variable may be used in a query, a constant may also be used. Thus `sub_type(synsem,forward)` is a valid query, as are `sub_type(synsem,X)` and `sub_type(Y,forward)`. The first argument is the more general type, with the second argument being the subtype.

Type unifications are handled by the following form of query:

```
| ?- unify_type(T1,T2,T).
```

The interpretation here is that `T1` unified with `T2` produces `T3`. As before, any subset of the three variables may be instantiated for the test and the remaining variables will be solved for.

The following query will indicate whether given features have been defined and can also be used to iterate through the features if the argument is uninstantiated:

```
| ?- feature(F).
```

Feature introduction can be tested by:

```
| ?- introduce(F,T).
```

which holds if feature `F` is introduced at type `T`.

Type constraints can be tested using:

```
| ?- show_cons(Type).
```

which will display the description of the constraint assigned to the type, `Type`.

Finally, the inherited appropriateness function can be tested by:

```
| ?- approp(Feat,Type,Restr).
```

A solution indicates that the value for feature `Feat` for a type `Type` structure is of type `Restr`. As usual, any of the variables may be instantiated, so that it is possible to iterate through the types appropriate for a given feature or the features appropriate for a given type, the restrictions on a given feature in a fixed type, and so on.

There is one higher-level debugging routine for the signature that outputs a complete specification for a type, including a list of its subtypes and supertypes, along with the most general feature structure of that type (after all type inference and constraint satisfaction has been performed). An example of the `show_type/1` query is as follows:

```
| ?- show_type functional.

TYPE: functional
SUBTYPES: [forward,backward]
SUPERTYPES: [synsem]
MOST GENERAL SATISFIER:
    functional
    ARG synsem
    RES synsem
```

If `synsem` had any appropriate features, these would have been added, along with their most general appropriate values.

## 2.2 Evaluating Descriptions

Descriptions can be evaluated in order to find their most general satisfiers. ALE provides the following form of query:

```
| ?- mgsat tl:e_list.

ne_list_quant
HD quant
  RESTR proposition
  SCOPE proposition
  VAR individual
TL e_list

ANOTHER? n.

yes
```

Note that there must be whitespace between the `mgsat` and the description to be satisfied. The answer given above is the most general satisfier of the description

`tl:e_list` using the signature in the categorial grammar in the appendix. It is important to note here that type inference is being performed to find most general satisfiers. In the case at hand, because lists in the categorial grammar are typed to have quantifiers as their HD values, the value of the HD feature in the most general satisfier has been coerced to be a quantifier.

Satisfiable non-disjunctive descriptions always have unique most general satisfiers as a consequence of the way in which the type system is constrained. But a description with disjunctions in it may have multiple satisfiers. Consider the following query:

```
| ?- mgsat hit,hitter:(j;m).
```

```
hit
HITTEE individual
HITTER j
```

```
ANOTHER? y.
```

```
hit
HITTEE individual
HITTER m
```

```
ANOTHER? y.
```

```
no
```

After finding the first most general satisfier to the description, the user is prompted as to whether or not another most general satisfier should be sought. As there are only two most general satisfiers of the description, the first request for another satisfier succeeds, while the second one fails. Failure to find additional solutions is indicated by the `no` response from Prolog.

Error messages will result if there is a violation of the type hierarchy in the query. For instance, consider the following query containing two type errors before a satisfiable disjunct:

```
| ?- mgsat hd:j ; a ; j.
```

```
add_to could not add incompatible type j to:
  quant
  RESTR proposition
  SCOPE proposition
  VAR individual
```

```
add_to could not add undefined type: a to
  bot
```

```
MOST GENERAL SATISFIER OF: hd:j;a;j
```

```
j
```

ANOTHER?

Here the two errors are indicated, followed by a display of the unique most general satisfiers. The problem with the first disjunct is that lists have elements which must be of the quantifier type, which conflicts with the individual type of *j*, while the second disjunct involves an undefined type *a*. Note that in the error messages, there is some indication of how the conflict arose as well as the current state of the structure when the error occurred. For instance, the system had already figured out that the head must be a quantifier, which it determined before arriving at the incompatible type *j*. The conflict arose when an attempt was made to add the type *j* to the quant type object.

To explore unification, simply use conjunction and `mgsat`. In particular, to see the unification of descriptions *D1* and *D2*, simply display the most general satisfiers of *D1*, *D2*, and their conjunction (*D1*, *D2*). To obtain the correct results, *D1* and *D2* must not share any common variables. If they do, the values of these will be unified across *D1* and *D2*, a fact which is not represented by the most general satisfiers of either *D1* or *D2*. Providing most general satisfiers also allows the user to test for subsumption or logical equivalence by visual inspection, by using `mgsat/1` and comparing the set of solutions. Future releases should contain mechanisms for evaluating subsumption (entailment), and hence logical equivalence of descriptions.

`mgsat` can also be used to test functional descriptions, e.g., for the functional append on p. 32:

```
| ?- mgsat append([bot],[bot,bot]).
```

```
ne_list
HD bot
TL ne_list
  HD bot
  TL ne_list
    HD bot
    TL e_list
```

ANOTHER?

The Prolog predicate `iso_desc/2` can be used to discover whether two descriptions evaluate to the same feature structure. This can be useful for testing extensional type declarations.

```
| ?- iso_desc(X,X).
```

```
X = _A-bot ?
```

```
yes
| ?- iso_desc((a_ atom),(a_ atom)).      % a_ atoms are extensional
yes
```

```

| ?- iso_desc(b,b).           % for b, intensional

no
| ?- iso_desc(a,a).           % for a, extensional, with feature f

no
| ?- iso_desc(f:(a_ at1),f:(a_ at1)). % f approp. to a_ atoms

yes
| ?- iso_desc(f:(a_ at1),f:(a_ at2)).

no

```

## 2.3 Hiding Types and Features

With a feature structure system such as ALE, grammars and programs often manipulate very large feature structures. To aid in debugging, two queries allow the user to focus attention on particular types and features by suppressing the printing of other types and features.

The following command suppresses printing of a type:

```
| ?- no_write_type(T).
```

After `no_write_type(T)` is called, the type *T* will no longer be displayed during printing. To restore the type *T* to printed status, use:

```
| ?- write_type(T).
```

If *T* is a variable in a call to `write_type/1`, then all types are subsequently printed. Alternatively, the following query restores printing of all types:

```
| ?- write_types.
```

Features and their associated values can be suppressed in much the same way as types. In particular, the following command blocks the feature *F* and its values from being printed:

```
| ?- no_write_feat(F).
```

To restore printing of feature *F*, use:

```
| ?- write_feat(F).
```

If *F* is a variable here, all features will subsequently be printed. The following special query also restores printing of all features.

```
| ?- write_feats.
```



## 2.4 Evaluating Definite Clause Queries

It is possible to display definite clauses in feature structure format by name. The following form of query can be used:

```
| ?- show_clause append.

HEAD: append(e_list,
             [0] bot,
             [0] )
BODY: true

ANOTHER? y.

HEAD: append(ne_list_quant
            HD [0] quant
              RESTR proposition
              SCOPE proposition
              VAR individual
            TL [1] list_quant,
              [2] bot,
              ne_list_quant
            HD [0]
            TL [3] list_quant)
BODY: append([1],
            [2],
            [3])

ANOTHER? y.

no
```

Note that this example comes from the categorial grammar in the appendix. Also note that the feature structures are displayed in full with tags indicating structure sharing. Next, note that prompts allow the user to iterate through all the clauses. The number of solutions might not correspond to the number of clause definitions in the program due to disjunctions in descriptions which are resolved non-deterministically when displaying rules. But it is important to keep in mind that this feature structure notation for rules is not the one ALE uses internally, which compiles rules down into elementary operations which are then compiled, rather than evaluating them as feature structures by unification. In this way, ALE is more like a logic programming compiler than an interpreter. Finally, note that the arity of the predicate being listed may be represented in the query as in Prolog. For instance, the query `show_clause append/3` would show the clauses for `append` with three arguments.

Definite clauses in ALE can be evaluated by using a query such as:

```
| ?- query append(X,Y,[a,b]).
```

```

append(e_list,
      [0] ne_list
      HD a
      TL ne_list
      HD b
      TL e_list,
      [0] )

```

```

ANOTHER? y.
append(ne_list
      HD [0] a
      TL e_list,
      [1] ne_list
      HD b
      TL e_list,
      ne_list
      HD [0]
      TL [1] )

```

```

ANOTHER? y.
append(ne_list
      HD [0] a
      TL ne_list
      HD [1] b
      TL e_list,
      [2] e_list,
      ne_list
      HD [0]
      TL ne_list
      HD [1]
      TL [2] )

```

```

ANOTHER? y.

```

```

no

```

The definition of `append/3` is taken from the syllabification grammar in the appendix. After displaying the first solution, ALE queries the user as to whether or not to display another solution. In this case, there are only three solutions, so the third query for another solution fails. Note that the answers are given in feature structure notation, where the macro `[a, b]` is converted to a head/tail feature structure encoding.

Unlike Prolog, in which a solution is displayed as a substitution for the variables in the query, ALE displays a solution as a satisfier of the entire query. The reason for this is that structures which are not given as variables may also be further instantiated due to the type system. Definite clause resolution in ALE is such that only the most general solutions to queries are displayed. For instance, consider the following query, also from the syllabification grammar in the appendix:

```
| ?- query less_sonorous(X,r).

less_sonorous(nasal,
              r)

ANOTHER? y.

less_sonorous(sibilant,
              r)

ANOTHER? n.
```

Rather than enumerating all of the `nasal` and `sibilant` types, ALE simply displays their supertype. On the other hand, it is important to note that the query `less_sonorous(s,r)` would succeed because `s` is a subtype of `sibilant`. This example also clearly illustrates how ALE begins each argument on its own line arranged with the query.

In general, the goal to be solved must be a literal, consisting only of a relation applied to arguments. In particular, it is not allowed to contain conjunction, disjunction, cuts, or other definite clause control structures. To solve a more complex goal, a definite clause must be defined with the complex goal as a body and then the head literal solved, which will involve the resolution of the body.

There are no routines to trace the execution of definite clauses. Future releases of ALE will contain a box port tracer similar to that used for Prolog. At present, the best suggestion is to develop definite clauses modularly and test them from the bottom-up to make sure they work before trying to incorporate them into larger programs.

## 2.5 Displaying Grammars

ALE provides a number of routines for displaying and debugging grammar specifications. After compile-time errors have been taken care of, the queries described in this section can display the result of compilation.

Lexical entries can be displayed using the following form of query:

```
| ?- lex(kid).

WORD: kid
ENTRY:
cat
QSTORE e_list
SYNSEM basic
      SEM property
        BODY kid
          ARG1 [0] individual
            IND [0]
          SYN n
```

```
ANOTHER? y.
```

```
no
```

As usual, if there are multiple entries, ALE makes a query as to whether more should be displayed. In this case, there was only one entry for `kid` in the categorial grammar in the appendix.

Another predicate, `export_words(Stream, Delimiter)`, writes an alphabetised list of all of the words in the lexicon, separated by *Delimiter*, to *Stream*. In SICStus Prolog, for example, `export_words(user_output, '\n')` will write the words to standard output (such as the screen), one to a line.

Empty lexical entries can be displayed using:

```
| ?- empty.
```

```
EMPTY CATEGORY:
```

```
cat
QSTORE ne_list_quant
  HD some
    RESTR [0] proposition
    SCOPE proposition
    VAR [1] individual
  TL e_list
SYNSEM forward
  ARG basic
    SEM property
    BODY [0]
    IND [1]
  SYN n
  RES basic
    SEM [1]
    SYN np
```

```
ANOTHER? no.
```

Note that the number specification was removed to allow the empty category to be processed with respect to the categorial grammar type system. As with the other display predicates, `empty` provides the option of iterating through all of the possibilities for empty categories.

Grammar rules can be displayed by name, as in:

```
| ?- rule forward_application.
```

```
RULE: forward_application
```

```
MOTHER:
```

```
cat
QSTORE [4] list_quant
```

```
SYNSEM [0] synsem
```

DAUGHTERS/GOALS:

```
CAT  cat
    QSTORE [2] list_quant
    SYNSEM forward
        ARG [1] synsem
        RES [0]
```

```
CAT  cat
    QSTORE [3] list_quant
    SYNSEM [1]
```

```
GOAL  append([2],
             [3],
             [4])
```

ANOTHER? n.

Rules are displayed as most general satisfiers of their mother, category and goal descriptions. It is important to note that this is for display purposes only. The rules are not converted to feature structures internally, but rather to predicates consisting of low-level compiled instructions. Displaying a rule will also flag any errors in finding most general satisfiers of the categories and rules in goals, and can thus be used for rule debugging. This can detect errors not found at compile-time, as there is no satisfiability checking of rules performed during compilation.

Macros can also be displayed by name, using:

```
| ?- macro np(X).
```

```
MACRO:
    np([0] sem_obj)
ABBREVIATES:
    basic
    SEM [0]
    SYN np
```

ANOTHER? n.

First note that the macro name itself is displayed, with all descriptions in the macro name given replaced with their most general satisfiers. Following the macro name is the macro satisfied by the macro description with the variables instantiated as shown in the macro name display. Note that there is sharing between the description in the macro name and the SEM feature in the result. This shows where the parameter is added to the macro's description.

Finally, it is possible to display lexical rules, using the following query:

```
| ?- lex_rule plural_n.
```

```

LEX RULE: plural_n
INPUT CATEGORY:
  n
  NUM sing
  PERS pers
OUTPUT CATEGORY:
  n
  NUM plu
  PERS pers
MORPHS:
  [g,o,o,s,e] becomes [g,e,e,s,e]
  [k,e,y] becomes [k,e,y,s]
  A,[m,a,n] becomes A,[m,e,n]
  A,B becomes A,B,[e,s]
    when fricative(B)
  A,[e,y] becomes A,[i,e,s]
  A becomes A,[s]

```

ANOTHER? n.

Note that the morphological components of a rule is displayed in canonical form when it is displayed. Note that variables in morphological rules are displayed as upper case characters. When there is sharing of structure between the input and output of a lexical rule, it will be displayed as such. As with the other ALE grammar display predicates, if there are multiple solutions to the descriptions, these will be displayed in order. Also, if there is a condition on the categories in the form of an ALE definite clause goal, this condition will be displayed before the morphological clauses. As with grammar rules, lexical rules are compiled internally and not actually executed as feature structures. The feature structure notation is only for display. Also, as with grammar rules, displaying a lexical rule may uncover inconsistencies which are not found at compile time.

## 2.6 Executing Grammars: Parsing

In this section, we consider the execution of ALE phrase structure grammars compiled for parsing. The examples shown in this section have been produced while running with the mini-interpreter off. The mini-interpreter will be discussed in the next section.

The primary predicate for parsing is illustrated as follows:

```
| ?- rec [john,hits,every,toy].
```

```
STRING:
0 john 1 hits 2 every 3 toy 4
```

```
CATEGORY:
cat
```

```

QSTORE e_list
SYNSEM basic
  SEM every
    RESTR toy
      ARG1 [0] individual
    SCOPE hit
      HITTEE [0]
      HITTER j
    VAR [0]
  SYN s

```

ANOTHER? y.

```

CATEGORY:
cat
QSTORE ne_list_quant
  HD every
    RESTR toy
      ARG1 [0] individual
    SCOPE proposition
    VAR [0]
  TL e_list
SYNSEM basic
  SEM hit
    HITTEE [0]
    HITTER j
  SYN s

```

ANOTHER? y.

no

The first thing to note here is that the input string must be entered as a Prolog list of atoms. In particular, it must have an opening and closing bracket, with words separated by commas. No variables should occur in the query, nor anything other than atoms. The first part of the output repeats the input string, separated by numbers (*nodes* in the chart) which indicate positions in the string for later use in inspecting the chart directly. ALE asserts one lexical item for every unit interval, with empty categories being stored as loops from every single node to itself. The second part of the output is a category which is derived for the input string. If there are multiple solutions, these can be iterated through by providing positive answers to the query. The final no response above indicates that the category displayed is the only one that was found. If there are no parses for a string, an answer of no is returned, as with:

```
| ?- rec([runs, john]).
```

```

STRING:
0 runs 1 john 2

```

no

Notice that there is no notion of “distinguished start symbol” in parsing. Rather, the recognizer generates all categories that it can find for the input string. This allows sentence fragments and phrases to be analyzed, as in:

```
| ?- rec [big,kid].

STRING:
0 big 1 kid 2

CATEGORY:
cat
QSTORE ne_list_quant
  HD some
    RESTR and
      CONJ1 kid
        ARG1 [0] individual
      CONJ2 big
        ARG1 [0]
    SCOPE proposition
      VAR [0]
      TL e_list
SYNSEM basic
  SEM [0]
  SYN np

ANOTHER? n.
```

There is also a two-place version of `rec` that displays only those parses that satisfy a given description:

```
| ?- rec([big,kid],s).

STRING:
0 big 1 kid 2

no
```

This call to `rec/2` failed because there were no parses of `big kid` of type `s`. Internally, the parser still generates all of the edges that it normally does — the extra description is only applied at the end as a filter.

Once parsing has taken place for a sentence using `rec/1`, it is possible to look at categories that were generated internally. In general, the parser will find every possible analysis of every substring of the input string, and these will be available for later inspection. For instance, suppose the last call to `rec/1` executed was `rec [john,hits,every,toy]`, the results of which are given above. Then the following query can be made:



```
| ?- edge(2,4).

COMPLETED CATEGORIES SPANNING: every toy

cat
QSTORE ne_list_quant
      HD every
      RESTR toy
          ARG1 [0] individual
          SCOPE proposition
          VAR [0]
      TL e_list
SYNSEM basic
      SEM [0]
      SYN np

Edge created for category above:
  index: 20
  from: 2 to: 4
  string: every toy
  rule: np_det_nbar
  # of dtrs: 2
Action(retract,dtr-#,continue,abort)?
|: continue.
```

```
no
| ?-
```

The possible replies in the action-line will be discussed in the next section. This tells us that from positions 2 to 4, which covers the string `every toy` in the input, the indicated category was found. Even though an active chart parser is used, it is not possible to inspect active edges. This is because ALE represents active edges as dynamic structures that are not available after they have been evaluated.

Using `edge/2` it is possible to debug grammars by seeing how far analyses progressed and by inspecting analyses of substrings.

There is also a predicate, `rec/4` that binds the answer to variables instead of displaying it:<sup>1</sup>

```
| ?- rec([kim,walks],Ref,SVs,Iqs).

Iqs = [ineq(_A,index(_B-third,_C-plur,_D-masc),_E,index(_B-third,_C-plur,
_D-masc),done)],
SVs = phrase(_F-synsem(_G-loc(_H-cat(_I-verb(_J-minus,_K-minus,_L-none,_M-
bool,_N-fin),_O-unmarked,_P-e_list),...)))?
```

`Ref` is an unbound variable that represents the root node of the resulting feature structure. `SVs` is a term whose functor is the type of the feature structure, whose

<sup>1</sup>Actually, this example is a bit of an improvisation — the sample HPSG grammar included in the ALE distribution does not use inequations.

arity is the number of appropriate features for that type, and whose arguments are the values at those appropriate features in alphabetical order. Each value, in turn, is of the form, *Ref-SVs*, another pair of root variable and type-functor term. *Iqs* is a list of the inequations that apply to the feature structure and its substructures. Each member of the list represents a disjunction of inequations, i.e., one must be satisfied, with the list itself being interpreted conjunctively, i.e., every member must be satisfied. Each member is represented by a chain of *ineq/5* structures:

```
ineq(Ref1,SVs1,Ref2,SVs2,ineq(. . . . .,done) . . .)
```

The first four arguments represent the *Ref-SVs* pairs of the two inequated feature structures of the first disjunct. The fifth argument contains another *ineq/5* structure, or *done/0*. These three structures are suitable for passing to *gen/3* or *gen/4*. These representations are not grounded; so if you want to assert them into a database, be sure to assert them all in one predicate to preserve variable sharing. For more details on ALE's internal representation, the reader is referred to Carpenter and Penn (1996).

There is also a *rec/5*, which works just like *rec/4*, but filters its output through the description in its last argument, just like *rec/2*:

```
| ?- rec([kim,walks],Ref,SVs,Iqs,phrase).
```

```
Iqs = [ineq(_A,index(_B-third,_C-plur,_D-masc),_E,index(_B-third,_C-plur,
_D-masc),done)],
SVs = phrase(_F-synsem(_G-loc(_H-cat(_I-verb(_J-minus,_K-minus,_L-none,_M-
bool,_N-fin),_O-unmarked,_P-e_list),...)))?
```

The call succeeds here because the given answer is of type, *phrase*.

*rec\_list/2* iteratively displays all of the solutions for each string in a list of list of words that satisfy a description:

```
| ?- rec_list([[kim,sees,sandy],[sandy,sees,kim]],phrase).
```

```
STRING:
0 kim 1 sees 2 sandy 3
CATEGORY:
```

```
phrase
QRETR e_list
QSTORE e_set
SYNSEM synsem...
```

```
ANOTHER? y.
```

```
STRING:
0 sandy 1 sees 2 kim 3
CATEGORY:
```

```
phrase
QRETR e_list
```

```
QSTORE e_set
SYNSEM synsem
```

```
ANOTHER? y.
```

```
no
```

If no filtering through a description is desired, the description, `bot`, which is trivially satisfied, can be used. When `rec_list/2` runs out of solutions for a string, it moves on to the next string.

`rec_best/2` iteratively displays all of the solutions satisfying a given description for the *first* string in a list of list of words that has a solution satisfying that description.

```
| ?- rec_best([[kim,sees,sandy],[sandy,sees,kim]],phrase).
```

```
STRING:
0 kim 1 sees 2 sandy 3
CATEGORY:
```

```
phrase
QRETR e_list
QSTORE e_set
SYNSEM synsem...
```

```
ANOTHER? y.
```

```
no
```

When `rec_best/2` runs out of solutions for a string that had at least one solution, it fails. It only tries the strings in its first argument until it finds one that has solutions.

There is also a three-place `rec_list/3` that collects the solutions to `rec_list/2` in a list of terms of the form `fs(Tag-SVs,Iqs)`.

## 2.7 Executing Grammars: Generation

The generator can be used with the predicate `gen/1,3` predicate. It can take a single description argument:

```
| ?- gen((sentence,
          sem:(pred:decl,
              args:[(pred:call_up,
                    args:[pred:mary,pred:john])]))).
```

```
CATEGORY:
```

```
sentence
SEM sem
```

```

ARGS arg_ne_list
  HD sem
    ARGS arg_ne_list
      HD sem
        ARGS arg_list
        PRED mary
      TL arg_ne_list
        HD sem
          ARGS arg_list
          PRED john
        TL e_list
      PRED call_up
    TL e_list
  PRED decl

```

```

STRING:
mary calls john up

```

```

ANOTHER? y.

```

```

STRING:
mary calls up john

```

```

ANOTHER? y.

```

```

no

```

Notice the extra set of parentheses necessary to make the whole description a single argument for `gen/1`.

`gen` can also take three arguments:

```

gen(Ref,SVs,Iqs)

```

where `Ref`, `SVs` and `Iqs` are the three parts of ALE's internal representation of a feature structure as defined in the last section. This alternative is most useful when the feature structure has been generated before by another process, like parsing, or retrieved from a database.

```

| ?- rec([john,calls,mary,up],Ref,SVs,Iqs),gen(Ref,SVs,Iqs).

```

```

CATEGORY:

```

```

sentence
SEM sem
  ARGS arg_ne_list
    HD sem
      ARGS arg_ne_list
        HD sem

```

```

        ARGS e_list
        PRED john
    TL arg_ne_list
        HD sem
            ARGS e_list
            PRED mary
        TL e_list
    PRED call_up
    TL e_list
    PRED decl

```

STRING:

john calls mary up

ANOTHER? n.

```

Iqs = [],
SVs = sentence(_O-sem(_N-arg_ne_list(_M-sem(_L-arg_ne_list(_K-
sem(_J-e_list,_I-john),_H-arg_ne_list(_G-sem(_F-e_list,_E-mary),
_D-e_list)),_C-call_up),_B-e_list),_A-decl)) ?

```

yes

In both cases, ALE will print the input feature structure and then will generate and display all possible string solutions through backtracking.

It is also possible to bind the string to a variable, using `gen/4`:

```
gen(Ref,SVs,Iqs,Ws).
```

`Ws` will non-deterministically be bound to the word lists that constitute valid solutions to the generation problem. This can be used as input to `rec/1`, for example.

```
| ?- rec([john,calls,mary,up],Ref,SVs,Iqs),gen(Ref,SVs,Iqs,Ws).
```

```

Iqs = [],
SVs = sentence(_O-sem(_N-arg_ne_list(_M-sem(_L-arg_ne_list(_K-
sem(_J-e_list,_I-john),_H-arg_ne_list(_G-sem(_F-e_list,_E-mary),
_D-e_list)),_C-call_up),_B-e_list),_A-decl)),
Ws = [john,calls,mary,up] ? ;

```

```

Iqs = [],
SVs = sentence(_O-sem(_N-arg_ne_list(_M-sem(_L-arg_ne_list(_K-
sem(_J-e_list,_I-john),_H-arg_ne_list(_G-sem(_F-e_list,_E-mary),
_D-e_list)),_C-call_up),_B-e_list),_A-decl)),
Ws = [john,calls,up,mary] ? ;

```

```

Iqs = [],
SVs = s(_L-finite,_K-sem(_J-arg_ne_list(_I-sem(_H-e_list,_G-john),

```

```
_F-arg_ne_list(_E-sem(_D-e_list,_C-mary),_B-e_list)),_A-call_up)),
Ws = [john,calls,mary,up] ? ;
```

```
Iqs = [],
SVs = s(_L-finite,_K-sem(_J-arg_ne_list(_I-sem(_H-e_list,_G-john),
_F-arg_ne_list(_E-sem(_D-e_list,_C-mary),_B-e_list)),_A-call_up)),
Ws = [john,calls,up,mary] ? ;
```

no

The last two solutions in the example above are generated because the input string, `john calls mary up`, can be parsed both as a sentence type and as an `s` type.

## 2.8 Mini-interpreter (parsing only)

ALE contains a mini-interpreter that allows the user to traverse and edit an ALE parse tree. By default, the mini-interpreter is off when ALE is loaded. To turn the mini-interpreter on, simply type:

```
| ?- interp.

interpreter is active

yes
| ?-
```

To turn it off again, use `nointerp`. Any parse automatically stores the following information on the edges added to ALE's chart:

- Spanning nodes
- Substring spanned
- Creator
- Daughters (if any)

The spanning nodes are the nodes in the chart that the edge spans. The substring spanned is the concatenation of lexical items between the spanning nodes. If an edge was formed by the application of an ALE grammatical rule, its creator is that rule, with the daughters being the daughters of the rule, i.e., the `cat>`, and `cats>` of the rule). If the rule was created by EFD closure, the mini-interpreter will treat it as the user's rule it was created from, displaying the empty categories that had matched its daughters during EFD closure in the same way as daughters matched at run-time. If an edge represents an empty category, its creator is normally `empty`; but empty categories created during EFD closure will show the rules that created them, along with their empty category daughters. If an edge represents a lexical item, its creator is `lexicon`. In the case of empty categories not created by EFD closure and all lexical items, there are no daughters.

The status of the mini-interpreter has no effect on compilation. The same code is used regardless of whether the mini-interpreter is active or inactive. The mini-interpreter only has an effect on the run-time commands `rec/1,2,4,5` and `drec/1`.

When the mini-interpreter is active, `rec/1` operates in one of two modes: query-mode or go-mode. When the mini-interpreter is active, `rec/1` always begins in *query-mode*. In query-mode, the user is prompted just before any edge is added. Because ALE parses from right to left, the edges are encountered in that order. The prompt consists of a display of the feature structure for the edge, followed by the mini-interpreter information for that edge, followed by an *action-line*, which lists the options available to the user. For example (from the HPSG grammar included in the ALE distribution):

```
| ?- rec([kim,sees,sandy]).

STRING:
0 kim 1 sees 2 sandy 3

word
QRETR list_quant
QSTORE e_set
SYNSEM synsem
  LOC loc
    CAT cat
      HEAD noun
        CASE case
        MOD none
        PRD bool
        MARKING unmarked
        SUBCAT e_list
      CONT nom_obj
        INDEX [0] ref
        GEN gend
        NUM sing
        PER third
      RESTR e_set

Edge created for category above:
  from: 2 to: 3
  string: sandy
  rule: lexicon
  # of dtrs: 0

Action(add,noadd,go(-#),break,dtr-#,abort)?
|:
```

We see, in this example, the main action-line for `rec`. If the user selects `add`, the edge is added, and `rec` proceeds, in query-mode, as usual. If `noadd` is selected, the edge is not added, and `rec` proceeds in query-mode. In every ALE action-line,

the first option is the one that ALE would have chosen if the mini-interpreter were disabled.

`go` puts the mini-interpreter into *go-mode*. In *go-mode*, `rec` proceeds as it would if the mini-interpreter were inactive, or to think of it another way, it functions as if the user always chose the first option on every action-line, but it does not stop to ask. As it adds the edges, it displays them, along with their mini-interpreter information. `go` suffixed with a number, e.g., `go-1`, puts the mini-interpreter into *go-mode* until it encounters an edge whose left node is that number, and then, beginning with that edge, automatically switches back into *query-mode*. With ALE's current parsing strategy, `go-N` will remain in *go-mode* until it encounters the first edge corresponding to the  $(N+1)$ st lexical item in the string being parsed.

`break` simply invokes the Prolog `break` command, placing the user into a Prolog interpreter with a new `break-level`. Edges that have been added so far can be examined and retracted at this time. When the user pops out of the break, the current prompt is redisplayed.

`dtr-N` displays the  $N$ th daughter, its mini-interpreter information, and the action-line for `dtr`:

```
Action(retract,dtr-#,parent,abort)?
|:
```

`retract` removes the displayed edge (in this case, the daughter) from the chart. When the parse continues, ALE grammar rules will not be able to use that edge. The current edge (the parent of this daughter), however, can still be added. `dtr-N` has the same effect as in the `rec` action-line. `parent` returns to the current edge's parent and its action-line (either `rec` or `dtr`).

The mini-interpreter will not display any edge that has already been retracted. Note that if edges are retracted, there may be gaps in the sequence of chart-edge indices.

If `abort` is selected, the parse is aborted. All of the edges added so far remain in memory until the next `rec` statement. The edge that was displayed when `abort` was chosen is discarded.

Mini-interpreter information is also available through the run-time commands, `edge/2` and `edge/1`, e.g.:

```
| ?- edge(2,4).
```

```
COMPLETED CATEGORIES SPANNING: every toy
```

```
cat
QSTORE ne_list_quant
      HD every
      RESTR toy
          ARG1 [0] individual
          SCOPE proposition
          VAR [0]
      TL e_list
SYNSEM basic
```



```

SEM [0]
SYN np

Edge created for category above:
  index: 20
  from: 2 to: 4
  string: every toy
  rule: np_det_nbar
  # of dtrs: 2
Action(retract,dtr-#,continue,abort)?
| :
```

Every edge that is actually asserted into the chart is assigned a unique number, called an *index*, which `edge/2` displays also. `retract` and `dtr` behave the same as in the `dtr` action-line. `continue` tells the mini-interpreter that the user is done traversing the parse tree rooted at the current edge, and to find more.

`edge/1` works exactly as `edge/2` does, except that its input is the unique index number that ALE stores with every edge.

## 2.9 Subsumption Checking (parsing only)

ALE can perform subsumption checking on edges during parsing. By default, it does not. To enable it, use the command:

```

| ?- subtest.

edge subsumption checking active

yes
| ?-
```

To turn it, off, use `nosubtest`.

When subsumption checking is enabled, ALE will only add a new feature structure to the chart between nodes  $n$  and  $m$  if there is no other edge currently spanning  $n$  and  $m$  whose feature structure subsumes the new one. If, instead, the new feature structure subsumes an existing one's, then the existing edge is retracted and replaced with the new one.

Note that if edges are retracted, there may be gaps in the sequence of chart-edge indices.

Extra compiled code is required in order to make subsumption checking more efficient. If subsumption checking is enabled when a grammar is compiled, this code will be compiled also. If it is disabled, the subsumption checking code will be compiled when the `subtest` command is given. Only in the former case, however, will subsumption checking be used during EFD closure to reduce the number of empty categories to consider at run-time. The empty categories and phrase structure rules must be recompiled (with `compile_rules`) if subsumption checking is enabled after the initial compilation of a grammar for empty categories to be tested for subsumption.

Our experience has been that subsumption checking is not required in most unification-based grammars, and should therefore be left disabled. It is useful only for those grammars which have true spurious ambiguity or redundancy. Grammars that incorporate some notion of thematic or functional structure for representing the meaning of a sentence normally realise structural ambiguities as semantic ambiguities that should be retained in the chart.

If both subsumption checking and the mini-interpreter are enabled, then the user may override either of these behaviours. In the former case, before the new feature structure is discarded, it will be displayed along with the discard action-line:

```
Action(noadd,continue,break,dtr-#,existing,abort)?
|:
```

`continue` instructs the parser to look for more subsuming edges. If no more are found, the new feature structure is added to the chart. `existing` displays the existing chart edge that subsumes the new one with the `edge/2` action-line. The rest behave as described above.

In the latter case, before an existing edge is retracted, its feature structure will be displayed along with the retract action-line:

```
Action(retract,continue,break,dtr-#,incoming,abort)?
|:
```

`retract` retracts the existing, subsumed edge and asserts the incoming feature structure. `incoming` displays the incoming feature structure, along with the incoming action-line:

```
Action(noadd,dtr-#,existing,abort)?
|:
```

the functions of whose options have already been described.



```

:- nomrs.                % no output (default)

:- fs.                   % switches feature structure output on (default)
:- nofs.                 % switches feature structure output off

```

The macros that are given as arguments to `root_symbol`, `imp_symbol`, `decl_symbol`, and `que_symbol` have to be specified in your grammar.

You may customize the chart display yourself or use one of the `dot.chartdisplay` files supplied in `trale_home(chart_display)`. They should be moved to `~/chartdisplay`.

If you put the following line in your `.emacs`, the prompt of the command `go` will be recognized by SICStus Prolog and you can type `Ctrl-C Ctrl-P` to get to the previous input and parse it (or call it) again.

```
(setq sicstus-prompt-regexp ">>> *\\|| [ ?][- ] *")
```

If your cursor is in the line of a previously parsed utterance, you may simply hit `return` and the sentence is parsed again.

### T2.1.2 Working with the Chart Display

Typing `go.` brings you to an interactive mode. You can type in a sentence and you will get parsing results displayed either with `grisu` or to `stdout` depending on whether you use `grisu` (strongly recommended). If not open the chart display will pop up (after a parse) and the chart will be displayed. The left mouse button gives you actions you can apply to the edge you point to and the right mouse button gives you actions you can perform on the whole chart. The rule names at the left hand side are also clickable. Empty elements are not clickable yet.

While in the interactive mode, you can execute simple commands directly, provided the command does not correspond to a lexical item in your grammar. For instance reloading of a grammar can be done in the interactive mode by typing `c.` However, if your grammar contains the letter `c` as a lexical object, it will be parsed rather than executed.

`fs` and `nofs` only affect the feature structure output for sentences parsed in the interactive mode. Parses initiated with `rec` directly are not affected by this switch.

### Debugging

You can use the chart display for debugging rule applications: First select a rule by clicking at the rule and choosing the menu item 'Select rule'. Then select an edge from the chart. All goals that are at the first position of the rule (specified with `goal>` in the rule) will be executed. Then the selected edge will be unified with the first daughter of the rule. If the unification succeeds and the blocked constraints that are attached to the rule and to the edge are satisfied, the next goals specified in the rule will be executed. If this succeeds it is checked whether there are further daughters. If this is not the case we have a passive edge which is shown to the user. If there are further daughters an active edge is stored. This active edge can be displayed by clicking somewhere in the display and pressing the right mouse

button. In the following other passive edges from the chart can be combined with the active edge. If the combination succeeds, the rule and the successfully combined daughters are marked green. Otherwise the edge that could not be combined is marked red. In case of a failure another passive edge can be tested against the active edge.

If you select another rule, the active edge is deleted.

The debugging of constraints that fire due to instantiations during a unification is difficult. To enable the debugging of single constraints a flag is set if the `chartdisplay` is used for debugging. For instance if you want to check whether a certain constraint fires, you may include debug information in the constraint. In the following example the flag `chart_debug/1` is tested and if its value is `on`, a debug message is printed.

```
undelayed_accusative([E1|T1]) if
  prolog((chart_debug(on) ->
        write(user_error,'Trying to assign accusative.'),
        nl(user_error)
        ;true
        )),
  assign_acc(E1),
  accusative(T1).
```

Instead of printing a message you can also call the debugger or do other things. Since this constraint may be called during the lexicon compilation as well it would be very difficult to debug without the flag, since you would enter debug modus thousand times before you loaded the grammar completely.

### Generation

Chart edges can be used as input for generation. You have to specify the path in your feature structures that yields the semantic information. (`SYNSEM|LOC|CONT` is the predefined path.) The semantic contribution of a selected chart edge will be taken as input for generation and all generation results will be displayed.

## 2.10 Source-Level Debugger

ALE also provides an XEmacs-based source-level debugger. This can only be used for parsing or definite clause resolution, and only with SICStus 3.8.6 or higher, and XEmacs 20.3 or higher. In future releases, a debugger with a more restricted functionality will be made available for users of SWI Prolog.

The ALE source-level debugger is implemented on top of the SICStus source-level debugger. The debugger provided with ALE has the complete functionality of the SICStus source-level debugger with ordinary Prolog programs; so you only need this one if you will need to debug both. SICStus debugger commands that are not explicitly mentioned in this section are not supported in ALE debugging.

ALE source code must occur in a single file in order to be debugged. To debug Prolog source code, please refer to the SICStus documentation. For both ALE and Prolog debugging, the prolog flag, `source_info`, needs to be turned on, using the command:

```
| ?- prolog_flag(source_info,_,on).
```

The SICStus XEmacs interface should do this automatically.

When a Prolog hook is encountered while debugging an ALE grammar, the SICStus debugger is automatically invoked. The hook will be embedded in a Prolog `call/1` statement. If the `leap` option of the SICStus debugger is used, the leap ends at the end of the hook — ALE will creep when it resumes control.

To install the ALE source-level debugger, follow the directions in the distribution file, `debugger/INSTALL`.

### 2.10.1 Running without XEmacs

To run the debugger without XEmacs, simply run SICStus Prolog from the directory with `debugger.pl` and type:

```
| ?- compile(debugger).
```

This assumes that `ale.pl` and the `debugger` subdirectory are located in the same directory. If ALE.PL has not been loaded already, this will compile `ale.pl` as well. There will be a warning message about buffer-mode when it loads, which can be ignored. Then type:

```
| ?- noemacs.
```

You can add a `noemacs/0` directive to the end of `debugger.pl` to do this automatically.

### 2.10.2 Running with XEmacs

To run the debugger with XEmacs, you must run SICStus Prolog and ALE within XEmacs as an inferior process. To do this:

1. set the `EPROLOG` environment variable to the command that runs SICStus Prolog in the shell that you will run XEmacs from (this is not necessary if the command is 'sicstus'),

2. run XEmacs from the directory with `debugger.pl`,
3. load the file to be debugged in Prolog major mode,
4. Use the XEmacs command, `M-x run-prolog`, to run SICStus prolog as an inferior process,
5. From the SICStus Prolog prompt, type:

```
| ?- compile(debugger).
```

There is also a command:

```
| ?- emacs.
```

that will turn on the XEmacs interface, if it has been disabled by `noemacs/0`.

### 2.10.3 Debugger Commands

There are seven basic commands in the ALE debugger, four of which are variants of normal ALE commands. These four, `dcompile_gram/1`, `drec/1`, `dquery/1` and `dgen/1`, are the debugger variants of `compile_gram/1`, `rec/1`, `query/1` and `gen/1`, respectively. The other three, `dleash/1`, `dskip/1`, and `dclear_bps/0`, will be explained below.

Currently, the ALE source-level debugger can only debug grammars down to the level of feature structure unification, i.e., feature structure unification is treated as an atomic operation. Constraints and procedural attachments on types using `cons` and `goal` cannot be debugged either, nor can inequation enforcement, edge subsumption checking, or extensionalization. These will be possible in a future version. For now, `dcompile_gram/1` compiles a grammar in exactly the same way that ALE normally does to the point where it can be debugged. It then reopens the grammar file and uses the token-stream directly to index those parts of the grammar source code that it can debug by line number. `dcompile_gram/1` also requires that all of the ALE source code for a grammar be in a single file, i.e., unlike ALE without the debugger, you cannot load other auxiliary files from within a grammar file. This restriction will also be lifted in a future version.

`dcompile_gram/1` also allows for tracing through the EFD closure algorithm, again down to the level of feature structure unification.

After a grammar has been compiled for debugging with `dcompile_gram/1`, `drec/1`, `dquery/1` and `dgen/1` can be used for parsing, definite clause resolution or generation, respectively, with the debugger. `dquery/1` works exactly as `query/1` does: it first finds most general satisfiers of the arguments and then searches for a solution with Prolog-style SLD-resolution using the clauses provided by the user. `drec/1` and `dgen/1` work exactly as `rec/1` and `gen/1` do, except that lexical rules are not compiled out, so that they can be debugged. Also remember that ALE parses right-to-left with EFD-closed rules and empty categories, and that it uses a semantic-head-driven generator. With parsing, the debugger can also be used in conjunction with the chart mini-interpreter described above for extra control of the chart. With generation, the debugger successively shows the construction of the pivot template, pivot matching, pivot checking and the linking of the pivot

with the root. Generation of non-semantic-head daughters is performed recursively.

#### 2.10.4 Debugger Ports and Steps

The ALE debugger is loosely based on the procedural box model of execution that many Prolog debuggers use. There are four kinds of ports, *call*, *exit*, *redo*, and *fail*. The ALE debugger does not support exception ports. A call port occurs at every initial invocation of a step that ALE takes in parsing or definite clause resolution, a list of which is given below. An exit port occurs at the successful completion of such a step; a redo port occurs when a subsequent step has failed and ALE backtracks into the current step to find more solutions, and a fail port occurs when a step has failed to produce any or any more solutions.

Consider, for example, what happens when ALE applies the following description to a feature structure that occurs in the lexical entry for the word, `foo`

```
foo --->(a
        ;f:b),
        g:c.
```

The first port we encounter is when ALE tries to add the type `a`. This is a call port:

```
Call: add type, a, to lex entry?
```

If this succeeds, then an exit port occurs:

```
Exit: add type, a, to lex entry?
```

and processing moves on to `g:c`. If this fails, then we must backtrack through adding `a` for more solutions (for example, if there is a disjunctive constraint on that type):

```
Redo: add type, a, to lex entry?
```

If there is another solution, then another exit port occurs. Otherwise, the next port is fail port:

```
Fail: add type, a, to lex entry?
```

and processing continues with the other disjunct, `f:b`. Certain steps in ALE, notably the depth-first rule application of the chart parser, are failure-driven loops. To indicate that these “failures” are actually a normal part of execution, they are displayed as, e.g.:

```
Finished:close chart edge under rule application?
```

Enforcing descriptions at a feature also counts as a step:

```
Call: enforce description on f value of lex entry?
```

If we agree to this, then the next step would be to add the type `b`

```
Call: add type, b, to value at f?
```



Step	Kind	Example Message
Adding types	desc	add type, a, to lex entry
Adding a_/1 atoms	desc	add atom, foo(X), to lex entry
Feature selection	desc	enforce description on f value of lex entry
Adding path equations	desc	path equate [f,g] with [f,h]
Adding inequations	desc	inequate ...with description
Unification (from shared variables)	desc	unify ...with <i>Var</i>
Macro substitution	desc	substitute macro description for np/1
Functional description evaluation	desc	evaluate functional description, append/2

Figure B.1: Description-level Steps

and so on.

The steps in a description whose ports the ALE debugger keeps track of are given in Figure B.1: There are other kinds of steps besides description-level ones, that pertain to ALE's built-in control for parsing and definite clause resolution. These steps, along with their *kind*, are given in the table in Figure B.2. The steps for generation are given in Figure B.3. Almost all of these involve sub-steps that enforce descriptions. Some, such as chart-edge closure, involve other sub-steps such as rule selection. Lexical rule application includes input and output morph application, as well as morph condition (given in a *when* clause) application. The one generation step of kind, *lr*, takes place when the pivot template is matched against a (base or derived) lexical entry. Unlike the compiled generator, the debugger does not compile lexical rules into the non-chain-rule selection step, so that the user can see their application. Instead, a base lexical entry is first selected without reference to the pivot template, then zero or more lexical rules are applied bottom up, until a derived entry is eventually unified with the pivot template. Lexical rules are, thus, treated as a special kind of unary chain rule. The length of these special chains is still controlled by `lex_rule_depth/1`, not `chain_length/1`.

### 2.10.5 Leashing

With so many steps, and four possible ports, stepping through an entire parse in a large grammar would be a very trying experience. The ALE debugger provides three basic ways to filter through the steps to find points of interest in a parse or definite clause query.

The first is leashing. Leashing allows the user to distinguish at which steps information is simply displayed and at which steps the debugger stops and asks the user what to do. Unlike the SICStus debugger, leashing in the ALE debugger is a property of steps, not ports. The command to control leashing is `dleash/1`. The argument to `dleash/1` consists of a sign, + or -, plus the kind of step. A + sign indicates that the debugger should stop and ask what to do at steps of that kind; and a - sign indicates that it should simply display the port and proceed. For

Step	Kind	Example Message
Empty category derivation	empty	empty category
EFD Closure	empty	close empty categories under rules
EFD matching	empty	apply daughter 1 of rule sentence1 to empty category
Lexical entry derivation	lex	derive seen from base entry: see
Lexical rule application	lr	apply lexical rule, passive, to see
Input morph application	lr	apply morph to input: see
Output morph application	lr	apply morph to input: see, output: seen
Morph condition application	lr	apply morph condition
Functional description clause selection	fun	evaluate functional clause for append/2
Definite clause selection	rel	evaluate relational clause for head.feature_principle/2
Definite clause resolution	rel	resolve goal, append/3
Negated goal (\+) resolution	rel	resolve negated goal
Shallow cut (-> ;) execution	rel	execute shallow cut
Extensional identity (@=) check	rel	resolve extensional identity
Prolog hook call	rel	resolve prolog hook: (num(N),write(N))
Closing new chart edge under rules as leftmost daughter	rule	close chart edge under rule application
Rule selection	rule	apply rule, schema1.

Figure B.2: Parsing and Definite Clause Resolution Steps

example, to turn leashing off for empty categories, type:

```
| ?- dleash(-empty).
```

```
yes
```

There is also a special kind, `all`, that allows the user to turn on and off leashing on all kinds of steps at once. The default leashing at start-up is `+all`.

If a kind of step is leashed, then the debugger will stop at every port for every step of that kind, and ask what to do. The possible responses are given in Figure B.4: Not all responses are available at all ports. The kind of port (call, fail, etc.) is what determines the possible responses. The responses, `?` and `h` are always available, and list the other legal responses at the current port.

### 2.10.6 Skipping

Even leashing may not be enough for very large parses or queries because of the sheer number of ports displayed. The ALE debugger also provides a facility for *auto-skipping*. Whereas turning leashing off at a kind of step is like au-

Step	Kind	Example Message
Build semantic index	gen	build semantic index from root
Build pivot template	gen	build pivot template from index
Find pivot	gen	find pivot
Match pivot (lexical entry)	gen	match pivot against entry derived from see
Match pivot (empty category)	gen	match pivot against empty category
Match pivot (mother of non-chain rule)	gen	match pivot against mother of non-chain rule, sentence1
Apply chain rule	gen	apply chain rule, s, to, pivot
Pivot check	gen	check for link from pivot to root
Generate from pivot	gen	recursively generate from pivot
Generate pre-head daughters	gen	recursively generate pre-head daughters from pivot
Generate post-head daughters	gen	recursively generate post-head daughters from pivot
Connect pivot to root	gen	connect pivot to root
Connect chain node to root	gen	connect new chain node to root
Unify pivot template with lexical entry	lr	unify pivot template and see
Unify chain node	gen	unify mother of chain rule, s, with root

Figure B.3: Generation Steps

tomatically answering `c` (advance to next port) at those steps, auto-skipping is like automatically answering `s`, which advances to the next exit or fail port of the current step without stopping at or even displaying the ports in between. The command for this is `dskip/1`, and its argument is of the same form as the argument to `dleash/1`. The signs have a different meaning, of course. For example, `dskip(+empty)` means that you want the debugger to auto-skip steps of kind `empty`, i.e., not stop and ask, whereas `dleash(+empty)` means that you want to leash steps of kind `empty`, i.e., stop and ask. When a step where auto-skipping is set is encountered, it is displayed with an automatic reply without stopping, e.g.:

```
Call: empty category? <auto-skip>
Exit: empty category? <auto-skipped>
Edge added: Number:0, Left:1, Right:1, Rule:empty
Call: close chart edge under rule application?
```

### 2.10.7 Breakpoints

The final kind of filtering is the breakpoint. In the ALE debugger, breakpoints are a property of lines in a grammar source file, not steps or ports. For a finer grain of resolution, it would be necessary to give each potential breakable step its own line in the input. By setting breakpoints and then using the `1` response, the debugger will advance to the next step whose line has a breakpoint without displaying any

Input	Description	Ports
?	show available commands at current port	c,e,r,f
h	same as ?	
a	abort processing	c,e,r,f
f	fail at this step (go to fail port)	c,e,r
r	retry this step (go from fail to call port)	f
c	advance to next port	c,e,r,f
LF	same as c	
s	advance to next exit/fail port of this step	c,r
n	advance to first port on new line of grammar file	c,e,r,f
l	advance to next breakpoint	c,e,r,f
+	set breakpoint at current line	c,e,r,f
-	clear breakpoint at current line	c,e,r,f
@ <i>PrologGoal</i>	pass a goal to Prolog	c,e,r,f
i	toggle chart mini-interpreter	c,e,r,f
d	display current structure	c,e,r,f

Figure B.4: Possible Responses at Debugger Ports

steps in between. If that step is not leashed or has auto-skipping set, the debugger acts accordingly after displaying it.

There are currently two ways to set a breakpoint. One is to use the + response from within the debugger at a step at whose line you wish to set a breakpoint. The other is only available when the debugger is used with an installation of XEmacs that supports XPM resources. In this case, when a source file is compiled a small glyph will be displayed at the left edge of every breakable line. Clicking on this glyph once with the left mouse button sets a breakpoint. Clicking again clears it. A breakpoint can also be cleared with the - response.

It is often the case that a line will have several breakable steps on it, for example, feature paths:

```
synsem:local:cat:head:verb,
qretr:e_list
```

If a breakpoint were set at the first line, then leaping from the call port for SYNSEM would still advance to the call port for LOCAL:

```
Call: enforce description on synsem value of lex entry? l
Call: enforce description on local value of value at synsem?
```

To avoid this, the response n is provided, for leaping automatically to the first port on a different line in the source file. The combination of n and l can be used to leap more effectively in files that pack many steps, particularly description steps, into one line.

All breakpoints can be cleared at once using the command, `dclear_bps/0`.

# Appendix C

## [incr TSDB0]

There is code that generates output for the profiling and test suite management tool [incr TSDB0] (<http://www.coli.uni-sb.de/itsdb/>) developed by Stephan Oepen. If you want to implement larger grammar fragments it is recommended to use this tool. It provides the following functionality (and much more):

- storing time, number of passive edges, memory requirements, errors for every parsed item
- compare test runs (performance, coverage, overgeneration)
- detailed comparison on item basis of
  - number of readings/edges
  - the derivations (i.e. tree structure with rule names)
  - the MRSes
- test parts of test suites on a phenomenon basis or some other selection from your test items (restrictions may be formulated in SQL queries)

The [incr TSDB0] uses the Parallel Virtual Machine. If you have several CPUs idle you may distribute the processing of your test suite over several machines which enormously shortens the time needed for testing and speeds up grammar development.

### T3.1 Installing [incr TSDB0]

Get the itsdb package from <http://lingo.stanford.edu/ftp/> and install it. How to do this is described in the documentation which can be found also at this site.

#### T3.1.1 Setting up Shell Variables

Set the SICStus path. For instance, if you use tcsh, put the following in your .tcshrc:

```
setenv SP_PATH /usr/local/lib/sicstus-3.9.1/
```

[incr TSDB()] is called via foreign functions. The file is linked to libtsdb.so which has to be in the LD\_LIBRARY\_PATH.

Put something like the following in your ~/.tcshrc (or the respective file for the shell you are using):

```
setenv LD_LIBRARY_PATH /home/stefan/Lisp/lkb/lib/linux/
```

### T3.1.2 Setting up .tsdbrc

Create a file ~/.tsdbrc to include something like the following:

```
(setf *pvm-cpus*
  (list
    (make-cpu
      :host "laptop1"
      :spawn "/home/stefan/bin/trale"
      :options '("-s" "-c" "/home/stefan/Prolog/Trale/Bale/theory"
                "-e" "load_tsdb,itsdb:connect_tsdb")
      :class :bale :threshold 2)
    )
  )
)
```

In the options line you give a path to the grammar that should be loaded.

The item following :class is an identifier. You may have several calls to make-cpu, for instance if you want to use different machines or if you want to load different grammars.

If you want to use two cpus in parallel, you specify two cpus with the same class name:

```
(setf *pvm-cpus*
  (list
    (make-cpu
      :host "server1"
      :spawn "/home/stefan/bin/trale"
      :options '("-s" "-c" "/home/stefan/Prolog/Trale/Bale/theory"
                "-e" "load_tsdb,itsdb:connect_tsdb")
      :class :bale1 :threshold 2)
    (make-cpu
      :host "server2"
      :spawn "/home/stefan/bin/trale"
      :options '("-s" "-c" "/home/stefan/Prolog/Trale/Bale/theory"
                "-e" "load_tsdb,itsdb:connect_tsdb")
      :class :bale1 :threshold 2)
    )
  )
)
```

### T3.1.3 Setting up PVM

Furthermore you have to set up the parallel virtual machine (pvm): Create the file ~/.pvm\_hosts containing something like the following:

```

#
# list machines accessible to PVM; option fields are (see pvmd(8))
#
# - dx: path to 'pvmd3' executable (on remote host);
# - ep: colon-separated PATH used by pvmd(8) to locate executables;
# - wd: working directory for remote pvmd(8);
# - ip: alternate (or normalized) name to use in host lookup;
#
# sp=VALUE Specifies the relative computational speed of this host
# compared to other hosts in the configuration. VALUE is an inte
# ger in the range [1 - 1000000]

laptop1 dx=/home/stefan/Lisp/lkb/bin/linux/ \
        ep=/home/stefan/Lisp/lkb/lib/linux/ wd=/tmp sp=1004
server1 dx=/home/stefan/Lisp/lkb/bin/linux/ \
        ep=/home/stefan/Lisp/lkb/lib/linux/ wd=/tmp sp=1004
server2 dx=/home/stefan/Lisp/lkb/bin/linux/ \
        ep=/home/stefan/Lisp/lkb/lib/linux/ wd=/tmp sp=1004

```

The binaries and the man pages of pvm are part of the [incr TSDB0] distribution.

If you work on one host only, you will probably not have problems with PVM. If you want to work with several machines in parallel, chances are high that PVM does not work without debugging.

You can start the pvmd3 that comes with [incr TSDB0] by hand and then call pvm to check if the pvm demon initialized the other machines properly. pvm provides a console. If you have one host and type conf, you get:

```

pvmd already running.
pvm> conf
1 host, 1 data format
          HOST      DTID      ARCH      SPEED      DSIG
          laptop1  40000    LINUX     1000  0x00408841
pvm>

```

If you configured several hosts in .pvm\_hosts, you should see them here. If you do not, something went wrong in the initialization. If your PVM uses rsh (default) then you have to make sure that rexec is enabled on all hosts. You can test this by typing:

```
rsh server2 who
```

If this does not work, you have to consult your system administrator.

### T3.1.4 Initializing a CPU

Having loaded [incr TSDB0] you can initialize a CPU by typing:

```
(tsdb::tsdb :cpus :bale)
```

Where :bale is the identifier you have chosen for your CPU.

This should load the grammar and come back saying something like:

```
wait-for-clients(): 'laptop1' registered as tid <262179>.
```

After successful registration of a client (i.e. after loading your grammar) and after the creation of a test suite with the [incr TSDB()] podium you can process items for instance by using Process|All Items in the [incr TSDB()] podium. Processing can be interrupted by typing CTRL-G at the podium.

Initializing CPUs with (tsdb::tsdb :cpus :bale) redirects the output to the /tmp/ directory. This is reasonable if you work with several cpus. If you want to see the output you may use:

```
(tsdb::tsdb :cpus :bale :file t)
```

Looking at the load messages is also useful if you see something like:

```
wait-for-clients(): client exit for 'laptop1' <262150>
```

Please refer to the [incr TSDB()] manual (<http://www.coli.uni-sb.de/itsdb/publications/index>) for a description of example sessions and further documentation.

### T3.1.5 Creating a Test Suite

You can put together your own test suite by using the [incr TSDB()] import function (File|Import|Test Items). This function imports data from an ASCII text file. For example:

```
;;; intransitive
Karl schlft.      ;; this is a very simple example
Der Mann schlft.
;;; transitive
Liebt Karl Maria?
Karl liebt Maria.
;;; np
der Mann
der klug Mann
der Mann, der ihn kennt
;;; pronoun
Er schlft
Er kennt ihn.
;; subjless
Mich drstet.
;; particle_verbs
Karl denkt nach.
Karl denkt ber Maria nach.
*Karl nachdenkt ber Maria.
;; unaccusatives
Er flt mir auf.
;;; perfect
Er hat geschlafen.
Du wirst schlafen.
Du wirst geschlafen haben.
```



```

;;; free_rc
Wer schlft, stirbt.
Wen ich kenne, begre ich.
Was er kennt, it er.
Wo ich arbeite, schlafe ich.
ber was ich nachdenke, hast du nachgedacht.
Ich liebe, ber was du nachdenkst.
ber was du nachdenkst, gefllt mir.
;;; case
*Liebt ihn ihn?

```

Ungrammatical sentences are marked with a star. Everything that follows two ‘;’ is treated as a comment. The phenomenon is given on a separate line starting with three ‘;’. If you want [incr TSDB()] to display statistics on a phenomenon-based basis, you have to make [incr TSDB()] know these phenomena. This can be done by specifying them in the .tsdbrc file:

```

(setf *phenomena*
  (append
    (list "intransitive"
          "transitive"
          "np"
          "pronoun"
          "perfect"
          "free_rc"
          "case")
      *phenomena*))

```

The data base you created is in your data base root directory. If you want to use it as a general skeleton, you have to move it to the skeleton directory and update the Index.lisp file appropriately. You can find out about the location of these directories under Options in the [incr TSDB()] podium.

The following should be specified in your grammar file (theory.pl):

```

grammar_version('Trale GerGram 0.3').
root_symbol(@root).      % symbol for input that does not end with punctuation
imp_symbol(@imp).        % symbol for input that ends with '!'
decl_symbol(@decl).      % symbol for input that ends with '.'
que_symbol(@que).        % symbol for input that ends with '?'

```

The grammar version will be shown in the run relation. The root symbol is used for parsing. The macros that are given as arguments to root\_symbol, imp\_symbol, decl\_symbol, and que\_symbol have to be specified in your grammar.

If you use CVS for version control you may consider the following entry for your grammar version.

```

grammar_version('Trale GerGram $Revision: 1.1 $').

```

When you check in a new version the revision number is automatically updated by CVS.

The following specification can optionally be given in your theory.pl:

T101-6

*APPENDIX C. [INCR TSDB()]*

```
% before doing a parse all lexical descriptions  
% given in retract_before_parsing/1 are removed  
% this is used here to decrease the chart size  
% zero inflected elements do not interfere.  
retract_before_parsing(stem).
```

## Appendix D

# Pretty-printing Hooks

This section is intended for more advanced audiences who are proficient in Prolog and ALE .

ALE uses a data structure that is not so easily readable without pretty-printing or access predicates. In order to make pretty-printing more customizable, hooks are provided to portray feature structures and inequations. If these hooks succeed, the pretty-printer assumes that the structure/inequation has been printed and quietly succeeds. If the hooks fail, the pretty-printer will print the structure/inequation by the default method used in previous versions of ALE . The hooks are called with every pretty-printing call to a substructure of a given feature structure. It is, therefore, important that the user's hooks use the arguments provided to mark visited substructures if the feature structure being printed is potentially cyclic, or else pretty-printing may not terminate.

### 4.0.1 Portraying feature structures

The hook for portraying feature structures is:

```
portray_fs (Type, FS, KeyedFeats, VisIn, VisOut, TagsIn, TagsOut, Col,
           HDIn, HDOut)
```

FS is the feature structure to be printed. This is ALE's internal representation of this structure<sup>1</sup>. It is recommended that access to information in this structure be obtained by Type and KeyedFeats although the brave of heart may wish to work with it directly. FS is also used to check token identity with structures in the Vis and Tags trees, as described below. Type is the type of FS. KeyedFeats is a list of fval/3 triples:

```
[fval(Feat_1, Val_1, Restr_1), ..., fval(Feat_n, Val_n, Restr_n)]
```

where n is the number of appropriate features to Type. FS's value at Feat\_i is Val\_i, and the appropriate value restriction of Type at Feat\_i is Restr\_i.

VisIn, VisOut, TagsIn and TagsOut are AVL trees. They can be manipulated using the access predicates found in the library(assoc) module of SICStus Prolog. VisIn is a tree of the nodes visited so far in the current printing call, and

---

<sup>1</sup>The reader is referred to the ALE User's Guide for the structure of this representation.

`TagsIn` is a tree of the nodes with re-entrancies in the structure(s) currently being printed (of which `FS` may just be a small piece). Each node in an AVL tree has a key, used for searching, and a value. In both `Vis` and `Tags` trees, the key is a feature structure such as `FS`. For example, the call:

```
get_assoc(FS,VisIn,FSVal)
```

determines whether `FS` has been visited before. In the `Vis` tree, the value (`FSVal` in the above example) at a given key is not used by the default pretty-printer. The user may change them to anything desired. When the default pretty-printer adds a node to the `Vis` tree, it adds the current `FS` with a fresh unbound variable as the value.

In the `Tags` tree, the value at key `FS` is the numeric tag that the default pretty-printer would print in square brackets to indicate structure-sharing at that location. The user may change this value (using `get_assoc/5` or `put_assoc/3`), and the default pretty-printer will use that (with a `write/1` call) instead.

A hook must return a `TagsOut` and `VisOut` tree to the pretty-printer if it succeeds. At a minimum, this means binding `TagsOut` to `TagsIn` and `VisOut` to `VisIn`. If the structure being traversed is potentially cyclic, `VisOut` should, in general, update `VisOut` to record that the current node has been visited to avoid an infinite traversal.

`Col` is the number of columns that have already been indented before the hook was called. This is useful for formatting. `HDIn` and `HDOut` are hook data. They are not used by the default pretty-printer, and are simply passed around for the convenience of hooks to record information to pass to their next invocation. The initial top-level call to a `portray_fs` hook contains 0 (zero) as the value of `HDIn`. Thus, `HDIn` can also be used to distinguish the first call from subsequent calls provided that the 0 is replaced with something else in recursive calls.

The file `pphooks.pl` (discussed in subsection 4.0.3 below) shows a simple printing hook to produce output very much as the default pretty-printer would.

When a `portray_fs` hook prints one of `FS`'s feature values, it typically will call the pretty-printer rather than attempt to manipulate the data structure directly. This callback is provided by:

```
print_fs(VarType,FS,VisIn,VisOut,TagsIn,TagsOut,Col,HDIn,HDOut)
```

Note that the type and feature values of `FS` do not need to be supplied—those will be filled in by the system before control is passed to `portray_fs` or the default pretty-printer.

`VarType` is currently not used. The other positions are the same as in `portray_fs`.

## 4.0.2 Portraying inequations

The hook for inequations is:

```
portray_ineq(FS1,FS2,IqsIn,IqsOut,TagsIn,TagsOut,VisIn,VisOut,Col,
             HDIn,HDOut).
```

This is called when there is an inequation between FS1 and FS2 with IqsIn being the remaining inequations. The hook should return the remainder to consider printing next in IqsOut, which is normally bound to IqsIn. IqsIn can also be used to test whether FS1= $\backslash$ =FS2 is the last inequation.

Col is the number of columns that the default pretty-printer *would* print before printing FS1. This is different from the use of Col in portray\_fs where it is the number of columns *already* printed. The other arguments are the same as in portray\_fs.

Inequations are typically printed after all feature structures and their substructures in a top-level call to the pretty-printer have been printed. Likewise, portray\_ineq is only called once portray\_fs has been called on all feature structures and their substructures in a top-level call. Typically, the arguments to inequations will thus have been visited before—the only exceptions to this are inequated extensionally typed feature structures.

### 4.0.3 A sample pretty-printing hook

The following Prolog code shows how the default pretty-printer can be written as a hook. This code is available online on the ALE web site under the name pphooks.pl.

The file has two top-level predicates: portray\_fs/10 and portray\_ineq/11. As mentioned above, the former is responsible for printing feature structures and the latter, for printing inequations.

The first thing that portray\_fs/10 does is check whether the feature structure it is printing is tagged, i.e., whether it is structure-shared with another feature structure. This check is made using get\_assoc(FS,TagsIn,Tag). If so, the tag is printed between square brackets. Then using get\_assoc(FS,VisIn,\_), the system determines whether the feature structure has already been printed. Should this be the case, VisOut, TagsOut and HDOut are respectively bound with VisIn, TagsIn and HDIn, and the hook succeeds with nothing else printed.

If, on the other hand, the feature structure has not been printed already, and FS is a variable, then either Type, or in case Type has at least one appropriate feature, mgsat(Type), is printed. Then, put\_assoc/4 is used to update VisOut to include FS, and TagsOut and HDOut are bound to TagsIn and HDIn, respectively.

In case the feature structure is not a variable, its type is written, and a call to a recursive predicate is made to print each feature-value pair of FS in turn. Two important things to note are that (1) VisOut is updated to include FS *before* the call is made to avoid non-termination on cyclic structures, and (2) the feature values are actually printed with a callback to print\_fs/9 (which in turn calls portray\_fs/10 again).

The predicate portray\_ineq/11 works similarly. Note that Col spaces are added by the hook itself, unlike portray\_fs/10.

The source code of pphooks.pl is given below:

```
% pphooks.pl
```

```
% default printer written as hook (w/o type or feat hiding or
```

```

% FS expansion)

portray_fs(Type,FS,KeyedFeats,VisIn,VisOut,TagsIn,TagsOut,Col,
          HDIn,HDOut):-

% print Tag if shared
( get_assoc(FS,TagsIn,Tag)
-> write('['),write(Tag),write('] ')
; true),

% print structure if not yet visited
( get_assoc(FS,VisIn,_)
-> VisOut = VisIn,
   TagsOut = TagsIn,
   HDOut = HDIn      % already printed

% variable - use Type to print
; var(FS) -> ( approp(_,Type,_)
              -> write('mgsat('),write(Type),write(')')
              ; write(Type)
              ),
              put_assoc(FS,VisIn,_,VisOut),
              TagsOut = TagsIn,
              HDOut = HDIn

              % otherwise print Type and recurse
              ; write(Type),
              put_assoc(FS,VisIn,_,VisMid),
              print_feats(KeyedFeats,VisMid,VisOut,TagsIn,
                          TagsOut,Col,HDIn,HDOut)
).

print_feats([],Vis,Vis,Tags,Tags,_Col,HD,HD).
print_feats([fval(F,Val,Restr)|Feats],VisIn,VisOut,TagsIn,
            TagsOut,Col,HDIn,HDOut) :-
  nl,tab(Col),
  write_feature(F,LengthF), % capitalise, print and count
                        % characters
  NewCol is Col + LengthF + 1,
  print_fs(Restr,Val,VisIn,VisMid,TagsIn,TagsMid,NewCol,HDIn,
          HDMid),
  print_feats(Feats,VisMid,VisOut,TagsMid,TagsOut,Col,HDMid,
            HDOut).

portray_ineq(FS1,FS2,Rest,Rest,VisIn,VisOut,TagsIn,TagsOut,
            Col,HDIn,HDOut):-

```

```
tab(Col),
print_fs(bot,FS1,VisIn,VisMid,TagsIn,TagsMid,Col,HDIn,
         HDMid),
nl,write(' =\\= '), NewCol is Col + 7,
print_fs(bot,FS2,VisMid,VisOut,TagsMid,TagsOut,NewCol,
         HDMid,HDOut).
```

# Appendix E

## TRALE Output

*W. Detmar Meurers, Ohio State University*<sup>1</sup>

### T5.1 Saving of outputs

It can be useful to be able to save the output of a command like `mgsat`, `rec`, `lex` etc, e.g. in order to

- output it again without reperforming the actual task
- view it pretty printed using different pretty printers (e.g. `grisu` or `text`)
- run a diff of the two structures (see below)
- `save_results_on`. saves copies of the output for later use (during the same session)
- `save_results_off`. switches saving results off. What was saved is preserved.
- `show_saved(+<nr>)`. shows saved result number `nr`
- `show_all_saved`. shows all saved results
- `saved_id(-<nr>)`. returns number of saved results
- `save_results(+<filename>)`. save results in a file
- `load_results(+<filename>)`. loads results from a file

At system startup, saving of results is off.

---

<sup>1</sup>©2003, W. Detmar Meurers



## T5.2 Grisu interface support

- `grisu_off`. switches grisu output off and textual pretty printer on
- `grisu_on`. switches grisu output on, if the system had been started with `grisu`
- `grisu_debug`. sends the output normally sent to grisu to standard out instead
- `grisu_noddebug`. reverts to sending grisu output to the socket at which grisu listens

At system startup, grisu is on if trale has been started with `-g`

### T5.2.1 Unfilling

When running a grammar with a large signature (e.g., MERGE), you'll immediately notice how essential proper unfilling of uninformative features is. The code now assumes a node to be informative if

- it is of a type that is more specific than the appropriate value of the feature that led to it. For the top level, the appropriate value is taken to be `bot`.
- or
- it is structure shared with some other node
- or
- the value of one of its features is informative according to a), b) or c)

- `unfill_on`. switches on unfilling of uninformative nodes in the output
- `unfill_off`. switches it off

At system startup, unfill is on.

### T5.2.2 Feature ordering

To specify the order in which features are displayed by the pretty printer, include any number of statements of the form

- `f <<< g`. meaning: f will be ordered before g.

and at most one each of the following statements:

- `<<< h`. meaning: h will be ordered last.
- `>>> i`. meaning: i will be ordered first.

Currently only the grisu interface takes feature ordering into account, but this should be added to the default textual pretty printer at some point.

### T5.2.3 Diff of feature structures

- `diff(NrA,NrB)`.
- `diff(NrA,PathA,NrB,PathB)`.
  - NrA and NrB are numbers of saved results
  - PathA and PathB are paths of the form `f:g:h:i` or `[f,g,h,i]`. The empty path for both cases is `[]`.

Output is provided via the grisu interface.

## Appendix F

# ALE Keyword Summary

The following is a summary of keywords discussed in this manual, along with page references. A table of auxiliary keywords, those that only occur as arguments of other keyword operators, such as the `cat>` argument of a `rule`, will be provided in a future version.

A keyword of kind *Description* is one that occurs in an ALE description of a feature structure. One of kind *Def. Clause*, or *DCL*, is one that occurs in ALE's definite clause language. One of kind *Signature* is a declaration that occurs in an ALE signature. One of kind *Type* is an ALE type with special properties. One of kind *ALE* is a Prolog query (entered at the `| ?-` prompt) that can be used after ALE has been loaded (see p. 4). One of kind *Mini-interpreter* is a mini-interpreter command that appears in an interpreter action-line. One of kind *Debugger* is a Prolog query that can be used after the source-level debugger has been loaded. For debugger responses, the reader is referred to the table on page 101.

Keyword	Kind	Description	Page
,	Desc./DCL	Conjunction	23, 42
--->	Signature	Declare lexical entry.	47
-> (;)	Def. Clause	Shallow cut.	42
:	Description	Feature value.	23
;	Desc./DCL	Disjunction.	23, 42
!	Def. Clause	Cut.	42
\+	Def. Clause	Negation-by-failure.	42
==	Description	Path Equation.	23
=@	Def. Clause	Predefined token-identity definite clause predicate.	42
=\ =	Description	Inequation.	23
@	Description	Macro instantiation.	29
[. . .]	Description	Predefined list macro.	31
%	Prolog	Comment delimiter.	7

Keyword	Kind	Description	Page
a_	Description/Signature	Built-in extensional atom.	21
abort	Mini-interpreter	Abort parse.	92
add	Mini-interpreter	Add the current edge.	91
approp	ALE	Show value restriction on a feature at a type.	73
assert	Prolog	Add clause to Prolog database.	42
bot	Type	In an ALE signature, this type must appear, and must subsume all of the other types.	6
break	Mini-interpreter	Invoke Prolog break.	92
chain_length	ALE	Set limit on chain rule sequence length.	63
compile_gram	ALE	Compile ALE signature (or parts of it — see table, p. 66).	66
compile	Prolog	Compile a Prolog file.	4
cons	Signature	Declare type constraint.	33, 43
consult	Prolog	Load Prolog file (such as an ALE signature) into database.	4
continue	Mini-interpreter	Proceed to look for more (subsuming/subsumed) edges.	93,94
control-c	Prolog	Prolog interrupt.	5
control-z	Unix	Unix interrupt.	5
dclear_bps	Debugger	Clear all breakpoints in current grammar file.	101
dcompile_gram	Debugger	Compile grammar file for source-level debugging.	96
dgen	Debugger	Generate with source-level debugger.	96
dleash	Debugger	Set or remove leashing on a kind of step.	98
dquery	Debugger	Evaluate a definite clause with source-level debugging.	96
drec	Debugger	Parse with source-level debugger.	96
dskip	Debugger	Set or remove auto-skipping on a kind of step.	100
dtr- <i>N</i>	Mini-interpreter	Display <i>N</i> th daughter edge of current edge.	92
edge	ALE	Show a chart edge.	84
emacs	Debugger	Turn on XEmacs interface to source-level debugger.	96
empty	ALE	Show empty categories.	80
empty	Signature	Declare empty category.	49
export_words	ALE	Send list of words in lexicon to stream	80
existing	Mini-interpreter	Display edge that subsumes new feature structure.	94
ext	Signature	Declare extensional types.	20
feature	ALE	Test if feature exists.	72
generate	ALE	Tell compiler to produce code for generation only.	66

Keyword	Kind	Description	Page
gen	ALE	Generate a string using the compiled generator.	87,89
go	Mini-interpreter	Add current and all subsequent edges.	92
go- <i>N</i>	Mini-interpreter	Add current and all subsequent edges until node <i>N</i> is reached.	92
halt	Prolog	Exit from Prolog.	5
if	Def. Clause	Definite clause language equivalent of :-.	40
incoming	Mini-interpreter	Display incoming edge that subsumes existing edge.	94
interp	ALE	Turn on mini-interpreter.	90
intro	Signature	Declare appropriate features for type.	14
introduce	ALE	Test if a feature was introduced by a type.	72
iso_desc/2	ALE	Test whether two descriptions evaluate to the same feature structure.	75
lex	ALE	Show lexical entry.	79
lex_compile	ALE	Compile lexicon intermediate code (SICStus only)	68
lex_consult	ALE	Dynamically consult lexicon intermediate code	68
lex_rule	ALE	Show lexical rule.	81
lex_rule	Signature	Declare lexical rule.	51
lex_rule_depth	ALE	Set bound on lexical rule application.	51
list	Type	This type, along with types <code>e_list</code> and <code>ne_list</code> , and features <code>HD</code> and <code>TL</code> , must be defined in an ALE signature in order to use the predefined [...] macro in descriptions, or the <code>cats&gt; list</code> -argument operator in grammatical rules.	31, 59
macro	ALE	Show macro definition.	81
macro	Signature	Declare macro.	28
mgsat	ALE	Find most general satisfier(s) of a type.	73
no_write_feat	ALE	Hide a feature and its value.	76
no_write_type	ALE	Hide a type.	76
noadd	Mini-interpreter	Do not add the current edge.	91
noemacs	Debugger	Turn off XEmacs interface to source-level debugger.	95
nointerp	ALE	Turn off mini-interpreter.	90
nosubtest	ALE	Disable edge subsumption checking.	93
parent	Mini-interpreter	Return to parent edge.	92
parse	ALE	Tell compiler to produce code for parsing only.	66

Keyword	Kind	Description	Page
parse_and_gen	ALE	Tell compiler to produce code for parsing and generation.	66
prolog	Def. Clause	Definite clause hook to Prolog.	42
query	ALE	Evaluate a definite clause.	77
rec	ALE	Parse a string.	82,84–86
rec_best	ALE	Parse first parsable string in a list of strings	87
rec_list	ALE	Parse a list of strings	86,87
retract	Mini-interpreter	Retract currently displayed edge.	92
retract	Prolog	Remove clause from Prolog database.	42
retractall_lex	ALE	Retract all of a word's entries from lexicon	68
retract_lex	ALE	Retract a word's entry from lexicon	68
rule	ALE	Show grammatical rule.	80
rule	Signature	Declare grammatical rule.	55
semantics	ALE	Declares a semantics definite clause predicate.	62
show_clause	ALE	Show a definite clause.	77
show_cons	ALE	Show constraint for a type.	72
show_type	ALE	Show subtypes, supertypes, constraint and most general satisfiers for a type.	73
sub	Signature	Declare subtyping relationship.	6
subtest	ALE	Enable edge subsumption checking, and, if necessary, compile code for it.	93
sub_type	ALE	Test subsumption between two types.	72
true	Def. Clause	Definite clause that is always satisfied (also used to construct ground clauses in def. clause language).	41
type	ALE	Test if type exists.	71
unify_type	ALE	Unify two types.	72
update_lex	ALE	Add new entries to lexicon	68
when	Def. Clause	Co-routining	44
write_feat	ALE	Don't hide a feature.	76
write_feats	ALE	Don't hide any features.	76
write_type	ALE	Don't hide a type.	76
write_types	ALE	Don't hide any types.	76

## **Appendix G**

# **Flag Summary**

TO BE ADDED

# **HDRUG: A Graphical User Environment for Natural Language Processing in Prolog**

Hdrug is an environment to develop logic grammars / parsers / generators for natural languages. The package is written in Sicstus Prolog version 3 and uses library(tcltk) to implement its user interface. Tcl/Tk is a powerful script language to develop applications for the X-windows environment.

Hdrug offers various tools to visualize lexical entries, grammar rules, definite-clause definitions, parse trees, feature structures, lexical- rule- and type-hierarchies, graphs of the comparison of different parsers on a corpus of test sentences etc., in a Tk widget, LaTeX/DVI format, and the Clig system.

The package comes with a number of example grammars, including the grammars to be found in the distribution of the ALE system.

Hdrug allows for easy comparison of different parsers/generators; it has extensive possibilities to compile feature equations into Prolog terms; it can produce graphical (Tk), and ordinary Prolog output of trees, feature structures, Prolog terms (and combinations thereof), plotted graphs of statistical information, and tables of statistical information. Etc. Etc.

Using just menu's and buttons it is possible to parse sentences, generate sentences from logical form representations, view the parse trees that are derived by the parser or generator, change a particular version of the parser on the fly, compare the results of parsing the same sentence(s) with a set of different parsers, etc.

HDRUG was designed and implemented by Gertjan van Noord. The HDRUG home-page is <http://www.let.rug.nl/~vannoord/Hdrug/>.



# **Pleuk Grammar Development Environment**

For those using SICStus 2.1#9 under X windows, the Pleuk grammar development shell has been adapted for ALE. Pleuk provides a graphical user interface, facilities for maintaining and testing corpora, and an interactive, incremental derivation checker. Pleuk is available free of charge from:

`ftp.cogsci.ed.ac.uk:/pub/pleuk`

The file README contains instructions for downloading the system. Pleuk has been ported to Sun SPARC's SunOS 4.\* and HP-UX. For more information, send email to `pleuk@cogsci.ed.ac.uk`. Pleuk was developed by Jo Calder and Chris Brew of the Human Communication Research Centre at the University of Edinburgh, Kevin Humphreys of the Centre for Cognitive Science at the University of Edinburgh, and Mike Reape, of the Computer Science Department, Trinity College, Dublin.

As of this release, Pleuk will not work under SICStus 3.0 or later.

# Appendix H

## References

This collection of references only scratches the surface of the relevant literature. A much more complete survey of the historical perspective on typed unification grammars and programs can be found in Carpenter (1992), and in subsequent papers in *ACL*, *EACL*, *COLING*, etc.

Ait-Kaci, H. (1991). *The WAM: A (Real) Tutorial*. MIT Press, Cambridge, Massachusetts.

The best available introduction to Prolog compiler technology, focusing on Warren's Abstract Machine for Prolog.

Ait-Kaci, H. (1986a). An algebraic semantics approach to the effective resolution of type equations. *Theoretical Computer Science*, 45:293–351.

Seminal work in sorted feature structures, based on Ait-Kaci's 1984 University of Pennsylvania dissertation. Focuses on general constraint resolution.

Ait-Kaci, H., and Nasr, R. (1986b). LOGIN: A logical programming language with built-in inheritance. *Journal of Logic Programming*, 3:187–215.

The first application of feature structures to logic programming. Includes sorted, but not typed feature structures. Also includes good details on the Martelli and Montanari (1984) unification algorithm applied to feature structures.

Ait-Kaci, H. (1984). *A Lattice-Theoretic Approach to Computation based on a Calculus of Partially Ordered Type Structures*. Univ. of Pennsylvania dissertation.

Ait-Kaci's introduction of sorted  $\psi$ -terms, which are like our feature structures only without the appropriateness conditions, inequations and extensionality. An appendix contains a coding of the Zebra Puzzle, a benchmark logic puzzle for constraint resolution.

Carpenter, B. (1992) *The Logic of Typed Feature Structures*. Cambridge Tracts in Theoretical Computer Science 32, Cambridge University Press, New York.

Contains all the theoretical details behind the ALE feature structures, description language and applications. A must for fully understanding ALE and a number of related variations.

- Carpenter, B. and G. Penn (1996) *Compiling Typed Attribute-Value Logic Grammars*. In H. Bunt and M. Tomita, eds., *Recent Advances in Parsing Technology*. Kluwer.  
 A description of the theoretical underpinnings of ALE 2.0, including the data structures, type inference mechanism, description resolution, parsing, and inequation solving.
- Colmerauer, A. (1987). Theoretical model of prolog II. In van Canegham, M., and Warren, D. H., editors, *Logic Programming and its Application*, 1–31. Ablex, Norwood, New Jersey.  
 Describes unification with cyclic terms and inequations in a logic programming environment.
- Gazdar, G., and Mellish, C. S. (1989). *Natural Language Processing in Prolog*. Addison-Wesley, Reading, Massachusetts.  
 An introduction to computational linguistics using Prolog. Also contains a very general introduction to simple PATR-II phrase structure grammars, including simple implementations of unification and parsing algorithms. A version is also available using Lisp.
- Höfeld, M., and Smolka, G. (1988). Definite relations over constraint languages. LILOG-REPORT 53, IBM – Deutschland GmbH, Stuttgart.  
 Highly theoretical description of a constraint logic programming paradigm, including an application to feature structures similar to those used in LOGIN.
- Jaffar, J. (1984). Efficient unification over infinite terms. *New Generation Computing*, 2:207–219.  
 Unification algorithm for possibly cyclic terms in Prolog II. Includes quasi-linear complexity analysis.
- Kasper, R. T., and Rounds, W. C. (1990). The logic of unification in grammar. *Linguistics and Philosophy*, 13(1):35–58.  
 The details of Kasper and Rounds original feature structure description system and related theorems.
- Martelli, A., and Montanari, U. (1982). An efficient unification algorithm. *ACM Transactions on Programming Languages and Systems*, 4(2):258–282.  
 The Union/Find unification algorithm used by ALE, which was adapted to the cyclic case by Jaffar (1984).
- Mastroianni, M. (1993) *Attribute-Logic Phonology*. Carnegie Mellon University Laboratory for Computational Linguistics Technical Report CMU-LCL-93-4. Pittsburgh.  
 The description and motivation for an attribute-logic approach to phonology. Includes extensive discussion of its implementation in ALE, including syllable structure and morphologically conditioned effects such as epenthesis, harmony and assimilation.
- O’Keefe, R. A. (1990) *The Craft of Prolog*. MIT Press, Cambridge, Massachusetts.

Best text on advanced programming techniques using Prolog compilers. Should read Sterling and Shapiro's introduction as a pre-requisite.

Penn, G. (1993). A Utility for Typed Feature Structure-based Grammatical Theories. Technical Report. Laboratory for Computational Linguistics, Carnegie Mellon University, Pittsburgh.

This project served as the basis of the version 2.0 updates of ALE. The report details these updates, including the algorithms used to implement them and other efficiency issues. It also describes Penn's implementation of Head-Driven Phrase Structure Grammar (HPSG), as represented in the first eight chapters of (Pollard and Sag 1994).

Penn, G. (1999). A Parsing Algorithm to Reduce Copying in Prolog. Arbeitspapier des Sonderforschungsbereichs 340, Nr. 137.

A presentation of the Empty-First-Daughter closure algorithm, which can be used to reduce copying in Prolog-based parsers.

Penn, G., and Carpenter, B. (1993). Three Sources of Disjunction in a Typed Feature Structure-based Resolution System. *Feature Formalisms and Linguistic Ambiguity*, H. Trost ed. Ellis Horwood, New York.

A presentation of the principal sources of complexity in solving constraint puzzles, such as the Zebra Puzzle, a simplified version of which is presented in this manual; and an outline of steps taken to cope with them in the reversible general constraint resolver which was the precursor to ALE.

Penn, G. and Popescu, O. (1997). Head-Driven Generation and Indexing in ALE. *Proceedings of Workshop on Computational Environments for Grammar Development and Linguistic Engineering (ENVGRAM)*, 35th ACL / 8th EACL.

Describes the implementation of ALE's head-driven generator, and a simple indexing strategy for lexical entries during generation.

Popescu, O. (1996). Head-Driven Generation for Typed Feature Structures. Carnegie Mellon University MS Project.

The extension of semantic-head-driven generation to typed feature structures used in ALE.

Pereira, F. C. N., and Shieber, S. M. (1987). *Prolog and Natural-Language Analysis*. Volume 10 of *CSLI Lecture Notes*. Center for the Study of Language and Information, Stanford.

Excellent introduction to the use of term unification grammars in natural language. Includes a survey of Prolog, parsing algorithms and many sample grammar applications in syntax and semantics.

Pollard, C. J. (in press). Sorts in unification-based grammar and what they mean. In Pinkal, M., and Gregor, B., editors, *Unification in Natural Language Analysis*. MIT Press, Cambridge, Massachusetts.

Contains the original extension of Rounds and Kasper's logical language to sorts. Also motivates the use of sorts in natural language grammars.

Pollard, C. J., and Sag, I. A. (1994). *Head-driven Phrase Structure Grammar*. Chicago University Press, Chicago.

The primary grammar formalism which motivated the construction of the ALE system. Provides many examples of how typed feature structures and their descriptions are employed in a sophisticated natural language application.

Shieber, S. M. (1986). *An Introduction to Unification-Based Approaches to Grammar*. Volume 4 of *CSLI Lecture Notes*. Center for the Study of Language and Information, Stanford.

Best source for getting acquainted with the application of feature structures and their descriptions to natural language grammars.

Shieber, S. M., Uszkoreit, H., Pereira, F. C. N., Robinson, J., and Tyson, M. (1983). The formalism and implementation of PATR-II. In *Research on Interactive Acquisition and Use of Knowledge*. Volume 1894 of *SRI Final Report*, SRI International, Menlo Park, California.

Original document describing the PATR-II formalism.

Shieber, S. M., Pereira, C. N., van Noord, G., Moore, R. C. (1990). Semantic-Head-Driven Generation. In *Computational Linguistics*, Vol. 16(1):30–42.

The main reference for a description of the semantic- head-driven generation algorithm.

Smolka, G. (1988a). A feature logic with subsorts. LILOG-REPORT 33, IBM – Deutschland GmbH, Stuttgart.

An alternative logic to that of Rounds and Kasper, which includes sorts, variables and general negation.

Smolka, G. (1988b). Logic programming with polymorphically order-sorted types. LILOG-REPORT 55, IBM – Deutschland GmbH, Stuttgart.

An application of ordered term unification to typed logic programming.

Sterling, L., and Shapiro, E. Y. (1986). *The Art of Prolog: Advanced Programming Techniques*. MIT Press, Cambridge, Massachusetts.

Best general introduction to logic programming in Prolog.

van Noord, G. (1989). *BUG: A Directed Bottom-Up Generator for Unification Based Formalisms*. Working Papers in Natural Language Processing, Katholieke Universiteit Leuven, Stichting Taaltechnologie Utrecht.

Proposes the first semantic-head-driven generation algorithm.

# Appendix A

## Sample Grammars

### 1.1 English Syllabification Grammar

```
% Signature
% =====

bot sub [unit,list,segment].
  unit sub [cluster,syllable,word]
    intro [first:segment,
           last:segment].
  cluster sub [consonant_cluster, vowel_cluster]
    intro [segments:list_segment].
  consonant_cluster sub [onset,coda].
    onset sub [].
    coda sub [].
  vowel_cluster sub [].
  syllable sub []
    intro [syllable:list_segment].
  word sub []
    intro [syllables:list_list_segment].
  segment sub [consonant,vowel].
    consonant sub [sibilant,obstruent,nasal,liquid,glide].
      sibilant sub [s,z].
        s sub [].
        z sub [].
      obstruent sub [p,t,k,b,d,g].
        p sub [].
        t sub [].
        k sub [].
        b sub [].
        d sub [].
        g sub [].
      nasal sub [n,m].
        n sub [].
        m sub [].
```

```

liquid sub [l,r].
  l sub [].
  r sub [].
glide sub [y,w].
  y sub [].
  w sub [].
vowel sub [a,e,i,o,u].
  a sub [].
  e sub [].
  i sub [].
  o sub [].
  u sub [].
list sub [e_list,ne_list,list_segment,list_list_segment].
  e_list sub [].
  ne_list sub [ne_list_segment,ne_list_list_segment]
    intro [hd:bot,
           tl:list].
list_segment sub [e_list,ne_list_segment].
  ne_list_segment sub []
    intro [hd:segment,
           tl:list_segment].
list_list_segment sub [e_list,ne_list_list_segment].
  ne_list_list_segment sub []
    intro [hd:list_segment,
           tl:list_list_segment].

% Rules
% =====

word_schema_rec rule
(word,
 syllables:[Syllable|Syllables],
 first:First1,
 last>Last2)
==>
cat> (syllable,
      syllable:Syllable,
      first:First1,
      last>Last1),
cat> (word,
      syllables:Syllables,
      first:First2,
      last>Last2),
goal> (\+ less_sonorous>Last1,First2)).

word_schema_base rule
(word,

```

```

    syllables:[Syllable],
    first:First,
    last>Last)
===>
cat> (syllable,
      syllable:Syllable,
      first:First,
      last>Last).

v_syllable rule
(syllable,
 syllable:[Vowel],
 first:Vowel,
 last:Vowel)
===>
cat> (vowel,Vowel).

vc_syllable rule
(syllable,
 syllable:[Vowel|Segs1],
 first:Vowel,
 last>Last)
===>
cat> (vowel,Vowel),
cat> (coda,
      segments:Segs1,
      last>Last).

cv_syllable rule
(syllable,
 syllable:Segs,
 first:First,
 last:Vowel)
===>
cat> (onset,
      segments:Segs1,
      first:First),
cat> (vowel,Vowel),
goal> append(Segs1,[Vowel],Segs).

cvc_syllable rule
(syllable,
 syllable:Segs,
 first:First,
 last>Last)
===>
cat> (onset,
      segments:Segs1,
```



```

        first:First),
cat> (vowel,Vowel),
cat> (coda,
        segments:Segs2,
        last:Last),
goal> append(Segs1,[Vowel|Segs2],Segs).

consonant_cluster_base rule
(consonant_cluster,
  segments:[Consonant],
  first:Consonant,
  last:Consonant)
===>
cat> (consonant,Consonant).

onset rule
(onset,
  segments:[Consonant1|Consonants],
  first:Consonant1,
  last:Consonant3)
===>
cat> (consonant,Consonant1),
cat> (onset,
  segments:Consonants,
  first:Consonant2,
  last:Consonant3),
goal> less_sonorous(Consonant1,Consonant2).

coda rule
(coda,
  segments:[Consonant1|Consonants],
  first:Consonant1,
  last:Consonant3)
===>
cat> (consonant,Consonant1),
cat> (coda,
  segments:Consonants,
  first:Consonant2,
  last:Consonant3),
goal> less_sonorous(Consonant2,Consonant1).

% Lexicon
% =====

p ---> p.
t ---> t.
k ---> k.

```

```

b ----> b.
d ----> d.
g ----> g.
s ----> s.
z ----> z.
n ----> n.
m ----> m.
l ----> l.
r ----> r.
y ----> y.
w ----> w.
a ----> a.
e ----> e.
i ----> i.
o ----> o.
u ----> u.

```

```

% Definite Clauses
% =====

```

```

less_sonorous_basic(sibilant,obstruent) if true.
less_sonorous_basic(obstruent,nasal) if true.
less_sonorous_basic(nasal,liquid) if true.
less_sonorous_basic(liquid,glide) if true.
less_sonorous_basic(glide,vowel) if true.

```

```

less_sonorous(L1,L2) if
  less_sonorous_basic(L1,L2).
less_sonorous(L1,L2) if
  less_sonorous_basic(L1,L3),
  less_sonorous(L3,L2).

```

```

append([],Xs,Xs) if true.
append([X|Xs],Ys,[X|Zs]) if
  append(Xs,Ys,Zs).

```

## 1.2 Categorical Grammar with Cooper Storage

```

% Signature
% =====

bot sub [cat,synsem,syn,sem_obj,list_quant].
  cat sub []
    intro [synsem:synsem,
           qstore:list_quant].
  synsem sub [functional, basic].
    functional sub [forward,backward]
      intro [arg:synsem,
             res:synsem].
    forward sub [].
    backward sub [].
  basic sub []
    intro [syn:syn, sem:sem_obj].
  syn sub [np,s,n].
    np sub [].
    s sub [].
    n sub [].
  sem_obj sub [individual, proposition, property].
    individual sub [j,m].
      j sub [].
      m sub [].
    property sub []
      intro [ind:individual,
             body:proposition].
  proposition sub [logical,quant,run,hit,nominal].
    logical sub [and,or].
      and sub []
        intro [conj1:proposition,
               conj2:proposition].
      or sub []
        intro [disj1:proposition,
               disj2:proposition].
    quant sub [every,some]
      intro [var:individual,
             restr:proposition,
             scope:proposition].
      every sub [].
      some sub [].
    run sub []
      intro [runner:individual].
    hit sub []
      intro [hitter:individual,
             hittee:individual].
    nominal sub [kid,toy,big,red]

```

```

        intro [arg1:individual].
    kid sub [].
    toy sub [].
    big sub [].
    red sub [].
list_quant sub [e_list, ne_list_quant].
    e_list sub [].
    ne_list_quant sub []
        intro [hd:quant,
            tl:list_quant].

% Lexicon
% =====

kid --->
    @ cn(kid).

toy --->
    @ cn(toy).

big --->
    @ adj(big).

red --->
    @ adj(red).

every --->
    @ gdet(every).

some --->
    @ gdet(some).

john --->
    @ pn(j).

runs --->
    @ iv((run,runner:Ind),Ind).

hits --->
    @ tv(hit).

% Grammar
% =====

forward_application rule

```

```

(synsem:Z,
 qstore:Qs)
===>
cat> (synsem:(forward,
      arg:Y,
      res:Z),
      qstore:Qs1),
cat> (synsem:Y,
      qstore:Qs2),
goal> append(Qs1,Qs2,Qs).

```

```

backward_application rule
(synsem:Z,
 qstore:Qs)
===>
cat> (synsem:Y,
      qstore:Qs1),
cat> (synsem:(backward,
      arg:Y,
      res:Z),
      qstore:Qs2),
goal> append(Qs1,Qs2,Qs).

```

```

s_quantifier rule
(synsem:(syn:s,
        sem:(Q,
             scope:Phi)),
 qstore:QsRest)
===>
cat> (synsem:(syn:s,
             sem:Phi),
      qstore:Qs),
goal> select(Qs,Q,QsRest).

```

```

% Macros
% =====

```

```

cn(Pred) macro
  synsem:(syn:n,
          sem:(body:(Pred,
                    arg1:X),
              ind:X)),
  @ quantifier_free.

```

```

gdet(Quant) macro
  synsem:(forward,
    arg: @ n(Restr,Ind),
    res: @ np(Ind)),
  qstore:[@ quant(Quant,Ind,Restr)].

```

```

quant(Quant,Ind,Restr) macro
  (Quant,
  var:Ind,
  restr:Restr).

```

```

adj(Rel) macro
  synsem:(forward,
    arg: @ n(Restr,Ind),
    res: @ n((and,
      conj1:Restr,
      conj2:(Rel,
        arg1:Ind)),
      Ind)),
  @ quantifier_free.

```

```

n(Restr,Ind) macro
  syn:n,
  sem:(body:Restr,
  ind:Ind).

```

```

np(Ind) macro
  syn:np,
  sem:Ind.

```

```

pn(Name) macro
  synsem: @ np(Name),
  @ quantifier_free.

```

```

iv(Sem,Arg) macro
  synsem:(backward,
    arg: @ np(Arg),
    res:(syn:s,
      sem:Sem)),
  @ quantifier_free.

```

```

tv(Rel) macro
  synsem:(forward,
    arg:(syn:np,
      sem:Y),
    res:(backward,
      arg:(syn:np,
        sem:X),

```

```
        res:(syn:s,
            sem:(Rel,
                hitter:X,
                hittee:Y))))),
    @ quantifier_free.

quantifier_free macro
  qstore: [].

% Definite Clauses
% =====

append( [], Xs, Xs) if
  true.
append( [X|Xs], Ys, [X|Zs]) if
  append(Xs, Ys, Zs).

select( [Q|Qs], Q, Qs) if
  true.
select( [Q1|Qs1], Q, [Q1|Qs2]) if
  select(Qs1, Q, Qs2).
```

### 1.3 Simple Generation Grammar

```

% An implementation in {\sc ale} of the grammar in Shieber & al,
% "Semantic-Head-Driven Generation", CL 16-1, 1990.

% Signature
% =====

bot sub [pred, list, sem, form, agr, sign].
  pred sub [decl, imp, love, call_up, leave, see, john, mary, mark,
            friends, often, friend, up, you, i].
  decl sub []. imp sub [].
  leave sub []. love sub []. call_up sub []. see sub [].
  john sub []. mary sub []. mark sub [].
  friends sub []. friend sub [].
  often sub []. up sub [].
  you sub []. i sub [].
list sub [e_list, ne_list, arg_list, subcat_list].
  e_list sub [].
  ne_list sub [arg_ne_list, subcat_ne_list]
    intro [hd:bot, tl:list].
  arg_list sub [e_list, arg_ne_list].
    arg_ne_list sub [] intro [hd:sem, tl:arg_list].
  subcat_list sub [e_list, subcat_ne_list].
    subcat_ne_list sub [] intro [hd:sign, tl:subcat_list].
sem sub [] intro [pred:pred, args:arg_list].
form sub [finite, nonfinite].
  finite sub [].
  nonfinite sub [].
agr sub [sg1, sg2, sg3, pl1, pl2, pl3].
  sg1 sub []. sg2 sub []. sg3 sub [].
  pl1 sub []. pl2 sub []. pl3 sub [].
sign sub [sentence, verbal, np, adv, p]
  intro [sem:sem].
  sentence sub [].
  verbal sub [s, vp] intro [form:form].
    s sub [].
    vp sub [] intro [subcat:subcat_list].
  np sub [det, n]
    intro [agr:agr, arg:sem].
    det sub [] intro [np_sem:sem].
    n sub [].
  adv sub [] intro [varg:sem].
  p sub [].
ext([sg1,sg2,sg3,pl1,pl2,pl3]).

% Lexicon
% =====

```



```

love --->
  vp, form:nonfinite,
  subcat: [(np,sem:Obj), (np,sem:Subj)],
  sem: (pred:love, args: [Subj,Obj]).

call --->
  vp, form:nonfinite,
  subcat: [(np,sem:Obj), (p,sem:(pred:up,args: [])), (np,sem:Subj)],
  sem: (pred:call_up, args: [Subj,Obj]).

call --->
  vp, form:nonfinite,
  subcat: [(p,sem:(pred:up,args: [])), (np,sem:Obj), (np,sem:Subj)],
  sem: (pred:call_up, args: [Subj,Obj]).

leave --->
  vp, form:nonfinite,
  subcat: [(np,sem:Subj)],
  sem: (pred:leave, args: [Subj]).

see --->
  vp, form:nonfinite,
  subcat: [(np,sem:Obj), (np,sem:Subj)],
  sem: (pred:see, args: [Subj,Obj]).

see --->
  vp, form:nonfinite,
  subcat: [(s,form:finite,sem:Obj), (np,sem:Subj)],
  sem: (pred:see, args: [Subj,Obj]).

john --->
  np, agr:sg3, sem: (pred:john, args: []).

mary --->
  np, agr:sg3, sem: (pred:mary, args: []).

mark --->
  np, agr:sg3, sem: (pred:mark, args: []).

friends --->
  np, agr:pl3, sem: (pred:friends, args: []).

friend --->
  n, agr:sg3, arg:X, sem: (pred:friend, args: [X]).

i --->
  np, agr:sg1, sem: (pred:i, args: []).

```

```

you --->
  np, agr:sg2, sem:(pred:you,args:[]).

often --->
  adv, varg:VP, sem:(pred:often,args:[VP]).

up --->
  p, sem:(pred:up,args:[]).

% Lexical Rules
% =====

sg3 lex_rule (vp, form:nonfinite, subcat:Subcat, sem:Sem) **>
  (vp, form:finite, subcat:NewSubcat, sem:Sem)
  if add_sg3(Subcat,NewSubcat)
  morphs (X,y) becomes (X,i,e,s),
  X becomes (X,s).

non_sg3 lex_rule (vp, form:nonfinite, subcat:Subcat, sem:Sem) **>
  (vp, form:finite, subcat:NewSubcat, sem:Sem)
  if add_nong3(Subcat,NewSubcat)
  morphs X becomes X.

% Grammar Rules
% =====

sentence1 rule
  (sentence,sem:(pred:decl,args:[S])) ==>
  cat> (s,form:finite,sem:S).

sentence2 rule
  (sentence,sem:(pred:imp,args:[S])) ==>
  cat> (vp,form:nonfinite,
  subcat:[(np,sem:(pred:you,args:[]))],sem:S).

s rule
  (s,form:Form,sem:S) ==>
  cat> Subj,
  sem_head> (vp,form:Form,subcat:[Subj],sem:S).

vp1 rule
  (vp,form:Form,subcat:Subcat,sem:S) ==>
  sem_head> (vp,form:Form,subcat:[Compl|Subcat],sem:S),
  cat> Compl.

```

```
vp2 rule
  (vp,form:Form,subcat:[Subj],sem:S) ==>
  cat> (vp,form:Form,subcat:[Subj],sem:VP),
  sem_head> (adv,varg:VP,sem:S).

% Semantics Directive
% =====

semantics sem1.

% Definite Clauses
% =====

sem1(sem:S,S) if true.

add_sg3([(np,sem:Sem)],[(np,agr:sg3,sem:Sem)]) if !, true.
add_sg3([Cat|Cats],[Cat|NewCats]) if add_sg3(Cats,NewCats).

add_nong3([(np,sem:Sem)],[(np,agr:(=\sg3),sem:Sem)]) if !, true.
add_nong3([Cat|Cats],[Cat|NewCats]) if add_nong3(Cats,NewCats).
```

## Appendix B

# Error and Warning Messages

### 2.1 Error Messages

`a_/1` atom declared subsumed by type  $T$

Subsumption over `a_/1` atoms has a fixed definition. Subtyping specifications with `a_/1` atoms are not allowed.

`add_to` could not unify  $FS_1$  and  $FS_2$

The unification between feature structures  $FS_1$  and  $FS_2$  failed.

`add_to` could not unify paths  $\pi$  and  $\phi$  in  $FS$

The unification between paths  $\pi$  and  $\phi$  failed in feature structure  $FS$  failed.

`add_to` could not inequante  $FS_1$  and  $FS_2$

The inequation between features structures  $FS_1$  and  $FS_2$  could not be satisfied.

`add_to` could not add feature  $F$  to  $FS$

The value of feature  $F$  is not unifiable with the similar value in feature structure  $FS$ .

`add_to` could not add undefined macro  $M$  to  $FS$

Macro  $M$  used in feature structure  $FS$  is undefined.

`add_to` could not add incompatible type  $T$  to  $FS$

Type  $T$  is not compatible with the type of feature structure  $FS$ .

`add_to` could not add undefined type  $T$  to  $FS$

Type  $T$  in feature structure  $FS$  is undefined.

`add_to` could not add ill formed complex description  $D$  to  $FS$

Description  $D$  is ill formed and could not be unified with feature structure  $FS$

appropriateness cycle following path  $\pi$  from type  $T$

There is a sequence of features  $\pi$  which must be defined for objects of type  $T$  where the value must be of type  $T$ .

bot has appropriate features

The most general type,  $\perp$ , cannot have any appropriate features.

bot has constraints

The most general type,  $\perp$ , must not have cons constraints.

cats> value with sort  $S$  is not a valid list argument

An argument of cats> was detected at run-time, which is not of a type subsumed by list.

consistent  $T_1$  and  $T_2$  have multiple mgus  $T_s$

Types  $T_1$  and  $T_2$  have the non singleton set  $T_s$  as their set of most general unifiers.

constraint declaration given for atom

a\_/1 atoms must not have cons constraints.

description uses unintroduced feature  $F$

A description uses the feature  $F$  which has not been defined as appropriate for any types.

edge/2: arguments must be non-negative

The arguments to edge/2 represent nodes in a parsing chart, and thus must be non-negative integers.

edge/2: first argument must be < second argument

The arguments to edge/2 represent nodes in a parsing chart. Edges only span from one node to an equal or greater valued node. edge/2 shows edges where the other node has a greater value. empty/0 shows edges where the node has an equal value.

extensional type  $E$  is not maximal

Type  $E$  is declared extensional but does not observe the maximality restriction.

feature  $F$  multiply introduced at  $T_s$

The feature  $F$  is introduced at the types in  $Ts$ , which are not comparable with one another.

illegal variable occurrence in  $T$  sub  $Ss$  (intro  $FRs$ )

In subtype/feature specifications, neither  $T$  or any of the types in  $Ss$  or  $FRs$ , or any of the features in  $FRs$  can be unbound variables. If a value restriction is an  $a_/1$  atom, that atom can be unbound, or contain unbound variables, but the  $a_/1$  operator must still appear.

incompatible restrictions on feature  $F$  at type  $T$  are  $Ts$

The inherited restrictions, consisting of types  $Ts$ , on the value of  $F$  at type  $T$  are not consistent.

invalid line  $\phi$  in rule

A line of a grammar rule is neither a goal nor a category description.

lexical rule  $LR$  lacks morphs specification

The obligatory morphs part of lexical rule  $LR$  is missing.

multiple constraint declaration error for  $T$

More than one cons declaration exists for type  $T$ .

multiple feature specifications for type  $T$

The appropriate features of  $T$  can be introduced along with subtyping or by themselves, but there can only be one declaration of appropriate features.

multiple specification for  $F$  in declaration of  $T$

More than one restriction on the value of feature  $F$  is given in the definition of type  $T$ .

no lexical entry for  $W$

Expression  $W$  is used, but has no lexical entry.

pathval: illegal path specified -  $\pi$

Path  $\pi$  is not a valid path specification for the given feature structure.

rule  $R$  has multiple `sem_head` specifications

More than one semantic head declaration was found in grammar rule  $R$ .

rule  $R$  has no `cat>` `cats>` or `sem_head>` specification

The grammar rule named  $R$  is empty in that it does not have any daughter specification.

rule  $R$  has wrongly placed `sem_goal>` specifications

A `sem_goal>` specification occurs somewhere other than immediately before or immediately after a `sem_head>` specification in rule  $R$ .

subtype/feature specification given for `a_/1` atom

Subsumption over `a_/1` atoms has a fixed definition, and they can have no features. Subtype or feature specifications for them are not allowed.

subtyping cycle at  $T$

The subsumption relation specified is not anti-symmetric. It can be inferred that the type  $T$  is a proper subtype of itself.

subtype  $T_1$  used in  $T_2$  undeclared

Undefined type  $T_1$  declared as subtype in definition of  $T_2$ .

$T$  multiply defined

There is more than one definition of type  $T$ .

$T$  subsumes bot

$T$  is declared as subsuming the most general type,  $\perp$ .

$T_1$  used in appropriateness definition of  $T_2$  undeclared

Undefined type  $T_1$  used as value restriction in definition of  $T_2$ .

undefined macro  $M$  used in description

A description uses a macro which is not defined.

undefined type  $T$  used in description

A description uses a type  $T$  which is not defined.

undefined feature  $F$  used in path  $\pi$

A path  $\pi$  of features uses undefined feature  $F$  in a description.

unsatisfiable lexical entry for  $W$

Word  $W$  has a lexical entry which has no satisfying feature structure.

upward closure fails for  $F$  in  $S1$  and  $S2$

$S1$  subsumes  $S2$ , but the value restriction for  $F$  at  $S1$  does not subsume the value restriction for  $F$  at  $S2$ .

## 2.2 Warning Messages

=@ accessible by procedural attachment calls from constraint for  $T$

The built-in =@ predicate (extensional identity check) is non-monotonic, so its use should be avoided in constraints attached to types.

atom a\_/1 *Atom* is ground in declaration of  $T$

One of the appropriate features of  $T$  has a value restriction that is a ground a\_/1 atom. Every feature structure of type  $T$  will have the same value at this feature.

homomorphism condition fails for  $F$  in  $T_1$  and  $T_2$

It is not the case that the appropriateness restriction on the type  $T = T_1 + T_2$  is the unification of the appropriateness restrictions on  $T_1$  and  $T_2$ .

lexical description for  $W$  is unsatisfiable

Incompatibilities in the lexical description for word  $W$  could not produce a satisfying feature structure.

no chain rules found

All the grammar rules in the program were non-chain rules (no semantic heads).

no definite clauses found

There were no definite clause rules specified in the program.

no features introduced

There are no appropriate features for any types.

no functional descriptions found

There are no functional description definitions in the program.

no lexical rules found

There were no lexical rules specified in the program.

no lexicon found

There were no lexical entries specified in the program.

no non\_chain rules found



All the grammar rules in the program were chain rules (with semantic heads).

no  $P$  definite clause found

A definition for definite clause predicate  $P$ , which was declared as the semantics predicate, was not found in the program.

no phrase structure rules found

There were no phrase structure rules specified in the program.

no semantics specification found

There was no specification of the semantics definite clause predicate in the program.

no types defined

There were no sub or intro declarations found in the program.

unary branch from  $T_1$  to  $T_2$

The only subtype of  $T_1$  is  $T_2$ . In this situation, it is usually more efficient to eliminate  $T_1$  if every instance of  $T_1$  is a  $T_2$ .

## Appendix C

# BNF for ALE

```
<desc> ::= <type>
         | <variable>
         | (<feature>:<desc>)
         | (<desc>,<desc>)
         | (<desc>;<desc>)
         | @ <macro_spec>
         | <func_spec>
         | a_ <prolog_term>
         | <path> == <path>
         | =\= <desc>

<type> ::= <prolog_functor>

<feature> ::= <prolog_atom>

<path> ::= list(<feature>)

<macro_def> ::= <macro_head> macro <desc>.

<macro_head> ::= <macro_name>
                | <macro_name>(<seq(var)>)

<macro_spec> ::= <macro_name>
                | <macro_name>(<seq(desc)>)

<func_def> ::= <func_spec> +++> <desc>.

<func_spec> ::= <func_name>
                | <func_name>(<seq(desc)>)

<clause> ::= <literal> if <goal>.

<literal> ::= <pred_sym>
            | <pred_sym>(<seq(desc)>)
```

```

<cut_free_goal> ::= true
                  | <literal>
                  | prolog(<prolog_goal>)
                  | (<cut_free_goal>,<cut_free_goal>)
                  | (<cut_free_goal>;<cut_free_goal>)
                  | (<desc>=@ <desc>)
                  | (<cut_free_goal> -> <cut_free_goal>)
                  | (<cut_free_goal> -> <cut_free_goal>
                    ; <cut_free_goal>)
                  | (\+ <cut_free_goal>)

```

```

<goal> ::= true
         | <literal>
         | prolog(<prolog_goal>)
         | (<goal>,<goal>)
         | (<goal>;<goal>)
         | (<desc>=@ <desc>)
         | (<cut_free_goal> -> <goal>)
         | (<cut_free_goal> -> <goal> ; <goal>)
         | !
         | (\+ <goal>)
         | when(<cond>,<goal>)

```

```

<cond> ::= <variable>^(<cond>)
         | <quantified_cond>

```

```

<quantified_cond> ::= <quantified_cond>, <quantified_cond>
                   | <quantified_cond> ;<quantified_cond>
                   | <variable>=<cond_desc>

```

```

<cond_desc> ::= <variable>
              | <type>
              | max(<type>)
              | <feat>:<cond_desc>
              | <path> == <path>
              | <cond_desc>, <cond_desc>
              | <cond_desc> ;<cond_desc>

```

```

<lex_entry> ::= <word> ---> <desc>.

```

```

<rule> ::= <rule_name> rule <desc> ==> <rule_body>.

```

```

<rule_body> ::= <sem_less_rule_body>
              | <sem_less_rule_body>,<sem_rule_body>,
                <sem_less_rule_body>

```

```

<sem_less_rule_body> ::= <rule_clause>
                        | <rule_clause>, <sem_less_rule_body>

<rule_clause> ::= cat> <desc>
                | cats> <desc>
                | goal> <goal>

<sem_rule_body> ::= sem_head> <desc>
                  | sem_goal> <goal>, sem_head> <desc>
                  | sem_head> <desc>, sem_goal> <goal>
                  | sem_goal> <goal>, sem_head> <desc>,
                    sem_goal> <goal>

<lex_rule> ::= <lex_rule_name> lex_rule <lex_rewrite>
              morphs <morphs>.

<lex_rewrite> ::= <desc> **> <desc>
                | <desc> **> <desc> if <goal>

<morphs> ::= <morph>
            | <morph>, <morphs>

<morph> ::= (<string_pattern>) becomes (<string_pattern>)
           | (<string_pattern>) becomes (<string_pattern>)
             when <prolog_goal>

<string_pattern> ::= <atomic_string_pattern>
                   | <atomic_string_pattern>, <string_pattern>

<atomic_string_pattern> ::= <atom>
                           | <var>
                           | <list(<var_char>>)

<var_char> ::= <char>
             | <var>

<seq(X)> ::= <X>
           | <X>, <seq(X)>

<empty_prod> ::= empty <desc>.

<type_spec> ::= <type> sub <list(<type>>)>.
              | <type> sub <list(<type>>)>
                intro <list(<frestr_spec>>)>.
              | <type> intro <list(<frestr_spec>>)>.

<frestr_spec> ::= <feature>:<type>
                | <feature>: a_ <prolog_term>

```

```
<cons_spec> ::= <type> cons <desc>.  
              | <type> cons <desc>  
                goal <goal>.
```

```
<ext_spec> ::= ext(list(<type>)).
```

```
<prog> ::= <prog_line>  
          | <prog_line> <prog>
```

```
<prog_line> ::= <type_spec>  
               | <ext_spec>  
               | <cons_spec>  
               | <macro_def>  
               | <empty_prod>  
               | <clause>  
               | <rule>  
               | <lex_entry>  
               | <lex_rule>
```

## T3.1 BNF for TRALE

The same as ALE except:

```

<cons_spec> ::= <type> cons <desc>
              | <type> cons <desc>
                  goal <goal>
              | <cond_desc> * <desc>
              | <cond_desc> * <desc>
                  goal <goal>

<macro_def> ::= <macro_head> macro <desc>.
              | <macro_head> := <desc>.

```

In addition, when the topological parser is loaded:

```

<rule> ::= <topo-rule>
          | <linking-rule>
          | <matches-rule>
          | <matched-by-rule>
          | <bidirectional-matching>
          | <covers-rule>
          | <covered-by-rule>
          | <bidirectional-covering>
          | <compaction-rule>
          | <tecto-rule>

<topo-rule> ::= <region> topo <list(<field>)>.

<linking-rule> ::= <region> <<-- <field>.
                  | <region> <<-- (<field-disj>).
                  | <field> -->> <region>.
                  | <field> -->> (<region-disj>).

<field> ::= <atom>

<region> ::= <atom>
            | matrix

<field-disj> ::= <field>
                | <field> ; <field-disj>

<region-disj> ::= <region>
                | <region> ; <region-disj>

<matches-rule> ::= <conj-desc> matches <field-or-region>.
                  | <conj-desc> matches (<fields-or-regions>).

```

```

<conj-desc> ::= <type>
              | <variable>
              | (<feature>:<desc>)
              | (<desc>,<desc>)
              | @ <macro_spec>
              | <func_spec>
              | a_ <prolog_term>
              | <path> == <path>
              | =\= <desc>

<field-or-region> ::= <field>
                    | <region>

<fields-or-regions> ::= <field-or-region>
                       | <field-or-region> ; <fields-or-regions>

<matched-by-rule> ::= <field-or-region> matched_by <desc>.

<bidirectional-matching> ::= <conj-desc> <==> <field-or-region>.

<covers-rule> ::= <conj-desc> covers (<fields-or-regions>).

<covered-by-rule> ::= <field-or-region> covered_by (<desc>).

<bidirectional-covering> ::= <conj-desc> <--> <field-or-region>.

<compaction-rule> ::= compacts( <list(<desc>)> ).

<tecto-rule> ::= <desc> *--> <list(<desc>)>.
              | <desc> *--> [ <seq(<desc>)> , {<conditions>}].

<conditions> ::= <condition>
                | (<conditions>,<conditions>)
                | (<conditions>;<conditions>)

<condition> ::= <local-covers>
                | <local-matches>
                | <precedence>
                | <immediate-precedence>
                | <local-compaction>
                | <prolog_goal>

<local-covers> ::= <non-zero-index> covers <field-or-region>

<local-matches> ::= <non-zero-index> matches <field-or-region>

<precedence> ::= <non-zero-index> < <non-zero-index>

```

`<immediate-precedence> ::= <non-zero-index> << <non-zero-index>`

`<local-compaction> ::= compacts(<index>)`



# Appendix D

## Reference Card

Grammar

---

Type sub Subtypes (intro [F1:R1,...,Fn:Rn]).	(Subtyping)
Type intro [F1:R1,...,Fn:Rn].	(Appropriateness)
ext([Type1,...,Typen]).	(Extensionality)
(Types not otherwise declared, but used in the above definitions are assumed to be maximal and/or immediately subsumed by bot)	
Type cons Desc (goal Goal).	(Type Constraint)
Word ---> Desc.	(Lexical Entry)
empty Desc.	(Empty Category)
RuleName lex_rule DescIn **> DescOut (if Goal) morph X becomes [X,e,s] when PrologGoal.	(Lexical Rule)
RuleName rule Desc ==> cat> Desc/ cats> ListDesc/ sem_head> Desc/ goal> Goal/ sem_goal> Goal.	(Phrase Structure Rule) [Daughter] [List of Daughters] [Semantic Head] [Procedural Attachment] [P.A. to Semantic Head]
f(Desc1,...,Descn) if g(Desc1,...,Descn)/ Desc1 =@ Desc2/ prolog(PrologGoal).	(Definite Clause) (Extensionality Check) (Prolog Hook)
fun(Desc1,...,Descn) +++> DescResult.	(Function Declaration)
macro(X1,...,Xn) macro Desc.	(Macro Declaration)
semantics Pred.	(Semantic Pred. Declaration)

### Compile-time Options

---

?- parse/generate/parse_and_gen.	(Compilation Mode)
?- lex_consult/lex_compile.	(Lexicon Compilation Mode)
?- (no)adderrs.	(Toggle Description Errors)
?- (no)subtest.	(Toggle Subsumption Check)
?- chain_length(Num).	(Chain Rule Length Bound)
?- lex_rule_depth(Num).	(Lexical Rule Depth Bound)

### Compilation

---

?- (d)compile_gram(GramFile).	(Grammar Compilation)
?- update_lex(File).	(Incremental Lexicon Update)
?- retract(all)_lex.	(Incremental Lexicon Retraction)

### Grammar Inspection

---

?- (no_)write_type(s).	(Type Hiding/Showing)
?- (no_)write_feat(s).	(Feature Hiding/Showing)
?- show_type Type.	(Signature Inspection)
?- unify_type(Type1,Type2,LUB).	(Type Unification)
?- approp(Feat,Type,Restr).	(Appropriateness Inspection)
?- introduce(Feat,Type).	(Feature Introduction)
?- iso_desc(Desc1,Desc2).	(Extensional Identity)
?- mgsat Desc.	(Most General Satisfier)
?- show_clause PredName(/Arity).	(Definite Clause)
?- lex Word.	(Lexical Entry)
?- rule Rulename.	(Phrase Structure Rule)
?- empty.	(Empty Category)
?- macro Macro.	(Macro Definition)
?- lex_rule Lexrulename.	(Lexical Rule)
?- export_words(Stream,Delim).	(Lexicon Export)

### Execution

---

?- (d)rec WordList.	(Bottom-up Parsing)
?- (d)query Query.	(SLD Resolution)
?- (d)gen Desc.	(Head-Driven Generation)

### Run-time Options

---

?- (no)interp.	(Toggle Mini-Interpreter)
?- edge(Left,Right).	(Chart Edge Inspection)
?- dleash(+/-Kind).	(Port Leashing)
?- dskip(+/-Kind).	(Port Auto-Skipping)
?- dclear_bps.	(Breakpoint Clearing)

# Bibliography

Meurers, W. D. (2001). On expressing lexical generalizations in HPSG. *Nordic Journal of Linguistics* 24, 161–127.

Meurers, W. D. and G. Minnen (1997). A computational treatment of lexical rules in HPSG as covariation in lexical entries. *Computational Linguistics* 23(4), 543–568.

Penn, G. and M. Haji-Abdolhosseini (2002). Topological Parsing: A linearisation parser specification submitted to Project MiLCA/A4 of the BNBF. Technical report, University of Toronto.

Pollard, C. and I. Sag (1994). *Head-Driven Phrase Structure Grammar*. Studies in Contemporary Linguistics. Chicago: CSLI.