

Model Generation for Natural Language Interpretation and Analysis

Dissertation zur Erlangung des Grades Doktor der Ingenieurwissenschaften
(Dr.-Ing.) der Technischen Fakultät der Universität des Saarlandes,
Saarbrücken

Karsten Konrad

30. Januar 2000

Short Abstract

Model generation refers to the automatic generation of mathematical structures that prove the satisfiability of logical theories. The research documented in this thesis investigates the use of model generation in the analysis and interpretation of formal semantic representations of natural language. Based on standard techniques for first-order model generation, we develop a model generation technique for a restricted higher-order logic and show how this method can be used to investigate the criteria that distinguish valid natural-language interpretations from interpretations that do not correspond to the intended meaning of the represented sentences. In particular, we investigate the analysis of singular definite descriptions and reciprocal sentences and show that model generation gives a computational method for describing theories of preference for natural-language interpretations.

For Daisy

Contents

1	Motivation	13
1.1	The Subject of this Thesis	13
1.2	Interpretation, Analysis, Computation	14
1.2.1	Interpretation	15
1.2.2	Analysis	16
1.2.3	Computation	17
1.3	Acknowledgments	18
I	Logics	20
2	Model Generation	21
2.1	Introduction	21
2.2	Preliminaries	22
2.3	Topics	23
2.3.1	Models and Decidability	23
2.3.2	Herbrand Models	23
2.3.3	Minimality	25
2.3.4	Enumeration	27
2.4	Methods	28
2.4.1	Analytical Tableaux	29
2.4.2	Positive Unit Hyper-Resolution	32
2.4.3	A Method Complete for Finite Satisfiability	33
2.4.4	The Davis-Putnam Procedure	35
2.5	Related Work	38
3	Higher-Order Model Generation	40
3.1	The λ -Calculus in Linguistics	41

3.1.1	Composition of Meaning	41
3.1.2	Quantifiers as Higher-Order Expressions	42
3.1.3	First-Order Limitations	45
3.1.4	A Motivation for a New Kind of Logic	46
3.2	Higher-Order Logic	48
3.2.1	Syntax	48
3.2.2	Semantics	50
3.2.3	Model Generation for Generalised Frames?	54
3.3	A Fragment of Higher-Order Logic	57
3.3.1	Syntax	58
3.3.2	Semantics	59
3.3.3	An \mathcal{MQL} Logic	62
3.4	Constructing Models	67
3.4.1	Determining Models Intelligently	67
3.4.2	Translating Formulas into Constraints	70
3.4.3	An Example	71
3.4.4	Properties of the Translation	73
3.5	Minimal Model Generation	77
3.5.1	Preliminaries	77
3.5.2	Decidability of Local Minimality	77

II Linguistics

79

4	The Analysis of Definites	80
4.1	Introduction	80
4.1.1	The Semantics of Definite Descriptions	81
4.1.2	Definites and Deduction	82
4.1.3	How Models Interpret Sentences	83
4.1.4	Discourse Models	85
4.1.5	Models for Definites	86
4.1.6	Uniqueness and Lots of Rabbits	88
4.2	Some Representations	90
4.2.1	Simple Cases	90
4.2.2	Donkeys, Context Sets, and Anaphoric Use	92
4.2.3	Restrictions with Knowledge	95
4.2.4	Simple Cases Revisited	98
4.3	What We Have Learned so Far	102

5	Reciprocity	104
5.1	Introduction	104
5.2	Exploring the Meaning of <i>Each Other</i>	106
5.2.1	Reciprocals for Larger Groups	107
5.2.2	Classifying Reciprocal Meaning	108
5.2.3	Parameterisation	112
5.2.4	Interpreting Reciprocals	115
5.3	Inference to Best Reciprocal Meaning	117
5.3.1	To Strong Meaning through Minimality	118
5.3.2	A First Attempt at Computation	121
5.3.3	Conservative Minimality	125
5.4	Experiments	128
5.4.1	Pitchers and Pearls	128
5.4.2	Measles	130
5.4.3	Marriages	132
5.5	Loose Ends	133
5.6	How We Can Understand Each Other	134
6	Abduction	136
6.1	What is Abduction?	136
6.2	Models for Anaphora Resolution	138
6.2.1	Chasing The Criminal	139
6.2.2	Explaining Resolutions	140
6.2.3	Discussion	141
6.2.4	An Alternative by Conservative Minimality	143
6.3	Weighted Abduction	146
6.3.1	Logic Programming and Abduction	146
6.3.2	Applications	151
6.3.3	Discussion	155
7	Implementation	159
7.1	Introduction	159
7.2	System Architecture	160
7.3	The Syntax	162
7.3.1	Logical Constants	162
7.3.2	Formulas	163
7.3.3	Problem Specifications	164
7.3.4	Definitions	166

7.4	The Semantics	166
7.4.1	Logic Definition Structures	167
7.4.2	Propagator Procedures	167
7.4.3	The Translation	174
7.5	Proof Engines and Controlling Search	175
7.5.1	Proof Engines	175
7.5.2	Search	179
7.6	System Performance	181
7.6.1	Identifying Single Solutions	182
7.6.2	KIMBA as a Propositional Theorem Prover	184
7.6.3	Generating Minimal Models	185
8	Conclusion	188
8.1	Why Inference is Worth the Effort	188
8.2	Contributions	190
8.3	Models as Meaning	192
A	Some Example Problems	194
A.1	The Job Puzzle	194
A.2	Reciprocals: The Boston Pitchers	196

Zusammenfassung in Deutscher Sprache

Der Forschungsgegenstand der formalen linguistischen Semantik ist die Bedeutung von natürlichsprachlichen Äusserungen und die Werkzeuge, die diese Bedeutung durch Repräsentation und Berechnung zugänglich machen. Die in dieser Arbeit dokumentierte Forschung befaßt sich mit Modellgenerierung, einem in der formalen linguistischen Semantik neuartigen Werkzeug, das die Auswertung logisch-semantischer Repräsentationen erlaubt.

Modellgenerierung bezeichnet Methoden des automatischen Beweisens, mit denen die Erfüllbarkeit einer logischen Spezifikation dadurch gezeigt wird, daß man Modelle der Spezifikation berechnet. Im Bereich des automatischen Beweisens gilt Modellgenerierung schon seit längerem als eines der erfolgreichsten Teilgebiete, und die dort entwickelten Systeme haben praktische Anwendung außerhalb der Sprachforschung gefunden. Bis vor kurzem war Modellgenerierung allerdings noch kein Thema in der formalen Semantik. Während deduktive Beweissysteme bereits als Schlußfolgerungsmodule in natürlichsprachlichen Systemen eingesetzt werden, sind die Möglichkeiten der Modellgenerierung in diesem Bereich noch weitgehend unerforscht. Hierfür sind im besonderen zwei Gründe zu nennen.

Zum einen handelt es sich bei der Modellgenerierung um ein überraschend kleines Forschungsgebiet. In den über vierzig Jahren, in denen man Verfahren des automatischen Beweisens nun bereits untersucht, lag der Schwerpunkt des Interesses stets auf deduktiven Widerlegungskalkülen, die die Un erfüllbarkeit von logischen Spezifikationen zeigen. Das Hauptziel des automatischen Beweisens war ursprünglich das Beweisen von allgemeingültigen mathematischen Sätzen. Der Beweis eines solchen Satzes ist äquivalent zur Widerlegung seines Negats. Modelle, d.h., die mathematisch-logischen Strukturen, die Erfüllbarkeit zeigen, wurden daher oft nur als Gegenbeispiele un-

wahrer Behauptungen gebraucht. Modellgenerierungsverfahren konnten zwar eine Reihe von offenen Existenzsätzen in der endlichen Mathematik beantworten, die Erfolge in diesem etwas exotischen Gebiet haben jedoch kaum dazu beigetragen, Modellgenerierung allgemein bekannter zu machen. Ein Linguist, der sich für Methoden des logischen Schließens interessiert, kann auf ungleich mehr Material zum Themenbereich Deduktion zurückgreifen als auf solches, das sich mit Modellgenerierung beschäftigt.

Zum anderen ist das Gebiet der formal-linguistischen Semantik bis heute von der Frage beherrscht, welche Repräsentation für die Bedeutung von Sprache die beste ist. Auf dem Weg zu einer formalen Methode, mit der sich nicht zuletzt die Verarbeitung von komplexen Diskursen brauchbar darstellen läßt, entwerfen Semantiker immer neue Repräsentationssprachen, jedoch meist ohne die zahlreichen, für ihre Verarbeitung notwendigen Berechnungsverfahren mitzuliefern. Die Fähigkeit, logische Schlüsse zu ziehen, ist zwar allgemein als zentraler Punkt des Sprachverstehens anerkannt, die Umsetzung dieser Erkenntnis wird jedoch aufgrund des damit verbundenen sehr hohen Aufwands meist gescheut. Erst in letzter Zeit setzt sich langsam die Idee durch, daß die im automatischen Beweisen entwickelten, sehr leistungsstarken Systeme einen vergleichsweise einfachen Zugang zum logischen Schließen in der Sprachverarbeitung darstellen könnten. Es ist jedoch immer noch so, daß nur sehr wenige Semantiker den zahlreichen Arbeiten im Gebiet der Semantik-konstruktion und -repräsentation gleichwertiges im Bereich der semantischen Auswertung hinzuzufügen.

Das Ziel dieser Arbeit ist es, einige der Anwendungsmöglichkeiten der Modellgenerierung im Bereich der formalen linguistischen Semantik darzustellen und einem breiteren Leserkreis zugänglich zu machen. Das Thema dieser Arbeit ist die Verwendung von automatisch generierten Modellen in der semantischen Interpretation und Analyse, und die Berechnungsverfahren, die man benötigt, um solche Modelle zu erhalten.

Interpretation, Analyse, Berechnung

Eines der Hauptziele der modernen linguistischen Semantik ist die Entwicklung einer formalen Theorie über die Wahrheitsbedingungen natürlichsprachlicher Äußerungen. Hinter diesem Ziel steht die Erkenntnis, daß die Bedeutung einer Äußerung im großem Maße dadurch bestimmt wird, unter welchen Bedingungen man sie für wahr halten kann. Natürlich erfaßt eine semantische

Theorie nicht alle Bedeutungen, wenn sie sich nur auf solche Wahrheitsbedingungen konzentriert. Hinter einer Äußerung stehen oft auch pragmatische Ziele des Sprechers, und diese liegen immer außerhalb der Beziehung zwischen der Äußerung und ihrer Wahrheit. Dessen ungeachtet verstehen viele moderne semantische Theorien die Bedeutung von Diskursen im wesentlichen als eine Abbildung von sprachlicher Struktur (Syntax) in eine wahrheitserhaltende Form der Repräsentation.

Interpretation

In logik-basierten semantischen Formalismen ist die gewählte wahrheitserhaltende Repräsentation für eine natürlichsprachliche Aussage oft die einer logischen Formel. Die Bedeutung einer Formel wird durch Interpretation bestimmt. Hierbei erhalten die Symbole, die in der Formel vorkommen, eine Interpretation als mathematische Objekte des geeigneten Typs. Diese Auffassung von Interpretation bestimmt die mathematisch-logische Sichtweise in der Modelltheorie seit der Zeit, als Tarski eine in sich widerspruchsfreie und vollständige Definition einer Semantik entwickelte.

In der Sprachforschung wird der Begriff Interpretation allgemeiner verwendet. Hier versteht man darunter im weiteren Sinn jede Abbildung von Sprache in Bedeutung. In der Semantik der natürlichen Sprache meint dies oft die Berechnung der Bedingungen, unter denen eine zu interpretierende Aussage wahr wird. Eine Aussage ϕ in einem Zusammenhang Γ zu interpretieren heißt daher, die Annahmen Δ anzugeben, aus denen die Wahrheit von ϕ in Bezug zu Γ folgt. Handelt es sich bei ϕ , Γ und Δ um geeignet gewählte Formelmengen, so ist deren Beziehung wie folgt:

$$\Gamma \cup \Delta \models \phi$$

Die Annahmengen Δ , für die die obige Beziehung gilt, sind formale *Erklärungen* für die Gültigkeit von ϕ im Zusammenhang Γ . Interpretation heißt, diejenigen Erklärungen zu finden, die der Bedeutung von ϕ entsprechen.

Wenn wir annehmen, daß ϕ konsistent zu Γ ist, dann hat die logische Theorie $\Gamma \cup \{\phi\}$ zumindest ein Modell \mathcal{M} . Dieses Modell ist eine Interpretation im logischen Sinn, die $\Gamma \cup \{\phi\}$ als wahre Formelmenge interpretiert. Wenn wir \mathcal{M} als eine Formelmenge $\Delta_{\mathcal{M}}$ darstellen könnten, so könnten wir ohne große Schwierigkeiten beweisen, daß folgender Zusammenhang gilt.

$$\Gamma \cup \Delta_{\mathcal{M}} \models \phi$$

Dies bedeutet, daß die Modelle einer logischen Repräsentation unter bestimmten Voraussetzungen als Interpretationen in der linguistischen Semantik verwendet werden können. Tatsächlich existieren Klassen von Modellen, die selbst wieder in einer logischen Sprache ausdrückbar sind. Aus der Prädikatenlogik erster Stufe kennen wir Herbrandmodelle, die sich in ihrer einfachsten Form als Mengen atomarer Annahmen darstellen. Diese Klasse von Modellen hat zudem die Eigenschaft, daß jede erfüllbare logische Theorie des Prädikatenkalküls auch mindestens ein Herbrandmodell besitzt.

Leider ist nicht jedes Herbrandmodell einer Spezifikation $\Gamma \cup \{\phi\}$ auch eine geeignete Interpretation im natürlichsprachlichen Sinn. Es gibt im allgemeinen sehr viel weniger akzeptable Erklärungen, die den Wahrheitsbedingungen einer repräsentierten Aussage entsprechen, als es Modelle gibt. Dies hängt damit zusammen, daß der Begriff des Modells allein nicht die vielfältigen Konventionen erfaßt, die mit einer kommunikativen Handlung verbunden sind.

Die Form des logischen Schließens, die wir hier eigentlich durchführen wollen, ist die Abduktion. Abduktion bezeichnet das zielgerichtete Schließen hin zu einer bestmöglichen Erklärung. In unserem Fall sind bestmögliche Erklärungen diejenigen Modelle, die die Bedeutung einer Äußerung in ihrem Zusammenhang erfassen. Modellgenerierung wird dann zur Abduktion, wenn nur bestimmte Modelle aufgrund ihrer Eignung als Lösung eines Problems aus der Menge aller Modelle ausgewählt werden. Semantische Interpretation ist eine abduktive Aufgabe. Die vorliegende Arbeit zeigt, daß diese Aufgabe durch geeignete Modellgenerierungsverfahren angegangen werden kann, die die Prinzipien der Kommunikation als Bedingungen auf den Modellen der logischen Form verwirklichen. Als Beispiel hierfür stelle ich eine neuartige Analyse definiter Beschreibungen vor, in der die empirisch nachweisbare Präferenzierung von anaphorischer Auflösung vor der Annahme neuer Diskursteilnehmer durch eine spezielle Form von minimalen Modellen simuliert werden kann.

Analyse

Semantische Analyse ist eine Aufgabe, die sehr stark mit der Interpretation zusammenhängt. Sie wählt geeignete semantische Repräsentationen aus einer Menge von Möglichkeiten durch Anwendung von Weltwissen aus. Auch hier sind die Modelle $\Delta_{\mathcal{M}}$ unserer Spezifikationen $\Gamma \cup \{\phi\}$ von Wert. Tatsächlich ist es so, daß die Existenz eines Modells allein schon eine sehr

wichtige Information sein kann. “Sei widerspruchsfrei!” ist eine der wichtigsten Richtlinien menschlicher Kommunikation, und eine Repräsentation ϕ , die in Bezug zu einem Zusammenhang Γ widersprüchlich ist, kann im allgemeinen nicht als geeignete Bedeutungsrepräsentation angesehen werden.

Von einem rein theoretischen Standpunkt aus betrachtet ist semantische Analyse eine unentscheidbare Aufgabe. Sie setzt voraus, daß wir feststellen, welche Repräsentationen innerhalb des durch den Diskurs gegebenen Zusammenhangs konsistent sind und welche nicht. Diese Frage läßt sich jedoch im Rahmen von Logiken erster und höherer Stufe allgemein nicht beantworten, und wird auch praktisch mit der Anzahl der zu betrachtenden semantischen Repräsentationen zu aufwendig.

So versuchen die bekannten linguistischen Theorien zu englischen Rezipropronomen (*each other* und *one another*), die Bedeutung von Reziproken durch eine Klassifizierung der möglichen Bedeutungen des Pronomens zu erfassen. Dies führt dazu, daß jedes logisch-semantische Berechnungsverfahren, das die Bedeutung einer Aussage bestimmen soll, notwendigerweise sehr aufwendig wird und eine ganze Menge von logischen Formen betrachten muß. Hier bietet nun Modellgenerierung einen Ausweg, da sie die Konventionen, die die Auswahl unter den Semantiken bestimmt, als berechenbare Bedingung auf den Modellen statt als unentscheidbare Bedingung auf den logischen Formen umsetzen kann. Im Fall von Reziproken heißt das, daß man ausgehend von einer sehr schwachen logischen Form durch eine gesteuerte Modellgenerierung diejenigen Modelle berechnen kann, die den passenden natürlichsprachlichen Interpretationen entsprechen. Der in dieser Arbeit vorgestellte Ansatz ist dabei meines Wissens der einzige überhaupt, der in der Lage ist, die von den linguistischen Theorien geforderte Interaktion zwischen Semantik, Weltwissen und Präferenz zu implementieren.

Berechnung

Der technische Beitrag der hier vorliegenden Arbeit ist die Entwicklung eines Modellgenerierungsverfahrens, das als formales Werkzeug in der semantischen Interpretation und Analyse eingesetzt werden kann. Herkömmliche Verfahren zur Modellgenerierung haben in diesem Zusammenhang den Nachteil, das sie meist auf konventionelle Prädikatenlogik erster Stufe zugeschnitten sind. Wie das gesamte Gebiet des automatischen Beweisens, so ist auch Modellgenerierung von Verfahren erster Stufe dominiert. Dies liegt nicht zuletzt an theoretischen Eigenschaften, denn die Logiken erster Stufe sind bekannt-

termaßen die ausdrucksmächtigsten, für die sich sowohl vollständige als auch korrekte deduktive Beweisverfahren angeben lassen.

Die Sprachen der Logik erster Stufe sind als semantische Repräsentationssprachen jedoch nur bedingt geeignet, da sich mit ihnen das Prinzip der Kompositionalität nicht ohne weiteres verwirklichen läßt. Dieses Prinzip besagt, daß sich die Bedeutung eines zusammengesetzten Ausdrucks durch die Komposition der Bedeutungen seiner Teile ermitteln lassen sollte. Seit den frühen 70er Jahren, als dem amerikanischen Logiker und Linguisten Richard Montague erstmals mithilfe von Churchs einfach getypten λ -Kalkül eine uniforme und kompositionelle Behandlung von natürlichsprachlichen Quantoren gelang, gebrauchen Semantiker oft Sprachen höherer Stufe, um die Bedeutung sprachlicher Konstituenten zu definieren. Dies hat auch ästhetische Gründe, da der λ -Kalkül Repräsentationen erlaubt, deren logische Form sehr nahe an der syntaktischen Struktur der natürlichen Sprache bleibt.

Es gibt Modellgenerierungsverfahren, die auf der Idee beruhen, mithilfe von festgelegten Universen aus Individuen erster Stufe eine logische Theorie erster Stufe zunächst in eine aussagenlogische Theorie umzuwandeln, um deren Modelle dann mit einem der hocheffizienten Verfahren der Aussagenlogik aufzuzählen. Indem man die Größe der verwendeten Universen iterativ erhöht, ist es prinzipiell möglich, alle endlichen Modelle der Eingabe modulo einer Umbenennung von Individuen zu generieren. Da die Fähigkeit, endliche Modelle aufzählen zu können, im Bereich der semantischen Analyse offensichtlich sehr wertvoll sein könnte, wäre ein vergleichbares Verfahren für semantische Repräsentationen höherer Stufe wünschenswert. Das Problem der Modellgenerierung ist jedoch für Logiken höherer Stufe sehr viel schwerer als im Fall der ersten Stufe.

In diesem Zusammenhang stellt sich die Frage, ob die Art von logischen Interpretationsverfahren, die man für die natürlichsprachliche Semantik benötigt, nicht tatsächlich in einem Fragment der Logik höherer Stufe angesiedelt werden kann, für das es in der Praxis noch ausreichend effiziente Berechnungsverfahren gibt. Der Versuch einer Antwort führte mich zur Entwicklung einer speziellen Klasse von Logiken höherer Stufe, *MQC*, die einen brauchbaren Kompromiß aus Ausdrucksstärke und Berechenbarkeit darstellen. Der Modellgenerierer KIMBA, der speziell für Experimente im Bereich der formalen linguistischen Semantik entwickelt wurde, ist in der Lage, logische Spezifikationen der Sprache *MQC* in kombinatorische Probleme über ganzzahligen Variablen mit beschränktem Wertebereich umzurechnen. Die Lösung dieser Probleme ergibt Sequenzen von Modellen der Eingabe.

Wie bereits erwähnt ist die Anzahl der Modelle viel größer als die Anzahl der Erklärungen, die brauchbare Interpretationen der logischen Form im linguistischen Sinn sind. Würde KIMBA lediglich endliche Modelle aufzählen, so wäre sein Nutzen in der semantischen Auswertung sehr begrenzt. Es ist jedoch vergleichsweise einfach, über Eingriffe in die Suche und das Berechnungsverfahren verschiedene Formen von Minimalität als notwendige Eigenschaft aller generierten Modelle zu sichern. Wie meine Arbeit zeigt, kann man die Bedingungen, die sprachliche Konventionen und Präferenzen an die Interpretationen einer Aussage stellen, auch als berechenbare Minimalitätskriterien von erzeugten Modellen annähern und implementieren. Oft ist es so, daß erst durch die Experimente mit Modellgenerierung explizit gemacht wird, was überhaupt notwendig ist, damit die Interaktion von Weltwissen und semantischer Repräsentation eine Korrespondenz von Bedeutung und Wahrheit herstellt. Modellgenerierung ist daher auch ein ideales experimentelles Werkzeug, um eben genau diejenigen Kriterien zu erforschen, die die Wahl einer bestimmten Interpretation bestimmen.

Chapter 1

Motivation

To understand a proposition means to know what is the case if it is true.

(Wittgenstein, *Tractatus*)

1.1 The Subject of this Thesis

Computational semantics describes and determines the meaning of human language using computational tools. The research reported in this thesis investigates the use of a new such tool in the interpretation and analysis of logic-based semantic representations: model generation.

Model generation refers to the automatic construction of interpretations that satisfy logical specifications. Model generation is a successful area in automated theorem proving, and systems for generating models have found various applications outside of linguistics. However, model generation has, until recently, not been a topic of research in the computational semantics community. While conventional automated theorem provers enjoy some interest as inference modules in natural-language processing, the potential of systems that prove satisfiability by generating models is still widely unknown. Two reasons for this can be identified.

First, model generation is a surprisingly small research area. In the forty years that automated theorem proving has been investigated, work has concentrated on deductive methods that prove unsatisfiability and entailment in classical first-order logics. These methods have been developed originally for proving mathematical theorems without human interaction. In this context,

model generation has been considered only as a specialised tool for computing counterexamples of dubious theorems, or for proving the existence of certain small structures in finite mathematics. Hence, a linguist who wants to apply some techniques developed in automated theorem proving is likely to find much more material on first-order deduction than on model generation.

Second, natural-language semantics concentrates on representation rather than on inference. With the explicit goal of developing formal representations that are suitable for language-oriented inference, semanticists define and investigate formal representation languages for the meaning of human language. However, semanticists almost never implement inference systems for these languages. Only recently, there has been a growing interest in using the well-developed inference techniques from automated theorem proving for building natural-language systems. Still, only a small group of semanticists actually work on inference in semantics, and there is too little effort for exploring new techniques that have been developed elsewhere.

This thesis aims at presenting model generation to a larger audience of researchers who are interested in logic-oriented computational semantics. The subject of this thesis is the use of automatically generated models in the interpretation and analysis of semantic representations, and the computational methods that are needed for generating such models.

1.2 Interpretation, Analysis, Computation

One of the principal goals of modern linguistic semantics, as initiated by Frege and Russell, has been to develop a computational theory of the conditions under which a statement can be uttered truthfully. Behind this interest lies the insight that the meaning of a statement uttered in a situational context can often be captured completely by such truth conditions. The correspondance theory of meaning and truth is not a universal one because utterances have a wide range of purposes that cannot be conceived by knowing their truth value alone. Nevertheless, many modern semantic theories determine the meaning of natural-language utterances by a relationship between linguistic form and the truth conditions of some formal semantic representations.

The relation of form and meaning can be given as a translation from syntactical structures into semantic representations that preserves truth and entailment with regard to the original utterance. In semantic formalisms such as Discourse Representation Theory (DRT) by Kamp and Reyle [56],

the representations computed from natural-language sentences are logical formulas. The process that determines the meaning of such formulas is called interpretation.

1.2.1 Interpretation

Since Tarski formulated the first mathematically sound truth-conditional semantic for a logic, interpretations are conceived as recursively defined mathematical structures that assign truth values to the well-formed formulas of a logical language. Interpreting a formula in a logic simply means to apply such an interpretation.

In logic-based natural-language semantics, interpretation is understood as a related, but different concept. Here, interpretation refers to processes that compute the truth conditions for a natural-language sentence L uttered in a given context consisting of situational and world knowledge.

More formally, interpreting a semantic representation ϕ for L within in context Γ means to compute facts Δ that must be assumed to be true if ϕ is true in Γ . The assumptions Δ that we look for formulate “explanations” why L is true in the situation described by Γ . Computing semantic interpretations is an inference process that identifies suitable sets of facts Δ for ϕ and Γ such that the following condition holds:

$$\Gamma \cup \Delta \models \phi$$

Suppose that there exists a logical interpretation \mathcal{M} such that $\mathcal{M}(\Gamma \cup \{\phi\})$ is true, i.e., that there is a model \mathcal{M} that satisfies both Γ and the representation ϕ . If we could represent \mathcal{M} as a set of facts $\Delta_{\mathcal{M}}$, we could prove immediately that the following condition holds:

$$\Gamma \cup \Delta_{\mathcal{M}} \models \phi$$

This implies that one can use models \mathcal{M} as semantic interpretations if there is a way to represent \mathcal{M} within the logical language as a set of facts $\Delta_{\mathcal{M}}$. For an important class of models, the Herbrand models, this is indeed possible. A Herbrand model can be represented unambiguously as a set of ground literals. Furthermore, every first-order specification $\Gamma \cup \{\phi\}$ that is satisfiable at all must also be satisfiable by at least one Herbrand model.

Unfortunately, not every Herbrand model \mathcal{M} of $\Gamma \cup \{\phi\}$ is suitable as an explanation for the truth of a semantic representation ϕ in context Γ . There

are, in general, many models that interpret a logical specification $\Gamma \cup \{\phi\}$ to be true even though they do not have a correspondence to the truth conditions that we look for in a semantic interpretation.

What we actually want to perform here is abduction, i.e., that kind of inference that finds the best explanation. One of the topics of this thesis is to determine the relationship between abduction and the generation of models in the interpretation of semantic representations.

1.2.2 Analysis

Semantic analysis identifies the valid semantic representations within a set of possible readings by common-sense reasoning. For this, too, our models \mathcal{M} for ϕ and Γ are interesting. In fact, the existence of a model \mathcal{M} for a specification is a valuable information on its own. If a semantic representation ϕ is inconsistent within a context Γ , then $\Gamma \cup \{\phi\}$ will have no model at all. “Be consistent!” is one of the most important constraints for human communication and most discourses obey this maxim. If a semantic representation has no model within its situational context, we can assume that either the representation itself is faulty, or there is a discourse anomaly. Consider for instance the discourse *Katja’s husband knows her sister. She is not married.* A natural-language parser might provisionally propose two different resolutions for the pronoun *She*:

(2.1) ***Katja*₁**’s husband knows her sister. ***She*₁** is not married.

(2.2) *Katja*’s husband knows her ***sister*₁**. ***She*₁** is not married.

Only the second reading makes sense since the first reading leads to an inconsistency—we know that *Katja* is married because this information was implicitly given. A parser which has no access to common-sense reasoning will not be able to distinguish a valid from an inconsistent reading.

Model generation is useful for distinguishing valid and invalid semantic representations. Given a sentence L , a listener will in general have an intuition for the assumptions Δ that must hold if L is to be true in a context Γ . A computational tool that enumerates sets of assumptions Δ' for the semantic representations ϕ of L can be used to investigate whether the assumption sets Δ' meet the listener’s expectations or not. If not, the sets Δ' can help us to detect flaws in the original representation.

For this form of semantic analysis, we need inference for proving the satisfiability of semantic representations in context. In this thesis, I will discuss some existing model generation methods that have been developed originally for other applications than linguistics. As we will see, there is no off-the-shelf method available that is sufficient for the task at hand. For semantic analysis, we need specialised technical machinery, and I will show how to tailor the existing tools to our needs.

1.2.3 Computation

Technically, this thesis is about developing a model generation method that is suitable for the interpretation and analysis of logical semantic representations. One of the stumbling stones for conventional model generation in this context is higher-order logic. Model generation, like automated deduction, has concentrated on first-order predicate logic or fragments thereof. The standard language of classical first-order predicate logics can be very awkward in places, especially for representing natural language. Since the early 1970s, when Montague [70] initiated the use of the simply typed λ -calculus in natural-language semantics, semanticists use, among other means, higher-order logics with specific linguistic quantifications as a more convenient form of representation. The following are two possible representations of the sentence *Two women love John*, one with a Montague-style syntax and linguistic quantification, and one in a standard first-order form.

$$(2.3) \quad \text{Two}(w)(\text{love}(j))$$

$$(2.4) \quad \exists x_1 \exists x_2 w(x_1) \wedge w(x_2) \wedge x_1 \neq x_2 \wedge \text{love}(j)(x_1) \wedge \text{love}(j)(x_2)$$

It is easy to see that the Montague-style representation (2.3) is both closer and more natural with respect to the original sentence than the formula in standard syntax (2.4). The use of such Montague-style formalisations can have both theoretical and practical advantages for first-order inference, as has been shown by McAllester and Givan [66]. Linguists use higher-order logics because of its expressivity and because the semantics of natural language utterances can be constructed in a compositional way.

All model generation methods expect their input in a form similar (2.4). As natural-language semantics makes use of logical languages that are more expressive than classical first-order predicate logic, a model generation method

that is suitable for natural-language processing must be able to deal with such languages too. However, the use of more expressive logics may change the computational tractability of inference. Although the extra expressive power is often worth the price, we would like to have efficient techniques that can be used in practice and as the basis of scientific evaluation in computational semantics. The mutually exclusive design goals of expressivity and computational tractability can be dealt with in systems that use the more expressive power of the richer formalism, but in fact inhabit a fragment that has better computational properties. One of the contributions of this thesis is the development of an efficient model generation method for Montague-style higher-order logical specifications that has the same complexity as first-order model generation.

1.3 Acknowledgments

Starting at the beginning, I would like to thank Michael Kohlhase and Jörg Siekmann for introducing me to the exciting interdisciplinary research in automated reasoning and formal semantics. They gave me the opportunity to be a part of a very active group which, among other privileges, provided me with access to knowledge of logical systems and how to implement them. The knowledge that I needed desperately was donated by several people, but I have to thank especially my colleagues Michael Kohlhase and Christoph Benzmüller who spend a lot of their valuable time by initiating me into the mysteries of higher-order logics, both with an amount of patience that I still find incredible.

My other colleagues at the computer science department, Lassaad Cheikhrourou, Detlef Fehrer, Armin Fiedler, Helmut Horacek, Andreas Meier, Erika Melis, Martin Pollet, and Volker Sorge not only let me be a part of their workload sometimes, but also accepted my “exotic” interest in natural language semantics—they are after all in a group whose primary goal is to develop techniques for planning and presenting mathematical proofs. I suspect that most of them believe that my research interests were only a clever scheme of mine to avoid doing any real work, but they did not let it show too much.

In the computational linguistics department, my colleague Claire Gardent provided the linguistic background that was essential for the present work. Our discussions and close cooperation strongly influenced what I did and how I did it. Some of my “best” ideas died an instant death as soon as I gave

her the opportunity to examine them from an experienced linguist's point of view. I am therefore not only grateful for the work that we did together, but also for a lot of unnecessary work that she prevented me from doing.

Johan Bos kindly let us plug KIMBA into the DORIS system as an inference module and wrote the translation routine that was necessary for this. Of our students at the computer science department, I would like to thank Andreas Franke for giving me more than a hand with KIMBA's integration into MATHWEB such that it could be used as part of the DORIS system. Carsten Brockmann at the computational linguistics department was equally helpful for the modification of KIMBA into a web-based applet.

David Wolfram of the computer science department at The Australian National University, Canberra, gave me the opportunity to further develop the idea of higher-order model generation as a guest of his university for ten weeks. François Bry and his group at the Ludwig-Maximilian University in Munich also were several times my hosts. They showed me the beauty of "lean" model generation and have answered some of my questions about the connections of model generation and abduction. Martin Müller and Christian Schulte of the Programming Systems Lab, Saarbrücken, were the ones I could always turn to for technical support with Oz.

My girlfriend Katja often had to notice that whatever she had just told me did not find a way into my mind: it was still busy with such important questions as whether *giving measles to each other* should be interpreted by IAO or IAR reciprocity. There is also no question that she had to endure bad moods and inexplicable behavior beyond whatever can be expected from a PhD student's girlfriend. Sorry love, I promise that I will never write a dissertation again.

Michael Kohlhase, Christoph Benz Müller, Jörg Siekmann, and Manfred Pinkal had a closer look at parts of this thesis at various stages of its development. I am more than grateful for the constructive criticism and helpful remarks they gave me. Needless to say, all remaining errors are my own.

Finally, I always had and have the support of my family, although I think my dad would like me to stop doing research now. He does not consider it to be proper work at all because there is too much fun involved.

Part I

Logics

Chapter 2

Model Generation

Approach your problem from the right end and begin with the answers. Then one day, perhaps you will find the final question. (van Gulik)

Overview: This chapter presents research topics in model generation that will be relevant for later chapters. It also gives an introduction to some of the most popular methods for generating models.

2.1 Introduction

Model generation is important for the automation of reasoning. Its applications are in many different fields such as finite mathematics [69, 83], diagnosis [7, 35], planning [58], scheduling [48], deductive databases [18], software verification [52], expert systems [26], and program synthesis [36]. For this, and other reasons, model generation developed into a rich field of research of its own.

In Section 2.3, we give an introduction to the topics of model generation that are especially interesting for computational semantics and that will play some rôle in later chapters. Such topics are for instance minimal model generation and model enumeration. On the other hand, the methods discussed here are not those that we will actually use for natural-language analysis and interpretation, but survey the standard technical machinery of model generation. The staple techniques of model generation include the use of (hyper-)tableaux calculi and that of efficient propositional decision procedures. In Section 2.4, we discuss both PUHR tableaux and the Davis-Putnam

α	α_1	α_2	β	β_1	β_2
$A \wedge B$	A	B	$\neg(A \wedge B)$	$\neg A$	$\neg B$
$\neg(A \vee B)$	$\neg A$	$\neg B$	$A \vee B$	A	B
$\neg(A \Rightarrow B)$	A	$\neg B$	$A \Rightarrow B$	$\neg A$	B
$A \Leftrightarrow B$	$A \Rightarrow B$	$B \Rightarrow A$	$\neg(A \Leftrightarrow B)$	$\neg A \wedge B$	$A \wedge \neg B$

Figure 2.1: α - and β -Formulae and Components

γ	$\gamma(t)$	δ	$\delta(t)$
$\forall xF$	$[t/x]F$	$\neg\forall xF$	$[t/x]\neg F$
$\neg\exists xF$	$[t/x]\neg F$	$\exists xF$	$[t/x]F$

Figure 2.2: γ - and δ -Formulae and Instantiations

procedure, because these two methods exemplify the two most common approaches to model generation. Section 2.5 gives some pointers to selected works that investigate other methods and topics in model generation.

2.2 Preliminaries

We assume the reader to be familiar with first-order logics and inference procedures such as presented, e.g., in Fitting's textbook on first-order logic and automated theorem proving [34].

A **specification** is a finite set of closed formulas, a **theory** a possibly infinite one. Unless indicated otherwise, I use "formula" in lieu of closed formula, i.e., formulas F do by default not contain free variables.

An **atom** a is a formula that does not contain any logical connectives. A **literal** l is either an atom a or its negation $\neg a$. If a literal l is an atom a , then its **complement** \bar{l} is $\neg a$, otherwise its complement is a .

Following Smullyan [84], we classify formulas as implicitly conjunctive (α), disjunctive (β), existentially quantified (δ) or universally quantified (γ). Figure 2.1 shows the components of α - and β -formulae, while figure 2.2 shows the relationship between γ - and δ -formulae and their instantiations. We denote a substitution that instantiates a variable x with a term t with $[t/x]$.

2.3 Topics

In the following, some selected concepts and research topics in model generation are introduced.

2.3.1 Models and Decidability

Models are interpretations of the symbols of a logical theory T such that T becomes true. For a first-order theory, each model is associated with a nonempty domain \mathcal{D} of individuals. This domain sometimes is called the universe of discourse.

First-order **interpretations** are pairs $\mathcal{I} = \langle \llbracket \cdot \rrbracket, \mathcal{D} \rangle$ of an interpretation function $\llbracket \cdot \rrbracket$ that maps predicate and constant symbols to predicates and constants of the appropriate type, and a domain \mathcal{D} of individuals that provides entities for the interpretation of constants. Every interpretation of a term t , $\llbracket t \rrbracket$, must be an element of \mathcal{D} .

A model \mathcal{M} for a classical propositional logic is simply a mapping $\llbracket \cdot \rrbracket$ from Boolean constants c into the truth values. We call a problem **combinatorial** if it is equivalent to finding an assignment for finitely many variables with finite domains, and verifying in polynomial time whether the assignment is a solution or not. All combinatorial problems are decidable, because we can use a brute-force algorithm that simply tries out all possible variable assignments. Satisfiability of specifications in classical propositional logic is a combinatorial problem.

The predicate calculus, on the other hand, is undecidable. We know that there is no method that will be able to prove satisfiability in general. Hence, there can also be no procedure that is always able to compute a model for an arbitrary satisfiable first-order specification. However, there exist methods that are complete for finite satisfiability, i.e., that are decision procedures for specifications that either are unsatisfiable or can be satisfied by a finite model (see Section 2.4.3).

2.3.2 Herbrand Models

An important class of models are the **Herbrand models**. Within these, constants and other terms are interpreted as themselves, i.e., $\llbracket t \rrbracket = t$ for all terms. This implies that the universe of discourse is a set of the terms that can be built from the signature of the interpreted theory. In a classical

first-order logic, a Herbrand model \mathcal{M} and its associated domain \mathcal{D} can be specified completely by the set $\text{GROUND}(\mathcal{M})$ of all ground literals that are satisfied by \mathcal{M} . The domain \mathcal{D} of \mathcal{M} are the terms that occur in this set of literals. We identify Herbrand models \mathcal{M} with their representation as sets of ground literals.

Herbrand models are an important class of models because each satisfiable first-order theory must have at least one Herbrand model. Methods in automated deduction make use of this property when proving unsatisfiability. If a theory is inconsistent, it suffices to show that it cannot have a Herbrand model. This, in general, is simpler to show than the non-existence of arbitrary models. Unless indicated otherwise, we use the term “model” for Herbrand models.

Finite Models

Not all first-order Herbrand models for a formula may have a *finite* representation as a finite set of ground literals. For instance, the formula $p(f(a)) \wedge \forall x p(f(x)) \Rightarrow p(f(f(x)))$ is satisfiable, but only by an infinite model of the form $\mathcal{M} = \{p(f(a)), p(f(f(a))), p(f(f(f(a))))\dots\}$. The universe of discourse \mathcal{D} of this model is infinite. A **finite model** is a model with a finite universe of discourse. Most model generation methods are restricted to the generation of such finite models (see Section 2.5).

A serious problem for completeness in model generation is Skolemisation. A Skolem term is a term that represents an equivalence class of δ -formulas up to a renaming of variables [11, 39]. Skolemisation is a convenient method for eliminating existential quantifiers in theorem proving, but for model generation purposes it has serious drawbacks. Skolem terms often introduce new function symbols which extend the domain of the Herbrand models of the input specification. While Skolemisation does not turn a satisfiable specification into an unsatisfiable one, it may transform a finitely satisfiable specification into one that has only infinite Herbrand models. For instance, the formula $p(a, a) \wedge \forall x (p(x, x) \Rightarrow \exists y p(x, y))$ has a finite Herbrand model $\{p(a, a)\}$, while a Skolemised form $p(a, a) \wedge \forall x (p(x, x) \Rightarrow p(x, f(x)))$ has not. As mentioned before, infinite models are a problem because there can be no method that detects their existence in general.

Representations

A finite representation $\text{GROUND}(\mathcal{M})$ still can be awkward in places because of their size. Herbrand models are therefore represented more compactly by their positive part alone, i.e., by the set of positive literals $\mathbf{Pos}(\mathcal{M})$. The interpretation function I that is derived from \mathcal{M} simply interprets all literals P as false if P is not in \mathcal{M} . Note that the universe of discourse \mathcal{D} of a positive model might actually be smaller than the terms that occur in the complete model. Throughout this thesis, we will usually present models as sets of atoms that leave out all negativ literals.

2.3.3 Minimality

An important research topic in model generation is the generation of models with certain properties, the most important property of which is some form of minimality. In the following, we summarize the forms of minimality that have been investigated in the literature.

Subset Minimality

A model \mathcal{M} for a specification Φ is called **subset-minimal** iff for all models \mathcal{M}' of Φ the following holds: $\text{Pos}(\mathcal{M}') \subseteq \text{Pos}(\mathcal{M}) \Rightarrow \text{Pos}(\mathcal{M}') = \text{Pos}(\mathcal{M})$. Subset-minimality is the most important form of minimality and has been the focus of a great deal of attention in the literature. Minimal model reasoning is at the heart of the circumscriptive approach to common sense reasoning. Circumscription, as introduced by McCarthy [67], is one way to formalise non-monotonic logics, and it has been applied to reasoning about actions, planning, and the semantics of logic programs. The logical consequence relation \models_c of circumscriptive reasoning can be formalised as follows: $\Gamma \models_c \Phi$ iff Φ is true in all subset-minimal models of Γ . This consequence relation is stronger than the classical consequence relation \models . For instance $A \wedge B \models_c \neg C$, but obviously $A \wedge B \not\models \neg C$. It is non-monotonic because $A \wedge B \models_c \neg C$ holds, but $A \wedge B \wedge C \models_c \neg C$ does not.

Circumscription needs methods that can verify the subset-minimality of models. Several such methods have been proposed in the literature on circumscriptive reasoning [40, 71, 72]. There are also other methods from model generation to generate subset-minimal models [20, 81]. Yet, there is no method available that enumerates the minimal models of a specification

both in a space- and time-efficient way. Most solutions proposed are not able to handle larger examples because of exponential worst-case complexity. For instance, the MM-SATCHMO system of Bry and Yahya [20] needs a potentially expensive lookup mechanism for models that have been computed earlier (see Section 2.4.2).

The minimal model reasoning method proposed by Niemelä [73] seems to be the only one on the market that has no exponential space requirements by design. In a tableaux setting, the method is able to locally verify the minimal model property of an open saturated tableau branch, but does not prevent that minimal models are sometimes generated multiple times in different branches. Also, the method is restricted (in its original form) to propositional logics.

Domain Minimality

A second definition of model minimality is domain minimality. A Herbrand model \mathcal{M} for a specification Φ is **domain-minimal** if the size $|\mathcal{D}(\mathcal{M})|$ of its universe of discourse is minimal, i.e., there is no model \mathcal{M}' of Φ such that $|\mathcal{D}(\mathcal{M}')| < |\mathcal{D}(\mathcal{M})|$. Reasoning with domain minimal models has been investigated by Hintikka [49] and Lorenz [63] as an alternative form of circumscription where the focus lies on minimizing the domain size rather the facts. Unlike subset-minimality, verifying domain-minimality is computationally tractable for every model generation method that is guaranteed to find a domain-minimal model first. A method that enumerates models in increasing size of the domain can be derived from the EP calculus presented in Section 2.4.3.

Predicate-Specific Minimality

Let p be a certain predicate symbol, and $p(\mathcal{M})$ be the subset of atoms in $\text{Pos}(\mathcal{M})$ whose head predicate symbol is p . We call a model \mathcal{M} of a specification Φ **p-minimal** if there is no model \mathcal{M}' of Φ such that $|p(\mathcal{M}')| < |p(\mathcal{M})|$.

Applications such as diagnosis and planning often employ some forms of p -minimal models. For instance, logical specifications in model-based diagnosis are used for describing the correct behavior of a device, while a certain predicate ab is used for abnormal behavior. The following specification may describe an AND-gate in a circuit:

$$on(in_1) \wedge on(in_2) \Rightarrow on(out_1) \vee ab(AND_1)$$

The formula specifies that $on(out_1)$, the output of our AND-gate AND_1 , must hold if both inputs in_1 and in_2 are on as well, otherwise the device AND_1 is abnormal. A set of such rules could be used to describe the in/out-behavior of an integrated circuit. A **diagnosis** is then a model of such a description Γ together with a set of observations Φ that indicate a faulty device. The observations Φ correspond to the in/out-states of parts of the circuit. The models that are interesting are those that explain the truth of $\Gamma \cup \Phi$ which refer to a minimum of devices in the predicate ab . These minimal models show how the error indicated by Φ can be caused without assuming more faulty devices than necessary. The specific task in diagnosis is to compute models that minimize not all atoms, but only the atoms that depend on the predicate ab . Other parts of the model are of no interest. We call models that are minimal with respect to a certain set of predicates **predicate-specific minimal models**.

2.3.4 Enumeration

The minimal models of a logical problem specification usually represent different solutions of the problem. As a consequence, we are often not interested in just computing one single model for a specification, but want to enumerate models instead. This implies that the generation method is complete with regard to the solutions we are interested in. For many applications, the form of completeness that is asked for is completeness for finite satisfiability [19].

Model Enumeration with Theorem Provers

Some theorem provers decide fragments of first-order logics. For instance, the theorem proving system BLIKSEM implements a decision procedure for the guarded fragment [30]. Within such decidable fragments, the related proof procedure can sometimes be modified such that a model is constructible from a failed refutation proof attempt. Yet, high-speed theorem proving calculi are not designed for enumerating the models of a specification. The general use of Skolemisation makes it impossible for conventional automated theorem provers to be complete for finite satisfiability.

On the other hand, calculi for model generation theorem proving can simultaneously be refutation complete and complete for finite satisfiability if we do not Skolemise (see Section 2.4.3). The FINFIMO system by Bry and Torge [19] implements this idea. It is the only theorem prover currently available whose underlying hyper-tableaux calculus is provably complete for finite satisfiability. The original SATCHMO [65] system uses Skolemisation, and therefore shows the same enumeration incompleteness with regard to finite satisfiability as other theorem provers.

Enumeration with Finite Model Generators

Unlike model generation theorem provers, finite model generators are programs whose only purpose is the generation of finite models and not the refutation proof of a theorem. One of the best known systems for finite model generation is Slaney's FINDER [82].

Finite model generators apply propositional decision procedures to first-order specifications that have been translated into propositional logic using a given universe of discourse. Programs such as FINDER are frequently used for applications that require the exhaustive search for finite models within very large search spaces [83]. By using efficient decision procedures, finite model generators can solve combinatorial problems that are far beyond the capabilities of first-order methods.

As a welcome side-effect, the use of complete propositional decision procedures results in a finitely complete form of model generation where all finite models of a specification up to a renaming of constant symbols are generated. Every propositional model is also a unique finite model of the input first-order specification. By iterative deepening over the size of the domain, all finite models can be successively enumerated in a space- and time-efficient way.

2.4 Methods

Most of the calculi developed in automated deduction that prove unsatisfiability can also sometimes be used for proving satisfiability. For many non-classical logics such as a variety of modal logics, the proof procedures available are not only refutation complete, but actually decide theories in these logics. Also, efficient procedures for testing satisfiability (SAT) in (classical)

propositional logics are a major research topic in Automated Reasoning and Artificial Intelligence [22, 31, 37, 43, 53].

Model generation as a field distinguishes itself from research on SAT and decision procedures in two aspects. First, the logics considered are undecidable in general and have some unrestricted notion of quantification over individuals. This means that the logical languages used are not tailored to be a decidable fragment of first-order logics such as those used in knowledge representation [74]. Second, the methods generate objects from which models can be derived.

The border line between model generation and other research on satisfiability is not a strict one. For instance, propositional decision and SAT procedures have been used successfully for first-order model generation in planning [58] and finite mathematics [68].

Model generation as a field of research is associated with a variety of tableaux methods, the most prominent being Positive Unit Hyper-Resolution (PUHR). We present this tableaux method in Section 2.4.2 as a standard device in model generation. An important improvement of PUHR with regard to our intended application is completeness for finite satisfiability, which is exemplified in the Extended Positive (EP) tableaux calculus in Section 2.4.3. The core method for model generation that we will use throughout this thesis is constraint-based finite model generation. This approach shares many properties with the Davis-Putnam procedure, a propositional decision procedure that we will discuss in Section 2.4.4. Constraint-based model generation is then the main topic of the next chapter.

2.4.1 Analytical Tableaux

Many calculi for model generation are based on analytical tableaux, although conventional tableaux calculi with an unrestricted syntax are rarely used for model generation. In the following, we show why standard tableaux methods are not suitable for finite model generation in general.

Ground Tableaux

The set of expansion rules shown in Figure 2.3 defines a sound and refutation complete theorem proving calculus for first-order predicate logics called **ground tableaux** [10].

$$\begin{array}{ccccc}
 \frac{\beta}{\beta_1 \mid \beta_2} & \frac{\alpha}{\alpha_1 \\ \alpha_2} & \frac{\neg\neg\mathbf{F}}{\mathbf{F}} & \frac{\gamma}{\gamma(t)} & \frac{\delta}{\delta(\text{sko}(\delta))}
 \end{array}$$

Figure 2.3: Tableaux expansion in conventional tableaux

$$\begin{array}{c}
 (\forall x \neg p(x)) \wedge (p(a) \vee p(b)) \\
 \forall x \neg p(x) \\
 p(a) \vee p(b) \\
 \neg p(a) \\
 \neg p(b) \\
 p(a) \mid p(b) \\
 \star \quad \mid \quad \star
 \end{array}$$

Figure 2.4: A tableau proof for refuting $(\forall x \neg p(x)) \wedge (p(a) \vee p(b))$

In conventional analytical tableaux, an arbitrary input formula is analysed by applying expansion rules, thus creating the well-known tree-structure of tableaux proofs. A branch in such a tableau is closed if it contains an elementary contradiction, such as an atom \mathbf{a} and its negation $\neg\mathbf{a}$. A closed tableau is a tableau where all branches are closed. A closed formula proves that the input formula is unsatisfiable. Figure 2.4 shows a closed tableau for the formula $(\forall x \neg p(x)) \wedge (p(a) \vee p(b))$. The tableau proof instantiates the γ -formula $\forall x \neg p(x)$ twice, creating two new formulas, $\neg p(a)$ and $\neg p(b)$. These are used for showing that the disjunction $p(a) \vee p(b)$ can not be satisfied.

Branches are frequently identified with the set of formulas they contain. A branch is called **saturated** if no application of an expansion rule can add new elements to this set. In analytical tableaux, the atoms in a saturated open branch define a positive model of the input formula.

The γ -rule of ground tableaux is a serious source of inefficiency. When applying the rule, we must “guess” the right terms and number of instantiations that are necessary in order to close all branches. There is, in principle, no limit for the number of instantiations that might be necessary to close a branch. For model generation, we must use all different terms that occur in the same branch as the γ -formula while Skolemisation in the δ -rule successively might introduce new terms. This quickly leads to infinite branches.

$$\frac{\gamma}{\gamma(x)}$$

Figure 2.5: Free variable tableaux expansion.

Free Variable Tableaux

The free variable tableaux is a basis of many deduction systems. It is particularly easy to implement, and some systems consist only of a few lines of Prolog code [12]. In a free variable tableaux, the instantiation of γ -formulas by ground terms is replaced by an instantiation with some variable x that is new to the tableau and whose instantiation has to be determined later (see Figure 2.5). A branch containing two formulas E and $\neg F$ can be closed iff there is a variable assignment, i.e., a unifier σ , such that $\sigma(E) = \sigma(F)$. The unifier σ can be computed efficiently by first-order **unification**, and σ is then applied to the whole tableau in order to remove all occurrences of the free variables in E and F .

By using unification, the selection of a term that instantiates a γ -formula can be delayed until this instantiation actually helps closing a branch. This is a considerable advantage in practice. Still, the selection can be wrong, because there might be several formula pairs E and F in a branch that are suitable for closing it. Applying a unifier has a global effect on the whole tableau, and a wrong choice in one part of a tableau might require additional γ -instantiations in another part. In order to have a fair expansion of γ -formulas, implementations often use iterative deepening over the number of γ -instantiations that are allowed in one branch.

The quantifier rules in a free variable tableaux still are a source of model incompleteness. Once a branch contains one γ -formula, it can be extended by an unlimited number of successive applications of the γ -rule. Hence, even if a branch has a model, we might not be able to detect it because one or several γ -formulas prevent the finite saturation of a branch. In practice, there are nevertheless methods that can be used for detecting models in potentially infinite branches. A formula that contains a variable represents an infinite set of ground Herbrand terms, and if a branch provably cannot be closed, such methods are even able to generate infinite models. However, infinite model generation requires a considerable amount of technical machinery and still cannot prevent model incompleteness because of Skolemisation. For a

					a_1	
					\vdots	
					a_m	
b_1	b_2	\dots	b_{n-1}	b_n	$a_1 \wedge \dots \wedge a_m \Rightarrow b_1 \vee \dots \vee b_n$ is a ground instance of a clause in the input Φ	
$\neg b_2$	$\neg b_3$	\dots	$\neg b_n$			
\vdots	\vdots					
$\neg b_{n-1}$	$\neg b_n$					
$\neg b_n$						

Figure 2.6: The PUHR tableaux expansion rule

survey of the state-of-the-art techniques in this area, we refer to Klengenbeck's thesis [60].

2.4.2 Positive Unit Hyper-Resolution

Positive Unit Hyper-Resolution (PUHR) tableaux is the theoretical basis of the SATCHMO theorem prover [17,65]. Introduced by Bry and Yahya [20], this calculus addresses the efficiency problem that the blind instantiation of γ -formulas poses for conventional analytical tableaux calculi. PUHR tableaux is a ground calculus, i.e., the formulas in a branch do not contain free variables.

PUHR tableaux are constructed from an initially empty branch by applying the inference rule in Figure 2.6. The single PUHR tableaux rule is a sound and refutation complete calculus for clausal specifications.

If a branch θ contains all body atoms $a_1 \dots a_m$ of some ground instance of a clause $C = a_1 \wedge \dots \wedge a_m \Rightarrow b_1 \vee \dots \vee b_n$ in the input specification Φ , then θ is split into n branches, each extended by some atom b_i . In addition, each branch also contains the complements of the atoms b_j with $j > i$. If C has no atoms b_i , i.e., $n = 0$, then the branch is closed. A branch is also closed if it contains an atom and its negation. PUHR tableaux satisfy a regularity condition that forbids the application of a ground clause C if one of the atoms b_i already occurs in the current branch.

The input to the PUHR calculus is a set of **range-restricted rules** $F_1 \wedge \dots \wedge F_m \Rightarrow E_1 \vee \dots \vee E_n$, i.e., first-order clauses with free variables where all variables in $E_1 \vee \dots \vee E_n$ also occur in $F_1 \wedge \dots \wedge F_m$. By using a special domain predicate *dom* for introducing variables on the left-hand

side, all first-order specifications can be translated into this form. A rule is applied only if the left-hand side as a whole can be instantiated with the atoms that occur in the current branch. PUHR tableaux can avoid the effect of the blind instantiation of γ -formulas because implicative rules need only be instantiated if all their preconditions are met.

A branch of a PUHR tableaux is frequently identified with the set of atoms it contains. These atoms define a partial interpretation. A saturated PUHR tableau Θ for a clausal specification Φ has, among others, the following properties:

- Every open branch of Θ is a positive model of Φ .
- Every subset-minimal model of Φ appears as a branch of Θ .
- If two models \mathcal{M} and \mathcal{M}' appear in Θ and $\text{Pos}(\mathcal{M}) \subseteq \text{Pos}(\mathcal{M}')$, then the branch containing \mathcal{M} appears left of the branch of \mathcal{M}' .

PUHR tableaux can be used for the generation of subset-minimal models if non-minimal models are eliminated. In order to do so, we traverse a tableau in a depth-first left-to-right manner and check for every generated model whether it is not a superset of some earlier computed model. This requires a potentially expensive lookup mechanism. We either have to keep track of all models that have been computed earlier, or we must repeat parts of the earlier tableau proof.

PUHR tableaux, like other clausal calculi, must employ Skolemisation in order to have a clausal input that is equivalent to the original first-order specification. As mentioned before, Skolemisation may cause infinite Herbrand models. A saturated PUHR tableau always has a branch for each minimal model, but this theoretical property is often of no relevance in practice because a tableau that contains infinite models cannot be saturated by finitely many expansion steps. In this case, a branch will be expandable infinitely many times, and we will not be able to verify the existence of its model in general.

2.4.3 A Method Complete for Finite Satisfiability

The simple PUHR calculus presented in the last section is in many ways interesting for both theoretical and practical research and has inspired many related works. One of the various derivations based on the PUHR tableaux

$$\frac{\frac{\beta}{\beta_1 \mid \beta_2} \quad \frac{\alpha}{\alpha_1 \mid \alpha_2} \quad \frac{\forall \bar{x} (C(\bar{x}) \rightarrow D)}{D[\bar{c}/\bar{x}]}}{\frac{\exists x F(x)}{F[c_1/x] \mid \dots \mid F[c_k/x] \mid F[c_{new}/x]}}$$

Figure 2.7: The EP calculus

addresses the question of how Skolemisation can be eliminated in order to get completeness for the generation of minimal models. The Extended Positive (EP) tableaux calculus by Bry and Torge [19] uses a special δ -rule which avoids Skolemisation. Fig. 2.7 displays EP's expansion rules.

EP accepts so-called Positive Restricted Quantification (PRQ) rules as its input. These rules are rules of the form $\forall \bar{x}(C \rightarrow \exists \bar{y} F)$ where C is a conjunction of atoms and F is a disjunction of atoms and PRQ rules. The rules do not contain function symbols. Every first-order specification in classical logic can be translated into a PRQ specification.

In the \forall -rule in Fig. 2.7 above, \bar{c} is a tuple of constants occurring in the expanded node. These constants are determined by evaluating C , a conjunction of atoms, against the already constructed interpretation, i.e., the Herbrand model determined by the set of ground positive literals occurring in the current branch. The \forall -rule handles both universally and implicative formulas, and, like the original PUHR rule from which it is derived, only extends the current branch if there is a complete match of the implication's left side with the already constructed interpretation.

The \exists -rule (originally called the δ^* -rule in [19]) instantiates existentially quantified formulas x with constants $\{c_1, \dots, c_k\}$ that occur in the current branch and that form the current domain \mathcal{D} . It also extends the tableau with a new branch that instantiates x with a new constant c_{new} . By this, the use of a Skolem term with a potentially dangerous introduction of function symbols is avoided. The reuse of constants also ensures that EP constructs models with minimal domains first when using a left-to-right traversal of the tableau.

With the elimination of Skolem terms, and function symbols as a whole, EP is complete for finite satisfiability. Like PUHR tableaux, EP is refutation complete and simultaneously searches for models and proofs.

2.4.4 The Davis-Putnam Procedure

All tableaux methods presented in Sections 2.4.1–2.4.3 are refutation complete proof calculi for first-order logics. They not only compute models and refutation proofs simultaneously, but also determine the domain of the models they compute. A quite different approach to model generation is the use of propositional decision procedures where the universe of discourse must be chosen in advance.

Given a domain of individuals \mathcal{D} , one translates the input specification Φ into ground formulas or ground clauses by translating δ -formulas into disjunctions and γ -formulas into conjunctions. The result are propositional theories for which efficient decision procedures are available. A propositional model determines the truth value of all ground literals and can be translated back into a model of the first-order input specification. If the specification is unsatisfiable, one simply extends \mathcal{D} by an additional entity and starts the translation into propositional logics again. By iterative deepening on the size of \mathcal{D} , all finite models of the input can be enumerated.

A simple, yet efficient decision procedure is the Davis-Putnam procedure¹. In its original form, as presented in [29], it was intended to be a theorem-proving technique suitable for automation, covering both classical propositional and first-order logics. The first-order version was not as efficient as resolution, which was introduced soon after, but the propositional version is still among the fastest [34].

Calculus and Procedure

The Davis-Putnam procedure works on ground clauses $\mathbf{a}_1 \vee \dots \vee \mathbf{a}_m \vee \neg \mathbf{b}_1 \vee \dots \vee \neg \mathbf{b}_n$ that we will write in the following as **literal sets** $\{l_1, \dots, l_{n+m}\}$. Like all sets, a literal set does not contain the same elements twice. Additionally, we remove all clauses from our input that contain both some literal l and its complement \bar{l} , i.e., that are tautologies.

Like a tableaux calculus, the Davis-Putnam procedure manipulates sequences of formulas that we organise into branches. Unlike tableaux, however, we do not have a tree-like proof structure, but each branch is kept separately. A branch can be modified destructively by removing formulas. Like in a tableaux, a saturated open branch proves the satisfiability of the

¹It is more precisely referred to as the Davis-Logeman-Loveland procedure [28].

$$\begin{array}{ccc}
\frac{\{l_1, \dots, l_n\}}{\{l_1\}} & \frac{\{l_1, l_2, \dots, l_n\}}{\{l_1\}} & \frac{\{l_1, \dots, l_n\}}{\{l_1\} \mid \{l'_1\}} \\
\text{remove } \{l_1, \dots, l_n\} & \begin{array}{l} \text{remove } \{l_1, \dots, l_n\} \\ \text{add } \{l_2, \dots, l_n\} \end{array} &
\end{array}$$

Figure 2.8: The Davis-Putnam calculus

input. A branch is closed if it contains the empty clause $\{\}$, or a unit clause $\{l_1\}$ together with the complementary clause $\{\bar{l}_1\}$.

Figure 2.8 shows the rules of the Davis-Putnam calculus. The first rule eliminates clauses $\{l_1, \dots, l_n\}$ if there is a unit clause $\{l_1\}$ in the same branch. This is sound because every interpretation that validates $\{l_1\}$ must also validate $\{l_1, \dots, l_n\}$. The second rule is propositional unit resolution. If both $\{l_1, l_2, \dots, l_n\}$ and $\{\bar{l}_1\}$ are in the same branch, we can simplify $\{l_1, l_2, \dots, l_n\}$ to $\{l_2, \dots, l_n\}$ because we know already that literal l_1 must be interpreted as false in the current branch.

The final rule splits a branch if there is a literal l_1 that occurs both positively and as its complement in some non-unit clauses. This results in two new branches, one that contains the unit clause $\{l_1\}$ and one that contains the unit clause $\{\bar{l}_1\}$.

The Davis-Putnam procedure defines an effective order of rule applications for the calculus. The first two rules are always applied exhaustively on a branch until no further simplification is possible. Then, a literal is selected for splitting, and the process of simplification that is defined by the first two rules starts again. As a heuristic, the procedure selects literals that occur in as many other formulas as possible in order to maximize the effect of the simplification rules.

Branches as Models

The output of the Davis-Putnam procedure is a finite set of branches. If at least one of them is open, then the input is satisfiable. Unlike the tableaux calculi discussed before, a single branch can actually represent several models. Consider for instance the set of literal sets $\{\{a, b\}, \{d, e, c\}, \{\neg c\}\}$. This set can be simplified by applying unit resolution to the last two formulas, and we obtain the saturated open branch $\{\{a, b\}, \{d, e\}, \{\neg c\}\}$. The models represented in a branch are defined by all atoms that occur in paths through

$$\frac{\{l_1, \dots, l_n\}}{\{l_1\} \mid \{\bar{l}_1\}}$$

Figure 2.9: The modified splitting rule for enumeration

the set of clauses that select one literal in each clause. In our example, we obtain the following positive models:

$$\{a, d\}, \{a, e\}, \{b, d\}, \{b, e\}$$

As in the case of first-order models, the interpretation I derived from a positive model \mathcal{M} meets the following condition: $I(c) = \text{true}$ iff $c \in \mathcal{M}$ for all atoms c .

With a slight variation of the splitting rule, the Davis-Putnam procedure actually enumerates all models of its input in separate branches. Figure 2.9 shows the rule that can now split a branch on literal l_1 if this literal occurs in any non-unit formula in the branch.

Efficiency

The Davis-Putnam-procedure is the basis of a family of model generation programs (see Section 2.3.4) that all have been used with some success in finite mathematics. The Davis-Putnam programs have been applied to combinatorial problems with very large search spaces that are beyond the capabilities of first-order methods.

At a first glance, it is not obvious why the Davis-Putnam procedure is more effective for propositional decision problems than so many other proof procedures. The splitting rule for instance requires that the set of clauses is split into independent partitions that can be manipulated destructively while a tableaux allows us to keep common parts of different branches in the same branch the tree-like proof structure. In fact, if the modified splitting rule would be applied for all literals before simplification takes place, a Davis-Putnam proof would have the same complexity than a truth-table method, i.e., there would always be 2^n branches with n being the number of different literals in the input problem.

The success of the Davis-Putnam procedure can be explained by the potential of simplification due to the the first two rules. Any unit clause in the current branch is used for simplifying all other clauses in which it occurs.

If a literal l occurs in the same polarity in a clause C in the branch, then C is removed immediately, which shortens the branch. If l occurs as \bar{l} in some non-unit clause C' , then C' can be shortened at least by one literal. These shortened clauses are possibly unit clauses, and can be used immediately for simplification. Splitting on some literal introduces two new branches, each of these simplifiable. As an effect, many potential splittings on literals may never become necessary, either because the literals involved are eliminated from the branch entirely, or because the literal or its complement remain as a unit clause in the branch. In this case, the interpretation of the literal is already given.

2.5 Related Work

There currently is no overview article or an annotated bibliography available that gives a throughout introduction to the various topics, methods, and applications in model generation. Hasegawa [44] focuses on model generation theorem proving such as the PUHR tableaux approach. For finite model generation, the article by Slaney, Fujita and Stickel [83] sketches the techniques of different programs, while Zhang and Stickel [91] explain in detail how the Davis-Putnam procedure can be implemented efficiently.

There are several model generation programs available that have been developed in the automated deduction community. A family of model generators is based on the Davis-Putnam procedure, for instance LDPP, SATO [91], and MACE [68]. Both SEM [93] and its predecessor FALCON [92] use techniques from constraint solving instead or in addition to those that have been pioneered by FINDER [82].

MGTP [45] is a model generation theorem prover that is based on a PUHR-like tableaux calculus. The system implements some interesting refinements of PUHR, for instance an optimised proof procedure for horn clauses.

While PUHR tableaux as used in SATCHMO and MGTP is a ground calculus, the hyper-tableaux derivation of PUHR by Baumgartner et al. uses free variables [6, 8] in order to suppress unnecessary ground instantiations when proving theorems. Also for theorem proving applications, variants of SATCHMO have been developed that apply rules in a more goal oriented way by ordering them with respect to relevancy measure [64]. Abdennadher and Schütz [1] present a PUHR tableaux calculus that can handle exis-

tentially quantified variables whose domain of instantiation is restricted by constraints.

This thesis concentrates on the generation of *finite* Herbrand models for natural-language semantic representations. There are some methods available in other application areas that generate certain classes of infinite models. The methods proposed by Klingenbeck [60] and Peltier [75] investigate languages with exponents that have finite representations. For an overview of techniques in infinite model generation, we refer to Klingenbeck [60].

There are many model generation methods that are related to some form of first-order (hyper)resolution (see [33] for a summary).

Unlike proof presentation in theorem proving, the computation of an appropriate representation of models or parts thereof is a relatively new research topic in model generation. Horacek and Konrad [51] propose linguistically motivated techniques for presenting finite Herbrand models in a human-oriented way.

Chapter 3

Higher-Order Model Generation

By keenly confronting the enigmas that surround us, and by considering and analysing the observations I had made, I ended up in the realms of mathematics. (M.C. Escher)

Overview: The simply typed λ -calculus can be taken as the basis of a formalism in which we assign meaning to the basic expressions of a natural language and explain the meaning of larger constituents by the composition of the meaning of their parts. However, under a conventional higher-order semantic for the λ -calculus, we cannot give a model generation method based on the well-known techniques of first-order model generation. By weakening the semantic of higher-order logic, we formulate a logical language that has the compositional expressivity of a higher-order logic but the finitely representable models of a first-order one. A generic model generation technique for this language is presented and discussed, and a refinement for minimal model generation is introduced.

3.1 The λ -Calculus in Linguistics

The simply typed λ -calculus has been invented by Church [23] with the goal of providing a uniform language with which to describe functions. In the 1950s, its untyped variant was the starting point for the functional programming paradigm in computer science, notably as the theoretical basis of the programming language LISP. It has become a standard tool of computational semantics at least since the 1970s when Montague [70] first introduced his theory of quantification in natural language.

The simply typed λ -calculus is an expressive, elegant, and uniform method of composing functions out of more primitive functions. The tools for building its complex expressions are function application and λ -abstraction. The operation of β -reduction is the essence of computation, whereas λ -abstraction is the essence of function definition.

In computational semantics, the λ -calculus can be taken as the basis of a formalism in which we assign a formal meaning to the basic expressions of a natural language and explain the meaning of larger constituents by the composition of the meaning of their parts.

3.1.1 Composition of Meaning

In a standard first-order predicate logic, we can represent the meaning of a sentence like *a man loves a woman* as a formula, e.g., as $\exists x \exists y \text{man}(x) \wedge \text{woman}(y) \wedge \text{love}(x, y)$ where the non-logical constants *man*, *woman* and *love* are given a postulated meaning that corresponds to the meaning of the associated words in the natural language.

What we cannot represent directly as one first-order logical expression is the contribution that the meaning of the verb phrase has to the meaning of the whole sentence. This means that some constituent like *loves a woman* has no corresponding first-order logical form, at least not without additional pieces of machinery. In this sense, standard first-order predicate logic lacks **compositionality**.

One of the primary goals of semantics construction is to develop a truly compositional method for constructing meaning representations that define the meaning of complex constituents in terms of the meaning of its parts. The stumbling stone for giving any such compositional construction mechanism for natural-language semantics is quantification.

Quantification in Natural Language

Russell [79] notes the disturbing variety among the logical contributions of subject constituents to simple sentences. Consider for instance the following sentences.

Sentence	Logical Form
Peter died	$died(peter)$
A man died	$\exists x man(x) \wedge died(x)$
The man died	$died(\iota x man(x))$
One man died	$\exists x man(x) \wedge died(x) \wedge \forall y man(y) \wedge died(y) \Rightarrow x = y$
Every man died	$\forall x man(x) \Rightarrow died(x)$

As we can see, structurally small changes in the subject constituent of the natural language sentences lead to very different logical forms. The variety of logical representations seemingly prohibits a compositional method for constructing the semantics of sentences.

The contemporary solution to this problem is based on the use of **generalised quantifiers** and **generalised determiners**. A generalised quantifier applies to a property and produces a truth value. Such quantifiers in natural language are proper names like *Peter* that in the sentence *Peter died* takes a property, *died*, in order to produce the truth of the sentence *Peter died*. A generalised determiner can then be taken to be a relation between two or more sets. In the case of diadic determiners we have two sets, one contributed by the restriction from the noun and one contributed by the predicate that supplies the scope of quantification. In the sentence *every man died*, the word *every* is a diadic determiner that formulates a relation between the set of men and the set of died entities. Like *Peter*, the phrase *every man* is a quantifier, because it applies to a property. What is puzzling in natural language is the infinite variety of such quantifiers that can be constructed easily from the set of available determiners and nouns, and the variety of logical contributions they show when sentences are translated into first-order logic. A compositional treatment of linguistic quantification requires some uniform way of constructing complex generalised quantifiers from their parts. Such a uniform way exists for the λ -calculus.

3.1.2 Quantifiers as Higher-Order Expressions

The simply typed λ -calculus can represent generalised quantifiers as second-order functions that take a set and return a truth value. Indeed, we can even

calculate suitable semantic representations for quantifiers by solving higher-order equations. In the following, the quantifiers *every man* and *some man* are treated as higher-order variables EVERYMAN^* and SOMEMAN^* . Typical equations that we have to solve are

$$(1.1) \quad \text{EVERYMAN}^*(died) = \forall x \text{ man}(x) \Rightarrow died(x)$$

$$(1.2) \quad \text{SOMEMAN}^*(died) = \exists x \text{ man}(x) \wedge \neg died(x)$$

The right-hand side of the equation is the logical form that we would like to have for the whole sentence. Solving the equations gives us the two results

$$(1.3) \quad \text{EVERYMAN}^* = \lambda P \forall x \text{ man}(x) \Rightarrow P(x)$$

$$(1.4) \quad \text{EVERYMAN}^* = \lambda P \forall x \text{ man}(x) \Rightarrow died(x)$$

for the first equation and

$$(1.5) \quad \text{SOMEMAN}^* = \lambda P \exists x \text{ man}(x) \wedge P(x)$$

$$(1.6) \quad \text{SOMEMAN}^* = \lambda P \exists x \text{ man}(x) \wedge died(x)$$

for the second equation. Of these four solutions, only two, namely (1.3) and (1.5), are linguistically valid. The other two are vacuous abstractions that do not correspond to the semantics of the natural-language quantifiers. The application of result (1.4) to the representation of a different verb shows why we must discard such solutions.

$$(1.7) \quad \begin{array}{ll} \text{EVERYMAN}^*(sleep) & =_{def} \\ (\lambda P \forall x \text{ man}(x) \Rightarrow died(x))(sleep) & =_{\beta} \\ \forall x \text{ man}(x) \Rightarrow died(x) & \end{array}$$

Clearly, the resulting first-order logical form is not the intended semantic representation of *every man sleeps*. Vacuous abstractions can be dealt with automatically by extending higher-order equational reasoning with syntactic filter mechanisms [38].

By discarding the linguistically invalid solutions, we obtain some suitable semantics for the linguistic quantifiers *every man* and *some man* in the form of λ -terms. But how are such semantics composed from their parts? Now that

we have the logical form of the quantifiers at our hands, we simply use higher-order equational reasoning again for computing the semantics of their parts, namely of the determiners *every* and *some*. Their semantic representations, when applied to the representation of the noun *man*, should yield the non-vacuous quantifiers that we have computed above. Hence, the higher-order equations that we must solve now are as follows.

$$(1.8) \quad \text{EVERY}^*(\mathit{man}) = \lambda P \forall x \mathit{man}(x) \Rightarrow P(x)$$

$$(1.9) \quad \text{SOME}^*(\mathit{man}) = \lambda P \exists x \mathit{man}(x) \wedge P(x)$$

By solving these, and discarding the vacuous solutions, we obtain definitions in higher-order logic for our determiners.

$$(1.10) \quad \text{EVERY} \equiv \lambda Q \lambda P \forall x Q(x) \Rightarrow P(x)$$

$$(1.11) \quad \text{SOME} \equiv \lambda Q \lambda P \exists x Q(x) \wedge P(x)$$

With these definitions, we can represent simple sentences like *every man drinks* or *no Greek lies* in a compositional way that is closer to the syntax of the original sentence than a standard first-order logical form. At the same time, we can always reduce the representation to a purely first-order logical form by expanding the definition and performing β -reduction.

$$(1.12) \quad \begin{array}{ll} \text{EVERY}(\mathit{man})(\mathit{drink}) & =_{def} \\ (\lambda Q \lambda P \forall x Q(x) \Rightarrow P(x))(\mathit{man})(\mathit{drink}) & =_{\beta} \\ (\lambda P \forall x \mathit{man}(x) \Rightarrow P(x))(\mathit{drink}) & =_{\beta} \\ \forall x \mathit{man}(x) \Rightarrow \mathit{drink}(x) & \end{array}$$

One of the central ideas of Montague's approach is the treatment of proper names such as *Peter*. As mentioned earlier, proper names act as quantifiers in sentences. At the same time, they appear as individual constants in the logical form. This discrepancy is resolved by introducing quantifiers for proper names that, when applied to their scope of quantification, introduce the first-order constant into the logical form. For instance, the proper name *Peter* is represented by a quantifier $\text{PETER} \equiv \lambda P P(\mathit{peter})$ where *peter* is an individual constant. By raising the type of proper names to quantifiers, we can treat them as required by a compositional approach.

$$(1.13) \quad \begin{array}{ll} \text{PETER}(\mathit{drink}) & =_{def} \\ (\lambda P P(\mathit{peter}))(\mathit{drink}) & =_{\beta} \\ \mathit{drink}(\mathit{peter}) & \end{array}$$

3.1.3 First-Order Limitations

It seems that we now have the best of two worlds. While a compositional higher-order logic is used at the representation level, we actually still have a standard first-order logical form that is immediately computable by definition expansion and β -reduction. However, the approach has its limits in that sentences with certain quantifiers cannot be reduced to first-order formulas at all. Consider the following.

$$(1.14) \quad \mathbf{More} \text{ boys than girls drink.}$$

MORETHAN is a three-place generalised determiner that takes three sets instead of two. Its definition can be derived from the following equation.

$$(1.15) \quad \text{MORETHAN}^*(\text{boy})(\text{girl})(\text{drink}) = |\text{boy} \cap \text{drink}| > |\text{girl} \cap \text{drink}|$$

The right-hand side of the equation states that the set of boys that drink has a greater cardinality than the set of girls that drink. The linguistically valid solution of the equation is

$$(1.16) \quad \text{MORETHAN} \equiv \lambda P \lambda Q \lambda R |P \cap R| > |Q \cap R|$$

Unlike as in previous examples, we cannot give a higher-order formalisation of cardinality comparison such that a definition expansion of MORETHAN will produce a first-order formula. The reason for this is a result of the compactness of first-order logics which prohibits us to formulate conditions which distinguish finite and infinite sets. From this, it follows that comparing the cardinality of two first-order predicates that might denote arbitrary sets is not directly expressible in first-order predicate calculus. There is, however, an indirect approach that makes use of a first-order axiomatisation of mathematical set theory. Such an axiomatisation¹ treats sets as first-order entities, and specifies their properties in terms of first-order formulas. Axiomatic set theory can be taken as the basis of a formalisation of mathematics as a whole. Naturally, we would therefore be able to give an expression that captures the semantic of the sentence (1.14). Unfortunately, all models that depend on a first-order axiomatisation of set theory have an infinite universe of discourse

¹A first-order formalisation of set theory that has been developed specifically for automated deduction can be found in Boyer et al. [16].

because they must have a distinct entity in the universe for every set. Hence, for our purposes, such an approach is not practicable.

What is really disturbing about our discovery here is that the statement (1.16) itself does not look more “difficult” to model in any way than previous examples. Given a situation in which we have n boys and m girls who either drink or do not drink, it should be very easy to verify whether the statement holds or not. Yet, we cannot give a self-contained first-order formalisation that expresses the simple cardinality constraint of the set of boys and the set of girls that we have in our example.

3.1.4 A Motivation for a New Kind of Logic

The simply typed λ -calculus gives us a formalism for representing not only the semantics of complete sentences, but also of the semantics of their parts. Functional application and β -reduction provide a simple, yet powerful compositional construction mechanism for semantic representations. By using higher-order definitions for quantifiers and determiners, we obtain compositional representations that still are reducible to a standard first-order format in many cases.

Nevertheless, certain forms of quantification do not have first-order formalisations that are suitable for the application of conventional model generation methods. First-order compactness prohibits even simple cardinality constraints such as needed for the determiner *more*. This is disturbing for our intended application of model generation in the interpretation of natural language.

What we would like to have for natural-language semantics is an expressive language in which we can experiment with higher-order definable concepts, such as generalised determiners, in the usual way. For this, as we have seen, we cannot stay in the domain of first-order logic in all cases. On the other hand, there is, to my knowledge, no model generation method at all that can deal with higher-order logic in any form. The technical difficulties of model generation in the context of higher-order semantics will be discussed in Section 3.2.3. Here, I will sketch only one, namely that conventional higher-order models can often not be reduced to the interpretation of predicate symbols over *small* universes of discourse. The performance of all model generation methods presented in Chapter 2 depends on having small domains of individuals, whether they are derived from the Herbrand interpretation of term occurrences or being given adjacently as in the case of finite

model generators. First-order models are often small and simple structures because they must only define the interpretation of predicates and functions over such small domains. In higher-order logics, the argument of a function symbol can be any expression of appropriate type, and quantification can range not only over individuals, but also over entities of higher type. Even in cases where we only have a small number of individuals in our universe of discourse, the number of distinct functions that can be derived from them grows exponentially with both the number of individuals and the arity of the functions involved. Any model computation that must deal with complete domains of functions instead of just a domain of individuals will inevitably face intractable combinatorial problems.

Our intended application requires that we identify an interesting fragment of higher-order logic that captures the linguistically motivated forms of quantification that we have mentioned earlier as well as being open to efficient methods of model generation. This is possible by considering a higher-order logic where the domains of all entities of arbitrary type can be restricted to finite (and small) subsets of the full domain that is given by the standard semantics. In other words, we aim at a higher-order logic whose notion of quantification is not only explicitly limited to finite domains, but whose range of quantification can further be restricted such that higher-order forms of quantification stay tractable. In a model where the denotations of predicate symbols are guaranteed to be finite sets, we will be able to formalise the truth conditions that we need for certain forms of natural-language quantification, and find computational means to deal with them when generating models.

There is a price that we will have to pay for this additional expressivity: we will no longer be able to formally capture the meaning of infinite concepts or discourse situations with infinitely many participants. For the overwhelming majority of natural-language utterances, this restriction plays no rôle. Our working hypothesis is that a concentration on finiteness is even a necessity for exploring practicable model generation methods in computational semantics.

3.2 Higher-Order Logic

In the following, the syntax and semantics of higher-order logics are presented in order to provide a compact reference. Apart from the notational conventions, most of the introduced concepts are as usual, and are discussed in more detail for instance in Barendregt's textbook on the λ -calculus [5].

3.2.1 Syntax

The simply typed λ -calculus is a formal language whose expressions, i.e., whose terms, are composed from constants, variables, function applications and λ -abstractions. Each of these terms carries a type, i.e., a symbolic annotation which makes sure that the functions defined in our language can only be applied to arguments of the appropriate domain.

Types

We have a set `BASETYPE` of **base types** that consists of the **type ι of individuals** and the **type o of truth values**. The set `TYPE` of **types** then is defined as the smallest set such that

- `BASETYPE` \subseteq `TYPE`, and
- if $\alpha \in \text{TYPE}$ and $\beta \in \text{TYPE}$, then $\alpha \rightarrow \beta \in \text{TYPE}$.

The **type constructor** \rightarrow is right-associative, and a type of the form $\iota \rightarrow (\iota \rightarrow o)$ for instance will be written as $\iota \rightarrow \iota \rightarrow o$.

Unless indicated otherwise, we will use α and β for denoting arbitrary types, and Φ_α denotes any syntactic or semantic structure Φ whose type is α . The basic expressions of the simply typed λ -calculus are typed constants and typed variables. For these symbols we will use the following notational conventions.

Symbol	Type	Denotation
$c, jon, peter, \dots$	ι	individual constants
x, y, z, \dots	ι	individual variables
$p, love, man, \dots$	$\alpha_1 \rightarrow \dots \rightarrow \alpha_n \rightarrow o$	predicate constants
P, Q, R, \dots	$\alpha_1 \rightarrow \dots \rightarrow \alpha_n \rightarrow o$	predicate variables
$C_\alpha, D_\alpha, \dots$	α	constants of arbitrary type
$X_\alpha, Y_\alpha, \dots$	α	variables of arbitrary type

Terms

We assume a countable signature Σ of constant symbols, and a countable set of variables ∇ . The set $\text{TERM}_{\nabla}^{\Sigma}$ of well-formed typed λ -terms with respect to Σ and ∇ is the smallest set such that

- if $X_{\alpha} \in \nabla$, then $X_{\alpha} \in \text{TERM}_{\nabla}^{\Sigma}$
- if $C_{\alpha} \in \Sigma$, then $C_{\alpha} \in \text{TERM}_{\nabla}^{\Sigma}$
- if $\mathbb{T}_{\alpha \rightarrow \beta} \in \text{TERM}_{\nabla}^{\Sigma}$ and $\mathbb{U}_{\alpha} \in \text{TERM}_{\nabla}^{\Sigma}$, then $(\mathbb{T}_{\alpha \rightarrow \beta} \mathbb{U}_{\alpha})_{\beta} \in \text{TERM}_{\nabla}^{\Sigma}$
- if $\mathbb{T}_{\beta} \in \text{TERM}_{\nabla}^{\Sigma}$ and $X_{\alpha} \in \nabla$, then $(\lambda X_{\alpha} \mathbb{T}_{\beta})_{\alpha \rightarrow \beta} \in \text{TERM}_{\nabla}^{\Sigma}$

If the type of a well-formed λ -term is given by the context or implied by our notation, we avoid its explicit mention. We will use \mathbb{T} , \mathbb{U} , and \mathbb{W} for denoting arbitrary λ -terms. A term of the form $(\mathbb{T} \mathbb{U})$ is called a **function application**, and one of the form $\lambda X \mathbb{T}$ a **λ -abstraction**. To ease readability, we follow the usual convention for λ -terms and leave out brackets in every case where the construction of an expression is uniquely determined.

In a λ -abstraction $\lambda X \mathbb{T}$, the variable X is a **bound variable**. The set $\text{free}(\mathbb{T})$ denotes the free variables in \mathbb{T} , i.e., those variables in \mathbb{T} which are not bound. A λ -term \mathbb{T} is a **closed λ -term** iff $\text{free}(\mathbb{T})$ is empty, otherwise it is open.

Some λ -terms are in a structural equality relation that is induced by **$\beta\eta$ -reduction**:

$$(\lambda X \mathbb{T})\mathbb{U} \longrightarrow_{\beta} [\mathbb{U}/X]\mathbb{T} \quad (\lambda X \mathbb{T}X) \longrightarrow_{\eta} \mathbb{T}$$

where X is not free in \mathbb{T} , and $[\mathbb{U}/X]\mathbb{T}$ denotes the substitution of all free occurrences of X in \mathbb{T} by \mathbb{U} . It is well known that the reduction relations β , η , and $\beta\eta$ are terminating and confluent, so that we have unique normal forms for λ -terms. We chose the $\beta\eta$ -normal form as the standard syntactical form of all λ -terms.

Depending on which notational variant is more elegant or clear, we will sometimes write $((\dots (\mathbb{T} \mathbb{U}_n) \mathbb{U}_{n-1}) \dots) \mathbb{U}_1$ either as $\mathbb{T}(\overline{\mathbb{U}_n})$, as $\mathbb{T}(\mathbb{U}_1, \dots, \mathbb{U}_n)$, or as $\mathbb{T}(\mathbb{U}_n) \dots (\mathbb{U}_1)$. Function application is considered to be left-associative.

3.2.2 Semantics

The simply typed λ -calculus can be taken as the basis of different logics. The semantic of each such logic is defined by the postulated meaning of its logical constants. In what follows, we first describe the interpretation of general λ -terms without considering a logical content.

Functional Interpretations

A **frame** $\{\mathcal{D}_\alpha \mid \alpha \in \text{TYPE}\}$ is a collection of domains for types α . A **function domain** $\mathcal{D}_{\alpha \rightarrow \beta}$ is a domain that consists of functions $f : \mathcal{D}_\alpha \rightarrow \mathcal{D}_\beta$. It is common practice to assume that the elements of function domains are total functions. Alternatively, frames can also be based on partial functions, as has been done for instance by Farmer [32].

An **interpretation of the simply typed λ -calculus** with respect to a signature Σ is a pair $\mathcal{I} = \langle \llbracket \cdot \rrbracket, \mathcal{D} \rangle$ where \mathcal{D} is a frame and $\llbracket \cdot \rrbracket$ is an evaluation that assigns to each constant $C_\alpha \in \Sigma$ an object in \mathcal{D}_α .

A **variable assignment** σ for the set of variables ∇ is a mapping from variables $X_\alpha \in \nabla$ into a domain \mathcal{D}_α of objects of appropriate type. A variable assignment $\sigma \cup \{X_\alpha := a\}$ denotes a mapping σ' derived from σ where $\sigma'(Y) = a$ if $Y = X$, and $\sigma'(Y) = \sigma(Y)$ otherwise.

The **denotation** $\llbracket \mathbb{T} \rrbracket_{\mathcal{I}}^\sigma$ for an arbitrary λ -term $\mathbb{T} \in \text{TERM}_{\nabla}^\Sigma$ with regard to an interpretation $\mathcal{I} = \langle \llbracket \cdot \rrbracket, \mathcal{D} \rangle$ and a variable assignment σ is recursively defined as follows.

1. $\llbracket C \rrbracket_{\mathcal{I}}^\sigma = \llbracket C \rrbracket$ if $C \in \Sigma$
2. $\llbracket X \rrbracket_{\mathcal{I}}^\sigma = \sigma(X)$ if $X \in \nabla$
3. $\llbracket (\mathbb{U} \mathbb{W}) \rrbracket_{\mathcal{I}}^\sigma = \llbracket \mathbb{U} \rrbracket_{\mathcal{I}}^\sigma (\llbracket \mathbb{W} \rrbracket_{\mathcal{I}}^\sigma)$
4. $\llbracket (\lambda X_\alpha \mathbb{U}_\beta) \rrbracket_{\mathcal{I}}^\sigma = f \in \mathcal{D}_{\alpha \rightarrow \beta}$ such that $f(a) = \llbracket \mathbb{U} \rrbracket_{\mathcal{I}}^{\sigma \cup \{X := a\}}$ for all $a \in \mathcal{D}_\alpha$

We assume in the definition of denotations that we are dealing with well-typed terms, e.g., that $(\mathbb{U} \mathbb{W})$ denotes the application of a function term $\mathbb{U}_{\alpha \rightarrow \beta}$ to a term \mathbb{W}_β . The denotation of a closed term does not depend on the initial variable assignment σ , and we will simply write $\llbracket \mathbb{T} \rrbracket_{\mathcal{I}}$ for the denotation $\llbracket \mathbb{T} \rrbracket_{\mathcal{I}}^\sigma$ of a closed term \mathbb{T} with respect to some arbitrary assignment σ .

Logical Constants

A **logic** based on the simply typed λ -calculus is a set of definitions that determines the denotation of the set of **logical constants** in the signature Σ .

In a higher-order logic, we will have at least the logical constants $\bigvee_{o \rightarrow o}$, $\neg_{o \rightarrow o}$, and $\Pi_{(\alpha \rightarrow o) \rightarrow o}$ for all types α . All logical constants are given a postulated meaning, i.e., a meaning that fixes their denotation in all interpretations.

For the purpose of this thesis, we assume a family of logical constants in Σ that are distinguished by their type. Depending on our linguistic applications, we will later extend our set of logical constants as needed. For now, we have

- **unary logical connectives** of type $o \rightarrow o$, e.g., \neg ,
- **binary logical connectives** of type $o \rightarrow o \rightarrow o$, e.g., \vee , \wedge , \Rightarrow , \Leftrightarrow ,
- **monadic quantifier constants** of type $(\alpha \rightarrow o) \rightarrow o$, e.g., \forall and \exists , for all types α ,
- **diadic quantifier constants** of type $(\alpha \rightarrow o) \rightarrow (\alpha \rightarrow o) \rightarrow o$, e.g., EVERY, SOME, etc., for all types α , and
- an **equality sign** = of type $\alpha \rightarrow \alpha \rightarrow o$ for all types α .

The somewhat exotic diadic quantifiers are inspired from the linguistic theory of quantification in natural language where they are known as *generalised determiners*. We will later use the diadic quantifiers for exemplifying how linguistic quantification can be encoded as constraints over finite models.

All constants that are not logical constants are called **parameter constants**, short **parameters**.

Defining a Logic

In order to define a logic, we must first fix a domain of truth values \mathcal{D}_o that appears in each frame \mathcal{D} of the logic. For a classical two-valued logic, we chose the domain \mathcal{D}_o to be $\{0, 1\}$ where 1 denotes truth and 0 denotes falsity.

As mentioned earlier, different logics may be based on varieties of logical constants and denotations for them. The semantic of a logical constant for

a certain logic can be defined as a constraint that describes how a formula governed by the constant is to be interpreted relative to the formula's parts or instantiations. For instance, a basic higher-order logic \mathcal{HOL} can be defined by giving the constraints of a minimal set of logical constants \neg , \vee , and $\Pi_{(\alpha \rightarrow \mathfrak{o}) \rightarrow \mathfrak{o}}$ as follows.

- $\llbracket \neg \mathbb{T} \rrbracket_{\mathcal{I}}^{\sigma} = 1 - \llbracket \mathbb{T} \rrbracket_{\mathcal{I}}^{\sigma}$
- $\llbracket \vee(\mathbb{T}, \mathbb{U}) \rrbracket_{\mathcal{I}}^{\sigma} = \max(\llbracket \mathbb{T} \rrbracket_{\mathcal{I}}^{\sigma}, \llbracket \mathbb{U} \rrbracket_{\mathcal{I}}^{\sigma})$
- $\llbracket \Pi_{(\alpha \rightarrow \mathfrak{o}) \rightarrow \mathfrak{o}}(\mathbb{T}) \rrbracket_{\mathcal{I}}^{\sigma} = 1$ iff $\llbracket \mathbb{T} \rrbracket_{\mathcal{I}}^{\sigma}(u) = 1$ for all $u \in \mathcal{D}_{\alpha}$

To ease our notation, we will from now on use infix notation whenever a logical constant denotes a binary connective, and adopt the common notation for quantification that is known from first-order predicate calculus. For instance, we will write $\mathbb{T} \vee \mathbb{U}$ instead of $\vee(\mathbb{T}, \mathbb{U})$, and $\forall X \rho(X)$ instead of $\forall(\lambda X \rho(X))$ or $\forall(\rho)$.

The set of standard logical constants that are usually found in higher-order logics can be derived from the basic set above by using **higher-order definitions**, i.e., by λ -terms that replace all occurrences of the defined constants and that make use only of other logical constants whose denotation is already given. The definitions are as follows.

- $\wedge \equiv \lambda X \lambda Y \neg(\neg X \vee \neg Y)$
- $\Rightarrow \equiv \lambda X \lambda Y \neg X \vee Y$
- $\Leftrightarrow \equiv \lambda X \lambda Y (X \Rightarrow Y) \wedge (Y \Rightarrow X)$
- $\forall \equiv \lambda P_{\alpha \rightarrow \mathfrak{o}} \Pi(P)$
- $\exists \equiv \lambda P_{\alpha \rightarrow \mathfrak{o}} \neg \Pi(\neg P)$
- $\text{EVERY} \equiv \lambda P_{\alpha \rightarrow \mathfrak{o}} \lambda Q_{\alpha \rightarrow \mathfrak{o}} \forall X_{\alpha} P(X) \Rightarrow Q(X)$
- $\text{SOME} \equiv \lambda P_{\alpha \rightarrow \mathfrak{o}} \lambda Q_{\alpha \rightarrow \mathfrak{o}} \exists X_{\alpha} P(X) \wedge Q(X)$
- $= \equiv \lambda Q_{\alpha} \lambda R_{\alpha} \forall P_{\alpha \rightarrow \mathfrak{o}} P(Q) \Rightarrow P(R)$

A **formula** in \mathcal{HOL} is a λ -term of type o . Unless indicated otherwise, all formulas that we consider are closed. An **atom** in \mathcal{HOL} is a formula $h(\overline{u}_n)$ in $\beta\eta$ -normal form that does not have a logical constant as its head h . A **literal** in \mathcal{HOL} is a \mathcal{HOL} atom or its negation.

Our definition of the equality signs $=_{\alpha \rightarrow \alpha \rightarrow o}$ is a higher-order formalisation of equality that has been proposed by Leibniz: two things are equal iff they have the same properties.

Standard Frames and Generalised Interpretations

So far, we have not given a precise notion of the content of the function domains $\mathcal{D}_{\alpha \rightarrow \beta}$ in our interpretations for the simply typed λ -calculus. The standard assumption is that every total function from \mathcal{D}_α to \mathcal{D}_β is an element of $\mathcal{D}_{\alpha \rightarrow \beta}$. The frames that are given by this construction are called **standard frames**.

In a higher-order logic that is based on standard frames, the denotation f of a λ -abstraction $\lambda X_{\alpha \rightarrow \beta} U_\gamma$ is a total function from the complete set of total functions $\mathcal{D}_{\alpha \rightarrow \beta}$ into the domain \mathcal{D}_γ . Standard frames for higher-order logic suffer from the drawback that there can be no calculus that axiomatises logical consequence. In other words, there is no complete proof theory for a higher-order logic based on standard frames. A method for demonstrating this incompleteness of higher-order logic is to give an encoding of arithmetic and applying Gödel's incompleteness theorem [3].

Henkin [47] introduced a weaker notion of higher-order semantics where the domains $\mathcal{D}_{\alpha \rightarrow \beta}$ may consist only of a subset of all total functions from \mathcal{D}_α to \mathcal{D}_β . In order to have a denotation for each λ -term, the lower bound for each of these subsets is the set of all functions of type $\alpha \rightarrow \beta$ that are expressible as λ -terms, i.e., as computable functions. The frames obtained in this way are called **generalised frames**, and a **generalised interpretation** is an interpretation $\mathcal{I} = \langle \llbracket _ \rrbracket, \mathcal{D} \rangle$ whose frame \mathcal{D} is a generalised one.

In a higher-order logic that is based on generalised interpretations, sound and complete proof theories for the classical consequence relation can be given. This even leads to mechanisable calculi, e.g., the variant of extensional higher-order resolution by Benzmüller and Kohlhasse [13]. In recent years, Henkin's semantic for higher-order logic has become the standard theoretical basis of higher-order automated deduction.

3.2.3 Model Generation for Generalised Frames?

We call an interpretation $\mathcal{I} = \langle \llbracket \cdot \rrbracket, \mathcal{D} \rangle$ a **model of a set of formulas** Φ if $\llbracket F \rrbracket_{\mathcal{I}} = 1$ for all formulas F in Φ .

It would be very desirable to have a model generation method for a higher-order logic whose semantic is based on Henkin's generalised interpretations. A model generator for higher-order logic would take a finite set Φ of \mathcal{HOL} formulas as its input and determine interpretations $\mathcal{I} = \langle \llbracket \cdot \rrbracket, \mathcal{D} \rangle$ such that $\llbracket F \rrbracket_{\mathcal{I}} = 1$ for all F in Φ . Finite model generation for higher-order logic is possible at least in theory because all function domains $\mathcal{D}_{\alpha \rightarrow \beta}$ that are based on finite basic domains \mathcal{D}_ι and \mathcal{D}_o are finite themselves, as has been shown by Andrews [2]. The functions f from a finite domain \mathcal{D}_α into a finite domain \mathcal{D}_β have a finite representation, for instance as a table. Hence, we could enumerate all models for a specification Φ simply by enumerating all possible interpretations \mathcal{I} of the constants C that occur in the input and verifying in finite time whether the interpretation is a model or not. Of course, such a method would be intractable in general for all except trivial input specifications.

If we could design a more efficient method, higher-order model generation could be put to use for instance to compute counterexamples in higher-order automated theorem proving. Such a method, if it exists, would also be very handy for exploring the meaning of natural-language semantic representations that have higher-order logical forms, i.e., for the topic of research that this thesis is about. Unfortunately, there is some evidence that model generation for generalised models is not mechanisable by using the technical machinery that is available to us.

Equivalency for Higher-Order Atoms

In first-order logics, finite model generation methods can be reduced to methods that compute the interpretation of predicate symbols over some finite domains of first-order entities. A model is determined completely by the interpretation of a set of atoms of the form $p(a_1, \dots, a_n)$ where the a_n are symbolic, i.e., are constants, ground terms or entities of a finite domain. In the case of SATCHMO-like model generation theorem proving, the a_j can be complex terms t , but these are interpreted as themselves and one does not have to consider the interpretation of function symbols. Hence, it is trivial to decide whether two ground literals of the form $p(a_1, \dots, a_n)$ and

$\neg p(b_1, \dots, b_n)$ are contradictory. The efficiency of first-order model generation methods relies on having a simple way for detecting such elementary inconsistencies while they construct partial interpretations.

In the simply typed λ -calculus, the arguments of predicates can be arbitrarily complex terms or even formulas. Unlike as in first-order logics, the equality of embedded terms cannot be reduced to some syntactical equality relation. For instance, in a higher-order logic \mathcal{HOL} that is based on generalised frames, the following formulas are **logically equivalent**, i.e., must have the same denotation in all models.

$$(2.1) \quad p_{\sigma \rightarrow \sigma}(a_\sigma)$$

$$(2.2) \quad p_{\sigma \rightarrow \sigma}((\lambda x a_\sigma)(c_i))$$

$$(2.3) \quad p_{\sigma \rightarrow \sigma}(a \wedge (b \vee \neg b))$$

$$(2.4) \quad p_{\sigma \rightarrow \sigma}((b \vee \neg b) \wedge a)$$

$$(2.5) \quad p_{\sigma \rightarrow \sigma}(\exists x ((q(x) \vee a) \vee \neg q(x)))$$

The formulas (2.1) and (2.2) have the same $\beta\eta$ -normal form. Their logical equivalency is therefore given by the syntactically determined equivalency of formulas induced by $\beta\eta$ -reduction. In contrast to this, a purely syntactical approach is not sufficient for proving the equivalency of (2.1) to all other formulas. The equivalency is given only semantically by the denotations of the logical constants. In our examples, the proof problems are quite simple, but as embedded terms can be arbitrarily complex formulas, the resulting proof problems easily become undecidable. If we have two literals $p(\top)$ and $\neg p(\perp)$, we cannot decide in general whether they form an elementary contradiction or not.

There is an analogous problem in higher-order automated theorem proving where the first-order variant is efficiently computable, while the higher-order instance of the same operation is undecidable. The operation we refer to is *unification* which is essential for finding non-trivial proofs in many proof procedures. In higher-order automated theorem proving, unification problems are frequently not solved by simply calling a procedure, but are treated as constraints of the proof problem and thus become part of the overall proof search. By this, one avoids to call an external unification procedure which,

in general, might not terminate. In the same way, one could treat the equivalency problem in higher-order model generation as a set of constraints: if two formulas $p(a_1, \dots, a_n)$ and $\neg p(b_1, \dots, b_n)$ have been computed as part of a partial interpretation, we could simply add some additional formulas to the input specification that make sure that the tuples $\langle a_1, \dots, a_n \rangle$ and $\langle b_1, \dots, b_n \rangle$ do not have the same denotation. For instance, if we have two formulas $p_{a \rightarrow o}(a_o)$ and $\neg p_{a \rightarrow o}(b_o)$ in a partial interpretation, we prevent a and b from having the same denotation by adding the formula $\neg(a \Leftrightarrow b)$ to our input. In any model that we obtain for the extended specification, a and b must denote different truth values. In cases where our additional conditions can not be satisfied, we will not be able to compute a model, as required.

Nevertheless, a treatment of equivalence in higher-order models still is intractable in general. If a partially determined interpretation contains n positive occurrences of a formula with head p and m negative occurrences, we have $n \times m$ pairs that generate additional formulas for the input problem. These formulas might further add to the complexity of the model generation problem, because the arguments of a predicate could be formulas that introduce new equivalency conditions themselves. It seems that a simple and efficient treatment of the basic formulas of a logic such as we have in first-order finite model generation is not possible for higher-order logics.

Function Domains and Quantification

One of the criteria in first-order finite model generation for the complexity of a model generation task is the number of ground atoms whose interpretation must be determined. This number is roughly the same for all methods based on propositional satisfiability procedures such as the Davis-Putnam procedure. While only being a crude method for classifying model generation problems, a smaller number of atoms generally indicates problems that are easier to solve. Experiments in propositional planning show that the state-of-the-art exhaustive SAT procedures can handle some hundreds of atoms while methods that employ local probabilistic search sometimes find propositional models with several thousand atoms whose interpretation must be determined [59].

The number of ground atoms that are derived from a first-order model generation problem depends on two orthogonal parameters. First, the number of atoms usually grows with the complexity of the input. A specification that consists of a large number of formulas will in general be harder to treat

than one that only consists of a few, although it is quite simple to find notoriously hard problems that have small formalisations. Second, the number of ground atoms derivable from a first-order specification can grow exponentially in the size of the universe of discourse. For certain problems in quasi-group theory, the search space grows with such speed that even particularly efficient systems like FINDER are able to search exhaustively only in a domain size of a dozen elements.

If the number of atoms that must be considered in an interpretation is taken as an indicator for a problem's complexity, then higher-order model generation for generalised frames will usually be much harder than first-order model generation. Even in cases where we have a small domain \mathcal{D}_l of individuals, the number of functions that can be defined from them is quite large. In the presence of higher-order quantification, we can easily formulate statements over the function domains $\mathcal{D}_{\alpha \rightarrow \beta}$ that are intractable. As an example, consider a first-order domain \mathcal{D}_l of size n , and the domain of diadic relations $\mathcal{D}_{l \rightarrow l \rightarrow o}$. The number of different relations in $\mathcal{D}_{l \rightarrow l \rightarrow o}$ is 2^{n^2} . In order to build a model for a formula of the form $\forall R_{l \rightarrow l \rightarrow o} P(R)$, we will have to make sure that all the 2^{n^2} instantiations of $P(R)$ are true. In a worst-case scenario, P itself could be a formula that quantifies over some function domain.

3.3 A Fragment of Higher-Order Logic

As discussed in Section 3.2.3, a higher-order logic that is based on Henkin semantics cannot be given a practically useful model generation method that uses the same techniques that work for first-order model generation. The size of the function domains $\mathcal{D}_{\alpha \rightarrow \beta}$ in generalised frames are an obstacle for a computational treatment of quantification, and an unrestricted syntax will leave us with formulas where even basic contradictions in atoms cannot easily be detected. The solution that I propose for the first problem is to quantify only over small subsets of the function domains $\mathcal{D}_{\alpha \rightarrow \beta}$ instead of considering the whole function space. The second problem is attacked by a restricted syntax where the ground atoms that must be considered in the computation of a denotation have a simple syntactic structure. By furthermore adopting a Herbrand-like interpretation for constant symbols, we obtain a linguistically motivated logic *MQC* (Montague-style Quantification Language) that combines the compositionality of a higher-order logic with the compactly

1	Natural Language	\mathcal{MQL} Representation
2	loves Ken	$love(ken)$
3	loved by Ken	$\lambda x love(x)(ken)$
4	loves him/herself	$\lambda x love(x)(x)$
5	Every man loves Ken	$EVERY(man)(love(ken))$
6	Ken (as quantifier)	$\lambda P P(ken)$
7	Ken loves Ken	$(\lambda P P(ken))(love(ken))$
8	Some apples are red	$SOME(apple)(red)$
9	Red is a color	$color_{(\iota \rightarrow o) \rightarrow o}(red)$
10	Some colors are primary colors	$SOME(color_{(\iota \rightarrow o) \rightarrow o})(primary)$

Figure 3.1: Examples for well-formed \mathcal{MQL} expressions

representable interpretations of a first-order logic.

3.3.1 Syntax

A formula F of the simply-typed λ -calculus is called **function-free quantified** iff for all atoms $h(\bar{u}_n)$ that occur in F , each u_j is either a variable or a constant. A well-formed formula in \mathcal{MQL} is a closed function-free quantified formula of the simply-typed λ -calculus.

The function-free quantified format does not have complex λ -expressions as arguments of non-logical constants. This means that the $\beta\eta$ -normal form of an argument u_j in an atom $h(\bar{u}_n)$ must always be a symbol. Note that while every atom itself must of course be well-typed, the arguments may be of arbitrary type.

Figure 3.1 gives some examples for well-formed \mathcal{MQL} formulas and expressions that might occur as semantic representations of natural language. The logical encodings are straightforward. In examples 7–10, we have \mathcal{MQL} formulas that formalise properties of higher-order objects. The constant $color$ denotes a second-order predicate of type $(\iota \rightarrow o) \rightarrow o$, and $color(red)$ specifies that the constant red denotes a first-order set has the property of being a color. In a first-order logic, one would have to reify colors as first-order entities, and could not use them at the same time as predicate symbols and arguments of other predicates.

The syntactic restriction that we have here is very similar to that of the EP calculus (cf. 2.4.3) that also uses a function-free quantified input format

where no complex term may occur as an argument of a predicate symbol.² EP’s syntax is motivated primarily by the need to avoid function symbols in the context of Herbrand interpretations. Our main motivation instead is to have only atoms $h(\overline{u}_n)$ in the recursive computation of a denotation that have a simple syntactical structure. We do not have to care for embedded function expressions whose denotation must be considered for instance when we check for basic contradictions in a partial interpretation.

3.3.2 Semantics

We cannot restrict the function domains $\mathcal{D}_{\alpha \rightarrow \beta}$ in a frame at will without potentially violating the **Denotatpflicht**, i.e., the requirement that each expression that is necessary for computing a denotation has a denotation itself.

In order to meet the Denotatpflicht in Henkin’s construction of generalised frames, each function domain includes all functions of appropriate type that can be expressed as λ -terms—this property is provided by the so-called **comprehension axioms** [47]. In principle, all functions that are expressible must be made available, as Henkin’s semantics must assign a meaning to arbitrary terms. In our simpler syntax, we may not encounter complex functional expressions that must be given an interpretation. In the recursive interpretation of a closed \mathcal{MQL} formula, the Denotatpflicht is already met if we can assign a meaning to the formula’s components and instantiations down to the level of simple, ground atoms.

The intuition behind the construction below is that we can further simplify the interpretation if all u_i in each atom $h(\overline{u}_n)$ are interpreted “as themselves”. The interpretations that we obtain in this way are term interpretations, i.e., a constant c_α that occurs as an argument will be interpreted as some “individual” c of type α rather than some object of the domain \mathcal{D}_α . Higher-order quantification and the Denotatpflicht for our term interpretations require that there are appropriate sets of individuals of type α available. These sets are provided in our formal framework by constant frames.

²This format is not to be confused with the function-free *fragment* of first-order logic, i.e., the set of formulas without function symbols whose prenex form starts with an initial sequence of universal quantifiers followed by a sequence of existential ones [34]. EP’s input normal form is equivalent to full first-order logics, while the function-free fragment is decidable.

Constant Frames

A **constant frame** \mathcal{C} is a collection $\{\mathcal{C}_\alpha | \alpha \in \text{TYPE}\}$ of sets \mathcal{C}_α that obeys the following conditions.

- $\mathcal{C}_t \in \mathcal{C}, \mathcal{C}_o \in \mathcal{C}$
- \mathcal{C}_t is not empty
- $\{0, 1\} \subseteq \mathcal{C}_o$

A **constant frame** \mathcal{C}_Φ for an \mathcal{MQL} specification Φ is a constant frame where all parameter constants C_α that occur in Φ are elements of $\mathcal{C}_\alpha \in \mathcal{C}$. A constant frame \mathcal{C}_Φ may also contain additional constants C_α that do not occur in Φ . A constant frame is an **initial constant frame** \mathcal{C}_Φ^0 if it contains only the minimum number of constants with regard to Φ . A **finite constant frame** is a constant frame where all sets \mathcal{C}_α are finite. Obviously, all initial constant frames are finite.

Interpretations and Denotations

What we will do in the following is to define a classical two-valued semantic for our logic \mathcal{MQL} that is based on the concept of a finite constant frame \mathcal{C}_Φ .

Let \mathcal{C}_Φ be a finite constant frame for a logical specification Φ . An \mathcal{MQL} interpretation $\mathcal{I} = \langle \llbracket \cdot \rrbracket, \mathcal{C}_\Phi \rangle$ for Φ is an evaluation of all constants in \mathcal{C}_Φ where the evaluation function $\llbracket \cdot \rrbracket$ obeys the following conditions:

- $\llbracket 1 \rrbracket = 1, \llbracket 0 \rrbracket = 0$
- $\llbracket c_o \rrbracket \in \{0, 1\}$ for all Boolean parameters c_o
- $\llbracket h_{\alpha_1 \rightarrow \dots \rightarrow \alpha_n \rightarrow o} \rrbracket$ is a function $f : \langle \mathcal{C}_{\alpha_1}, \dots, \mathcal{C}_{\alpha_n} \rangle \rightarrow \{0, 1\}$ for all n-ary predicate parameters h

An interpretation \mathcal{I}_Φ for a specification Φ determines the denotation of all parameter constants in \mathcal{C}_Φ as needed for defining a denotation of an \mathcal{MQL} formula later on. As usual, we furthermore assume that the logic itself determines an evaluation $\llbracket C \rrbracket$ of logical constants C in all interpretations.

Unlike conventional interpretations in higher-order logics, a \mathcal{MQL} interpretation does not give a denotation for individual constants or function

constants that are not predicates. This is actually not necessary because we will adopt a Herbrand-like interpretation of constant symbols in cases where constants appear as arguments of predicates. A Herbrand-like interpretation of constants is particularly attractive for applications in linguistics, as has been noted for instance by Baumgartner and Kühn [9]. In logical representations of natural language discourse, we often have the convention that individual constants are interpreted as names for distinct individual entities (e.g., *jon*, *mary*, etc.). The **unique-name assumption**, i.e., the practice of treating different individual constants as different individuals, is one of the basic constraints of interpretation in many logic-based linguistic formalisms. Our logic \mathcal{MQL} simply generalises this idea to constants of higher type.

Let $\mathcal{I} = \langle \llbracket \cdot \rrbracket, \mathcal{C}_\Phi \rangle$ be an interpretation for a specification Φ . The **denotation** $\llbracket F \rrbracket_{\mathcal{I}}$ of a \mathcal{MQL} formula in Φ is recursively defined as follows.

- $\llbracket c_o \rrbracket_{\mathcal{I}} = \llbracket c_o \rrbracket$ for all Boolean constants c_o in \mathcal{C}_Φ
- $\llbracket h(\overline{u}_n) \rrbracket_{\mathcal{I}} = \llbracket h \rrbracket(\overline{u}_n)$ for all predicate parameter constants h in \mathcal{C}_Φ
- $\llbracket \circ F \rrbracket_{\mathcal{I}} = \llbracket \circ \rrbracket(\llbracket F \rrbracket_{\mathcal{I}})$ for all unary logical connectives \circ
- $\llbracket F_1 \circ F_2 \rrbracket_{\mathcal{I}} = \llbracket \circ \rrbracket(\llbracket F_1 \rrbracket_{\mathcal{I}}, \llbracket F_2 \rrbracket_{\mathcal{I}})$ for all binary logical connectives \circ
- $\llbracket Q(T_{\alpha \rightarrow \sigma}) \rrbracket_{\mathcal{I}} = \llbracket Q \rrbracket(\langle \llbracket (T p_1) \rrbracket_{\mathcal{I}}, \dots, \llbracket (T p_n) \rrbracket_{\mathcal{I}} \rangle)$ for all monadic quantifier constants Q and all $p_j \in \mathcal{C}_\alpha$.
- $\llbracket Q(T_{\alpha \rightarrow \sigma})(U_{\alpha \rightarrow \sigma}) \rrbracket_{\mathcal{I}} = \llbracket Q \rrbracket(\langle \llbracket (T p_1) \rrbracket_{\mathcal{I}}, \dots, \llbracket (T p_n) \rrbracket_{\mathcal{I}} \rangle)(\langle \llbracket (U p_1) \rrbracket_{\mathcal{I}}, \dots, \llbracket (U p_n) \rrbracket_{\mathcal{I}} \rangle)$ for all diadic quantifier constants Q and all $p_j \in \mathcal{C}_\alpha$.

The recursion is a bit different from that of conventional higher-order logics. The crucial points are as follows.

First, the denotation of an atom $h(\overline{u}_n)$ is that of the interpretation $\llbracket h \rrbracket$ of the head predicate symbol h applied directly to the arguments \overline{u}_n . In a recursive computation of a denotation, we now have only atoms $h(\overline{u}_n)$ where all u_j are constants that actually are interpreted as themselves. The situation is analogous to a first-order Herbrand interpretation where embedded terms t that denote individuals are interpreted as themselves, i.e., $\llbracket t \rrbracket = t$. Unlike first-order logics, however, our constants u_j can have any type and might be interpreted as predicates when they occur in a predicate position.

Second, the denotation of quantified formulas is that of the interpretation $\llbracket Q \rrbracket$ of the quantifier constant Q applied to tuples of truth values that come from instantiations over the constant sets \mathcal{C}_α . Quantification in \mathcal{MQL} does not range over domains \mathcal{D}_α , but only over sets of constants \mathcal{C}_α that name entities of type α . Hence, the truth of a quantified formula is determined by the denotations of the formula's instantiations using constants of appropriate type. A set of constants \mathcal{C}_α that is used for quantification can be small, as long as it contains at least the constants of type α that are part of the input specification Φ that we are about to interpret. Informally, a quantified formula $\prod P_{\alpha \rightarrow \alpha} Q(P)$ is interpreted as a specification of properties Q over a set of predicates P that have been given names in the specification Φ . Additionally, $Q(P)$ may also influence the interpretation of other named properties P that do not occur in Φ . We may choose the domains of quantification \mathcal{C} as large as we want to as long as they remain finite. However, we can always start with an initial constant frame \mathcal{C} that is much smaller than a generalised frame \mathcal{D} , and whose domains of quantification only includes a basic set of distinguished entities. Their existence is implied by giving them distinct names in the specification Φ .

3.3.3 An \mathcal{MQL} Logic

Our logical language \mathcal{MQL} contains the usual set of logical constants that we find in the same form in conventional higher-order logics such as \mathcal{HOL} . Additionally, we may add logical constants that are motivated by our linguistic application, namely certain universal determiners whose semantic cannot be given conveniently in the form of a higher-order definition.

Like the λ -calculus, \mathcal{MQL} can be taken as the basis of different actual logics, all of them derived from different sets of logical constants and interpretations for them. As usual, we assume that \mathcal{MQL} determines for each logical constant one denotation that is mandatory in all interpretations.

Connectives

Because unary and binary connectives are predicates, they must be interpreted now as functions from sets of Boolean constants $\mathcal{C}_{type\ o}$ into the set $\{0, 1\}$. This has no real consequences because we already include the truth values 0 and 1 into the set \mathcal{C}_o of each constant frame. We can therefore keep the interpretation of unary and binary logical connectives as is, i.e., the

functions f that denote the logical connectives in a conventional higher-order logic \mathcal{HOL} work analogously in \mathcal{MQL} .

For \mathcal{HOL} , we have only given some arithmetic constraints for the semantics of the basic logical connectives \neg and \vee . The full set of constraints that defines the denotation of the standard set of logical connectives in all \mathcal{MQL} interpretations is as follows.

- $\llbracket \neg \rrbracket(x) = 1 - x$
- $\llbracket \vee \rrbracket(x, y) = \max(x, y)$
- $\llbracket \wedge \rrbracket(x, y) = \min(x, y)$
- $\llbracket \Rightarrow \rrbracket(x, y) = \min(1 - x, y)$
- $\llbracket \Leftrightarrow \rrbracket(x, y) = 1 - |x - y|$

The x and y are the truth values 0 or 1 that are provided by the recursive computation of the denotation of a formula.

Quantifiers

The interpretation of quantifier constants in \mathcal{MQL} must be adapted slightly in comparison to \mathcal{HOL} . The denotation of a quantifier constant is now a function from one or two tuples of truth values into a truth value. The tuples of truth values come from the denotation of formulas, and the size of the tuples is always finite. Our notion of quantification is obviously different from that of higher-order logic and first-order logic because quantification in \mathcal{MQL} is restricted to finite sets. However, due to this restriction we are more free to define forms of quantification that are hard to capture in first-order logics. Quantifier constants can now be given an operational semantic by defining their denotations as computable functions over finite sets of truth values. For instance, the semantic of the basic quantifier constants $\Pi_{(\alpha \rightarrow \omega) \rightarrow \omega}$ could be given as follows.

$$(3.1) \quad \llbracket \Pi \rrbracket(\langle k_1, \dots, k_n \rangle) = \min(\langle k_1, \dots, k_n \rangle)$$

Here, \min denotes the function that selects from a tuple of integers the smallest one. Given a set of truth values $k_j \in \{0, 1\}$, we obviously can compute

the value of $\llbracket \Pi \rrbracket(\langle k_1, \dots, k_n \rangle)$. The definition (3.1) expresses a constraint in the sense that it relates $\llbracket \Pi \rrbracket(\langle k_1, \dots, k_n \rangle)$ to a set of variables k_j .

From Π , we can derive the full set of standard quantifier constants³ by higher-order definitions as usual. Alternatively, we can give constraints over finite tuples of truth values to define denotations such as the following.

$$(3.2) \quad \llbracket \forall \rrbracket(\langle k_1, \dots, k_n \rangle) = \min(\langle k_1, \dots, k_n \rangle)$$

$$(3.3) \quad \llbracket \exists \rrbracket(\langle k_1, \dots, k_n \rangle) = \max(\langle k_1, \dots, k_n \rangle)$$

$$(3.4) \quad \llbracket \text{EVERY} \rrbracket(\langle k_1, \dots, k_n \rangle)(\langle l_1, \dots, l_n \rangle) = 1 \text{ iff } k_j \leq l_j \text{ for all } i \leq n$$

$$(3.5) \quad \llbracket \text{SOME} \rrbracket(\langle k_1, \dots, k_n \rangle)(\langle l_1, \dots, l_n \rangle) = 1 \text{ iff } k_j = l_j = 1 \text{ for some } i \leq n$$

The encoding of truth conditions into computable functions allows us to define the operational semantics of quantifiers that otherwise could not be formalised. For instance, we want a formula $F = \text{MORE}(\mathsf{T}_{\alpha \rightarrow \circ})(\mathsf{U}_{\alpha \rightarrow \circ})$ to be true iff the denotation of T has more elements than that of U . We know that F is true in an interpretation if more formulas $\mathsf{T}(C)$ than formulas $\mathsf{U}(C)$ are true for the C that are in \mathcal{C}_α . Hence, the sum of the denotations $k_j = \llbracket \mathsf{T}(C_j) \rrbracket_{\mathcal{I}}$ must be larger than that of the denotations $l_j = \llbracket \mathsf{U}(C_j) \rrbracket_{\mathcal{I}}$ for all $C_j \in \mathcal{C}_\alpha$. This condition is expressed by the following constraint.

$$(3.6) \quad \llbracket \text{MORE} \rrbracket(\langle k_1, \dots, k_n \rangle)(\langle l_1, \dots, l_n \rangle) = 1 \text{ iff } \sum_{j=1}^n k_j > \sum_{j=1}^n l_j$$

In first-order and higher-order logics, constraints such as (3.6) do not define computable denotations for quantifiers, because the domains of quantification can be infinite, and infinitely many truth values k_j and l_j may have to be considered. The MORE quantifier in \mathcal{MQL} has a computable denotation for all interpretations, but only because we have sufficiently weakened the semantic of our logical language to finite domains.

³We use the standard first-order quantification in order to improve readability.

Definitions

The collection of quantifier constants in \mathcal{MQL} can be extended simply by defining new ones as λ -terms. A definition expansion with subsequent $\beta\eta$ -conversion eliminates all occurrences of the defined constants which therefore do not need to have a denotation of their own in an interpretation.

A motivation for having definable logical constants is that we would for instance like to experiment with alternative formalisations of linguistic quantifiers in the context of semantic analysis. For instance, we have a diadic linguistic quantifier **ONE** where for instance $\text{ONE}(\textit{man})(\textit{drink})$ is true iff there is exactly one $x \in \mathcal{C}_i$ such that both $\textit{man}(x)$ and $\textit{drink}(x)$ are true. Hence, its denotation in \mathcal{MQL} must be as follows.

$$(3.7) \quad \llbracket \text{ONE} \rrbracket (\langle k_1, \dots, k_n \rangle) (\langle l_1, \dots, l_n \rangle) = 1 \text{ iff there is exactly one } i \text{ such that } k_j = l_j = 1$$

The quantifier **ONE** can be implemented simply by a definition that formalises the truth conditions in (3.7). The higher-order definition scheme for quantifiers **ONE** of arbitrary type $(\alpha \rightarrow o) \rightarrow (\alpha \rightarrow o) \rightarrow o$ looks as follows.

$$(3.8) \quad \text{ONE} \equiv \lambda P \lambda Q \exists X (P(X) \wedge Q(X)) \wedge \forall Y (P(Y) \wedge Q(Y)) \Rightarrow X \doteq Y$$

For reasons of simplicity, \mathcal{MQL} 's logical constants only include monadic and diadic quantifier constants, but the device of definitions allows us to formalise other forms of quantifications over finite sets as well. For instance, in Section 3.1.3, we considered the example (3.9) where the cardinality of two sets is part of the truth condition of a three-place quantifier.

$$(3.9) \quad \text{MORETHAN}(\textit{boy})(\textit{girl})(\textit{drink}) = |\textit{boy} \cap \textit{drink}| > |\textit{girl} \cap \textit{drink}|$$

The quantifier constant **MORETHAN** has a simple definition in \mathcal{MQL} that makes use of the **MORE** quantifier.

$$(3.10) \quad \text{MORETHAN} \equiv \lambda P \lambda Q \lambda R \text{ MORE}(\lambda x P(x) \wedge R(x))(\lambda x Q(x) \wedge R(x))$$

We will sometimes approximate the meaning of generalised determiners whose truth conditions cannot be exactly captured in our logic. An example is the universal determiner **Most** whose meaning is influenced by pragmatic considerations. That is, the choice of which truth conditions actually hold if

a sentence like *most men lie* is true lies to some extent with the speaker. However, a formula $\text{MOST}(P)(Q)$ in general induces at least that the majority of elements in P will also have the property Q . This truth condition is formalised in the following higher-order definition, again with the help of the MORE quantifier.

$$(3.11) \quad \text{MOST} \equiv \lambda P \lambda Q \text{ MORE}(\lambda x P(x) \wedge Q(x))(\lambda x P(x) \wedge \neg Q(x))$$

While definitions often allow us to formulate complex truth conditions in a compact and human-oriented way, the technical machinery that later deals with the expanded logical forms can be kept very simple by considering only a basic set of logical constants from which others are derived.

Equality

The equality signs $=_{\alpha \rightarrow \alpha \rightarrow \circ}$ in our higher-order logic \mathcal{HOL} have been given a semantic via the following definition scheme.

$$(3.12) \quad = \equiv \lambda Q_{\alpha} \lambda R_{\alpha} \forall P_{\alpha \rightarrow \circ} P(Q) \Rightarrow P(R)$$

In \mathcal{MQL} , the same formalisation (3.12) defines a form of equality that is determined only by those properties P in $\mathcal{D}_{\alpha \rightarrow \circ}$ that have a representative constant h in the set $\mathcal{C}_{\alpha \rightarrow \circ}$ with $\llbracket h \rrbracket = P$. This implies that the content of our constant frame \mathcal{C}_{Φ} plays an important rôle in the way in which equality is interpreted. Syntactically, we cannot use the definition of $=$ to compare arbitrary terms because the function-free quantified format prohibits formulas $P(Q)$ where P is a parameter and Q is not symbolic. Hence, a formula $X = Y$ in \mathcal{MQL} is only well-formed if both X and Y are constants or bound variables.

For practical purposes, we will often use a simpler form of equality, the **Herbrand equality** \doteq , that identifies its two arguments X and Y as equal if their ground instantiations are syntactically identical. Herbrand equality obeys the rules of parameter constants, i.e., its arguments must be symbolic. Its semantic is as follows.

$$(3.13) \quad \llbracket \doteq \rrbracket(C_{\alpha})(D_{\alpha}) = 1 \text{ iff } C \text{ and } D \text{ are the same constant}$$

3.4 Constructing Models

An interpretation \mathcal{I} in \mathcal{MQL} depends on the adjacently given semantics of the logical constants and on the interpretation of the constants that occur as predicate symbols. Constants that occur as arguments of non-logical constants have a fixed interpretation as themselves as specified in Section 3.3.2. Variables play no rôle because we only consider closed formulas.

Given an instantiation of a logic as a set of constraints that determines the interpretation of logical constants, and an algorithm that instantiates the recursive denotation scheme in Section 3.3.2, a \mathcal{MQL} interpretation is determined completely by the interpretation of predicate symbols over the sets of constants in a finite constant frame \mathcal{C}_Φ . The number of different interpretations for a finite set of formulas Φ and a finite constant frame \mathcal{C}_Φ is finite itself. Hence, we could enumerate all possible models for Φ in a given constant frame by a brute-force algorithm that simply enumerates all interpretations \mathcal{I} and checks whether \mathcal{I} is a model for the formulas of Φ . Such a brute-force algorithm is bound to fail for efficiency reasons already for propositional theories, and we would certainly not be able to compute many models for interesting logical specifications that way.

3.4.1 Determining Models Intelligently

Formulas can be used for restricting the search space for valid interpretations. The intuition of this approach has already been hinted at when we defined the interpretations of logical constants as arithmetic constraints. The general idea is based on treating formulas as constraints over variables that range over truth values, and to use efficient techniques for solving such constraints.

Formulas as Constraints

Each complex formula can be seen as a constraint that defines how the denotation of a formula depends on the denotations of its sub-formulas or instantiations. Consider for instance the propositional formula $F = (a_o \vee b_o) \wedge c_o$. Following the recursive definition of a denotation that has been given in Section 3.3.2, each model \mathcal{M} for F and some arbitrary constant frame \mathcal{C} must obey the following conditions:

$$\begin{aligned} \llbracket (a \vee b) \wedge c \rrbracket_{\mathcal{M}} &= \min(\llbracket a \vee b \rrbracket, \llbracket c \rrbracket) \\ \llbracket a \vee b \rrbracket_{\mathcal{M}} &= \max(\llbracket a \rrbracket, \llbracket b \rrbracket) \end{aligned}$$

The interpretation of a , b , and c is constrained by the fact that each must be interpreted as a truth value. Hence, the following set of equations characterise all models \mathcal{M} of F :

$$\begin{array}{lll} \llbracket a \rrbracket_{\mathcal{M}} \geq 0, & \llbracket a \rrbracket_{\mathcal{M}} \leq 1, & \llbracket (a \vee b) \wedge c \rrbracket_{\mathcal{M}} = 1, \\ \llbracket b \rrbracket_{\mathcal{M}} \geq 0, & \llbracket b \rrbracket_{\mathcal{M}} \leq 1, & \llbracket (a \vee b) \wedge c \rrbracket_{\mathcal{M}} = \min(\llbracket a \vee b \rrbracket_{\mathcal{M}}, \llbracket c \rrbracket_{\mathcal{M}}), \\ \llbracket c \rrbracket_{\mathcal{M}} \geq 0, & \llbracket c \rrbracket_{\mathcal{M}} \leq 1, & \llbracket a \vee b \rrbracket_{\mathcal{M}} = \max(\llbracket a \rrbracket_{\mathcal{M}}, \llbracket b \rrbracket_{\mathcal{M}}) \end{array}$$

Solving Constraints

In order to enumerate the models \mathcal{M} for F , we could enumerate all evaluation functions $\llbracket \cdot \rrbracket$ of the boolean constants in F and check whether the set of equations is satisfied. Alternatively, we could instead try to solve the set of (in)equations that we have computed by expanding the semantic of the logical constants. Each solution to this set corresponds to a valid interpretation of the Boolean constants and hence denotes a model.

Efficiently solving sets of linear (in)equations over finite-domain integer variables or, more generally, any computable set of relations between finite sets of variables, is the realm of constraint solving. Decision procedures for classical propositional logics can be seen as specialised constraint solvers for variables that can only have two values. For instance, the Davis-Putnam procedure is an efficient constraint solver that computes the instantiation of a finite set of boolean variables, i.e., atoms, whose dependencies are given by ground clauses.

The set of inequations given in the last section is an instance of the so-called Integer Programming (IP) Problem. IP has applications in operations research and economics [80], and solving inequations over finite-domain integer variables enjoys a great deal of interest in the constraint solving community. There are some off-the-shelf constraint solvers available that can deal with large sets of inequations in practice. Some of these come as modules for programming languages such as the finite-domain integer package of the constraint programming language Oz [78].

The general idea of a constraint solver is to have a two-part decision procedure. The first part is responsible for constraint propagation. Propagation refers to a process that first partially determines or restricts the value of variables out of the information that is encoded in a set of constraints and then uses this partly determined instantiation for inferring new constraints or simplifying old ones. For instance, in the Davis-Putnam procedure presented

in Section 2.4.4, the rule for unit resolution is a typical propagation rule in that it uses the information given by a unit clause for simplifying all clauses in which this unit occurs. The variable that is determined here is the interpretation of the unit, and this determined value is used for simplifying the constraints given by other ground clauses. For IP and other constraint systems, the propagation problem is a bit more complicated because the value of variables can range over many values instead of only two as in the case of propositional logic.

The second part of a constraint solving procedure is distribution. In general, distribution refers to a selection of a variable and a provisional restriction of its range in order to get new information that can be used for propagation. For completeness, a kind of backtracking mechanism must be able to reset the choice of restriction if it does not lead to a solution. In logical inference systems that have an analytical cut, the cut rule formulates a form of distribution. In the Davis-Putnam procedure we have such a cut-rule that selects a literal and splits into the case where the literal is interpreted as true and one in which it is interpreted as false.

An important paradigm of constraint solving is to delay distribution and to use as much information as possible for adding new constraints on the possible values of variables. The idea behind this is that it is very costly in general to backtrack a wrong decision while the efforts for propagation will pay off in the end by leading to more educated guesses. An important advantage of a larger set of constraints in practice is that we might be able to determine more precisely which variables are important and which are not. Distribution over a variable that occurs in many constraints will also help to simplify many more constraints in the end than determining a variable that only occurs in few or even only one constraint.

For the logical language \mathcal{MQL} , we do not have an off-the-shelf constraint solver that is comparable to the Davis-Putnam procedure for ground clause sets. We could give a translation that maps \mathcal{MQL} formulas over a given constant frame into sets of inequations whose variables are the interpretations of ground formulas. Such a set of equations can then be solved by a standard IP constraint solver. However, there are constraint solvers available that can already deal more efficiently with the kind of constraints that we find expressed in logical formulas, and whose constraint languages are more convenient than the sometimes awkward language of arithmetic inequations. We therefore will develop a generic translation from logics into constraints that does not directly depend on the actual constraint solver that is available.

3.4.2 Translating Formulas into Constraints

Our task in this section is to define a translation from a set of \mathcal{MQ} formulas into a set of constraints with regard to a given constant frame \mathcal{C} that defines the range of quantification for different types. Our translation must follow the definition of a \mathcal{MQ} semantic and map each complex formula into a constraint that reflects the relation that the interpretation of the formula has to the interpretation of its parts or instantiations. At the same time, we want to leave open which logical constants we have in our logic—we want to be able to extend the basic set as necessary—as well the actual constraint language that we are using. Our translation therefore is only a scheme that can be instantiated by an actual translation.

Fig. 3.2 shows the translation of \mathcal{MQ} specifications into constraints as a tableaux inference system. We use **signed formulas** $V_F: F$ where F is a formula and V_F is a 0/1-integer variable associated with the interpretation $\llbracket F \rrbracket$. For a specification $\Phi = \{F_1, \dots, F_n\}$, we start with an initial tableau $\Theta = \{V_{F_1}: F_1, \dots, V_{F_n}: F_n, V_{F_1} = 1, \dots, V_{F_n} = 1\}$. The top-level integer variables V_{F_j} all are restricted to 1 by the equations $V_{F_j} = 1$.

Each of the rules given in our Constraint Tableaux (CT) system can expand the current tableau branch by new signed formulas and constraints. Tableau expansion stops when the tableau is saturated, i.e., no rule can add new information to the tableau.

The *con2*-rule expands signed formulas $V_F: F_1 \circ F_2$ that are governed by a binary logical connective \circ . The tableau is extended by new signed formulas $V_{F_1}: F_1$ and $V_{F_2}: F_2$ for the components F_1 and F_2 . The rule adds a constraint $V_F = \llbracket \circ \rrbracket(V_{F_1}, V_{F_2})$ for the connective \circ . The constraint describes the relation between the variables V_F , V_{F_1} and V_{F_2} and is directly derived from the semantic of the logical connective \circ . In the case of a conjunctive formula and an IP constraint solver, this constraint could be the equation $V_F = \min(V_{F_1}, V_{F_2})$. The *con1*-rule corresponds to the *con2*-rule, but treats unary connectives and will add only one new signed formula.

The *uni*-rule adds an equality constraint for all integer variables that represent the same atom. This unification of variables ensures that we actually compute interpretations that are functions.

Finally, the *quan2*-rule and *quan1*-rules translates formulas governed by diadic and monadic quantifier constants Q . Given a constant frame \mathcal{C} , there is a tuple of constants $\mathcal{C}_\alpha = \langle C_1, \dots, C_n \rangle$, $n \geq 0$ for each type α that occurs in Φ . The order of constants may be arbitrary, but fixed. The *quan*-rules

build new formulas by applying each of the terms \top (and \cup) to all constants in \mathcal{C}_α of the appropriate type, and recursively translates these instantiations into constraints. The rule then generates a constraint that relates the integer variable of the whole quantified formula and the integer variables of its instantiations as defined by the denotation of the quantifier in question.

Unlike conventional tableaux systems, the CT calculus does not split the current tableau branch in any way during tableau construction.

3.4.3 An Example

We illustrate our tableaux system CT with a short example that uses arithmetic constraints for implementing the semantics of logical constants.

Consider a constant frame \mathcal{C} that contains a set $\mathcal{C}_i = \{john, pete, karl\}$ for individual constants and a set $\mathcal{C}_{i \rightarrow o} = \{drink, man\}$ for first-order predicates. A first-order specification $\{\forall x man(x) \Rightarrow drink(x)\}$, i.e., *all men drink*, would be expanded first into the following tableau:

$$\begin{array}{l}
 V_1: \forall x man(x) \Rightarrow drink(x) \\
 \star V_1 = 1 \\
 V_2: man(john) \Rightarrow drink(john) \\
 V_3: man(pete) \Rightarrow drink(pete) \\
 V_4: man(karl) \Rightarrow drink(karl) \\
 \star V_1 = \min(\langle V_2, V_3, V_4 \rangle)
 \end{array}$$

The tableau contains two constraints, indicated by \star . The first one, $V_1 = 1$, is part of the initial tableau and restricts the value of the truth variable V_1 to 1 because each top-level formula of a specification must be interpreted to 1 in a model. The second one, $V_1 = \min(\langle V_2, V_3, V_4 \rangle)$ relates V_1 to the truth variables of the three possible instantiations.

Each of the newly generated signed formulas is governed by a logical connective, and must therefore be expanded further by the *bin*-rule. This second stage of expansions will give us the following tableau:

$$\begin{array}{c}
\frac{V_F: F_1 \circ F_2}{V_{F_1}: F_1 \quad V_{F_2}: F_2} \text{ con2} \\
V_F = \llbracket \circ \rrbracket (V_{F_1}, V_{F_2})
\end{array}
\quad
\begin{array}{c}
\frac{V_{\neg F}: \neg F}{V_F: F} \text{ con1} \\
V_{\neg F} = \llbracket \neg \rrbracket (V_F)
\end{array}
\quad
\begin{array}{c}
\frac{V_1: A \quad V_2: A}{V_1 = V_2} \text{ uni (A is an atom)}
\end{array}$$

$$\frac{V_{QT}: Q(\top_{\alpha \rightarrow \beta})(U_{\alpha \rightarrow \beta})}{\begin{array}{c} V_{(\top C^1)}: (\top C^1) \\ \vdots \\ V_{(\top C^n)}: (\top C^n) \\ V_{(U C^1)}: (U C^1) \\ \vdots \\ V_{(U C^n)}: (U C^n) \end{array}} \text{ quan2}$$

$$V_{QT} = \llbracket Q \rrbracket (\langle V_{(\top C^1)}, \dots, V_{(\top C^n)} \rangle) (\langle V_{(U C^1)}, \dots, V_{(U C^n)} \rangle)$$

$$\frac{V_{QT}: Q(\top_{\alpha \rightarrow \beta})}{\begin{array}{c} V_{(\top C^1)}: (\top C^1) \\ \vdots \\ V_{(\top C^n)}: (\top C^n) \end{array}} \text{ quan1}$$

$$V_{QT} = \llbracket Q \rrbracket (\langle V_{(\top C^1)}, \dots, V_{(\top C^n)} \rangle)$$

Figure 3.2: The Constraint Tableaux (CT) expansion rules

- $$V_1: \forall x \text{ man}(x) \Rightarrow \text{drink}(x)$$
- ★ $V_1 = 1$
 - $V_2: \text{man}(\text{john}) \Rightarrow \text{drink}(\text{john})$
 - $V_3: \text{man}(\text{pete}) \Rightarrow \text{drink}(\text{pete})$
 - $V_4: \text{man}(\text{karl}) \Rightarrow \text{drink}(\text{karl})$
 - ★ $V_1 = \min(\langle V_2, V_3, V_4 \rangle)$
 - $V_5: \text{man}(\text{john})$
 - $V_6: \text{drink}(\text{john})$
 - ★ $V_2 = \min(1 - V_5, V_6)$
 - $V_7: \text{man}(\text{pete})$
 - $V_8: \text{drink}(\text{pete})$
 - ★ $V_3 = \min(1 - V_7, V_8)$
 - $V_9: \text{man}(\text{karl})$
 - $V_{10}: \text{drink}(\text{karl})$
 - ★ $V_4 = \min(1 - V_9, V_{10})$

After these expansions, no rule is applicable any more. Our translation has computed a set of constraints which can be solved by an IP constraint solver after all variables V_j have been restricted to integer values 0 and 1.

A solution is an instantiation $\sigma = \{V_1 = v_1, \dots, V_{10} = v_{10}\}$ for the integer variables V_j and truth values v_j that satisfies the constraints. For instance, a solution for our example is $\sigma = \{V_1 = 1, V_2 = 1, V_3 = 1, V_4 = 1, V_5 = 0, V_6 = 0, V_7 = 0, V_8 = 1, V_9 = 1, V_{10} = 1\}$. The variables that represent atomic formulas, i.e., the variables $\{V_5, \dots, V_{10}\}$, can be used for constructing the representation sets for evaluation functions. The solution σ represents the model $\mathcal{M} = \langle \{\text{drink}(\text{pete}), \text{man}(\text{karl}), \text{drink}(\text{karl})\}, \mathcal{C} \rangle$.

3.4.4 Properties of the Translation

A specification Φ is satisfiable with respect to a constant frame \mathcal{C} for Φ iff Φ has a model $\mathcal{M} = \langle \llbracket \cdot \rrbracket, \mathcal{C}_\Phi \rangle$. A constraint tableau Θ is satisfiable iff its set of constraints can be satisfied by a solution σ such that $\llbracket \mathbf{F} \rrbracket = \sigma(V_{\mathbf{F}})$ for an evaluation function $\llbracket \cdot \rrbracket$ and all signed formulas $V_{\mathbf{F}}: \mathbf{F}$ that occur in Θ .

Refutation Soundness

Theorem 1 *A constraint tableau for a specification Φ and a constant frame \mathcal{C}_Φ is satisfiable if Φ is satisfiable with respect to \mathcal{C}_Φ .*

Proof: If Φ is satisfiable with respect to \mathcal{C}_Φ , then Φ has a model $\mathcal{M} = \langle \llbracket \cdot \rrbracket, \mathcal{C}_\Phi \rangle$ such that $\llbracket F_j \rrbracket = 1$ for all $F_j \in \Phi$. An initial tableau Θ_0 for Φ and \mathcal{C}_Φ consists of pairs $V_{F_j}: F_j$ and equations $V_{F_j} = 1$. We chose $\sigma_0(V_{F_j}) = 1$ for all V_{F_j} . Then σ_0 obviously satisfies all constraints in Θ_0 and $\llbracket F_j \rrbracket_{\mathcal{M}} = \sigma_0(V_{F_j})$ holds for all signed formulas $V_{F_j}: F_j$ in Θ_0 . Hence, Θ_0 is satisfiable.

The set of constraints in a satisfiable tableau Θ can be satisfied by a solution σ such that $\llbracket F \rrbracket_{\mathcal{M}} = \sigma(V_F)$ for an evaluation function $\llbracket \cdot \rrbracket$ in \mathcal{M} and all signed formulas $V_{F_j}: F_j$ that occur in Θ . Let $V_F: F$ be a signed formula in Θ . We show that each expansion Θ' of $V_F: F$ in a satisfiable tableau Θ is satisfiable.

Let $F = \neg F'$. Then $\llbracket F \rrbracket_{\mathcal{M}} = \llbracket \neg \rrbracket \llbracket F' \rrbracket_{\mathcal{M}}$. The expansion Θ' of Θ by the new signed formula $V_{F'}: F'$ and the constraint $V_F = \llbracket \neg \rrbracket(V_{F'})$ is satisfied by a solution σ' such that $\sigma'(V) = \sigma(V)$ for all $V \neq V_{F'}$, and $\sigma'(V_{F'}) = \llbracket F' \rrbracket_{\mathcal{M}} = \llbracket \neg \rrbracket(V_F)$ else. If $F = F_1 \circ F_2$ then $\llbracket F \rrbracket_{\mathcal{M}} = \llbracket \circ \rrbracket(\llbracket F_1 \rrbracket_{\mathcal{M}}, \llbracket F_2 \rrbracket_{\mathcal{M}})$, and an expansion of Θ by the new signed formulas $V_{F_1}: F_1$ and $V_{F_2}: F_2$ and the constraint $V_F = \llbracket \circ \rrbracket(V_{F_1}, V_{F_2})$ is satisfiable by a modification of the solution σ to a solution σ' with $\sigma'(V_{F_1}) = \llbracket F_1 \rrbracket_{\mathcal{M}}$ and $\sigma'(V_{F_2}) = \llbracket F_2 \rrbracket_{\mathcal{M}}$.

If $F = Q(T^1) \dots (T_n)$, with Q being a quantifier constant, then an interpretation of F must formulate a constraint over all instantiations $T_i(p_j)$ of constants C_k of appropriate type for each T_i . Again, the expansion of Θ by the aforementioned constraint and the new signed formulas $V_{(TC_j)}: (TC_j)$ for the instantiations is satisfiable by a suitably extended solution σ' such that $\sigma'(V_{(Tp_j)}) = \llbracket (Tp_j) \rrbracket_{\mathcal{M}}$.

Let F be an atomic formula and Θ be a tableau containing two signed formulas $V_1: F$ and $V_2: F$. Because both are part of a satisfiable tableau, $\llbracket F \rrbracket_{\mathcal{M}} = \sigma(V_1) = \sigma(V_2)$ holds for an evaluation function $\llbracket \cdot \rrbracket$, and an extension Θ' of Θ by the constraint $V_1 = V_2$ is obviously satisfied by σ . \square

Theorem 1 proves a form of refutation soundness. Starting with a satisfiable specification Φ over a constant frame \mathcal{C} , each tableau expansion will create a satisfiable set of constraints. It is trivial to show that a solution σ that determines the interpretation $\llbracket A \rrbracket = \sigma(V_A)$ for each atom A in the tableau also determines the interpretations of all occurrences of complex formulas that depend on these. Additionally, a solution must interpret the atoms in a way such that the formulas of Φ become true. Hence, the part of a solution σ that determines the evaluation of all atoms A in the saturated tableau for Φ is equivalent to a model for the specification Φ relative to a given constant frame \mathcal{C}_Φ . Theorem 1 implies that a constraint translation of

a satisfiable specification over a constant frame will lead to a tableau that can be solved. Our method is complete in that we can be sure to find at least one solution that represents a positive model.

Completeness for \mathcal{MQL} Satisfiability

Our translation into constraints also gives us a decision procedure for the existence of models in a tableau. A saturated constraint tableau Θ may contain only finitely many signed formulas, and there can only be finitely many different evaluations σ for the atoms in Θ that represent models of Φ . Hence, we can decide whether a tableau Θ contains a model. In the following, we show that we can use this result for proving \mathcal{MQL} satisfiability.

Lemma 2 *Let Φ be satisfiable with respect to a constant frame \mathcal{C}_Φ . Let \mathcal{C}'_Φ be a constant frame for Φ with $|\mathcal{C}_\alpha| = |\mathcal{C}'_\alpha|$ for all $\mathcal{C}_\alpha \in \mathcal{C}_\Phi$ and all $\mathcal{C}'_\alpha \in \mathcal{C}'_\Phi$. Then Φ is also satisfiable with respect to \mathcal{C}'_Φ .*

Proof: Let Φ be satisfiable by a model $\mathcal{M} = \langle \mathcal{C}_\Phi, \llbracket \cdot \rrbracket \rangle$. Then $\llbracket \cdot \rrbracket$ has a finite representation as a set $\{\mathbf{A}_1, \dots, \mathbf{A}_k\}$ of ground atoms. We show that Φ then has a model $\mathcal{M}' = \langle \llbracket \cdot \rrbracket', \mathcal{C}'_\Phi \rangle$ with the representation set of $\llbracket \cdot \rrbracket'$ being equal to the set representation of $\llbracket \cdot \rrbracket$ up to a renaming of constant symbols.

Let Θ_0 be an initial constraint tableau for Φ and \mathcal{C}_Φ and Θ'_0 be an initial constraint tableau for Φ and \mathcal{C}'_Φ . It is easy to see that both tableaux are identical modulo a bijective mapping of constant symbols ρ from \mathcal{C}_Φ to \mathcal{C}'_Φ such that $\rho(h_\alpha) = h'_\alpha$ for all h_α that are in \mathcal{C}_Φ but not in \mathcal{C}'_Φ , and $\rho(h_\alpha) = h_\alpha$ else. We denote the renaming of all constant symbols in a formula F according to ρ as $\rho(F)$.

Every extension of a tableau Θ for \mathcal{C}_Φ has an equivalent extension in Θ' such that each signed formula $V : F$ in Θ has a corresponding signed formula $V : F'$ with $F = \rho(F')$. Because the constraints in both tableaux are isomorphic, each tableaux has the same solutions σ_j . We select one solution σ . Let $V_A : A$ be an atomic formula in Θ such that $\llbracket A \rrbracket = \sigma(V_A) = 1$. Then there is an atomic formula A' with $A' = \rho(A)$ in Θ' such that $\sigma(V'_A) = \sigma(V_A)$ and hence $\llbracket A \rrbracket = \llbracket A' \rrbracket'$. Because A is in the representation set of $\llbracket \cdot \rrbracket$, A' must be in the representation set of the evaluation function $\llbracket \cdot \rrbracket'$, and vice versa. Hence the representation sets for $\llbracket \cdot \rrbracket$ and $\llbracket \cdot \rrbracket'$ are equal up to a renaming ρ of constant symbols. \square

Theorem 3 (Satisfiability Completeness) *There is an algorithm that proves the satisfiability of an \mathcal{MQL} specification Φ .*

Proof: Let $\mathcal{M} = \langle \llbracket \cdot \rrbracket, \mathcal{C}_\Phi \rangle$ be an arbitrary \mathcal{MQL} model for Φ . Starting with a smallest initial constant frame \mathcal{C}_Φ^0 , we can enumerate models for extensions \mathcal{C}'_Φ of \mathcal{C}_Φ^0 by adding arbitrary new constants h_α to sets $\mathcal{C}_\alpha \in \mathcal{C}_\Phi^0$ and solving the constraint tableaux for Φ and \mathcal{C}'_Φ . As long as our extension strategy is fair and complete, we will eventually reach a constant frame \mathcal{C}'_Φ with $|\mathcal{C}_\alpha| = |\mathcal{C}'_\alpha|$ for all $\mathcal{C}_\alpha \in \mathcal{C}_\Phi$ and all $\mathcal{C}'_\alpha \in \mathcal{C}'_\Phi$. Lemma 2 shows that Φ has a model $\mathcal{M} = \langle \llbracket \cdot \rrbracket', \mathcal{C}'_\Phi \rangle$ that is equal to \mathcal{M} up to a renaming of constant symbols. Hence, the corresponding constraint tableau Θ is satisfiable and we can prove that Φ is satisfiable by producing a model of Φ relative to \mathcal{C}'_Φ . \square

Theorem 3 implies that we are bound to find a model for a satisfiable \mathcal{MQL} specification Φ when applying a fair iterative deepening strategy over the size of the constant frame \mathcal{C} . All models of \mathcal{MQL} are finite, and we therefore have a procedure for proving \mathcal{MQL} satisfiability.

Enumerating Models

Our translation from logics into constraints is not complete in the sense that we can guarantee to find a representative model \mathcal{M} for every model of the input Φ . The translation only considers formulas that can be expanded in some way from input formulas, and will not determine the interpretation of ground atoms that do not occur in the saturated tableau at all. The interpretation of such atoms can be chosen freely and therefore can be ignored when we are only interested in subset-minimal positive models. Still, we would like to have a method that truly enumerates the \mathcal{MQL} models of a specification.

The following theorem shows that a slight modification of our model generation method generates all isomorphic models of the input.

Theorem 4 (Model Completeness) *There is an algorithm that enumerates all models of a satisfiable specification Φ up to a renaming of constant symbols.*

Proof: Given a specification Φ and an initial constant frame \mathcal{C}_Φ^0 , we extend Φ to a specification Φ' in the following way: For each ground atom A that can be build from the constant frame \mathcal{C}_Φ^0 , we add the formula $A \vee \neg A$

to Φ . The newly added formula obviously is satisfiable under the usual interpretation of the logical constants \vee and \neg .

The solutions for a constraint tableau Θ_0 for Φ' and \mathcal{C}_Φ^0 determine the evaluation of all atoms A and represent all isomorphic models in the size of the constant frame \mathcal{C}_Φ^0 . By iterative deepening over the size of the constant frame, and by adding the tautologies $A \vee \neg A$ for all new atoms that we add, we can now enumerate all \mathcal{MQL} models up to isomorphism by enumerating the solutions σ for each resulting tableau. \square

3.5 Minimal Model Generation

We now define a new class of models for \mathcal{MQL} specification, the locally minimal models. They are an amalgamation of domain minimal models and subset-minimal models known from first-order model generation. The property of being a local minimal model in our logic \mathcal{MQL} is decidable.

3.5.1 Preliminaries

The individual domain size $|\mathcal{C}_i(\mathcal{M})|$ of a \mathcal{MQL} model \mathcal{M} is the number of individual constants c_i that occur in a finite representation of \mathcal{M} as a set of ground literals.

A model \mathcal{M} for a \mathcal{MQL} specification ϕ is **domain minimal** if its individual domain size $\mathcal{C}_i(\mathcal{M})$ is minimal, i.e., there is no model \mathcal{M}' of ϕ such that $|\mathcal{C}_i(\mathcal{M}')| < |\mathcal{C}_i(\mathcal{M})|$. A model \mathcal{M} for ϕ is a **locally minimal model** iff it is domain minimal and the condition $\text{Pos}(\mathcal{M}') \subseteq \text{Pos}(\mathcal{M}) \Rightarrow \mathcal{M} = \mathcal{M}'$ holds for all domain minimal models \mathcal{M}' of ϕ .

Local minimality is both stronger and weaker than the classical subset-minimal model property defined in Section 2.3.3. A locally minimal model always has a minimal individual domain while minimal models may have arbitrarily large domains. At the same time, a locally minimal model does not have to be minimal with regard to all other finite models of ϕ , but only in comparison with those that have the same domain size.

3.5.2 Decidability of Local Minimality

In this section, we prove that local minimality for finite models in \mathcal{MQL} is decidable. The method is based on the observation that the minimal

models \mathcal{M} up to a give size of the individual domain can be characterised by the literals that occur negatively within \mathcal{M} . Niemelä [73] originally used this characterisation for defining a method for propositional minimal model reasoning.

Lemma 5 *Let $\mathcal{M} = \{A_1, \dots, A_n, \neg B_1, \dots, \neg B_m\}$ be a local minimal model of a \mathcal{MQL} specification ϕ . Then there is no model \mathcal{M}' with $|\mathcal{C}_i(\mathcal{M}')| = |\mathcal{C}_i(\mathcal{M})|$ that satisfies $\phi' = \phi \wedge \neg B_1 \wedge \dots \wedge \neg B_m \wedge \neg(A_1 \wedge \dots \wedge A_n)$*

Proof: Assume that $\mathcal{M}' = \{A'_1, \dots, A'_k, \neg B'_1, \dots, \neg B'_l\}$ exists. Obviously, \mathcal{M}' is a domain minimal model of ϕ . \mathcal{M}' satisfies all negative literals that are satisfied by \mathcal{M} , but only a subset of the positive literals $\{A_1, \dots, A_n\}$. This implies that the set of positive literals of \mathcal{M}' , $\{A'_1, \dots, A'_k\}$, is a real subset of the positive literals $\{A_1, \dots, A_n\}$, while the set of negated literals of \mathcal{M} , $\{\neg B_1, \dots, \neg B_m\}$, is a real subset of $\{\neg B'_1, \dots, \neg B'_l\}$. Hence, \mathcal{M}' and \mathcal{M} must be different and \mathcal{M}' is a model of ϕ with $\text{Pos}(\mathcal{M}') \subseteq \text{Pos}(\mathcal{M})$ but $\mathcal{M}' \neq \mathcal{M}$. This contradicts that \mathcal{M} is a locally minimal model. \square

Theorem 6 (Decidability of Locally Minimal Models) *Given a model*

$$\mathcal{M} = \{A_1, \dots, A_n, \neg B_1, \dots, \neg B_m\}$$

of a \mathcal{MQL} specification ϕ , it is decidable whether \mathcal{M} is a locally minimal model of ϕ .

Proof: By lemma 5, there can be no model \mathcal{M}' that satisfies $\phi' = \phi \wedge \neg B_1 \wedge \dots \wedge \neg B_m \wedge \neg(A_1 \wedge \dots \wedge A_n)$ in the same domain size as \mathcal{M} if \mathcal{M} is locally minimal. Hence it suffices to show that ϕ' is unsatisfiable within the domain size $|\mathcal{C}_i(\mathcal{M})|$. We chose an initial constant frame \mathcal{C}_0 for ϕ' . If ϕ' is unsatisfiable, then a saturated constraint tableau Θ for ϕ' and our constant frame \mathcal{C}_0 cannot have a solution. Hence, the set of constraints in Θ is unsatisfiable, which itself is a decidable property. \square

Part II
Linguistics

Chapter 4

The Analysis of Definites

The use of *the* seems to suggest shared experience or knowledge: the listener/reader 'has been there too'.

(Michael Swan, *Practical English Usage*)

Overview: Model generation is a tool for natural-language interpretation in context. We apply a finite model generator to formalisations of singular definite descriptions and show that minimal model generation permits a computational treatment of definite noun phrases which directly reflects contemporary theories of definiteness. Model generation can be used as a uniform tool for analysing linguistic theories on definites.

4.1 Introduction

Given the semantic representation of a sentence and a specification of a context as a set of formulas, a finite model generator will enumerate a set of finite models satisfying both. The information in certain finite models encodes both the context change as well as the truth conditions of the sentence. In the truth-conditional sense of the word, model generation computes meaning. One of the things that we can use this computed meaning for is to (re)formulate and analyse linguistic theories.

The linguistic guinea pig that we¹ apply our analysis to are singular def-

¹The research presented here, as well as that presented in Chapter 5 originates in joint work with Claire Gardent.

inite descriptions. Our platform of experimentation is the finite model generator KIMBA (cf. Chapter 7) that implements a translation of \mathcal{MQL} logical specifications into constraints over finite-domain integer variables as well as the computation of minimal models as presented in Section 3.5.

We show that minimal model generation permits a computational treatment of singular definite noun phrases which can combine the theoretical insights from contemporary theories of definiteness with the processing of logically encoded knowledge about the situational context in which definite sentences are interpreted.

4.1.1 The Semantics of Definite Descriptions

Singular **definite descriptions** are expressions that refer to a particular individual entity without referring to its name, as in (1.1)– (1.3).

(1.1) *The President of Germany*

(1.2) *The first man on the moon*

(1.3) *Washington's mother*

The analysis of singular definite descriptions, in the following called **definites**, was one of the starting points of the logic-based approach to language philosophy. Both Russell and Frege investigated the problem how the principle of bivalence, i.e., the property of propositions to be either true or false, can be maintained if a definite refers to a non-existing entity, such as in the following famous example.

(1.4) ***The King of France** is bald.*

Russel's treatment of definites in predicate logics makes use of a special iota (ι) operator. The expression $\iota x P(x)$ denotes the x such that $P(x)$ holds. Thus, examples (1.4) can be formalised as follows.

(1.5) *bald(ιx kof(x))*

This approach fails in cases where such an x does not exist, or when there are too many x . Then, propositions like (1.5) can neither be interpreted as true nor as false. Russell's solution to this problem was a transformation from

what he calls the “misleading form” of the formula (1.5) into a logical form that refers only to standard quantifiers \exists and \forall . The transformation was first presented in his article “On Denoting” [79]. The result of the transformation for (1.5) into first-order predicate logic looks as follows.

$$(1.6) \quad \exists x \textit{kof}(x) \wedge \textit{bald}(x) \wedge \forall y \textit{kof}(y) \Rightarrow x = y$$

In higher-order logics, Russell’s formalisation is equivalent to a Montague-style linguistic quantification where THE is defined as in (1.7). This leads to the representation (1.8) for sentence (1.4) which nicely hides the discrepancy between the intended logical form and the “misleading” form of natural language.

$$(1.7) \quad \text{THE} \equiv \lambda P \lambda Q \exists x P(x) \wedge Q(x) \wedge \forall y P(y) \Rightarrow x = y$$

$$(1.8) \quad \text{THE}(\textit{kof})(\textit{bald})$$

4.1.2 Definites and Deduction

A formal representation such as (1.8) can be used for all kinds of inference. For instance, consider the consequence relation that holds for the sentences in example (1.9).

$$(1.9) \quad \textit{The King of France is bald} \models \textit{There is a King of France}.$$

This consequence relation can be verified mechanically by applying a first-order theorem prover to the theorem (1.10).

$$(1.10) \quad \text{THE}(\textit{kof})(\textit{bald}) \Rightarrow \exists x \textit{kof}(x)$$

Given a logically encoded context, deduction can sometimes be used to verify whether a provisional resolution of an anaphor is valid. By anaphor we mean any natural language expression whose referent can only be determined by relating it to an antecedent referent in the discourse. Example (1.12) is a typical proof task that is equivalent to the resolution of the anaphoric information given informally in (1.11). By proving (1.12), we show that the proposed resolution in (1.11) is correct.

$$(1.11) \quad \textit{Jon has a rabbit. The rabbit sleeps.} \models^? \textit{The rabbit that Jon has sleeps.}$$

$$(1.12) \quad \text{rbt}(r_1) \wedge \text{has}(\text{jon}, r_1) \wedge \text{THE}(\text{rbt})(\text{sleep}) \Rightarrow \\ \text{THE}(\lambda x \text{rbt}(x) \wedge \text{has}(\text{jon}, x))(\text{sleep})$$

Note that (1.12) has a standard first-order form when expanded, and proving the formula should present no difficulty to any state-of-the-art first-order theorem prover. However, a purely deductive approach to the resolution of definites will fail whenever the definite is not an anaphor. The interpretation of a definite sometimes requires the **accommodation** of new individuals. The term accommodation goes back to Lewis [62] and refers to the process by which the listener adjust her assumptions by adding just enough information to remedy the violation of some felicity condition. Consider the following sentence.

$$(1.13) \quad \textit{Jon's rabbit is cute.}$$

Here, we have no context in which we can resolve the definite *Jon's rabbit*. The interpretation of (1.13) requires that we accommodate the existence of some rabbit that has the property of both being the rabbit of Jon and sleeping. If we would insist on asserting the existence of a suitable referent in the context, then example (1.13) cannot be interpreted as true.

What we would like to have for examples such as (1.13) is a method that interprets a sentence relative to a given context and accommodates new information if necessary. The meaning of a sentences is then given by the change that the interpretative process imposes on the context. A method that provides this is minimal model generation.

4.1.3 How Models Interpret Sentences

A logical model can be conceptualised as a set of basic assumptions under which a given logical specification is true. These assumptions are what we are after when we interpret a sentence. In example (1.13), the minimal set of assumptions can be stated informally as follows.

- There is an individual named Jon who
- owns some rabbit r_1 and
- r_1 has the property of being cute.

The assumptions encode truth-conditional meaning. For natural-language interpretation, we need models whose entailed assumptions are either evident or accommodated as required. When accommodation is necessary, it should be restricted to sets of assumptions that at least are consistent with what we know about the situation at hand. Further, it should be obvious that the sets of assumptions must state all information that is required for validating the truth of the interpreted sentence: the assumptions should explain *completely* why a certain sentence can be true in the given context. Finally, we do not want over-explanation in the form of irrelevant or unjustifiable information. An interpretation must not state more information as is required by the task at hand.

It lies in the nature of a model that all assumptions that can be derived from it are true under the interpretation given by this model. If we use a background theory that contains the necessary world knowledge, a model of this background theory and a semantic representation will not entail logically inconsistent facts.

Further, a model always gives sufficient information for the truth of a theory. Every interpretation in the Tarskian sense of the word unambiguously defines the truth value of the specifications that it validates. If a model lacks any necessary information, it cannot be a model at all.

The elimination of irrelevant information is not an issue of model generation in general, but it is an issue for *minimal* model generation. Minimality constraints for models usually restrict the positive assumptions that are made by a model. The locally minimal models that we investigated in Section 3.5 are models that satisfy a specification both without referring to more individual entities than necessary and without making unnecessary assumptions relative to all other models that can be found in the smallest domain. In the sense of Occam's Razor, locally minimal models are the simplest, and therefore often the best explanations for the truth of a theory. In natural-language interpretation, the locally minimal models are essential because they minimise the accommodated individuals and the accommodated assumptions. For definite descriptions, the minimisation of individuals ensures that a resolution with referents in the context is preferred over the accommodation of new individuals. This preference has been advocated for on empirical grounds for instance by Strawson [85].

4.1.4 Discourse Models

In what follows, we understand the concept of a discourse model as a logical description of a context that is derived from a discourse. In general, the term *discourse model* denotes an abstraction of a situation in a real or hypothetical world. In cognitive psychology, such discourse models have been used to explain inferences that people draw in understanding text [24]. In Artificial Intelligence, Webber proposes them as describing the situation/state which the speaker is talking about [87]. In the dynamic/DRT (Discourse Representation Theory [55]) trend of natural language semantics, it represents the context created by previous discourse and against which subsequent utterances will be interpreted [55, 57].

Discourse models are abstractions in that we restrict some situation to a finitary or otherwise limited view that focuses on the information that is present in a discourse. In natural-language processing, the task of understanding an utterance is frequently modelled by the task of creating a discourse model by a hearer [24]. The crucial observation underlying the concept of a discourse model is that the situation a speaker is talking about influences the way in which discourse is interpreted. This intuition is made precise in contemporary dynamic theories of meaning which view meaning as a relation between contexts: a sentence is interpreted relative to a context and the interpretation of that sentence yields a new context, the context against which the next sentence will be interpreted. The context change that the models of a semantic representation induce is what will be in the focus of our interest, because this is where we expect to find a considerable part of the truth-conditional meaning.

As our formal framework is that of classical logics, we will focus on phenomena where we can widely ignore the effects of dynamic discourse structure. A typical problem that we will not attack is illustrated by example (1.14).

- (1.14) *Jon's rabbit is cute. The rabbit is white. Peter's rabbit is cute, too. The rabbit is black.*

Definite descriptions sometimes act as anaphors whose binding to referents is determined by the structure of the discourse. Hence, the association of the definite **The rabbit** to its referent dynamically changes in example (1.14) with the focus of the discourse. A convincing treatment of dynamic effects in discourse models would require suitable data structures that can deal

with structural properties such as dynamic accessibility of and focus. As these problems lie beyond the subject of this thesis, our discourse models are representations of situations and not that of discourses. In our analysis, we will concentrate on single sentences that are interpreted with respect to a given logical context.

4.1.5 Models for Definites

In \mathcal{MQL} , we can reduce the semantics of the standard semantic representation of the generalised determiner **THE** to a higher-order definition that relies only on the standard quantifiers \exists and \forall of predicate logic. As an alternative, definition (1.15) makes the two separate conditions in the semantic of Russellian definites more explicit. It is equivalent to the set-theoretical higher-order formulation (1.17) that is sometimes found in the literature [21].

$$(1.15) \quad \text{THE} \equiv \lambda P \lambda Q \text{ UNIQUE}(P) \wedge \text{EVERY}(P)(Q)$$

$$(1.16) \quad \text{UNIQUE} \equiv \lambda P \exists x P(x) \wedge \forall y P(y) \Rightarrow x \doteq y$$

$$(1.17) \quad \text{THE} \equiv \lambda P \lambda Q |P| = 1 \wedge P \subseteq Q$$

Russell's approach to definites leads to a first-order expressible form, and so does an expansion of the definitions (1.15)–(1.17).

We use the Herbrand equality symbol ' \doteq ' in the definition of the second-order unicity predicate **UNIQUE**. As discussed in Section 3.3.3, the symbol \doteq is interpreted as a two-place predicate on elements of a constant frame \mathcal{C} such that $C \doteq D$ is true iff C and D are identical constants. Under \mathcal{MQL} 's unique name assumption for all interpretations, two constants are considered equal iff they have the same name.

The finite model generator **KIMBA** for \mathcal{MQL} logics will be discussed in detail in Chapter 7. **KIMBA** computes, among other classes of models, the locally minimal models of an input specification. With **THE** now at our hands, we can apply **KIMBA** to sentences with singular definite descriptions and investigate whether our semantic representations have the locally minimal models that we expect. In the following, the logical specification of the discourse model for (1.18) is given in (1.19), the semantic representation of the sentence with the definite is the formula (1.20).

$$(1.18) \quad \textit{Jon has a rabbit. **The rabbit** is cute.}$$

$$(1.19) \quad \mathit{rbt}(r_1) \wedge \mathit{has}(\mathit{jon}, r_1)$$

$$(1.20) \quad \mathit{THE}(\mathit{rbt})(\mathit{cute})$$

We apply KIMBA to the logical specification (1.21), i.e., the conjunction of the logical description of the discourse model (1.19) and the semantic representation (1.20). A model for specification (1.21) must consider both the discourse model and the semantic representation. The model generation process generates the facts (1.22) under which both parts of the specification are true.

$$(1.21) \quad \mathit{rbt}(r_1) \wedge \mathit{has}(\mathit{jon}, r_1) \wedge \mathit{THE}(\mathit{rbt})(\mathit{cute})$$

$$(1.22) \quad \{\mathit{has}(\mathit{jon}, r_1), \mathit{rbt}(r_1), \mathit{cute}(r_1)\}$$

The minimal model clearly is an intuitive description of the meaning of the sentence (1.18). In order to make the semantic representation (1.20) true in the discourse model (1.19), the model generator adds the fact $\mathit{cute}(r_1)$ where r_1 has the property of being the only rabbit in the context. In other words, we have computed an interpretation of a the sentence *the rabbit is cute* in a given situational context. The example does not have any other locally minimal models.

Our analysis implements the interpretation of definites as local to a context which is restricted to a discourse. This focus is actually necessary for examples such as (1.23) where we cannot really expect to have a unique rabbit in London.

$$(1.23) \quad \textit{There once was a rabbit in London. } \mathbf{\textit{The rabbit}} \textit{ was } \\ \textit{Welsh.}$$

In our approach, quantification is restricted to the domain of discourse and therefore (1.23) is not taken to claim that there is a unique Welsh rabbit in London, but simply that there is a unique Welsh rabbit in London *which the speaker is talking about*. The analysis is here similar to that described by Groenendijk et al. [41]. Like them we relativise uniqueness to the domain of discourse, not to the world.

The definition of THE correctly expands to an unsatisfiable specification when we interpret the definite description in a discourse model where uniqueness does not hold even locally. In the following, we have a specification (1.24) where our model generator is unable to generate any model at all.

(1.24) *Jon has two rabbits. **The rabbit** is cute. (*)*

(1.25) $r_{bt}(r_1) \wedge has(jon, r_1) \wedge r_{bt}(r_2) \wedge has(jon, r_2) \wedge THE(r_{bt})(cute)$

KIMBA will run infinitely when applied to unsatisfiable specifications such as (1.25). For practical purposes, we restrict the search space to a limited, but reasonable number of extensions in the domain of individuals. The model generation problem then becomes decidable.

4.1.6 Uniqueness and Lots of Rabbits

The definition of the generalised determiner THE in sense of Russell's approach actually defines two separate truth conditions for a definite. Our formulation THE in higher-order logic makes these two conditions explicit.

(1.26) $THE \equiv \lambda P \lambda Q \text{ UNIQUE}(P) \wedge \text{ EVERY}(P)(Q)$

The first condition is that the set P provided by the noun must have exactly one member. The second condition is that this member of P is also a member of Q , i.e., a member of the set given by the scope of quantification of the verb phrase. We will call the first truth condition the **unicity condition**, and the second condition the **subset condition**. The unicity condition seems to be an essential part of the semantics of singular definite descriptions, but for a variety of linguistic material, we can show that unicity with respect to a given context is a constraint which is too strong.

First, as Heim [46] notes, quantifiers may weaken and even annihilate uniqueness. Thus in (1.27), the definite description *the carrot* is unique only per rabbit: there is one carrot per rabbit but there may be many rabbits and therefore many carrots. In (1.28), uniqueness completely disappears: there may even be several carrots per rabbit.

(1.27) *Most rabbits who see one carrot, eat **the carrot**.*

(1.28) *If a rabbit sees a carrot, **the rabbit** eats **the carrot**.*

Second, definites often have some implicit dependency on some other noun phrase. Examples like the following are discussed for instance by Asher and Wada [4] in the general context of anaphora resolution.

(1.29) *When Jon's rabbit dreams of carrots, **the tail** twitches.*

- (1.30) ***The tail** is perhaps the least known of the edible parts of a rabbit.*

Third, as example (1.30) indicates, not all uses of the singular definite article *the* really belong to a definite description. In cases like the following, the description is meant to be general.

- (1.31) ***The rabbit** is a vermin in Australia.*
- (1.32) ***The baseball cap** is an unlikely place to be for **the magician's rabbit**.*

Fifth, there are sentences such as (1.33) and (1.34) which explicitly deny that the description they contain give unique specifications of existing entities.

- (1.33) ***The best way to breed rabbits** does not exist.*
- (1.34) *There is no such thing as **the magical rabbit**.*

Finally, the uniquely identifying property may not be given by the definite description itself, but might have to be somehow inferred from the surrounding context. In (1.35) from Haddock [42], there is no unique hat in the context but the definite description *the hat* refers successfully to the hat that contains a rabbit. In (1.36), *the rabbit* actually refers to the rabbit in the hat, because to remove *x* from *y*, it must be the case that *x* is in *y*.

- (1.35) *A magician has two hats and two rabbits. One rabbit is in a hat. The magician says: "Now watch attentively, I will make **the rabbit in the hat** disappear."*
- (1.36) *Bugs and Bunny are rabbits. Bugs is in **the hat**. John removes **the rabbit from the hat**.*

In the face of such overwhelming evidence, it might seem best to give up uniqueness. There are a number of proposals however which manage to reconcile uniqueness with reality and on which we shall base our computational treatment [25, 41, 54]. In such proposals, uniqueness is not solely determined by the property denoted by the common noun occurring in the definite description. Additional contextual information also plays a rôle. In Cooper's approach [25] this property is a free variable whose value is determined by the context of use. Kadmon [54] asserts that it can be "accommodated, implicated or contextually supplied". Other approaches identify the property with the context set of the definite [41, 88]. In what follows, we adopt and experiment with a combination of these analyses.

4.2 Some Representations

We assume that the definite article *the* is assigned the semantic representation (2.1), which makes use of the definition of set intersection given in (2.2).

$$(2.1) \quad \text{THE} \equiv \lambda P_1 \lambda P_2 \lambda Q \text{UNIQUE}(P_1 \cap P_2) \wedge \text{EVERY}(P_1 \cap P_2)(Q)$$

$$(2.2) \quad \cap \equiv \lambda P \lambda Q \lambda x P(x) \wedge Q(x)$$

Here, P_2 is a first-order property that uniquely identifies the individual referent in the set P_1 which is given by the noun. We call the argument P_2 the **identifying property**. This form of representation requires that P_2 be given and therefore somehow determined. Determining such an identifying property is, as Kadmon remarks, an essentially pragmatic process which can involve accommodation, implicature and/or inference. A context will make many properties available, most of them are irrelevant for determining uniqueness. In what follows, we will manually determine and discuss arguments P_2 that give us correct analyses.

4.2.1 Simple Cases

For many definites, we do not need to give special identifying properties in order to obtain a correct analysis. In such cases, uniqueness holds for the context and we can use an unspecific identifying property $\top_\alpha = \lambda X_\alpha X \doteq X$ that holds for all entities in all domains \mathcal{C}_α of arbitrary type α . In what follows, \top denotes the unspecific identifying property \top_i for first-order individuals.

$$(2.3) \quad \textit{Jon has a rabbit. **The rabbit** is cute.}$$

$$(2.4) \quad \textit{rbt}(r_1) \wedge \textit{has}(\textit{jon}, r_1) \wedge \text{THE}(\textit{rbt})(\top)(\textit{cute})$$

The model generation problem for (2.4) is trivial. We already have a suitable referent for our definite at hand, namely the rabbit r_1 . The unicity condition constrains the resulting discourse model to one where r_1 is the only rabbit, and we therefore have no choice at all when generating a minimal model (2.4). But what about cases where such an individual is missing in the context?

$$(2.5) \quad \textit{Jon's rabbit is cute.}$$

$$(2.6) \quad \text{THE}(\lambda x \text{ rbt}(x) \wedge \text{has}(\text{jon}, x))(\top)(\text{cute})$$

Here, things are a bit more complicated. The minimal model (2.7) that we get could hardly be called intuitive.

$$(2.7) \quad \{\text{rbt}(\text{jon}), \text{has}(\text{jon}, \text{jon}), \text{cute}(\text{jon})\}$$

What is missing in (2.6) is some necessary world knowledge that prohibits an oversimplification. The possessive 's that we modelled by the *has* predicate has an implicit truth condition, namely that the possessor cannot be identical with the possession. Alternatively, we could add the constraint that rabbits cannot possess other rabbits, only humans can. In general, we will simply add additional knowledge about relations as formulas to the specification. When adding an irreflexivity axiom (2.8) for *has*, the model (2.9) that we generate is correct.

$$(2.8) \quad \forall x \neg \text{has}(x, x)$$

$$(2.9) \quad \{\text{rbt}(c_1), \text{has}(\text{jon}, c_1), \text{cute}(c_1)\}$$

The constant c_1 is an automatically generated constant that comes from the iterative extension of the first-order universe in our model generator KIMBA. There is no model that can satisfy (2.6) with a domain of individuals that only consists of the constant *jon*. Hence, KIMBA extends the universe and computes the necessary facts in (2.9). Without a restriction to minimal models, the following models could be generated as well. It should be clear from these examples why we prefer minimal models in general.

$$(2.10) \quad \{\text{rbt}(c_1), \text{has}(\text{jon}, c_1), \text{has}(c_1, \text{jon}), \text{cute}(c_1)\}$$

$$(2.11) \quad \{\text{rbt}(c_1), \text{has}(\text{jon}, c_1), \text{cute}(c_1), \text{cute}(\text{jon})\}$$

$$(2.12) \quad \{\text{rbt}(c_1), \text{has}(\text{jon}, c_1), \text{has}(c_1, \text{jon}), \text{cute}(c_1), \text{cute}(\text{jon})\}$$

The information that the noun phrase of a definite provides, together with some basic world knowledge, is usually sufficient for deciding between anaphor resolution and accommodation. In (2.13), the adjective *black* in the noun phrase *his black rabbit* makes it impossible to identify the white rabbit with Jon's rabbit in the discourse model—if we add the information that a white entity cannot be a black entity. As a result, KIMBA accommodates a new rabbit although we have already introduced a rabbit to the model.

$$(2.13) \quad \text{Jon loves his black rabbit. } \mathbf{The\ white\ rabbit\ is\ boring.}$$

4.2.2 Donkeys, Context Sets, and Anaphoric Use

In the following sections, we consider uses of definites where an analysis based on an unspecific identifying property and minimal models is not sufficient in general.

Quantifiers and Donkey Sentences

Quantifiers can weaken or annihilate uniqueness in definites. Thus in (2.14), the definite descriptions *the rabbit* and *the carrot* cannot be understood as referring to unique entities in the context. There may be many rabbits, and also many carrots. A suitable semantic representation should be equivalent to the first-order formula (2.15).

(2.14) *If a rabbit sees a carrot, **the rabbit** eats **the carrot**.*

(2.15) $\forall x \forall y \text{ rbt}(x) \wedge \text{crt}(y) \wedge \text{see}(x, y) \Rightarrow \text{eat}(x, y)$

Sentences such as (2.14) have been introduced by Geach and are known as “donkey sentences” in the literature. Donkey sentences are sentences where the anaphoric connection we perceive between an indefinite and a pronoun seems to conflict with the implicit existential quantification associated with the indefinite. In the following example, the indefinite in question is *some donkey*, the definite pronoun is *it*, and the

(2.16) *If Pedro owns some donkey, he beats it.*

The donkey sentences are the starting point for the dynamic/DRT approach to natural-language semantics. In Kamp and Reyle’s DRT, sentences are treated within dynamic semantic representations in which logical connectives and bound variables model to some extent the dynamic behaviour of pronouns in natural language. These Discourse Representation Structures (DRS) have a first-order relativisation, i.e., translation. A semantic construction for a donkey sentence in DRT yields a universal quantification in the antecedent, and our donkey sentence (2.16) has the following relativisation.

(2.17) $\forall x \text{ donkey}(x) \wedge \text{owns}(\text{pedro}, x) \Rightarrow \text{beat}(\text{pedro}, x)$

Example (2.18) shows a semantic representation in the spirit of a DRS for sentence (2.14) where we still use our standard formalisation for definites.

$$(2.18) \quad \forall x \forall y \text{ rbt}(x) \wedge \text{crt}(y) \wedge \text{see}(x, y) \Rightarrow \\ \text{THE}(\text{rbt})(\top)(\lambda x \text{ THE}(\text{crt})(\top)(\lambda y \text{ eat}(x, y)))$$

This representation clearly is faulty. When we apply KIMBA to (2.18) and a discourse model where we have several rabbits that each see a carrot, the specification becomes unsatisfiable. The uniqueness that is derived from the definition of THE and the unspecific identifying property \top is too strong. In what follows, we exchange the set \top with identifying properties that are more selective.

Context Set Restrictions

The class of identifying properties that we propose for definites in the scope of quantifiers is derived from a syntactical restriction which originates in Kadmon's B_K set [54]. The B_K set refers to the set of variables that are bound higher up than the variable representing the definite noun phrase. This corresponds to the set of variables that are in the accessibility relation of the definite in DRT. A similar restriction has been referred to as context sets for instance by Westerstahl and others [88]. What we do is to introduce the implicit equality relation that holds between the anaphoric definite and the variables that are bound higher up in the semantic representation.

$$(2.19) \quad \forall x \forall y \text{ rbt}(x) \wedge \text{crt}(y) \wedge \text{see}(x, y) \Rightarrow \\ \text{THE}(\text{rbt})(\lambda r r \doteq x \vee r \doteq y)(\lambda z \text{ THE}(\text{crt})(\lambda c c \doteq x \vee c \doteq \\ y)(\lambda u \text{ eat}(z, u)))$$

The idea is to relativise uniqueness to an equality relation over the set of individuals in the discourse mode, in this case, to the individuals related to the variables bound by a quantifier. The effect here is that uniqueness is relativised to arbitrary rabbit-carrot pairs and hence there is no uniqueness. Given such a specification and a discourse model with one rabbit, say *Bugs*, that sees two carrots, KIMBA will return a model where Bugs sees and eats both carrots. In contrast, given a discourse model where Bugs sees two carrots, but explicitly eats only one, KIMBA will find the specification unsatisfiable and no model will be generated.

The following example is very similar, which is hidden somewhat by the different quantifier. The semantic representation is given in (2.21).

$$(2.20) \quad \textit{Most rabbits that see a carrot, eat **the carrot**}.$$

$$(2.21) \quad \text{MOST}(\mathit{rbt})(\lambda x \forall y \mathit{crt}(y) \wedge \mathit{see}(x, y) \Rightarrow \\ \text{THE}(\mathit{crt})(\lambda z z \doteq x \vee z \doteq y)(\lambda u \mathit{eat}(x, u)))$$

Again, we have some quantifier variables that are used for identifying a referent in the definite. Because we use linguistic quantification with MOST, the quantifier variable x is the bound variable in the λ -abstraction of the scope given by the relative sentence *that sees a carrot* and the verb phrase. Apart from the fact that it is a challenging task to design a semantic construction method for such sentences, our treatment of the definite remains the same. With the identifying property $\lambda z z \doteq x \vee z \doteq y$, uniqueness actually disappears, and we can have for instance a majority of rabbits that see and eat more than one carrot in the discourse model without losing satisfiability.

The semantic representations that we have experimented with so far use identifying properties that are determined only by the context set as a whole. We use the set of all accessible variables because it would be very difficult in general to identify just the right variable automatically. In the examples above, we have some syntactic parallelism between the entities represented by the quantifier and the definite, e.g., *a carrot* and *the carrot*, but this syntactic parallelism is not necessarily sufficient for identifying the right variable in all cases. Example (2.22) is a variant of (2.20) where we would find it impossible to detect the dependency between the quantifier and the definite by syntactic means alone.

$$(2.22) \quad \textit{Every rabbit that sees a carrot, eats } \mathbf{the\ healthy\ veg-} \\ \mathbf{etable.}$$

The encoding (2.23) works for instance in the discourse model (2.24) where we have the minimal model (2.25). The formula $\text{NO}(\mathit{rbt})(\mathit{crt})$ gives the sortal information that rabbits are not carrots, and $\text{EVERY}(\mathit{crt})(\mathit{veg})$ adds the knowledge that carrots are vegetables.

$$(2.23) \quad \text{EVERY}(\mathit{rbt})(\lambda x \forall y \mathit{crt}(y) \wedge \mathit{see}(x, y) \Rightarrow \\ \text{THE}(\mathit{veg} \cap \mathit{hlthy})(\lambda z z \doteq x \vee z \doteq y)(\lambda u \mathit{eat}(x, u)))$$

$$(2.24) \quad \mathit{crt}(c_1) \wedge \mathit{rbt}(r_1) \wedge \mathit{see}(r_1, c_1) \wedge \text{NO}(\mathit{rbt})(\mathit{crt}) \wedge \text{EVERY}(\mathit{crt})(\mathit{veg})$$

$$(2.25) \quad \{\mathit{crt}(c_1), \mathit{hlthy}(c_1), \mathit{veg}(c_1), \mathit{rbt}(r_1), \mathit{see}(r_1, c_1), \mathit{eat}(r_1, c_1)\}$$

The encoding works just as well in cases where we have several rabbits and carrots. In each case, the entity that is both healthy and a vegetable is identified with the carrot mentioned in the antecedent clause for each rabbit-carrot pair.

The Treatment of Names

As we have discussed in the last section, some linguistic theories predict a dependency between an anaphoric definite and the quantifier variables in which scope it occurs. Such a dependency is actually a property of all singular anaphora. If an anaphor cannot be linked to a quantifier variable, it remains unresolved and the discourse becomes ill-formed. What we have not discussed so far is how names of individuals, which can be used just as well for anaphor resolution, fit into the picture.

In DRT, names are modelled by unary predicates and their occurrence introduces existential quantification over the whole representation. In \mathcal{MQL} , we have the names of individuals represented simply as constants because we have a Herbrand-like interpretation of constants in all \mathcal{MQL} interpretations. In the context set restriction that we use, names must be considered as potential referents for anaphoric definites. The following example illustrates the extension of the context set to names.

(2.26) *When Jon visits a doctor, **the man** is happy.*

(2.27) $\forall x \text{ doc}(x) \wedge \text{visit}(\text{jon}, x) \Rightarrow$
 $\text{THE}(\text{man})(\lambda u u \doteq x \vee u \doteq \text{jon})(\text{happy})$

The discourse is ambiguous because it is not entirely clear to whom *the man* refers to, Jon or the doctor. An identifying property $\lambda u u \doteq x \vee u \doteq \text{jon}$ that leaves the computation of the actual resolution to the model generator allows us to compute two minimal models, one where Jon is the man and is happy when he visits any doctor, and one where every (male) doctor is happy when Jon visits him.

4.2.3 Restrictions with Knowledge

What is interesting in our minimal-model-based approach to anaphor resolution is that additional contextual knowledge can be added which controls the resolution. For instance in example (2.26), if we add the information that Jon only visits female doctors, the computed models change accordingly. Note that an encoding of definites with context sets makes it impossible to accommodate new individuals for the definite—if we use context sets, we implicitly assume an anaphoric use of the definite. For some definites, the suitable identifying properties that are based on some pragmatically given restrictions still allow accommodation.

Implicit Knowledge and Accommodation

In example (2.28), the uniqueness in the definite *the rabbit* relates to *the rabbit in the hat*. The proposed logical specification is given in (2.29).

(2.28) *Bugs and Bunny are rabbits. Bugs is in the hat. Jon removes **the rabbit** from **the hat**.*

(2.29) $\text{THE}(\text{rbt})(\lambda v \text{ inhat}(v))(\lambda x (\text{THE}(\text{hat})(\top)(\lambda y \text{ rmv}(jon, x, y))))$

The context set restriction from the last section would not help here because both Bugs and Bunny are in the context set. Instead, we use an identifying property $\lambda x \text{ inhat}(x)$ that is a pragmatic restriction inferable only from what we know about the world: removing x out of y implies that x was in y^2 . Hence, the rabbit that is removed from the hat must have been in the hat. The locally minimal model generated by KIMBA for specification (2.29) is as expected (2.30).

(2.30) $\{\text{rbt}(\text{bugs}), \text{rbt}(\text{bunny}), \text{inhat}(\text{bugs}), \text{rmv}(jon, \text{bugs}, \text{hat})\}$

Now suppose there is no rabbit known to be in a hat. In that case, KIMBA accommodates the fact that one of the rabbit was indeed in the hat, and yields two possible minimal interpretations that satisfy the representation. In the first case (2.32), Bugs is in the hat, and in the second case (2.32), Bunny is in the hat.

(2.31) $\{\text{rbt}(\text{bugs}), \text{rbt}(\text{bunny}), \text{inhat}(\text{bugs}), \text{rmv}(jon, \text{bugs}, \text{hat})\}$

(2.32) $\{\text{rbt}(\text{bugs}), \text{rbt}(\text{bunny}), \text{inhat}(\text{bunny}), \text{rmv}(jon, \text{bunny}, \text{hat})\}$

In general, KIMBA will accommodate any fact that is necessary provided that it is consistent with the rest of the specification. An interesting question that would be worth exploring in this context is that of computable constraints on accommodation. We leave the question open for now and return to it later in Section 4.2.4.

Another variation³ on (2.28) is the following:

²For simplicity, we ignore temporal issues here.

³Thanks to Bonnie Webber for this particular version of the example. Zeevat [90] proposes a similar example namely: “A man died in a car accident last night. The Amsterdam father of four had been drinking.”

- (2.33) *Bugs and Bunny are rabbits. Bugs is in the hat. Jon re-moves **the birthday present** from the hat.*

In this case, we may infer that the birthday present is Bugs. And so does KIMBA in the locally minimal model.

Bridging

Bridging refers to the dependency of a definite to some other noun phrase such as in the following example.

- (2.34) *Jon's rabbit dreams. **The tail** twitches.*

As **The tail** here refers to the tail of Jon's rabbit, an acceptable interpretation of **The tail twitches** must resolve the definite to the tail of Jon's rabbit. A correct treatment of bridging requires that we make explicit the relationship of rabbits and tails by a formula such as (2.35). The formula states that for each rabbit we have a unique tail that is a part of that rabbit.

- $$(2.35) \quad \forall x \text{ rbt}(x) \Rightarrow \text{UNIQUE}(\lambda y \text{ tail}(y) \wedge \text{has}(x, y))$$

If we add formula (2.35) to the logically encoded context, the locally minimal model of the semantic representation $\text{THE}(\text{tail})(\top)(\text{twitch})$ represents the correct reading. The model identifies a tail which is introduced by formula (2.35) as that (only) tail that both belongs to Jon's rabbit and which also twitches. Note that the semantic representation does not give a specific identifying property.

However, in cases such as the following, it is reasonable that a suitable identifying property must somehow introduce the implicit dependency between the definite and the dependent noun.

- (2.36) *If a rabbit dreams, **the tail** twitches.*

Example (2.36) is a donkey sentence that cannot be dealt with by simply relating the definite to the context set because we only have a variable for rabbits, and not a variable for tails. Below, we give the semantic representations that do the trick, and the natural-language sentence that they actually stand for.

- (2.37) *When Jon's rabbit dreams, **the tail of Jon's rabbit** twitches.*

$$(2.38) \quad \text{THE}(\lambda x \text{rbt}(x) \wedge \text{has}(\text{jon}, x))(\top)(\text{dream}) \Rightarrow \\ \text{THE}(\text{tail})(\lambda u \text{THE}(\lambda v \text{rbt}(v) \wedge \text{has}(\text{jon}, v)) \\ (\top)(\lambda v \text{of}(u, v)))(\text{twitch})$$

$$(2.39) \quad \text{If a rabbit dreams, **the tail of it** twitches.}$$

$$(2.40) \quad \forall x \text{rbt}(x) \Rightarrow \text{THE}(\text{tail})(\lambda u \text{THE}(\text{rbt}) \\ (\top)(\lambda v \text{of}(u, v)))(\text{twitch})$$

As we can see, instead of an identifying property that simply relates two variables, we introduce an *of*-relation between two variables. As in the case of quantifier scoping, bridging introduces a relation, but this relation may be more complex than the simple equality of entities. The *of*-relation is a placeholder for a variety of actual relations that can hold between the definite and the dependent noun. For instance, in (2.41), the waiting room refers to the waiting rooms of the doctors' offices, not the waiting room of each doctor. Still, we can use *of* as long as we only want to make sure that an analysis with local uniqueness remains possible.

$$(2.41) \quad \text{If Jon visits a doctor, **the waiting room** empties.}$$

$$(2.42) \quad \forall x \text{doctor}(x) \wedge \text{visit}(\text{jon}, x) \Rightarrow \\ \text{THE}(\text{waitrm})(\lambda v \text{of}(x, v))(\text{empties})$$

4.2.4 Simple Cases Revisited

The identifying properties P_2 that we have used so far can be roughly classified into two categories.

- The identifying property is \top for simple cases.
- The identifying property relates the definite to a set of variables that are bound higher up in the semantic representation. A suitable relation cannot be determined from syntactical structure alone, but must sometimes be inferred from world knowledge.

It is not easy to classify the use of a definite as a simple case or as one where we need a non-trivial restriction on P_2 . The problem is that the simple cases with the identifying property \top sometimes work for sentences where a correct analysis should yield an unsatisfiable specification. This occurs mainly in connection with accommodation and the dynamic behaviour of definites that should be interpreted exclusively as anaphora.

(2.43) *Jon has no rabbit. The rabbit is cute. (*)*

(2.44) $\text{NO}(\text{rbt})(\lambda x \text{has}(\text{jon}, x)) \wedge \text{THE}(\text{rbt})(\top)(\text{cute})$

What we have here is a sentence (2.43) where a naive modelation of a discourse model in (2.44) will not prevent that we can generate a minimal model. Model generation implements an almost unconstrained form of accommodation. In our example, KIMBA simply accommodates an individual that is a cute rabbit but necessarily not Jon's one. Hence, formalisation (2.44) is at least odd as a semantic representation of (2.43). What is missing in our representation is that the definite should be treated as an anaphor.

Non-resolvable Anaphora in DRT

Solving the problem of anaphor resolution in the presence of dynamic quantifier scope and accessibility relations is the domain of dynamic logics. We might be tempted to say that our problem here is really that of a static logic representation. Arguably, we must be careful to use a suitable relativisation of intrinsically dynamic phenomena, such as the accessibility of discourse referents, in our logic \mathcal{MQL} . Consider for instance the following sentence.

(2.45) *Jon has no rabbit. He likes it.*

In DRT, the semantic construction process yields a DRT whose first-order form is given by (2.46). The pronoun *it* is resolved with a referent that is actually not accessible. The equation $y \doteq x$ which states that *it* refers to the non-existing rabbit results in a formula which is not well-formed. DRT on this ground rejects the semantic construction of (2.46).

(2.46) $\neg \exists y (\neg \exists x \text{rbt}(x) \wedge \text{has}(\text{jon}, x)) \wedge \text{like}(\text{jon}, y) \wedge y \doteq x$

If we use a DRT relativisation (2.47) for the example (2.43) where the definite is treated in the same spirit as an anaphor, we also have an ill-formed formula.

(2.47) $\neg(\exists x \text{rbt}(x) \wedge \text{has}(\text{jon}, x)) \wedge \text{THE}(\text{rbt})(\lambda u u \doteq x)(\text{cute})$

Definites are not Anaphora

The DRT-solution to non-resolvable anaphora relies on a semantic construction process that detects some syntactically ill-formed uses of anaphora before a semantic representation is presented. The ill-formedness in the representation comes from equations where one of the variable occurs free. The equations are introduced by the provisional instantiation of pronouns, i.e., by anaphor resolution.

For definites, there is no equivalent and generally useful filter mechanism does not exist because definites are not always anaphora. As mentioned earlier, definites sometimes presuppose the existence of discourse referents that have not been introduced yet (2.48), or trigger bridging (2.49) which only relates the definite to a previously mentioned referent in some indirect way. Finally, we have the use of singular *the*-phrases like (2.50) that are meant as general statements.

(2.48) *Jon's new rabbit is black.*

(2.49) *Jon has a new rabbit. **The tail** is black.*

(2.50) ***The rabbit** is a vermin in Australia.*

In any of these cases, if we use identifying properties that restrict the interpretation of definites to anaphoric uses, we will obtain unsatisfiable specifications. Unfortunately, linguistics so far has not provided us with computational methods for distinguishing anaphoric and non-anaphoric uses of definites.

Model generation for locally minimal models prefers anaphoric use over non-anaphoric use, because locally minimal models minimise the universe. This preference seems to match the empirical data. Apart from that, we must allow accommodation in all cases where we have consistency of the assumed facts in the model. In cases where only accommodation of new individuals yields a model, model generation on its own does not really have a means to distinguish valid and non-valid forms of accommodation.

Non-Existence

An interesting phenomenon are sentences which explicitly deny the uniqueness that the definite implicitly formulates.

(2.51) *The golden mountain does not exist.*

According to Russell, sentences such as (2.51) have two possible analyses, one where negation has a wide scope over the quantificational complex which represents the contribution of the definite, and one with narrow scope. The first analysis is equivalent to our formulation (2.52), while the second is equivalent to (2.53) for our example (2.51).

(2.52) $\text{THE}(\text{goldMountain})(\top)(\lambda x \neg \top(x))$

(2.53) $\neg \text{THE}(\text{goldMountain})(\top)(\top)$

Both analyses are not intuitive, as can be shown by investigating their models. In (2.52), we state that there is a unique golden Mountain that has the property of non-existence—note that \top denotes the property that all individuals in the discourse model have. This formalisation is unsatisfiable, because model generation must accommodate an individual on the one hand, and show that it does not belong to the discourse model at the same time. This is inconsistent.

As Kamp and Reyle argue [56], the second formalisation (2.53) says something which is closer to the intuitive meaning of (2.51). It says that there is no unique object which has the property of being the golden mountain that exist. Unfortunately, our model generator proves that the formalisation still is not correct. It has a non-minimal model (2.54) that instantiates another potentially valid interpretation of the formalisation.

(2.54) $\{\text{goldMountain}(c_1), \text{goldMountain}(c_2)\}$

Informally, Russell's semantic representation can be validated by assuming that there are two golden mountains c_1 and c_2 . The model (2.54) reveals the subtle problem that the Russellian encodings for definites do not always preserve entailment. In this special case, we cannot simply fix the problem by giving a suitable identifying property because an identifying property that determines an empty set is not possible.

The example shows that it sometimes pays to consider non-minimal models as well as minimal ones when we use model generation in the analysis of semantic representations.

4.3 What We Have Learned so Far

We have shown that model generation can provide interesting insights from existing semantic theories of definites with a computational interpretation that combines reasoning on linguistic and world knowledge. Given a suitable uniquely identifying property and a logical encoding of the context, the model generator either identifies the referent of a definite noun phrase with some already existing entity (coreference) or adds a new entity to the model (accommodation). We have determined some classes of identifying properties for our formalisation for some well-known examples from the literature. This by itself is not entirely trivial, and we know of no existing approach to the semantics of definites that can treat the range of examples that we have presented so far.

For the analysis of definites, the model generator KIMBA provides an adequate notion of minimality. Locally minimal models validate a specification within the smallest possible domain of individuals. As a result, our method for definites favours coreference over accommodation whenever coreferences can be assumed consistently. As has been stated first by Strawson [85], a correct computational treatment of definites asks for a method that can handle accommodation in this way.

Our approach often makes the meaning postulates explicit which are necessary for natural-language interpretation. A formalisation which leaves out this important part of a meaning quickly leads to oversimplification. We can use model generation as a software engineering tool and apply it to complex formalisations where we may not be sure that we actually have a correct semantic representation. For practical research in model-theoretic semantics, model generation helps to formulate and verify linguistic theories and allows us to experiment with their predictions.

Minimal model generation correctly treats the class of definites that we have called “simple cases”, i.e., where the identifying property can be safely chosen as \top , provided that we have a correct logic modelation of the context of an utterance and a suitable representation in which the definite is embedded. These simple case not only include some anaphoric uses of definites, but also some common forms of bridging and accomodation. These often require reasoning about the world which model generation provides as a built-in feature.

For other examples, we have shown that minimal model generation is general enough to allow for an implementation of contemporary linguistic

theories [25, 41, 46, 54, 88]. Our “identifying properties” can uniformly represent the various proposed constraints on the interpretation of definites.

Unfortunately, these theories do not yield specific methods for computing the identifying properties in general. For instance, as we have seen in Section 4.2.3, the use of a syntactic constraint such as a context set does block accommodation. If we want to employ the context set restriction correctly, we must somehow decide when to use it and when we should stay with a less restrictive formalisation that allows accommodation. It is unlikely that this problem has a sound computable answer. In Section 4.2.2, we have shown for instance that there can be no purely syntactic method that identifies anaphoric uses of definites. It remains unclear how we can then satisfactorily model an automatic context set restriction when we cannot even decide whether a definite is used as an anaphor or not. The contemporary theories state that the relation which determines the unicity of a definite must be ‘referred from the context’ or ‘is essentially a process of accommodation, implicature and/or inference’. The semantic theories that we know remain vague on the nature of the consequence relation that determines the inference that is referred to.

However, as Kamp and Reyle [56] note, it would be a non-trivial task to identify and describe the different purposes to which singular *the*-phrases can be put, let alone the actual constraints on their interpretation. Hence, our analysis remains incomplete in that we do not have more reliable computational means which automatically determine identifying properties from the linguistic data. What we have presented in this chapter is only a starting point for a more throughout analysis of definites.

Chapter 5

Reciprocity

Once, indeed, a problem was brought to me, and I solved it, obtaining very many solutions; I went into it fully, and found that there were 2678 valid answers. I marvelled at this, only to discover - when I spoke of it - that I was reckoned a simpleton or an incompetent, and strangers looked on me with suspicion.
(Abu Kamil)

Overview: Linguistic theories on reciprocal expressions identify a semantic principle, the Strongest Meaning Hypothesis, as the key concept for determining the logical contributions of reciprocals. Based on this theory, we present a new and simple method for computing reciprocal meaning by model generation.

5.1 Introduction

Research on the English reciprocal expressions *each other* and *one another* has uncovered a variety of meaning contributions that a reciprocal can provide. Consider the following sentences.

- (1.1) *The students like each other.*
- (1.2) *The students gave each other measles.*
- (1.3) *The students stare at each other in surprise.*

(1.4) *The students follow each other into the ballroom.*

In each case, the reciprocal implicitly formulates truth conditions for the interpretation of the relation that is expressed by the verb phrase. These truth conditions are different for each example. For instance, we can accept (1.1) to be true only if for each pair x and y of different students holds that x likes y . An analogous interpretation would be invalid in the case of (1.2)–(1.4) where not all pairs in the reciprocal group *the students* can consistently be in the scope relation.

Dalrymple et al. [27] argues that the variety of reciprocal meaning can be reduced to six different classes of reciprocal semantics. The correct choice amongst these classes is determined by a purely semantic principle, the Strongest Meaning Hypothesis. In its most general form, this principle can be stated as follows.

Strongest Meaning Hypothesis (SMH): *The meaning of an expression S in a context Γ is that meaning which corresponds to the interpretation of the logically strongest semantic representation ϕ available for S that is consistent with Γ .*

The SMH as formulated here lies at the border of semantics construction and natural-language interpretation. Each of the sentences (1.1)–(1.4) exemplifies a separate class of reciprocal semantics and the truth conditions expressed in the other examples are either too strong or too weak to cover the meaning of the reciprocal. The SMH identifies that class of reciprocity as the correct one whose contribution to the scope relation is the strongest one that still is consistent with the information provided by the context.

It can be argued that the variation of meaning in reciprocal expressions lies more in the way in which the scope relations can be interpreted over the antecedent groups rather than in some variation of the semantics of reciprocal expressions. We present a new theory of reciprocals where the SMH can be understood as a principle of interpretation alone. In our approach, we implement the preference for certain interpretations as a form of minimal model reasoning. Our version of the SMH maximises the logical contribution of the reciprocal to the scope relation in the set of logical models of a discourse. We use only one semantic representation for reciprocal expressions and model their complex behaviour by a refinement of a well-known minimal

model constraint that has found some use for instance in diagnosis applications. As usual, we will experimentally verify our theory with our model generator KIMBA.

5.2 Exploring the Meaning of *Each Other*

In the linguistic literature, the starting point for the semantic analysis of reciprocals are often sentences such as

$$(2.1) \quad \textit{Jon and Bill saw each other.}$$

where two discourse participants are in a reciprocal relation expressed by the verb phrase. Such sentences are particularly easy to analyse. The English reciprocal expressions *each other* and *one another* are frequently represented as diadic quantifiers over some first-order set called the **antecedent group**, and a binary first-order relation called the **scope relation**. In what follows, we will use the symbol RCP for such reciprocal quantifiers. In higher-order logics, example (2.1) is then represented as follows.

$$(2.2) \quad \text{RCP}(\{\textit{jon}, \textit{bill}\})(\lambda y \lambda x \textit{saw}(x, y))$$

Here, $\{\textit{jon}, \textit{bill}\}$ denotes the antecedent group, i.e., a first-order property that holds exactly for *jon* and *bill*, namely $\lambda x (x = \textit{jon} \vee x = \textit{bill})$.

When groups $P_{t \rightarrow o}$ of just two members are considered, each group member is required to stand in some scope relation $R_{t \rightarrow t \rightarrow o}$ to the other member, where R is provided by the verb phrase. This truth condition for reciprocals can be formalised by definition (2.3)¹. The formalisation uses the set exclusion operator $'/'$ which is specified by the λ -term in definition (2.4).

$$(2.3) \quad \text{RCP} \equiv \lambda P \lambda R \text{ EVERY}(P)(\lambda x \text{ EVERY}(P/{x})(\lambda y R(x, y)))$$

$$(2.4) \quad / \equiv \lambda P \lambda Q \lambda x P(x) \wedge \neg Q(x)$$

The minimal model (2.5) of the semantic representation (2.2) under the definition (2.3) is as expected, and represents nicely the natural-language interpretation of (2.1).

$$(2.5) \quad \{\textit{saw}(\textit{bill}, \textit{jon}), \textit{saw}(\textit{jon}, \textit{bill})\}$$

¹An equivalent definition is given by Carpenter [21].

5.2.1 Reciprocals for Larger Groups

The application of the semantic representation (2.3) to larger antecedent groups suggests for instance that

- (2.6) *House of Commons etiquette requires **legislators** to address only the speaker of the House and **refer to each other indirectly**.*

states that each legislator is required to refer to every other one indirectly. For example (2.6), this is indeed the intended meaning. However, the research on reciprocals has shown that the truth conditions implied by definition (2.3) turn out to be the wrong ones for many cases where the antecedent group has more than two members. Statement (2.7) from J.M. Barrie's *Peter Pan* exemplifies a meaning of the reciprocal where every member is claimed to relate to at least one other group member, but not necessarily to relate to *every* other one.

- (2.7) *"The captain!" said the **pirates, staring at each other in surprise**.*

It is impossible for the pirates to stare at each other such that each pirate stares at every other pirate at the same time. The truth conditions for the reciprocal that is implied by this example is obviously weaker than the one that we have for our first example. In other words, the binary relation R whose interpretation is partially determined by the semantics of the reciprocal can be interpreted more freely. As the following sentences further illustrate, R may instantiate a variety of relations.

- (2.8) *The blocks are stacked on top of each other.*
- (2.9) *Five Boston pitchers sat alongside each other.*
- (2.10) *Most people at the party are married to each other.*

Example (2.8) can be interpreted as true in contexts where the blocks mentioned are stacked in any way that is possible, including for instance a pyramid or a tower of blocks. Sentence (2.8) can never have the literal meaning that each block has the property of being stacked on top of each other one.

Sentence (2.9) describes a situation where five Boston pitchers sit on a bench. Naturally, we expect the ones in the middle each to have just two other pitchers sitting alongside, and the ones at the end to have just one pitcher on their side. The possible relation R for the situation at hand is more constrained than the previous one in that each pitcher must have another pitcher he sits alongside to. In sentence (2.8), the lowest layer of blocks may not on top of any other ones.

Finally, the reciprocal in (2.10) only means that we have a spouse for most persons that are at the party. The binary relation R must be one that assigns to each man in the group of most persons that is referred to a woman he is married to, and vice versa. The relation R in question does not even relate all married people with each other directly or indirectly.

5.2.2 Classifying Reciprocal Meaning

Dalrymple et al. [27] provides a classification system whose classes approximate the various forms of reciprocal meaning. Based on the earlier works of Langendoen [61], they propose six classes of reciprocal meaning: Strong Reciprocity, Strong Alternative Reciprocity, Intermediate Reciprocity, One-Way Weak Reciprocity, Intermediate Alternative Reciprocity, and Inclusive Alternative Ordering. Every occurrence of a reciprocal expression falls under one or several of these categories and each class can be given a precise semantic representation in higher-order logic. In what follows, we describe the aforementioned classes of reciprocal semantics and give their formalisations in higher-order logic.

Strong Reciprocity

The strictest form of reciprocity, **Strong Reciprocity** (SR) refers to reciprocals where the relation R holds for each member x of the antecedent set P and each other member y . As discussed above, the following statement (2.11) is an example for SR.

(2.11) *Legislators must refer to each other indirectly.*

The truth condition for such a reciprocal is twofold. First, the reciprocal can only be true in contexts where the antecedent set—in our case, the set of legislators—has at least two members. Second, we have the aforementioned

condition on the binary relation R . In (2.11), R is the property of referring indirectly, and its strong meaning is that $R(x, y)$ holds for each legislator x and each other legislator y . All this leads to the following formalisation for reciprocity.

$$(2.12) \quad \text{RCP}_{SR} \equiv \\ \lambda P \lambda R \text{CARD}^{\geq 2}(P) \wedge \\ \text{EVERY}(P)(\lambda x \text{EVERY}(P/\{x\})(\lambda y R(x, y)))$$

The quantifier $\text{CARD}^{\geq 2}$, i.e., $\text{CARDinality} \geq 2$, has a higher-order definition (2.13) that makes use of standard first-order quantification and equality.

$$(2.13) \quad \text{CARD}^{\geq 2} \equiv \lambda P \exists x \exists y \neg(x = y) \wedge P(x) \wedge P(y))$$

One-Way Weak Reciprocity

The example from *Peter Pan* exhibits a form of reciprocity where each member of the antecedent set must be the subject of the relation R and be related with at least one other member as the object. Strong Reciprocity does not hold.

$$(2.14) \quad \textit{“The captain!” said the pirates, staring at each other in surprise.}$$

The form of reciprocity in example (2.14) is called **One-Way Weak Reciprocity** (OWR). Its formalisation is as follows.

$$(2.15) \quad \text{RCP}_{OWR} \equiv \\ \lambda P \lambda R \text{CARD}^{\geq 2}(P) \wedge \\ \text{EVERY}(P)(\lambda x \text{SOME}(P/\{x\})(\lambda y R(x, y)))$$

One-Way Weak Reciprocity is a strictly weaker form than Strong Reciprocity, i.e., One-Way Weak Reciprocity does hold in all cases where Strong Reciprocity holds, but not necessarily vice versa. The weakening corresponds directly to the exchange of one of the second of the two universal determiners **EVERY** in **SR** by the existential determiner **SOME**.

Inclusive Alternative Ordering

A further weakening of the reciprocity conditions occurs in examples such as (2.16) where not every member of the antecedent set P must be a subject of the relation R . Instead, it suffices that every member of P is at least the object of the relation.

- (2.16) *He and scores of other inmates slept on **foot-wide wooden planks stacked atop each other**—like sardines in a can—in garage-sized holes in the ground.*

A slight modification of the definition of RCP_{OWR} suffices to obtain the formalisation of **Inclusive Alternative Ordering** (IAO). IAO is the weakest known form of reciprocity.

- (2.17) $\text{RCP}_{\text{IAO}} \equiv$
 $\lambda P \lambda R \text{CARD}^{\geq 2}(P) \wedge$
 $\text{EVERY}(P)(\lambda x \text{SOME}(P/\{x\})(\lambda y R(x, y) \vee R(y, x))$

Intermediate Reciprocity

Dalrymple et al. [27] cite the following statement from the *New York Times* for exemplifying a form of reciprocal meaning called **Intermediate Reciprocity** (IR). The statement can be true despite the impossibility of each group member sitting alongside.

- (2.18) *As the preposterous horde crowded around, waiting for the likes of Evans and Mike Greenwell, **five Boston pitchers sat alongside each other**: Larry Andersen, Jeff Reardon, Jeff Gray, Dennis Lamp and Tom Bolton.*

Informally, IR states that the the relation R relates all members of P directly or indirectly via a sequence of members of P . In other words, every two different elements x and y must be in the transitive closure of R with respect to P . The IR form of reciprocity is modelled by the following definition.

- (2.19) $\text{RCP}_{\text{IR}} \equiv$
 $\lambda P \lambda R \text{CARD}^{\geq 2}(P) \wedge$
 $\text{EVERY}(P)(\lambda x \text{EVERY}(P/\{x\})(\lambda y \text{TRANSCl}(P)(R)(x, y)))$

The formula $\text{TRANSCl}(R)(P)(x, y)$ denotes that there is a sequence z_0, \dots, z_n of elements of P such that $z_0 = x$, $z_n = y$ and $R(z_i, z_{i+1})$ for all $i < n$.

Intermediate Alternative Reciprocity

Examples (2.20)² and (2.21) exhibit yet another variant of reciprocity whose truth conditions are different from all other ones that we have investigated so far.

(2.20) *The students in Mrs. Smith's class gave each other measles.*

(2.21) *Instead, **countless stones**—each weighting an average of 300 pounds—**are arranged on top of each other** and are held in place by their own mass and the force of flying buttresses against the walls.*

The form of reciprocity here is called **Intermediate Alternative Reciprocity** (IAR) and is characterised by a relation R that connects all members of the antecedent set in the fashion of a strongly connected acyclic graph. It suffices that each member x is related to every other member y via a chain of R -relations where we ignore which way the pairs are connected. In the situation described by sentence (2.21), we have a cathedral which is built of stones arranged in a pattern like a brick wall. The relation we have here is not symmetric, but every brick is part of single connected structure. Dalrymple et al. claim that the sentence would be false in a context where the stones are arranged in a multiplicity of piles.

The formalisation of IAR is a variant of Intermediate Reciprocity where the concept of a transitive closure is weakened.

$$(2.22) \quad \text{RCP}_{IAR} \equiv \\ \lambda P \lambda R \text{CARD}^{\geq 2}(P) \wedge \\ \text{EVERY}(P)(\lambda x \text{EVERY}(P/\{x\}) \\ (\lambda y \text{TRANSCL}(P)(\lambda u \lambda v R(u, v) \vee R(v, u))(x, y)))$$

The formula $\text{TRANSCL}(P)(\lambda u \lambda v R(u, v) \vee R(v, u))$ denotes that there is a sequence z_0, \dots, z_n of elements of P such that $z_0 = x$, $z_n = y$ and $R(z_i, z_{i+1})$ **or** $R(z_{i+1}, z_i)$ for all $i < n$.

²It can be argued that this example illustrates IAO rather than IAR.

Strong Alternative Reciprocity

Strong Alternative Reciprocity is a weakened form of Strong Reciprocity and an exceptional class as there still are, as far as we know, no examples that can univocally be identified with it. All known examples of SAR also meet the truth conditions of SR. Hence, the SAR class owes its existence only to the parameterisation and classification scheme of Dalrymple et al. that we describe later in Section 5.2.3. SAR's definition in higher-order logic is as follows.

$$(2.23) \quad \text{RCP}_{SR} \equiv \\ \lambda P \lambda R \text{CARD}^{\geq 2}(P) \wedge \\ \text{EVERY}(P)(\lambda x \text{EVERY}(P/\{x\})(\lambda y R(x, y) \vee R(y, x)))$$

5.2.3 Parameterisation

The differences between the various definitions of RCP can be parameterised. This parameterisation was helpful to Dalrymple et al. to discover new forms of reciprocity that have not previously been considered in the literature.

The first parameter of variation determines how the scope relation R should cover the domain P . The examples we have seen each fall into the following categories.

- each pair of different individuals in P may be required to participate in the relation R directly (FUL).
- each pair of different individuals in P may be required to participate in the relation R directly or indirectly (LIN).
- each individuals in P may be required to participate in the relation R with another one (TOT).

The second parameter concerns how FUL, LIN, and TOT operate on the reciprocal's parameter R : whether the relation R that reciprocity requires between individuals in the domain is actually the extension of the reciprocal's scope, or the extension where we ignore the direction in which the relation R holds by adding inverse pairs R^{-1} . By R^{\vee} , we denote the parameter where we use the extension $R \cup R^{-1}$ in this way.

The Landscape of Reciprocity

The following table shows how the two parameters determine the various semantics of reciprocals. It presents the complete landscape of reciprocal meaning expressible by the variations of the two parameters. The landscape defines a partial order of the semantic contributions of the reciprocal meanings that is shown in Figure 5.1.

	FUL	LIN	TOT
R	SR	IR	OWR
R^\vee	SAR	IAR	IAO

The Strong Reciprocity (SR) form of reciprocity requires that each pair of different individuals directly participates in the scope relation R . Intermediate Reciprocity (IR) and One-Way Weak reciprocity (OWR) weaken the first parameter to LIN resp. TOT. SAR, i.e., Strong Alternative Reciprocity, is the weakening of SR that uses the notion R^\vee of directedness over the scope relation. Intermediate Alternative Reciprocity (IAR) and Intermediate Alternative Ordering (IAO) are the corresponding weakenings of IR and OWR resp. with regard to R^\vee .

Of the forms of reciprocity mentioned so far, SAR, OWR and IAR have been discovered by Dalrymple et al. and have not previously been considered in the literature.

Parameterised Definitions

The parameterisation in the last section gives us an elegant and uniform way to present the different formalisations for reciprocal quantification. The variations of the first parameter is given by the following sequence of definitions.

$$(2.24) \quad \text{FUL} \equiv \lambda P \lambda Q \text{ EVERY}(P)(\lambda x \text{ EVERY}(P/\{x\})(\lambda y Q(x, y)))$$

$$(2.25) \quad \text{LIN} \equiv \lambda P \lambda Q \text{ EVERY}(P) (\lambda x \text{ EVERY}(P/\{x\})(\lambda y \text{ TRANSCl}(Q)(P)(x, y)))$$

$$(2.26) \quad \text{TOT} \equiv \lambda P \lambda Q \text{ EVERY}(P)(\lambda x \text{ SOME}(P/\{x\})(\lambda y Q(x, y)))$$




Figure 5.1: The partial order on the logical contribution of reciprocals.

The first parameter of reciprocals is defined as a diadic quantifier over a first-order set P and a binary first-order relation Q . Every definition shows how each of the three types **FUL**, **LIN**, and **TOT** constrains the relation Q that holds between pairs of elements x and y in P . In the case of **FUL**, each pair must be directly in the relation Q , while **LIN** specifies that each pair of different elements x and y to be in the transitive closure of Q with regard to P . Finally, **TOT** states that every x must be in the relation R with some arbitrary y .

The second parameter of reciprocal meaning determines whether the argument Q in the definition of the first element is actually the relation R provided by the verb phrase of the reciprocal, or the relation R^\vee , i.e., the relation R extended by its inverse R^{-1} . We define an extension operator REL^\vee on binary relations which allows us to use the extension R^\vee as part of higher-order formalisations.

$$(2.27) \quad \text{REL}^\vee \equiv \lambda Q \lambda y \lambda x \ Q(x, y) \vee Q(y, x)$$

Every form of reciprocal semantics that Dalrymple et al. consider as valid can now be given a compact higher-order definition as follows.

$$(2.28) \quad \text{RCP}_{SR} \equiv \lambda P \lambda R \ \text{CARD}^{\geq 2}(P) \wedge \text{FUL}(P)(R)$$

$$(2.29) \quad \text{RCP}_{SAR} \equiv \lambda P \lambda R \ \text{CARD}^{\geq 2}(P) \wedge \text{FUL}(P)(\text{REL}^\vee(R))$$

$$(2.30) \quad \text{RCP}_{IR} \equiv \lambda P \lambda R \text{CARD}^{\geq 2}(P) \wedge \text{LIN}(P)(R)$$

$$(2.31) \quad \text{RCP}_{IAR} \equiv \lambda P \lambda R \text{CARD}^{\geq 2}(P) \wedge \text{LIN}(P)(\text{REL}^{\vee}(R))$$

$$(2.32) \quad \text{RCP}_{OWR} \equiv \lambda P \lambda R \text{CARD}^{\geq 2}(P) \wedge \text{TOT}(P)(R)$$

$$(2.33) \quad \text{RCP}_{IAO} \equiv \lambda P \lambda R \text{CARD}^{\geq 2}(P) \wedge \text{TOT}(P)(\text{REL}^{\vee}(R))$$

5.2.4 Interpreting Reciprocals

The interpretation of reciprocal expressions requires that we correctly identify the truth conditions that the reciprocal provides. As we have seen in the previous sections, reciprocals contribute different meanings in different contexts. The parameterisation in Section 5.2.3 identifies six candidate meanings, each of which will impose different constraints on the reciprocal's relation R . The actual selection amongst this set can be modelled by applying a certain semantic principle.

The Strongest Meaning Hypothesis

Many expressions in natural language, for instance homonyms, may ambiguously refer to different semantics where the choice of which one is actually meant relies with the speaker. As Dalrymple et al. argue, this is not the case for reciprocals where the literal meaning is determined only by the context in which the reciprocal is uttered. Hence, reciprocal meaning is independent from the speaker. The device for identifying the correct reciprocal semantic is a simple principle, the Strongest Meaning Hypothesis (SMH).

Strongest Meaning Hypothesis for Reciprocals: *The semantic representation of a reciprocal expression S in a context Γ is that representation ϕ whose logical contribution to the reciprocal relation is the strongest consistent one with respect to Γ .*

In the case of reciprocals, the semantic representations that are available in general are those which use any one of the six quantifiers SR, IR, OWR, SAR, IAR, and IAO. The “strongest” contribution to the scope relation is determined by the partial order shown in in Figure 5.1. Although this ordering is only a partial one, there are very few linguistic examples where a reciprocal can for instance be interpreted as both OWR and IAR, but not as




Figure 5.2: IR or OWR?

IR. Likewise, there is no example where both IR and SAR hold, but not SR, simply because SAR never has been attested so far. For practical purposes, we have a total order $SR > SAR > IR > IAR > OWR > IAO$ on the logical contributions of the reciprocal semantics.

It is easy to verify that most of the natural-language examples given earlier obey the rule that the correct interpretation of the scope relation in a given context is the logically strongest one possible. A weakening of the correct scope relation leads to interpretations that are not acceptable to a hearer, while a scope relation that comes from an overly strong reciprocal semantic even leads to inconsistencies with the contextual knowledge.

A Counter-Example

The linguistic theory sometimes predicts a reciprocal semantic that is too strong. The graphs in Figure 5.2 depict discourse situations for example (2.34) with four snipers. The two graphs show two ways in which they could train their rifles at each other.

(2.34) *The snipers train their rifles at each other.*

The example is problematic for the linguistic theory since the lefthand scope relation is an instance of IR. In the figure, each member of the reciprocal group is connected with each other one by a sequence of other group members. If we would accept IR as the reciprocal semantics, then we would have to discard the right-hand relation as a valid interpretation of the scope relation. Example (2.34) is an instance of OWR that is mistaken as IR reciprocity by the linguistic theory.

The SMH does not Compute (yet)

There is, to our knowledge, no computational method that implements an interpretation of reciprocal expressions and that makes use of the theoretical insights from the linguistic theories. This is not really surprising, as reciprocals seem to be one of the phenomena of natural language whose computational treatment requires an excessive amount of technical machinery. The most basic problem in this context is that the SMH can only select from a set of reciprocal semantics if it is already known which ones are consistent with the context and which ones are not. Apart from the fact that the problem is undecidable in general, even common reciprocal expressions are very problematic from a computational point of view. For instance, the semantics IR and IAR refer to a computation of transitive closure relations that can be a challenging task once the universe of discourse becomes larger than a few elements.

Fortunately, the interpretation of reciprocals can be seen from a different point of view where the SMH is understood as a constraint that strengthens the scope relation itself rather than one that chooses the right semantic representation for the reciprocal. As we will see, this alternative use of the SMH leads to an acceptable computational method for interpreting reciprocals by model generation.

5.3 Inference to Best Reciprocal Meaning

The goal of this section is to design a model generation method that can interpret reciprocal expressions. Our approach is based on some insights from the linguistic theory presented in the last section, namely that the reciprocal meaning ranges from a very weak form of reciprocity (IAO) to the strongest one (SR), and that the logical contribution of the reciprocal is determined by the Strongest Meaning Hypothesis. In contrast to the theory of Dalrymple et al., we understand the SMH as a principle of interpretation rather than a principle of semantics construction. The following alternative form of the SMH will be the basis of our computational treatment of reciprocals.

Maximise Meaning Hypothesis (MMH): *The valid interpretations of a reciprocal expression S in a context Γ are those which (a) are consistent both with the weakest form of reciprocal seman-*

tics and the context, and (b) whose logical contributions to the scope relation are the strongest.

The MMH refers only to the weakest semantic representation for reciprocals, namely IAO. The “strongest contribution” of reciprocal interpretations comes from a maximisation of the content of the scope relation over the antecedent set. We argue that this is a precise characterisation of the behavior of reciprocals. Consider again the following examples.

(3.1) *The men like each other. (SR)*

(3.2) *The snipers train their rifles at each other. (OWR)*

The two example sentences are compatible with the weakest form of reciprocity, IAO, which only requires that the antecedent group P that participates in the reciprocal expression has at least two members and that the reciprocal relation R relates each member of the group P with at least one other member. This is, however, not sufficient for a correct interpretation.

In both cases, we have a strengthening of the reciprocal’s semantic that is precisely predicted by the MMH. In the SR example, this strengthening includes all available pairs $\langle x, y \rangle$ to the scope relation where x and y are two different men. The reason for this strengthening is that the *liking* relation does not have truth conditions that further constrain how we interpret it on a set of men. In contrast to this, the truth conditions for *training a rifle* in the OWR example entail that no sniper x can be the subject of the scope relation for more than one other sniper y . Hence, the strengthening here only makes sure that that each x in P must be the subject of the scope relation R for exactly one other y .

As we can see in these first examples, our understanding of a reciprocal is that of a constituent whose semantic representation is unambiguous, but whose actual range of valid interpretations is determined by some additional constraint on the scope relation and the context. This implies that the truth conditions of the scope relation determine the logical contribution of the reciprocal rather than the other way around.

5.3.1 To Strong Meaning through Minimality

Based on a suitable semantic representation, we will later give the MMH an implementation as a minimal model constraint. This might seem paradoxical

at first because the MMH implies a *maximisation* of logical contribution, i.e., assumptions in a model, rather than a minimisation. However, as we will see, both tasks are actually equivalent.

Consider for instance three men a , b , and c in some context Γ . One of the locally minimal models for a semantic representation (3.3) of example (3.1) with respect to Γ is given by (3.4).

$$(3.3) \quad \mathit{man}(a) \wedge \mathit{man}(b) \wedge \mathit{man}(c) \wedge \mathit{RCP}_{IAO}(\mathit{man})(\mathit{like})$$

$$(3.4) \quad \{\mathit{man}(a), \mathit{man}(b), \mathit{man}(c), \mathit{like}(a, b), \mathit{like}(b, c)\}$$

The MMH implies that the weak IAO form of reciprocity is the only semantic representation of reciprocals. The model (3.4) satisfies the logical encoding of reciprocity that we have chosen, but does not represent a meaning that is predicted by the linguistic theory. However, we can turn model (3.4) into a model that corresponds to a valid natural-language interpretation of the reciprocal sentence by adding new consistent assumptions. A completion of (3.4) according to the MMH is given by (3.5).

$$(3.5) \quad \{\mathit{man}(a), \mathit{man}(b), \mathit{man}(c), \mathit{like}(a, b), \mathit{like}(b, c), \mathit{like}(a, c), \\ \mathit{like}(b, a), \mathit{like}(c, a)\}$$

By adding the new atoms to the set of atoms that is validated by the model, we have actually minimised those atoms that are *not* validated in (3.4). In other words, we have minimised the complement set of the predicate *like*. What we need is a form of minimality that can minimise the denotations of certain predicates or their complements in the set of generated models.

Predicate Minimisation

Predicate-specific minimality, as presented shortly in Section 2.3.3, minimises the occurrences of certain "costly" predicates. In applications of predicate-specific minimality, a model is generally considered to be an explanation for some kind of observation, and the predicates whose occurrences are to be minimised are those that are relevant for distinguishing good explanations from those that are not plausible or may not be attractive for economic reasons.

For instance, an assumption of the form $\mathit{ab}(d)$ in a diagnosis application could be used to encode that the behaviour of the part d of a circuit is

faulty. The model of a suitable logical encoding of the circuit and an observed error can then be taken as an explanation how the error is caused. The minimisation of assumptions $ab(d)$ in a model minimises the number of faulty parts that are needed to explain the error in the circuit. By considering only models that are ab -minimal, one implements the general principle of diagnosis that the cause of an error lies more probably in the failure of one part of a device than in a simultaneous failure of several parts.

Likewise, predicate-specific minimality can be used to optimise the movement of a robot in a planning application where the movement of a robot from a point x to a point y at a time t corresponds to an assumption $go(x, y, t)$. Each model of the plan task encodes how the robot must move in order to solve a certain task, and by minimising the go predicate, one could effectively optimise the (costly) movements of the robot.

In the case of reciprocal expressions, the observation that we want to explain is the truth of the reciprocal sentence, and the assumptions that we want to minimise correspond to those pairs x and y in the reciprocal group that are not in the scope relation. An explanation is best in the sense of the MMH if it explains the truth of a reciprocal sentence while assuming a maximum of pairs $\langle x, y \rangle$ in the scope relation. As indicated earlier, what we minimise here is actually the complement of some binary relation, but by using a suitable logical encoding, we can reduce this problem completely to predicate-specific minimal model generation.

A Logical Encoding of Less is More

Let ab be an arbitrary but fixed predicate symbol of type $\iota \rightarrow \iota \rightarrow o$. For each MQC model \mathcal{M} of a specification ϕ , the ab -index, in short index, is the number of atoms validated by \mathcal{M} whose head is ab . The index of a model refers to some number of “costly” assumptions that the model implies. The minimisation of the index predicate ab can implement the MMH if we can give a suitable semantic representation that makes explicit the connection between costly assumptions and pairs $\langle x, y \rangle$ that are not in the reciprocal relation,

As usual, we make use of our standard MQC logic presented in Section 3.3.3. Our higher-order definition of reciprocal semantics has two parts. The first part (3.6) defines the basic constraint on reciprocals that we know from the linguistic theory, namely that each reciprocal must at least meet the truth conditions of IAO reciprocity. We formalise IAO here exactly as

we will use it later for model generation.

$$(3.6) \quad \text{IAO} \equiv \lambda P \lambda R \text{CARD}^{\geq 2}(P) \wedge \text{EVERY}(P)(\lambda x \exists y \neg(x \dot{=} y) \wedge P(y) \wedge (R(x, y) \vee R(y, x)))$$

$$(3.7) \quad \text{CARD}^{\geq 2} \equiv \lambda P \exists x \exists y P(x) \wedge P(y) \wedge \neg(x \dot{=} y)$$

The definition is equivalent to the previously given definition of IAO reciprocity in higher-order logic, except that we use Herbrand equality instead of Leibniz equality. In \mathcal{MQC} semantics, we consider two individual constants to be different if they have different names, and Herbrand equality is a computationally very simple form of equality that implements this.

The second part (3.8) of reciprocal semantics formalises the connection between the index of a model and those pairs of group members that are not in the scope relation. Note that we exclude assumptions $\mathbf{ab}(x, x)$; group members that are in the scope relation with themselves do not play any rôle for the index of the model. The higher-order definition (3.9) combines the two parts of the reciprocal's representation.

$$(3.8) \quad \text{PRICE} \equiv \lambda P \lambda R \text{EVERY}(P)(\lambda x \neg \mathbf{ab}(x, x) \wedge \forall y (P(y) \wedge \neg(x \dot{=} y) \wedge \neg R(x, y)) \Leftrightarrow \mathbf{ab}(x, y))$$

$$(3.9) \quad \text{RCP} \equiv \lambda P \lambda R \text{IAO}(P)(R) \wedge \text{PRICE}(P)(R)$$

5.3.2 A First Attempt at Computation

Definition (3.9) allows us to represent reciprocal sentences according to our theory. Each model of the reciprocal sentence and the context will give us an essential information, the \mathbf{ab} -index. The index of a model gives us a means to compare models with respect to a relative satisfaction of the MMH. A model \mathcal{M} that has a smaller index than a model \mathcal{M}' is a better explanation for the truth of the reciprocal sentence according to the MMH. All finite model generation methods can be modified such that their output enumerates models with decreasing index, simply by filtering out all models that have predecessors with a smaller index.

In the case of a domain closure, i.e., whenever we focus only on models that are domain minimal or otherwise restricted in the size of the universe,

we can even decide which models have the smallest index. In our model generator KIMBA, three methods for the computation of such predicate-specific minimal models are particularly easy to implement.

First Method: Minimality by Proof

The first method is to prove for each generated model that there is no other model that has a smaller index. This method works similar to the computation of locally minimal models as presented in Section 3.5, but requires even less technical machinery. Suppose that ϕ is the logical specification, and \mathcal{M} is a generated model with fixed index n . If n is greater than 0 for \mathcal{M} , we simply start KIMBA again, but with a constraint on the interpretation of all *ab*-atoms that their sum must always be less than n . This constraint is actually a finite-domain integer constraint, which we can simply add to the other constraints that KIMBA generates while translating the logical input into a system of constraints. If KIMBA manages to generate a model \mathcal{M}' with an index $m < n$, then the previously computed \mathcal{M} is not *ab*-minimal, and we can discard it.

Second Method: Minimality by Bounded Search

The second method is to use branch-and-bound search. In this method, we use the index of each generated model as an upper bound for the index of all models that are generated later. In KIMBA, we can compute from each partial interpretation \mathcal{I} a preliminary index that defines a lower bound of the index of those models which are derived from \mathcal{I} . If this lower bound becomes greater than the upper bound that we have obtained from the previously computed models, then we know that a further exploration of the search space will not give us better explanations. By discarding those parts of the search space where we cannot expect to find better explanations, KIMBA efficiently enumerates models with decreasing index. If the index of a model is minimal, then all models that can be generated from that point on will also be *ab*-minimal.

Third Method: A Two-Stage Combination

The third method is a two-stage variation of the second one. In the first stage of the computation, one uses a branch-and-bound search where each computed model carries an index that is truly smaller than the upper bound

given by the previously computed models. That is, for each index n , we will compute only one representative model, and all later models must have an index $m < n$. This branch-and-bound constraint restricts the search space more quickly in general than that of the second method, and we will quickly compute a model \mathcal{M} whose index is some minimal index u . Then, in a second stage, we restart the model search again with the second method, but this time we restrict the search from the start to models whose index is u . This restriction will produce only *ab*-minimal models.

While all three methods for computing *ab*-minimal models are equally simple to implement, the second method is more efficient in many cases than the first because the search for the abductive explanations is implemented as a simple constraint that can effectively decrease the search space. The third then is more efficient than the second because it enumerates only the *ab*-minimal models in a much smaller search space that is already constrained by the most minimal index.

An Example

In order to illustrate the effect that the implementation of the MMH has, let us reconsider the example (3.10).

(3.10) *The snipers train their rifles at each other.*

As we know already, this is an example of the OWR form of reciprocity. We now interpret (3.10) in a simple discourse situation where we have a group of three snipers. The formalisation is as follows.

(3.11) $\text{sniper}(\text{peter}) \wedge \text{sniper}(\text{paul}) \wedge \text{sniper}(\text{mary})$

(3.12) $\text{NO}(\text{sniper})(\lambda x \text{CARD}^{\geq 2}(\lambda y \text{train}(x, y)))$

(3.13) $\text{RCP}(\text{sniper})(\text{train})$

Formula (3.11) formalises that there are three snipers, namely Peter, Paul, and Mary. Formula (3.12) expresses the world knowledge that no sniper trains his or her rifle at more than one other individual. Finally, formula (3.13) represents the reciprocal sentence.

We first apply our model generator KIMBA to the formalisation without using any model minimality constraints. A typical model generated by KIMBA is the following.

$$(3.14) \quad \{ab(mary, paul), ab(paul, peter), ab(peter, mary), \\ ab(peter, paul), sniper(mary), sniper(paul), \\ sniper(peter), train(mary, peter), train(paul, mary)\}$$

Figure 5.3 shows the search space of the example. Each diamond represents a model, while each circle is a branch point in the search. The search space has 25 such branch points, and leads to 26 different models. These solutions correspond to the different interpretations that one obtains by the IAO form of reciprocity.

In Figure 5.4, we have the same example, but this time we have used the branch-and-bound search of KIMBA for eliminating parts of the search space where we know that no *ab*-minimal models can be expected. There are now fifteen branch points in the search space, and sixteen models. Of these, the last eight are *ab*-minimal, and each minimal model contains exactly three *ab*-atoms. The eight minimal models correspond to the eight possible ways in which three snipers can train their weapons at each other.

Finally, in Figure 5.5, we have used the two-stage minimal model computation where those parts of the search space are discarded that lead to models whose index is higher than three. There are now only those models left that are *ab*-minimal. The red squares indicate a failure in the search where a locally computed index becomes too high. Below are the eight different scope relations that are found in the models.

$$(3.15) \quad \{train(mary, peter), train(paul, peter), train(peter, mary)\}$$

$$(3.16) \quad \{train(mary, peter), train(paul, peter), train(peter, paul)\}$$

$$(3.17) \quad \{train(mary, peter), train(paul, mary), train(peter, mary)\}$$

$$(3.18) \quad \{train(mary, peter), train(paul, mary), train(peter, paul)\}$$

$$(3.19) \quad \{train(mary, paul), train(paul, peter), train(peter, mary)\}$$

$$(3.20) \quad \{train(mary, paul), train(paul, peter), train(peter, paul)\}$$

$$(3.21) \quad \{train(mary, paul), train(paul, mary), train(peter, mary)\}$$

$$(3.22) \quad \{train(mary, paul), train(paul, mary), train(peter, paul)\}$$




Figure 5.3: The snipers example without any minimality constraint.




Figure 5.4: The snipers example with branch-and-bound search.

We have mentioned in Section 5.2.3 that the logical contribution of the reciprocal in our example sentence is not identified correctly by the linguistic theory. The scope relation for a larger group that is implied by the classification system is IR, because some of the interpretations meet the truth conditions of IR reciprocity. As we can see above, this does not pose a problem for our approach, as it distinguishes interpretations only by the number of assumptions they provide to the scope relation. Hence, the relation (3.19), an instance of IR reciprocity, occurs as well as a relation like (3.22), that is an example of OWR, in the set of our valid interpretations.

5.3.3 Conservative Minimality

The minimisation of a certain predicate *ab* so far seems to be an adequate form of minimality that identifies those models whose content each describes a valid logical contribution of the reciprocal. However, as discussed in some detail in Chapter 4, we have also argued for the use of *local minimality* for obtaining a correspondence of logical models and some valid natural-language




Figure 5.5: The snipers example with MMH minimality.

interpretations. This form of minimality must be considered as well in the interpretation of reciprocal expressions, for instance in the following sentence.

(3.23) *Peter, Paul, and Mary like each other.*

The *liking* relation implies SR reciprocity, and in this form of reciprocity, all members of the reciprocal group are in the reciprocal relation. SR reciprocals therefore have only one possible scope relation once the reciprocal group is determined. KIMBA's method for computing the *ab*-minimal models yields model (3.25) as one of the interpretations of the logical encoding (3.24), as required. Unfortunately, we also have the models (3.26)–(3.31).

(3.24) $\text{RCP}(\{\textit{peter}, \textit{paul}, \textit{mary}\})(\textit{like})$

(3.25) $\{\textit{like}(\textit{mary}, \textit{paul}), \textit{like}(\textit{mary}, \textit{peter}), \textit{like}(\textit{paul}, \textit{mary}),$
 $\textit{like}(\textit{paul}, \textit{peter}), \textit{like}(\textit{peter}, \textit{mary}), \textit{like}(\textit{peter}, \textit{paul})\}$

(3.26) $\{\textit{like}(\textit{mary}, \textit{paul}), \textit{like}(\textit{mary}, \textit{peter}), \textit{like}(\textit{paul}, \textit{mary}),$
 $\textit{like}(\textit{paul}, \textit{peter}), \textit{like}(\textit{peter}, \textit{mary}), \textit{like}(\textit{peter}, \textit{paul}),$
 $\textit{like}(\textit{mary}, \textit{mary})\}$

(3.27) $\{\textit{like}(\textit{mary}, \textit{paul}), \textit{like}(\textit{mary}, \textit{peter}), \textit{like}(\textit{paul}, \textit{mary}),$
 $\textit{like}(\textit{paul}, \textit{peter}), \textit{like}(\textit{peter}, \textit{mary}), \textit{like}(\textit{peter}, \textit{paul}),$
 $\textit{like}(\textit{paul}, \textit{paul})\}$

- (3.28) $\{ \textit{like}(\textit{mary}, \textit{paul}), \textit{like}(\textit{mary}, \textit{peter}), \textit{like}(\textit{paul}, \textit{mary}), \textit{like}(\textit{paul}, \textit{peter}), \textit{like}(\textit{peter}, \textit{mary}), \textit{like}(\textit{peter}, \textit{paul}), \textit{like}(\textit{peter}, \textit{peter}) \}$
- (3.29) $\{ \textit{like}(\textit{mary}, \textit{paul}), \textit{like}(\textit{mary}, \textit{peter}), \textit{like}(\textit{paul}, \textit{mary}), \textit{like}(\textit{paul}, \textit{peter}), \textit{like}(\textit{peter}, \textit{mary}), \textit{like}(\textit{peter}, \textit{paul}), \textit{like}(\textit{mary}, \textit{mary}), \textit{like}(\textit{paul}, \textit{paul}) \}$
- (3.30) $\{ \textit{like}(\textit{mary}, \textit{paul}), \textit{like}(\textit{mary}, \textit{peter}), \textit{like}(\textit{paul}, \textit{mary}), \textit{like}(\textit{paul}, \textit{peter}), \textit{like}(\textit{peter}, \textit{mary}), \textit{like}(\textit{peter}, \textit{paul}), \textit{like}(\textit{paul}, \textit{paul}), \textit{like}(\textit{mary}, \textit{mary}) \}$
- (3.31) $\{ \textit{like}(\textit{mary}, \textit{paul}), \textit{like}(\textit{mary}, \textit{peter}), \textit{like}(\textit{paul}, \textit{mary}), \textit{like}(\textit{paul}, \textit{peter}), \textit{like}(\textit{peter}, \textit{mary}), \textit{like}(\textit{peter}, \textit{paul}), \textit{like}(\textit{peter}, \textit{peter}), \textit{like}(\textit{paul}, \textit{paul}) \}$
- (3.32) $\{ \textit{like}(\textit{mary}, \textit{paul}), \textit{like}(\textit{mary}, \textit{peter}), \textit{like}(\textit{paul}, \textit{mary}), \textit{like}(\textit{paul}, \textit{peter}), \textit{like}(\textit{peter}, \textit{mary}), \textit{like}(\textit{peter}, \textit{paul}), \textit{like}(\textit{peter}, \textit{peter}), \textit{like}(\textit{paul}, \textit{paul}), \textit{like}(\textit{mary}, \textit{mary}) \}$

The last seven models are *ab*-minimal, but they contain information that is not implied by the natural-language sentence. This is because the model generator is free to assume a fact like *like(mary, mary)* unless it is not inconsistent. The unwanted assumptions are easy to identify because they do not occur in any locally minimal model of the input.

On the other hand, the model (3.25) that we want to obtain as an interpretation of the natural-language input is itself not a locally minimal model of our semantic representation. It contains assumptions that are unnecessary for validating the truth of the input with respect to local minimality. Our theory of reciprocal meaning explicitly requires the presence of such non-minimal assumptions.

In order to identify models for reciprocal sentences that correspond to natural-language interpretations, we arguably must identify those *ab*-minimal models that are locally minimal *relative* to all other *ab*-minimal models. In our example, the first model is locally minimal relative to all other predicate-minimal models, and hence is the only one that should be accepted. What we have in mind is a new form of minimality that we define formally as follows.

Conservative Minimality: *Let Φ be a set of satisfiable MQC formulas. Then there exists a nonempty set of models \mathbf{D} whose*

*domain \mathcal{D}_i of first-order individuals has a minimal size. Let \mathbf{C} be the set of those models in \mathbf{D} that are **ab**-minimal. A model $\mathcal{M} \in \mathbf{C}$ is conservative minimal iff $\text{Pos}(\mathcal{M}') \subseteq \text{Pos}(\mathcal{M})$ implies $\mathcal{M}' = \mathcal{M}$ for all models $\mathcal{M}' \in \mathbf{C}$.*

Conservative minimality is the desired combination of local minimality and **ab**-minimisation. We have chosen the term “conservative” because this form of minimality is a conservative extension of local minimality. If \mathbf{D} does not contain models with **ab**-atoms, then all conservative models are also locally minimal models. The new form of minimality is compatible for instance with our earlier approaches for computing the interpretations of singular definite descriptions.

But how exactly can we compute such models? This can be done by a variation of the standard search for locally minimal models. The models that we want to obtain are domain minimal models that (a) are predicate-specific minimal and (b) that are subset-minimal with respect to all other predicate-specific minimal models. The constraint (a) is met by identifying the predicate-specific models \mathcal{M} of the input that have some smallest possible domain \mathcal{D}_i . For (b), we use the same technique as for locally minimal models and prove for each model \mathcal{M} that there is no other model \mathcal{M}' that satisfies the input with a true subset of the assumptions in \mathcal{M} **and** that is also predicate-specific minimal. Hence, a predicate-specific minimal model is compared only to other predicate-specific minimal models. A modification of KIMBA’s minimal model search is straightforward and has been carried out for the following examples.

5.4 Experiments

In what follows, we will investigate how our approach computes the logical contributions of the reciprocal in the context for some selected examples from the literature. As we will see, a correct analysis of reciprocals sometimes requires a surprisingly complex specification of the properties of the scope relation.

5.4.1 Pitchers and Pearls

IR hold for scope relations R where each group member x is related to each other group member y by a sequence of members c_1, \dots, c_n such that $x = c_1$,




Figure 5.6: IR relations

$y = c_n$, and $R(c_i, c_{i+1})$ for all $i < n$. Some linguistic examples for IR are discourses such as (4.1) or (4.2) where the scope relation is equivalent to a linear or circular ordering of individuals.

(4.1) *The Boston pitchers sit alongside each other.*

(4.2) *The pearls in the necklace are separated from each other by semi-precious jewels.*

The truth conditions that IR formulates on its own are quite different from those that can hold in discourse (4.1) and (4.2). In Figure 5.6, we have four graphical representations of IR scope relations, each with four discourse participants. Only the first, leftmost one would be acceptable as a relation for (4.1) in a discourse situation with four Boston pitchers. The situation implied by sentence (4.1) requires a relation that is symmetric. The relations 2 and 3 do not meet this requirement. Relation 4 depicts a situation where one group member is related to three other ones, which conflicts with the world knowledge that a person cannot sit alongside more than two persons.

The Boston Pitchers

A correct interpretation of example (4.1) requires a suitable analysis of the *sitting alongside* relation. The properties of this relation can be captured by relating it to a *position*, i.e., a total ordering of individuals. In our case, this position could denote the position of the pitcher on the bench on which they sit. A formalisation of a *sitting alongside* relation *sital* must then express the following properties.

- For each two individuals x and y , if *sital*(x, y) holds, then x is either in a predecessor or a successor position of y .
- At each possible position we have at most one individual x .
- An individual cannot be at two different positions.
- For each position i , if some x is in a position i for a given relation, then there are exactly i individuals whose position p is smaller or equal to i .

The complete specification must also define concepts such as a successor relation and some total order on positions, but otherwise is straightforward. Section A.2 gives a logical encoding that can be taken as the input for our model generator KIMBA, and evaluates the results for a discourse situation with four Boston pitchers.

Pearls

In example (4.3), the logical specification of the properties of the scope relation is much easier than in the “Pitcher” example.

- (4.3) *The pearls in the necklace are separated from each other by semi-precious jewels.*

If we assume that we have a simple, round necklace, then it suffices to specify that no pearl can be separated by semi-precious levels from more than two other pearls, and that the relation itself is symmetric. There also are far fewer models than in the case of the pitchers—we have three models in a “necklace” with four pearls. All of them are instances of the relation depicted in Figure 5.7.

What is interesting here is that the interpretation of the relation provided by the verb phrase is heavily influenced by properties implied by the reciprocal group. The “pearls in the necklace” arguably imply a circular relation that is not part of the truth conditions of the scope relation on its own.

5.4.2 Measles

The linguistic theory describes IAR scope relations as relations where each group member is related directly or indirectly via the reciprocal relation without considering the direction of the relation.




Figure 5.7: The IR relation for pearls in a necklace

(4.4) *The students gave each other measles.*

The formalisation of IAR in Higher-Order Logics relies on a specification that uses the transitive closure relation of the scope relation. In our approach, it suffices to specify the characteristic properties of the scope relation for a discourse situation. For this, we again use a position relation that describes the distance of some person that has measles to the one who was a source of the measles infection in the discourse situation. The properties of the *giving measles* relation are then as follows.

- No one can be given measles by more than one other person.
- If someone has measles, then there must be at least one person that is the source of the measles, i.e., whose position is 1.
- If some person x gives measles to a second person y , then x 's position is the successor of the position of y .
- If some person x gives measles to a second person y , then both persons have measles.

Our specification of what *giving measles* means is quite different from what IAR formulates as the truth conditions of the scope relation. Indeed, it is easy to show that IAR insufficiently describes the reciprocal relation in (4.4). Figure 5.8 shows a selection of IAR relations over a group of four elements that all would not be acceptable as instances of the *giving measles* relation while those relations in Figure 5.9 are.

The model generator KIMBA generates 64 different conservative minimal models for the “Measle” example with four discourse participants. These




Figure 5.8: IAR relations that are not compatible with measles.




Figure 5.9: IAR relations that are compatible with measles.

models all are instances of the relations depicted in Figure 5.9. While relations 1 and 2 each have 24 instances, i.e., permutations, the relation 3 has twelve, and relation 4 only four instances.

5.4.3 Marriages

Example (4.4) exemplifies the OWR form of reciprocity.

(4.5) *The people in the room all are married to each other.*

A formalisation of the relation expressed by the verb phrase includes the following facts.

- If some person x is married to some person y , then y is also married to x .
- No one can be married to more than one other person.




Figure 5.10: The OWR relation for being married with each other.

In a discourse situation with four discourse participants, we have three models that all are an instance of the relation depicted in Figure 5.10.

Interestingly, if we add a fifth person to the discourse, then the model generation problem becomes unsatisfiable for this universe of discourse. The model generator solves this problem by accommodating a new individual who becomes the spouse of the “odd” person.

5.5 Loose Ends

There are some cases where the strong meaning predicted by the SMH as well as the MMH seems to be too strong. Consider the following example from Philip [76] that has also been discussed in detail in Winter [89].

(5.1) *The three boys tickle each other.*

Although it can happen that a boy tickles two objects at the same time, the natural interpretation of (5.1) is one where each boy tickles only one other boy. As Winter notes, this potential counter-example of meaning maximisation might come from some gap in our world-knowledge about the predicate *to tickle*. There seems to be a default assumption that one tickles only one other object in a situation. However, this “uniqueness” assumption can be overridden, i.e., is not a hard constraint on the interpretation of *to tickle*.

Winter also mentions a counter-example for the SMH where the truth conditions are actually too weak. This example is as follows.

(5.2) *Mary and Sue gave birth to each other. (*)*

In (5.2), both SMH and MMH predicts semantics/interpretations where either Mary gave birth to Sue or Sue gave birth to Mary. This is not the intended meaning of the clearly inconsistent statement. There seems to be a lower bound for the weakening of logical contribution at least in the case of the “giving birth” relation. It can be argued that sentence (5.2) implies a discourse anomaly as one of Grice’s maxims on conversation states that a contribution should not give redundant information. The weakening of the reciprocal’s contribution leads to such a redundancy, as the contribution of the reciprocal becomes nil.

5.6 How We Can Understand Each Other

The approach of Dalrymple et al. [27] to reciprocal meaning exemplifies a formal, theoretic solution that tries to classify ambiguity in meaning by some collection of semantic representations and a non-computational method that models how the correct reading is selected. In theory, the method is not computational because the criteria that must be verified for making a selection are undecidable.

In contrast to this, the approach by model generation strengthens the logical contribution of the weak semantic representation as a process based on *computational* constraints. Once any model can be computed at all for our logical form, the process of interpretation is an inference problem that is guaranteed to terminate. The algorithm may fail for complexity reasons, but in principle it should be able to give at least one linguistically acceptable interpretation if the reciprocal’s logical form is satisfiable with respect to its context and a restriction to finite domains.

Empirically, the predictions are the same in most cases because maximising the scope relation often results in yielding a logically stronger meaning on the entailment scale. In particular, the present approach captures the meanings postulated by Dalrymple et al. But the examples also show that there are cases where the predictions differ. Intuition dictates that the sentence *the snipers train their rifles at each other* has a natural OWR interpretation, but the classification scheme incorrectly gives us IR as the strongest meaning. Model generation, however, not only gives us the correct interpretations, it is also linguistically more appealing as it needs only one semantic representation. The model generator is the only computational tool that we need for determining preferable readings from this single representation.

One crucial feature of minimal model generation is that it permits comparing and ranking of natural-language interpretations against each other. In the case of reciprocals, the ranking is given mainly by the size of the scope relation. We have already found other useful ranking criteria as well. For instance, Chapter 4 shows that in the case of definite descriptions, the ranking defined by local minimality permits capturing the preference of binding over bridging over accommodation. Similarly the work of Baumgartner and Kühn [9], which we will discuss in Section 6.2, shows that a predicate minimisation together with a preference for logically consequent resolutions can be used to model the preferences in the interpretation of pronominal anaphora.

This suggests that one of the most promising application of model generators is as a device for developing and testing *preference systems* for the interpretation of natural language. Inference and knowledge based reasoning are needed in natural-language processing not only to check for consistency and informativity as illustrated for instance in Blackburn et al. [15], but also to express preferences between possible interpretations. For this, model generation seems to be a natural tool to use.

Chapter 6

Abduction

Overview: Abduction is inference to the best explanation. It is widely recognised as a form of inference that can conceptualise the problem of interpreting discourses. Model generation as we have used it is a form of abduction. We discuss the relation between our approach to abduction to others that have been applied to linguistics.

6.1 What is Abduction?

The term “**abduction**” was first used by Pierce [77] who defined it as follows.

The surprising fact, C , is observed; but if A were true, C would be a matter of course. Hence, there is reason to suspect that A is true.

In other words, abduction originally refers to an unsound inference rule that concludes A from the facts C and $A \Rightarrow C$. Abduction in this sense is inference to some *arbitrary* explanation A for C . Without some further restrictions, abductive inference would have few applications, as there usually are too many explanations that can be obtained from some rules which have C as their consequence. The actual method of inference is in general only secondary to what could be called the **overgeneration problem of abduction**. The usefulness of an explanation depends on additional application-dependent constraints of suitable explanations. The term abduction, as it is commonly used today, refers to all forms of inference that yield “best”, i.e. suitable, explanations for observed facts in some way.

Since Hobbs et al. presented their inspiring work, abduction is widely recognised as a form of inference that can conceptualise the problem of interpreting discourses. In this context, the best explanations sought for are sets of assumptions that explain the truth of a speaker's utterances in a way that is compatible with contextual information and the common ground of knowledge shared by speaker and hearer.

Model generation becomes a form of abduction whenever some models with desirable properties are identified as best solutions for a problem. Thus, in previous chapters, we have used model generation as a form of abduction where the best explanations are minimal models. For the interpretation of singular definite descriptions, we have argued for local minimality as that desirable property that characterises best explanations. In the case of reciprocal expressions, we have proposed conservative minimality which combines the minimal model constraints of logical minimality, domain minimality, and predicate-specific minimality.

The principal goal of this chapter is to discuss how minimal model generation fits exactly into the larger picture of abduction in natural-language processing, and how our approach compares to other work that uses inference in natural-language semantics.

A Formal Definition of Abduction

Let ϕ be a logical formula called **observation**, Γ be a set of formulas called **context** or **background theory**, and Ψ be a set of formulas called **allowable hypothesis**. An **abductive explanation** is a set Δ of assumptions such that the following holds [9].

- $\Delta \subseteq \Psi$
- $\Gamma \cup \Delta \models \phi$
- $\Gamma \cup \Delta$ is satisfiable

The term **abduction** refers to all forms of inference that generate best abductive explanations from specifications Γ , ϕ , and Ψ . The definition of an abductive explanation that has been given above is a general one that fits different forms of abduction without exactly capturing the characteristic properties that distinguish *best* abductive explanations from those explanations that are not considered suitable for the intended application.

For instance, in diagnosis applications, the allowable hypothesis Ψ is a set of assumptions that describes the cause a failure in some complex device described by Γ , and ϕ is a specification of an observable faulty behavior. An abductive explanation Δ is then a set that explains why Γ shows that behavior. A characteristic property of abductive explanations in diagnosis applications is a minimisation of Δ such that the failure of the device is explained by a minimal number of faulty device parts.

In our approach of semantic interpretation as minimal model generation, the formula ϕ consists of some semantic representation for a sentence, and Γ is the logical specification of a discourse model and the necessary world knowledge. The set of allowable hypothesis Ψ is the set of all atoms that can be built from a suitable signature. The set Δ of assumptions is a set of atoms that defines a model of $\Gamma \cup \{\phi\}$. We have characterised best explanations as minimal models whose forms of minimality have been found suitable as constraints of valid natural-language interpretations.

There is some other work that has investigated how different forms of abduction can be used for conceptualising the interpretation of discourses. Baumgartner and Kühn [9] presents an abductive solution to anaphora resolution in a model generation framework. Their approach uses predicate-specific minimality. In Hobbs et al. [50], we find **weighted abduction**, a variation of logic programming. A best explanation is an explanation where the set Δ of assumptions is minimal with respect to a numerical cost obtained from Δ itself and from the computation process that leads to Δ . Weighted abduction has been used for the interpretation of anaphora, compound nominals, syntactic ambiguities and metonymies.

In what follows, we will compare the results of this work with what can be achieved by the form of minimal model generation that we have developed.

6.2 Models for Anaphora Resolution

Baumgartner and Kühn [9] proposes model generation as a method for computing the resolution of anaphoric expressions. More specifically, they address the problem of establishing the links between pronominal and definite anaphora to their referents.

The model generation calculus that is used to attack the problem is based on hyper-tableaux as presented in Section 2.4.2. While a considerable part of the authors' efforts go into the design of a hyper-tableaux method that

incrementally computes best abductive explanations, the logical encoding and the abductive inference method is more general and can be used as well for other model generation approaches. In the following, we exemplify abductive anaphoric resolution as proposed by Baumgartner and Kühn.

6.2.1 Chasing The Criminal

Example (2.1) outlines the general problem of abductive anaphoric resolution. We have a definite description *the criminal* that has two possible referents, namely a politician and a gangster. Both referents are introduced by the antecedent sentence. It is not inconsistent per se to assume that a politician is a criminal, but we know that each gangster is a criminal. Hence, the reading where *the criminal* refers to the gangster in the context is preferred. The resolution of the anaphoric definite obviously is an abductive inference task: we have two linguistically consistent¹ interpretations, of which only one is the best, i.e., preferred one.

(2.1) *A politician chased a gangster. **The criminal died.***

For example (2.1), we use a logical encoding that formalises the information given in the first sentence by the first-order formula (2.2), the world knowledge by formula (2.3), and the sentence with the definite by formula (2.4).

(2.2) $pol(p) \wedge gan(g) \wedge chase(p, g)$

(2.3) $\forall x gan(x) \Rightarrow crim(x)$

(2.4) $\exists x crim(x) \wedge anaph_1(x) \wedge die(x)$

The predicate *anaph*₁ is used to indicate in the model which discourse participant has been used for resolving the anaphor expressed by the definite description.

The first-order Herbrand models of the formulas (2.2)–(2.4) are given by the tableaux in Figure 6.1. There are two branches, each containing the atoms that define an interpretation. While both branches are consistent and

¹Note that is a consequence of the weak logical encoding. The representation (2.4) does not formalise unicity which we have identified as one of the essential truth conditions of many definites in Chapter 4.

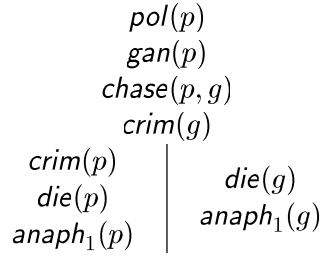


Figure 6.1: Model construction for discourse (2.1).

therefore are models of the input, only the second one corresponds to the preferred reading.

Baumgartner and Kühn investigate how resolutions can be eliminated that do not correspond to preferred interpretations. The anaphor resolution performed in the first branch leads to the assumption that the politician is a criminal. This is not a logical consequence of the information given by the discourse. In the second model, we have that the gangster is a criminal, a fact which logically follows from the world knowledge. This difference can be used to characterise the preference for the second model as a reading of the discourse.

6.2.2 Explaining Resolutions

Baumgartner and Kühn argue that the semantic content of the anaphoric expression should be *implied* by the context, and not only be consistent. For our example, the semantic content of the definite *the criminal* is given by formula (2.5).

$$(2.5) \quad \exists x \text{ crim}(x) \wedge \text{anaph}_1(x)$$

If we take (2.5) as an observation that must be explained, then we can instantiate the abduction scheme for our example as follows.

$$(2.6) \quad \Gamma = \{pol(p), gan(p), chase(p, g), \forall x \text{ gan}(x) \Rightarrow \text{crim}(x)\}$$

$$(2.7) \quad \phi = \exists x \text{ anaph}_1(x) \wedge \text{crim}(x)$$

$$(2.8) \quad \Psi = \{\text{anaph}_1(p), \text{anaph}_1(g)\}$$

There are three candidate explanations.

$$(2.9) \quad \Delta_1 = \{\}$$

$$(2.10) \quad \Delta_2 = \{\mathit{anaph}_1(p)\}$$

$$(2.11) \quad \Delta_3 = \{\mathit{anaph}_1(g)\}$$

$$(2.12) \quad \Delta_4 = \{\mathit{anaph}_1(p), \mathit{anaph}_1(g)\}$$

The abductive explanation (2.9) can be eliminated because it does not satisfy the condition that $\Gamma \cup \Delta_1 \models \phi$. Explanation (2.12) is eliminated by linguistic knowledge. It does not give a unique resolution of the anaphoric definite. Explanation (2.10) is not acceptable as well because it does not satisfy $\Gamma \cup \Delta_2 \models \phi$. The observation ϕ becomes a logical consequence of $\Gamma \cup \Delta_2$ only if we add the additional assumption $\mathit{crim}(p)$ which, however, is not part of the background theory Γ . Only (2.11) remains as a valid abductive explanation for the truth of observation ϕ .

Baumgartner and Kühn give an extension of the standard hyper-tableaux calculus that maintains a minimal model property such that each open branch is a anaph_i -minimal model of the input. Additionally, it can be verified at any time whether the semantic information of an anaphoric expression is implied by the background information. Thus, it can be tested whether a model provided by the tableaux meets the conditions of a best abductive explanation. This ensure that each model computed corresponds to a valid natural-language interpretation of the input. As their analysis of anaphors is based on closed, finite domains, Baumgartner and Kühn intuitively define an interesting new form of minimality for propositional logics that combines predicate minimisation and logical conditions on the consequence relation of models and input theories.

6.2.3 Discussion

The ambiguity that natural language shows on various levels leads to a combinatorial explosion when computing the candidate interpretations of discourses. In this context, a tableaux approach for model generation in general has the advantage that it can be space efficient. A tableaux expansion may restrict the search space to only one model candidate at a time. If an interpretation given by a tableaux branch is discarded by adding new information, then

backtracking can be used to compute a new consistent reading in a different branch. Additionally, a tableaux can be built such that already computed knowledge does not have to be recomputed completely when new data is added. Thus, model generation by tableaux, as proposed by Baumgartner and Kühn, offers both a solution to the combinatorial explosion of alternative readings and *incrementality*, the ability to deal with new discourse information.

Incremental Inference instead of Generate-And-Test

Baumgartner and Kühn identify those resolutions as preferable that explain the semantic content of the anaphoric expression by logical consequence from the context. Their approach works well for instance for singular pronouns as in example (2.13).

(2.13) *Peter loves Mary. He₁ sends her₂ letters every day.*

Pronouns provide almost no semantical information of their own. The observations that must be explained in (2.13) are the simple formulas below.

(2.14) $\exists x \textit{ male}(x) \wedge \textit{ anaph}_1(x)$

(2.15) $\exists x \textit{ female}(x) \wedge \textit{ anaph}_2(x)$

The background information includes the information that Peter is male and Mary is female. It is then no problem to identify the resolution of the pronouns with the best and only explanation $\Delta = \{\textit{ anaph}_1(\textit{ peter}), \textit{ anaph}_2(\textit{ mary})\}$. The method of how explanations are computed preserves one of the main advantage of the tableaux framework, namely that the computation of a model remains local with respect to the currently expanded branch. This is especially useful in cases where we have ambiguities.

(2.16) *Peter knows Mary's cousin. He₁ is married to her₂ sister.*

Here, we have one best explanation where the first pronoun refers to Peter and the second pronoun refers to Mary. This leads to the preferred reading where Peter is married to Mary's sister. Two alternative readings remain accessible in the tableau, one where Peter is married to the sister of Mary's cousin, and one where Mary's cousin is married to Mary's sister. These

two readings are not investigated further since the semantic content of the pronouns is not implied by the background knowledge. Hence, a further tableaux expansion can provisionally concentrate on one branch that contains the preferred reading. However, a sentence like (2.17) may later close the current branch by adding inconsistent background knowledge. The method then must backtrack and choose a new interpretation of the discourse from the available open branches. This new branch is then expanded up to the point where it is either closed or saturated.

(2.17) *Peter is now forty, and he still is a single.*

While the example appears to be trivial, it should be made clear that even many contemporary approaches to anaphora resolution use inference only for *verifying* that a provisional coreference of a pronoun to its referent is consistent with the context. The approach by Baumgartner and Kühn actually *computes* the resolutions of anaphoric expressions, and it does so in an incremental way that is more efficient than a generate-and-test approach which must in the end consider all possible mappings of pronouns to referents. The approach makes use of the ability of tableaux to deal with incremental information, which could be a serious advantage in comparison to those methods that must recompute all inferences when a discourse is extended.

6.2.4 An Alternative by Conservative Minimality

The linguistic theory behind abductive anaphora resolution considers preferable readings as abductive explanations for the semantic content of the anaphoric expression. This theory can be implemented by other forms of model generation as well, for instance by conservative minimality. The intuitive idea is to *minimise* the logical contribution that each anaphora provides to the model of the discourse.

In the case of discourse (2.17), we can use the formulas (2.19)–(2.21) as a logical specification that approximates its meaning. Further, we add the formula (2.22) that adds a *ab*-assumption for each discourse participant that meets the semantic content of the definite description *the criminal*.

(2.18) *A politician chased a gangster. **The criminal** died.*

(2.19) $pol(p) \wedge gan(g) \wedge chase(p, g)$

$$(2.20) \quad \forall x \text{ gan}(x) \Rightarrow \text{crim}(x)$$

$$(2.21) \quad \exists x \text{ crim}(x) \wedge \text{die}(x)$$

$$(2.22) \quad \forall x \text{ crim}(x) \Rightarrow \text{ab}(x)$$

There are two locally minimal models for the formulas (2.19)–(2.22).

$$(2.23) \quad \mathcal{M}_1 = \{\text{pol}(p), \text{gan}(g), \text{chase}(p, g), \text{crim}(g), \text{die}(g), \text{ab}(g)\}$$

$$(2.24) \quad \mathcal{M}_2 = \{\text{pol}(p), \text{gan}(g), \text{chase}(p, g), \text{crim}(g), \text{crim}(p), \text{die}(p), \text{ab}(g), \text{ab}(p)\}$$

Only the first model \mathcal{M}_1 is a conservatively minimal model that has a minimum *ab*-cost. Our encoding implies a penalty for each individual that has the property of being a criminal. Hence, a model where the definite is resolved with a discourse participant that already meets the properties described by the definite itself has necessarily a lower *ab*-cost than one where the properties for the referent must be accomodated. Intuitively, the definite description is resolved with that individual that “fits better” in the given discourse situation. This concept of a preferred resolution is analogous to that in the tableaux approach, only that it is implemented by a minimisation of *ab*-costs rather than additional inference steps in the model computation. The presence of *anaph_i*-predicates in the encoding is not necessary any more because the minimisation of anaphoric links is already implemented by the *ab*-minimisation.

We make use of these insights and propose a new universal determiner **THE** as an alternative to that known from Chapter 4. Instead of unicity, i.e., a strict constraint on the truth conditions of a definite, we now imply abductive resolution by *ab*-predicate minimisation. The definition in our *MQC* logic is as follows.

$$(2.25) \quad \text{THE} \equiv \lambda P \lambda Q (\exists x P(x) \wedge Q(x)) \wedge \forall x P(x) \Rightarrow \text{ab}(x)$$

The sentence *The criminal died* then has the representation **THE**(*crim*)(*die*), and conservative minimality ensures that we identify the correct preferred reading in the set of models.

Baumgartner and Kühn consider pronominal anaphora to be a special case of definite anaphora. The determiner **THE** can in this sense be used to

represent pronouns as well. The first argument of THE contains the semantic contribution of the anaphor. This semantic contribution is the gender information in the case of pronouns. Thus, we can have the logical representations (2.27)–(2.29) for the discourse (2.26). Formula (2.29) is the semantic representation of *He sends her letters every day.*

(2.26) *Peter loves Mary. He₁ sends her₂ letters every day.*

(2.27) $\text{love}(\text{peter}, \text{mary})$

(2.28) $\text{NO}(\text{male})(\text{female}) \wedge \text{male}(\text{peter}) \wedge \text{female}(\text{mary})$

(2.29) $\text{THE}(\text{male})(\lambda x \text{THE}(\text{female})(\lambda y \text{sendLet}(x, y)))$

The conservatively minimal model of the input contains the assumption $\text{sendLet}(\text{peter}, \text{mary})$, which corresponds to the correct resolution of the pronouns as required.

The alternative approach of abductive resolution by conservative minimality does not correspond exactly to the hyper-tableaux approach. Baumgartner and Kühn concentrate on anaphora. They do not consider definite descriptions where an accommodation of new discourse participants is necessary. We have discussed examples such as (2.30) in detail in Section 4. The tableaux approach is based on a domain closure assumption, i.e., all inferences are restricted to a given universe of discourse that is defined by the first-order constants in the current branch. While new linguistic data may add new discourse participants, the abductive inferences that establish the links between anaphora (or definites) and their referents will not. In contrast to this, the minimal model approach allows accommodation as needed for instance in (2.30), but on the other hand will not indicate a discourse anomaly when a pronoun does not have a referent such as in example (2.31).

(2.30) ***Jon's rabbit** is cute.*

(2.31) ***She** hates my house. (*)*

It is no technical problem to add a domain closure assumption to a model generator, either by a logical encoding or a simple modification in the search for models. However, if such a domain closure is used, an example like (2.29) cannot be treated correctly any more.

6.3 Weighted Abduction

Hobbs et al. [50] considers the interpretation of natural-language texts as a process that yields one preferable explanation of why a text is true. This process is modelled by **weighted abduction**, a form of abductive inference that is based on abductive logic programming.

6.3.1 Logic Programming and Abduction

Logic Programming refers to forms of computation where the programs and the input that define a particular computation are given as logical statements, and the process of executing a computation is made explicit by inference. The relation between logic and algorithms is summed up by Robert Kowalski's equation

$$\textit{Algorithm} = \textit{Logic} + \textit{Control}$$

Logic programming systems can be characterised as theorem provers where the way in which proofs are computed is controlled by some fixed control strategy that can be used by a programmer for defining algorithms. The logical language used usually is less expressive than first-order predicate logic, but may be extended by non-logical commands that control the inference process and the input-output of a computation. The best-known example of a logic programming language is Prolog whose logical part is restricted to first-order Horn clauses.

A **query** to a logic programming system is some formula ϕ that is refuted with respect to a **database** Γ which consists of logically encoded knowledge. An **answer** to a query is some instantiation σ of free variables in ϕ such that $\sigma(\neg\phi)$ is inconsistent with respect to Γ .

Abductive logic programming systems are logic programming systems where facts that are needed for a proof be assumed as hypothesis instead of being proved. Abductive logic programming systems answer queries ϕ not only by returning an instantiation, but also by some set of additional hypothesis Δ . The relation between the query ϕ , the background knowledge Γ and the assumptions Δ meet the formal definition of abductive explanations given in Section 6.1.

Weighted Abduction computes a set of assumptions that explains the truth of a semantic representation as a “best” answer of a query to an abductive logic programming system. In weighted abduction, the literals in

the query and in the clauses in the database carry numerical values called **weights** which determine a system of preference for which facts can be taken as allowable hypothesis. Weighted abduction is a method of abductive inference that is closely related to Prolog-style logic programming. That is, the query and the database are restricted to first-order Horn logic and answers are computed by the usual backward-chaining of a Prolog system. The Horn clauses in the database themselves carry numerical values which determine **inference costs** for deriving information when using them.

Abductive Explanations

Before we investigate the intuition behind weighted abduction, we first give a simple example for a set of explanations that can be computed by abductive logic programming.

Let (3.2) be the logical form of sentence (3.1) that we want to interpret in a context (3.3) where we have some car a and some red object b , and the world knowledge that every car is a vehicle.

$$(3.1) \quad \textit{The vehicle is red.}$$

$$(3.2) \quad \phi = \textit{vehicle}(x) \wedge \textit{red}(x)$$

$$(3.3) \quad \Gamma = \{ \textit{car}(a), \textit{red}(b), \forall x \textit{car}(x) \Rightarrow \textit{vehicle}(x) \}$$

Let us further assume that the set of allowable hypothesis is the set of atoms that can be built from all predicates in ϕ and Γ and the individuals in background theory Γ . Thus, the set of allowable hypothesis is given in (3.4). The answers (3.5)–(3.12) each give some abductive explanation why ϕ could be a logical consequence of Γ , and at the same time instantiate the query ϕ .

$$(3.4) \quad \Psi = \{ \textit{red}(a), \textit{red}(b), \textit{car}(a), \textit{car}(b), \textit{vehicle}(a), \textit{vehicle}(b) \}$$

$$(3.5) \quad \Delta_1 = \{ \textit{vehicle}(a), \textit{red}(a) \}$$

$$(3.6) \quad \Delta_2 = \{ \textit{vehicle}(b), \textit{red}(b) \}$$

$$(3.7) \quad \Delta_3 = \{ \textit{vehicle}(a), \textit{red}(a), \textit{car}(a) \}$$

$$(3.8) \quad \Delta_4 = \{ \textit{vehicle}(a), \textit{red}(a), \textit{car}(b) \}$$

$$(3.9) \quad \Delta_5 = \{\mathit{vehicle}(b), \mathit{red}(b), \mathit{car}(a)\}$$

$$(3.10) \quad \Delta_6 = \{\mathit{vehicle}(b), \mathit{red}(b), \mathit{car}(b)\}$$

$$(3.11) \quad \Delta_7 = \{\mathit{vehicle}(a), \mathit{red}(a), \mathit{car}(b), \mathit{vehicle}(b)\}$$

$$(3.12) \quad \Delta_8 = \{\mathit{vehicle}(b), \mathit{red}(b), \mathit{car}(b), \mathit{vehicle}(a)\}$$

For our purpose of natural-language interpretation, there are too many explanations. The situation is very similar to model generation where the number of models for a semantic representation is in general much higher than the set of valid natural-language interpretations. However, some explanations can be eliminated immediately as they use redundant assumptions—these explanations correspond to the models that are not minimal in model generation.

In many cases, redundant assumptions will never be considered as hypothesis by an abductive logic programming system. A fact can only become a candidate hypothesis if it either is used directly for refuting a literal of the query or can be reached by some chain of rule applications as a literal that must be refuted. For instance, Δ_8 will actually never be generated by a Prolog-style abductive logic programming system, as there is no sequence of backward-chaining applications of rules in Γ such that all assumptions Δ_8 can be reached. Only the explanations Δ_1 – Δ_4 may actually occur as answers.

For natural-language interpretation, we must find a way to eliminate all explanations that are not linguistically valid. For this, there are at least three approaches available.

The first method is to use a more appropriate logical encoding. The definite description in our example is actually modelled by an implicitly existentially quantified formula that does not include the uniqueness condition of singular definites. In a richer logical encoding that models unicity, such as the one discussed in Chapter 4, no explanations except Δ_3 and Δ_7 model the observation. Hence, all other explanations violate one of the conditions for abductive explanations.

A second approach to the problem is that by Baumgartner and Kühn that we have discussed in Section 6.2. Here, that explanation for the resolution of an anaphor is preferred that can be explained by logical consequence from the context. In our example, the definite *the vehicle* triggers an additional inference process that prefers the logically consequent link of the definite *the vehicle* to the car *a* instead of the accommodated link to the red object *b*.

The link is a logical consequence of an information given by the context Γ , namely that every car is a vehicle.

The third method is to encode the linguistic knowledge implicitly into a system of preferences for abductive explanations. We have implemented this idea for model generation in Section 6.2.4 where the *ab*-index of each minimal model is used to characterise best explanations. In weighted abduction, this system of preference is implemented by the weights and costs of the logical encoding.

Weights and Costs

The linguistic information that we must implicitly encode in our example is as follows. First, as Hobbs et al. argue, the determiner ‘the’ indicates an entity that is “*the most salient, mutually identifiable entity of that description*”. In other words, the part of the query that represents this entity should be difficult to assume, and rather be shown as some logical consequence of the contextual knowledge. We can encode this implicitly by giving this part of the query high assumption costs, say 100. Then, the cost of that part the query that does not depend on the definite description must be lower, say 20. We consider each fact in the background knowledge to be easily accessible, and give it a minimum cost of 1. The rule that formalises that each car is a vehicle must have some low weight, as it is a rule that encodes easily accessible world knowledge. We choose a weight of 5. If something is a vehicle, we may assume with some low probability that it also is a car. For this form of abductive inference, we assign an assumption cost of 70; it is less than the cost for assuming that some arbitrary object is a car, but still high. We write costs and weight as superscript values, and our query and the background knowledge is now as follows.

$$(3.13) \quad \phi = \mathit{vehicle}^{100}(x) \wedge \mathit{red}^{20}(x)$$

$$(3.14) \quad \Gamma = \{\mathit{car}^1(a), \mathit{red}^1(b), \forall x \mathit{car}(x)^{70} \Rightarrow^5 \mathit{vehicle}(x)\}$$

In order to answer the query ϕ , the weighted abduction system will refute the negation of ϕ using the background knowledge Γ . If an atom cannot be proven, it is taken as a hypothesis where its assumption cost goes into the overall costs of the proof. If a rule is used to derive new atoms that must be proven, then the rule’s weight goes into the cost of the whole proof. For

our example, the following explanations can be computed by abductive logic programming proof attempts.

$$(3.15) \quad \Delta_1 = \{\mathit{vehicle}^{100}(a), \mathit{red}^{20}(a)\}$$

$$(3.16) \quad \Delta_2 = \{\mathit{vehicle}^{100}(b), \mathit{red}^1(b)\}$$

$$(3.17) \quad \Delta_3 = \{\mathit{vehicle}^5(b), \mathit{car}^{70}(b), \mathit{red}^1(a)\}$$

$$(3.18) \quad \Delta_4 = \{\mathit{vehicle}^5(a), \mathit{car}^1(a), \mathit{red}^{20}(a)\}$$

Solution Δ_1 simply takes $\mathit{vehicle}(a)$ and $\mathit{red}(a)$ as two hypothesis that obviously are abductive explanations of the query. However, the assumptions costs are $100 + 20 = 120$, which makes this explanation the most expensive that can be computed by weighted abduction. Solution Δ_2 is slightly cheaper (101) since it only assumes $\mathit{vehicle}(b)$ and then uses the easily accessible knowledge that b is red. Then, Δ_3 includes some abductive inference where $\mathit{vehicle}(b)$ is “proven” by the assumption that b is a car. The derivation cost is 5, as our rule has a weight of 5, and the assumption cost for $\mathit{car}(b)$ is 70. Again, we make use of the fact that b is red. Hence, the cost of the answer Δ_3 is $5 + 70 + 1 = 76$. The best solution, however, is Δ_4 which derives that a is a red vehicle by proving that it is a vehicle at a very low cost of 6, and assuming that a is red with cost 20. The overall cost of 26 is the lowest possible for our example.

Our example illustrates the intuition behind weighted abduction as a method that uses numerical annotations in the logical specification to control the search for a best abductive explanation. By branch-and-bound-search, such a best explanation can often be determined quickly. Weighted abduction implements abductive inference in natural-language interpretation in a way that suppresses some effects of the combinatorial explosion caused by the large number of abductive explanations. However, as Hobbs et al. state themselves, the actual numerical values are ad hoc, and the high amount of interaction during computation it very difficult to engineer some universal and reliable system of annotations. There is, at present, no linguistic theory behind weights and costs. Nevertheless, weighted abduction has been used with some success in the TACITUS natural language system on a range of examples.

6.3.2 Applications

Hobbs et al. [50] gives a variety of applications for (weighted) abductive inference in natural-language interpretation. In the following, we will shortly present some of these applications.

Definite Reference

Weighted abduction can deal with different kinds of definite reference such as we have discussed in some detail in Chapter 4. Weighted abduction implements the empirically testified preference of anaphoric resolution over accommodation by high assumption costs of definite descriptions in the query. Examples of bridging, such as (3.19), are dealt with in a similar way as in our analysis. The existence of the referent for the definite *the engine* is proven by using² a rule (3.20) that has some low inference cost. In other words, from the existence of a car we can follow with low inference costs that there also is an engine that belongs to the car. This engine is used for proving the existence of the engine in sentence (3.19).

(3.19) *I bought a new car last week. **The engine** is already giving me trouble.*

(3.20) $\forall x \text{ car}(x) \Rightarrow \exists y \text{ engine}(y) \wedge \text{of}(x, y)$

Note that the logical form here is actually not Horn. Formula (3.20) must be translated into a Horn clause that approximates the semantics of (3.20) and that is acceptable by the abductive logic programming system. This may for instance require a special predicate that implements negation as failure.

Weighted abduction can correctly model examples such as (3.21) where the definite description is actually a determinative definite noun phrase that includes all information needed for its resolution. In examples such as this, there is no way to prove the existence of the definite referent from the context. Hence, the abductive explanation that simply assumes the existence of Jon's rabbit should be the cheapest one possible.

(3.21) ***Jon's rabbit** is cute.*

²Here, as in the following examples, we assume that suitable values for weights and costs have been chosen, and leave them out for having a simpler presentation.

Weighted abduction gives a clear preference for only one single solution, as there is not really some tractable way of determining “second best” solutions.

In both weighted abduction and minimal model generation, sentences such as (3.22) are a problem because there is an ambiguity that cannot be resolved by the usual heuristics. Does *the old man* refer to Konrad himself or maybe his supervisor? Domain-minimal model generation prefers the reading where Konrad is the old man, while weighted abduction may chose some arbitrary reading depending on how the weights and costs are chosen.

(3.22) *Konrad's PhD is terrible. **The old man** should really read it again.*

Composite Noun Phrases

Weighted abduction has been used in the analysis of complex noun phrases that are composed of given and new information. An example for such noun phrases is given in (3.23) in the form of the definite *the Boston office*.

(3.23) ***The Boston office** called.*

Ignoring the problem we must somehow expand the metonymy to “[Some person at] the Boston office called”, the primary interpretative problem in sentence (3.24) is to determine the implicit relation between the person that called and the office, and the relation between Boston and the office. The approach proposed by Hobbs et al. treats the sentence roughly³ as a formula (3.24).

(3.24) $\exists x \exists y \text{ call}(x) \wedge \text{person}(x) \wedge \text{rel}(x, y) \wedge \text{office}(y) \wedge \text{nn}(y, \text{boston})$

The relation *rel* is a placeholder for a relation that holds between an individual and a location, while *nn* stands for some relation that holds between a location and another one. The logical encoding informally states that some person called who somehow is related to some office that somehow is related to Boston. By giving *rel*(*x*, *y*) and *nn*(*y*, *boston*) high assumption costs, we enforce an explanation that uses some instances for relations that are derivable from the background knowledge. The following are examples for rules in the database that may be used for such a derivation.

³Again, we ignore temporal issues.

$$(3.25) \quad \forall x \forall y \text{ in}(x, y) \Rightarrow \text{nn}(x, y)$$

$$(3.26) \quad \forall x \forall y \text{ works-for}(x, y) \Rightarrow \text{rel}(x, y)$$

Suppose that we have given the additional information in the database that o_1 is an office in Boston and that a person *jon* works for office o_1 . Then (3.27) can be taken as an explanation for (3.24) where $\text{call}(\textit{jon})$ is the only new information that must be assumed. The explanation corresponds to a natural-language interpretation of sentence (3.24), namely that *jon* is the person who called and works for the office o_1 in Boston.

$$(3.27) \quad \Delta = \{ \text{call}(\textit{jon}), \text{person}(\textit{jon}), \text{rel}(\textit{jon}, o_1), \text{office}(o_1), \\ \text{nn}(o_1, \textit{boston}), \text{works-for}(\textit{jon}, o_1), \text{in}(o_1, \textit{boston}) \}$$

Examples such as (3.23) illustrate how some missing information in the semantic representation of a sentence can be filled in by abductive reasoning. Based on a similar logical encoding, the KIMBA model generator is able to reproduce the analysis of the example without using any heuristic control other than the world knowledge that persons are neither cities nor offices. Local minimality suffices here to identify the correct interpretation.

The weights and costs in weighted abduction are often used in a way that simulates the effects of local minimality. In our example, weighted abduction makes sure that a maximum of information that is easily accessible in the database is used, while $\text{call}(\textit{jon})$ remains the only true assumption. In model generation, we do not distinguish assumptions and provable facts, but the characteristic properties of a minimal model make sure that information which can be derived as a logical consequence of the context are preferred. The effect of the two approaches to abductive interpretation often is the same: the implementation of conversational conventions as computable constraints.

A difficulty with the abductive approach in general is that there may be many rules such as (3.25) and (3.26) that give possible instantiations for the placeholder relations. While weights may control which instances are chosen, it is not entirely clear how any system of weights can guarantee that the selection is correct. Likewise, model generation will produce a variety of models where we cannot decide which ones truly represent preferable interpretations. This is especially true when contextual information is missing in the background knowledge. Model generation seems to be somewhat more sensible to missing knowledge than weighted abduction which only makes assumptions that can be reached by some backward-chaining interpretation of the rules in the database.

Resolving Ambiguity

Common-sense reasoning helps listeners to associate words with the correct concepts. A typical example of word-level lexical ambiguities is given by example (3.28).

(3.28) *The plane taxied to the terminal.*

The interpretative process must identify which semantic concepts are behind such words as “plane”, “taxied”, and “terminal”, all of which carry more than one meaning. We can express a classification of concepts as rules in the database. The following rules give some background knowledge to the ambiguous words.

(3.29) $\forall x \text{ airport-terminal}(x) \Rightarrow \text{terminal}(x)$

(3.30) $\forall x \text{ computer-terminal}(x) \Rightarrow \text{terminal}(x)$

(3.31) $\forall x \forall y \text{ ride-cab}(x, y) \wedge \text{person}(x) \Rightarrow \text{taxi}(x, y)$

(3.32) $\forall x \text{ move-on-ground}(x) \wedge \text{airplane}(x) \Rightarrow \exists y \text{ taxi}(x, y)$

(3.33) $\forall x \text{ wood-smoother}(x) \Rightarrow \text{plane}(x)$

(3.34) $\forall x \text{ airplane}(x) \Rightarrow \text{plane}(x)$

Terminal is a hyponym of both airport-terminal and computer-terminal, as stated by the rules (3.29) and (3.30). Taxiing may refer to the process of a person riding a cab, or an air-plane moving on the ground, as stated by rules (3.31) and (3.32). Finally, the word “plane” may refer to a wood-smoother or an air-plane.

The task of weighted abduction in the example is to compute an answer for the query (3.35) that somehow makes explicit that the sentence (3.28) refers to an air-plane moving on the ground to an airport terminal. An essential hint that we give to the system is rule (3.36) which expresses the common knowledge that whenever we have an airport, we also have an air-plane and an airport-terminal.

(3.35) $\exists x \exists y \text{ plane}(x) \wedge \text{taxi}(x, y) \wedge \text{terminal}(y)$

$$(3.36) \quad \forall z \text{ airport}(z) \Rightarrow \exists x \exists y \text{ plane}(x) \wedge \text{airport-terminal}(y)$$

The analysis intuitively derives a minimal explanation (3.37) as follows. First, if $\text{plane}(x)$ in the query is explained by $\text{airplane}(c_1)$, then the fact $\text{airplane}(x)$ that is part of one of the rule (3.32) for $\text{taxi}(x, y)$ has already been assumed. Hence, the interpretation that uses this rule is cheaper than that where $\text{taxi}(x, y)$ requires more assumptions with respect to a person and some cab. If we further explain $\text{terminal}(y)$ by $\text{airport-terminal}(c_2)$, we can reduce the assumption costs of $\text{airport-terminal}(c_2)$ and $\text{airplane}(c_1)$ by using the rule (3.36) that assumes an airport c_3 . The cost of assuming a fact in weighted abduction is computed on the basis of the cheapest way to derive it as a candidate hypothesis. Any additional uses of an assumption in the computation of an explanation are free. As a result, the reuse of assumptions is preferred.

$$(3.37) \quad \Delta = \{\text{plane}(c_1), \text{taxi}(c_1, c_2), \text{terminal}(c_2), \text{airplane}(c_1), \\ \text{airport-terminal}(c_2), \text{airport}(c_3)\}$$

In model generation, we have no device analogous to the derivation costs in weighted abduction. In our example, the reuse of information makes the explanation (3.37) cheaper than other ones, even though it requires an assumption $\text{airport}(c_3)$ which is not required in other explanations. In comparison, the preference for a certain model is entirely determined by the atoms that constitute its representation. When we use minimal model generation on a similar logical encoding, the preferred interpretation that is computed is some model that corresponds to an explanation (3.38). The model uses less individuals as it does not require an airport. The correct interpretation is computed only if we leave out the “hint” that is necessary in weighted abduction. Domain minimality actually prevents us from identifying the most intended interpretation.

$$(3.38) \quad \Delta = \{\text{plane}(c_1), \text{taxi}(c_1, c_2), \text{terminal}(c_2), \text{airplane}(c_1), \\ \text{computer-terminal}(c_2)\}$$

6.3.3 Discussion

Apart from the applications presented in Section 6.3.2, weighted abduction has been put to work on such different phenomena as discourse structure, metonymy, redundancy, and various combinations of pragmatic problems.

The examples that have been investigated are real-world texts from equipment failure reports, naval operation reports, and newspaper articles on terrorist activities. For a more thorough presentation of the method and its applications, we refer to Hobbs et al. [50].

In the previous chapters, we have presented model generation as a reasoning method that can be seen as a competitor of weighted abduction as a computational tool for text interpretation. Before we discuss the differences of the methods, we would like to point out that apart from technicalities, the approaches themselves are very similar.

Similarities

First, weighted abduction and minimal model generation conceptualise text interpretation as a task that computes an explanation for truth. The intuition behind this can be taken as a universal approach to natural-language understanding. Even questions or statements with a high pragmatic content can be captured in principle by interpretation as abduction, since interpretation in general explains the truth of a speaker's intentions, emotions, and beliefs as well as the plain truth of an utterance. Given a suitable logical theory of intentions, emotions, and beliefs, abductive reasoning explains pragmatic phenomena as well as semantic ones—if we can solve the representational and complexity problems that are inevitable when reasoning about pragmatics. By generating abductive explanations of why a text has a conversationally sensible rhetoric structure, weighted abduction can even be used for discourse structure construction as well as for interpretation. There are few reasons to suspect that minimal model generation may not be useful as well for such inference tasks.

Second, the computation of truth is based on well-known methods in automated reasoning. Minimal model generation and weighted abduction are basically classical first-order inference methods whose purpose is the computation of best explanations. The use of standard tools, or rather of standard tools that have been tweaked to fit a special purpose, leads to a semantic representation language that often can only approximate the meaning of the natural-language input in the sense of a truth-conditional natural-language semantics. However, this disadvantage is paid for by the actual ability to investigate the phenomenon of inference by experimentation.

Third, both model generation and weighted abduction face the same principal problem of abductive inference, namely the combinatorial explo-

sion of explanations for an observation. As natural-language semantics is still dominated by the topic of representation, the criteria that distinguish a valid natural-language interpretation and the methods how these can be determined remain almost entirely unknown. Inference as abduction gives a practical tool for investigating these criteria, either in the form of a model generator or an abductive logic programming system.

Differences and Comparison

Apart from some minor technical details, the main difference of interpretation as weighted abduction to interpretation as model generation is the way in which the two approaches deal with the overgeneration problem of abductive explanations.

In weighted abduction, the main filter mechanism for unwanted explanation is actually the method with which answers for queries are computed. Given a query A and a database $\Gamma = \{B \Rightarrow A, B, C \Rightarrow D\}$, the abductive logic programming system will try to refute $\neg A$ either by assuming A as a hypothesis or by using the rule $B \Rightarrow A$ and the fact B by backward-chaining. The literals C and D cannot be reached at all with query A , and thus will never be part of an answer. In this sense, abductive logic programming is goal-oriented, as it only uses assumptions if they contribute something to the goal of explaining the query. Rules that do not contribute to the computation of an answer will never be executed. This can result in a considerable advantage in tractability as the knowledge base grows larger. However, the downside of the general approach is that without further technical machinery, there is no guarantee that the answers which are computed meet the formal conditions of abductive explanations. That is, abductive logic programming as implemented in weighted abduction sometimes gives sets of assumptions Δ which are not consistent with the database Γ of background knowledge.

In principle, it should be possible to augment weighted abduction with some proof method for finite satisfiability. However, as weighted abduction does not efficiently enumerate explanations, it seems to be a considerable technical problem to combine weighted abduction and a consistency check such that the resulting method still works in practice. The conversational maxim of consistency is not guaranteed by the current implementation. This has some serious disadvantage for natural-language interpretation: as long as consistency is not maintained by the proof procedure, there may be cases where a better logical encoding actually leads to inconsistent “explanations”.

There is little point in further developing truth-conditional theories in computational semantics if a better logical modelling negatively interacts with the inference process!

In model generation, there is no goal-orientedness or flow of control in principle. Model generation by itself has no operational filter mechanism that eliminates irrelevant literals such as C and D in our example. The answer to this problem is some form of minimality which prefers models that can explain the same observations with some subset of assumptions. The price that we have to pay for maintaining the formal properties of abductive explanations is a complexity problem. There seems to be no way to prevent a rule from interacting with the computation process once it is provided with the database, not even if it seems to be irrelevant for the task at hand. Furthermore, the computation of minimal models is considerably harder than that of simply proving satisfiability.

The weights in weighted abduction are a method for controlling the search for explanations. These weights can define a fine grained system of preferences that includes the plausibility of certain assumptions as well as some measure for the difficulty of inferring certain knowledge. However, weights currently have no theory that determines how they must be selected. As we have shown, our computational treatment of linguistic phenomena by minimal model constraints can be firmly based on contemporary linguistic theories.

Chapter 7

Implementation

Keep It Simple!
(Charles H. Moore)

Overview: This chapter describes the finite model generator KIMBA.

7.1 Introduction

Natural-language interpretation needs model generation methods that selectively enumerate the models of semantic representations. Finite model generation based on constraint solving provides such enumeration capabilities in an easily accessible way.

The finite model generator KIMBA¹ is based on the *MQC* specification language presented in Section 3.3 and its translation into constraints presented in Section 3.4. KIMBA implements an efficient enumeration procedure for finite models and can also enumerate models with certain minimal model properties.

The specification language *MQC* itself does not have an efficiently computable clause normal form, but KIMBA can avoid normalisation thanks to the properties of the translation. The basic two-valued logic *MQC* described in Section 3.3.3 can be generalised in the implementation to a collection of logics with finitely many truth values. The user may redefine the semantics of logical connectives and thus design her own logics.

¹The original Kimba, the Jungle Emperor, is a small white lion from a Japanese cartoon series in the 1970s. KIMBA is the small brother of LEO [14].

In the following, I start by sketching the general system architecture of KIMBA in Section 7.2. Then I discuss both the syntax of the specification language in Section 7.3, and, in Section 7.4, the way how semantics is assigned to it in the form of logic definition structures and propagator procedures. Finally, I outline the control mechanism of the search for models and the variety of proof procedures that are part of KIMBA in Section 7.5 and give some results about KIMBA's performance in Section 7.6.

7.2 System Architecture

The system architecture of KIMBA is depicted in figure 7.1. A logical input is first translated into a combinatorial constraint problem over finite-domain integer variables. This translation relates to one certain frame of constants that is determined by the problem itself and a current universe of first-order individuals. Together with a control strategy that describes how and which models are to be generated, the combinatorial problem is handed to the constraint solver.

The actual constraint solving process consists of two separate processes that are called iteratively. The first process is **propagation** which restricts the possible values of the finite-domain variables using the knowledge encoded in the constraints. Changing the value of a variable often imposes new constraints to the values of other variables as well, and propagation is applied until the set of constraints is stable, i.e., no further constraints can be derived and added. The second process, **distribution**, selects variables and restricts their values provisionally. The selection of the variable(s) and the actual form of their restriction is done heuristically. After a distribution step, constraints will propagate again, and so forth.

The iterative process of propagation and distribution continues until all variables are uniquely determined or a failure occurs. A model in KIMBA is represented internally as a list for the set of atoms that has been build from the constant frame that defines the range of quantification during model construction. The table's entries are pairs of atoms and finite-domain integer variables whose values each represent the truth value of the associated atom. Each change of a value in this table results in a different interpretation. By distributing over the variables' values, the constraint solver systematically enumerates all finite models with regard to the constant frame from which the set of atoms is derived.

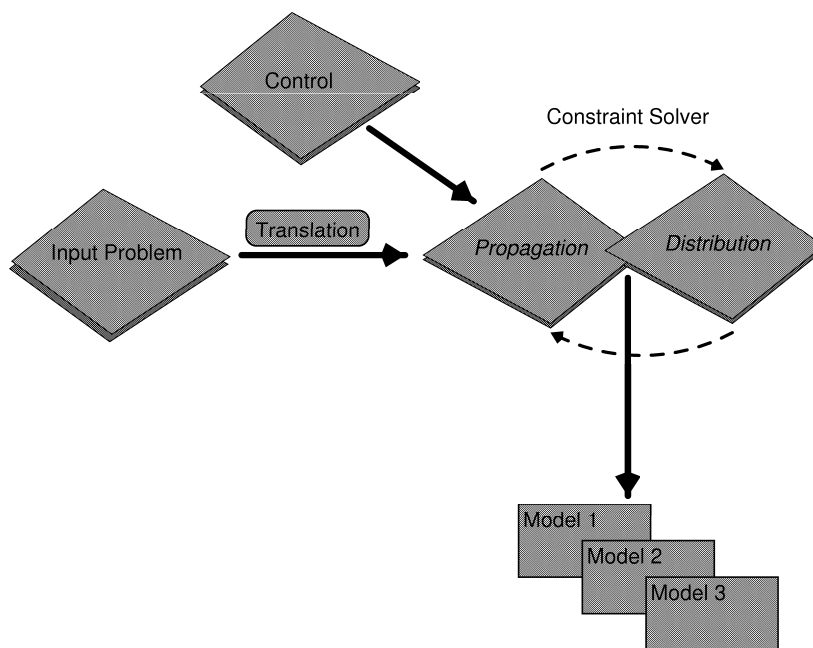


Figure 7.1: KIMBA's system architecture at a glance.

A failure, i.e., an assignment that leads to an unsatisfiable set of constraints, evokes backtracking, and the value of the variable that lead to the failure is constrained differently. The control strategy may both add constraints to the constraint problem itself, or start some additional inference process for each solution that is produced by the constraint solver. For instance, the constraint solver may produce some model for which the control strategy then proves whether it is a locally minimal model.

The constraint solving module of the finite model generator can be seen as independent from the logic and the problem specification considered. Once a suitable translation from the logical language into a constraint language is given, solving a model generation problem reduces to the application of any constraint solver that can deal with the constraint language in question. In KIMBA, the constraint solver used is provided by OZ [78], the constraint logic programming language in which KIMBA is implemented.

The main part of KIMBA is actually the translation from the input into a constraint language for finite-domain integer variables. The other parts of the system consist mainly of the implementation of the basic data structures and operations such as β -reduction. KIMBA has been implemented in the spirit of *lean deduction* as exemplified by the *lean^{AP}* tableaux prover [12]. The philosophy of lean deduction is it to have automated reasoning systems that are as small as possible. Lean automated reasoning aims at providing simple systems suitable for research and education rather than high-speed systems where efficiency is bought by complexity in the implementation. The whole KIMBA system is only about 30kBytes of OZ code, including a minimalistic graphical user interface for loading problems, controlling the parameters of the search, and visualising the results.

7.3 The Syntax

The syntax of KIMBA is derived from the *MLC* variant of Church's simply typed λ -calculus presented in Section 3.3.

7.3.1 Logical Constants

Figure 7.2 shows the basic set of logical connectives and quantifier constants. Other quantifiers, connectives and determiners can be defined as λ -terms in the problem specification, or permanently added to the system in the logical

\mathcal{MQL}	KIMBA	\mathcal{MQL}	KIMBA
\neg	not	\forall	forall
\vee	or	\exists	exists
\wedge	and	EVERY	every
\Rightarrow	implies	SOME	some
\Leftrightarrow	equiv	MORE	more

Figure 7.2: Standard Connectives and Quantifiers

Quantifier	Example	Truth condition
atLeast	[atleast man 5]	$ man \geq 5$
atMost	[atMost man 5]	$ man \leq 5$
exactly	[exactly man 5]	$ man = 5$

Figure 7.3: Cardinality Quantifiers

definition structures (see Section 7.4). The Herbrand equality ' \equiv ' from \mathcal{MQL} is denoted simply by the equality sign $=$.

Apart from the standard quantifier constants, KIMBA provides a set of cardinality quantifiers that can be used to constrain the cardinality of sets. The set of available cardinality quantifiers together with some examples of their use is given in Figure 7.3.

7.3.2 Formulas

Formulas in KIMBA follow the syntax of \mathcal{MQL} . Hence, the only terms allowed in a $\beta\eta$ -normalised formula are parameter constants or bound variables. Applications of the form $h(x_1, \dots, x_n)$ are written as [h $x_1 \dots x_n$]. Likewise, a quantified \mathcal{MQL} formula $Q(T)(U)$ is written as [Q T U].

A λ -abstraction $\lambda x_i. p(x)$ is written as lam(x#i [p x]). The bound variable of an abstraction must be given an explicit type, such as in this case by x#i.

We have basic types 'i' for individuals, 'o' for truth values, and the type 'n' for numbers. Numbers are used in general only for cardinality constraints. Higher-order types are written in a backward fashion such that the goal type of a function is always the first type. For instance, a KIMBA type [o [o i] [o i o]] denotes the \mathcal{MQL} type $(o \rightarrow o \rightarrow o) \rightarrow (i \rightarrow o) \rightarrow o$. KIMBA can up to some extent deal with polymorphic types, for instance

#	\mathcal{MQL} Formula	KIMBA
1	$\text{EVERY}(\text{man})(\lambda x_l \text{eat}(x) \wedge \text{sl}(x))$	<code>[every man lam(x#i [and [eat x] [sl x]])]</code>
2	$\text{NO}(\text{I}(\text{jon}))(\lambda x_l \text{I}(\text{jon}, x))$	<code>[no lam(x#i [l x jon]) lam(x#i [l jon x])]</code>
3	$\text{NO}(\lambda x_l \text{I}(\text{jon}, x))(\text{I}(\text{jon}))$	<code>[no [l jon] lam(x#i [l x jon])]</code>
4	$\text{MORE}(\lambda x_l \neg \text{man}(x))(\text{man})$	<code>[more lam(x#i [not [man x]]) man]</code>
5	$\forall x_l w(x) \Rightarrow \text{I}(\text{jon}, x)$	<code>forall(x#i [implies [w x] [l jon x]])</code>
6	$\text{EVERY}(w)(\lambda x \text{I}(\text{jon}, x))$	<code>[every w [l j]]</code>

Figure 7.4: Specification in KIMBA

in higher-order definitions of predicate constants whose arguments are polymorphic. A basic type 'x' is reserved for denoting polymorphic types α .

Fig. 7.4 shows some formulas and some of their possible translations into KIMBA's specification language. The redundancy in having a set of standard quantifiers as well as a set of linguistically motivated generalised determiners gives us some freedom in using whatever logical form we think is more elegant or appropriate. For instance, formula 5 is a formula in standard first-order notation whose semantics is equivalent to that of formula 6 which uses a generalised determiner in a more compact representation. Note that the order of arguments of n-ary predicates in KIMBA is different from the order of arguments in \mathcal{MQL} , for no especially good reason except for compatibility to some older first-order versions of the program.

7.3.3 Problem Specifications

A **problem specification** in KIMBA is a structure that consists at least of the following parts.

- a declaration of the signature for the non-logical constants
- a declaration of which logic used for formalising the problem
- a list of formulas

Optionally, we can have a set of higher-order definitions, a remark that describes the problem in natural language, and an entry that names a default constraint solving engine (cf. Section 7.5) for the problem.

```

p( remark: 'Two girls love jon.'
  logic: classical
  engine: 'local minimal'
  constants: are(girl:[o i] love:[o i i] jon:i)
  definitions:
    are(two: lam(p#[o i]
              lam(q#[o i]
                ['and' [exactly p 2] [every p q]])))
  formulas: [ ['not' [girl jon]]
             [two girl lam(x#i [love x jon])] ] )

```

Figure 7.5: A problem specification for KIMBA

A Small Example

Figure 7.5 shows a small problem specification. The specification starts with an entry at the label `remark` that describes the problem, in this case a formalisation of the sentence *two girls love jon*. The entry with the label `logic` specifies that the logic we used is the classical two-valued one. The label `engine` specifies the problem solving engine. For our problem, we use the 'local minimal' engine that enumerates locally minimal first-order models. This means that the solving engine is allowed to extend the initial universe of discourse if necessary, and that only those models are presented which meet the local minimality condition.

The initial constant frame is given by the problem specification. We have parameter constants `girl`, `love`, and `jon`, all of which have different types. We also have a higher-order definition for a new quantifier constant `two` such that $\text{two}(P)(Q)$ is true iff P has exactly two members and is a subset of Q .

Finally, the set of formulas specifies that Jon is no girl and that two girls love Jon. The problem specification can be loaded into KIMBA, and the only locally minimal model that is yielded by the model generation process has a representation as the following set of atoms.

$$(3.1) \quad \{ [\text{girl } C1], [\text{girl } C2], \\ [\text{love } C1 \text{ jon}], [\text{love } C2 \text{ jon}] \}$$

Here, the constants `C1` and `C2` are newly generated ones of type `i` that have

```

the: lam(p#[o i]
        lam(q#[o i]
            ['and' [exactly p 1] [every p q]]))

my: lam(p#[o i]
        lam(q#[o i]
            [the lam(x#i ['and' [p x] [isOf speaker x]]) q]))

```

Figure 7.6: Defining the determiners THE and MY by λ -terms

been added by the model generator while the first-order domain of constants was extended.

Some larger example specifications can be found in Appendix A.

7.3.4 Definitions

Definitions occur as part of a problem specification and are λ -terms that replace the occurrences of defined constants within the specification before the actual model generation process starts. KIMBA uses the standard β -normalisation to reduce expressions containing definitions into ones that have the function free quantified form required by the translation into constraints. Figure 7.6 shows two definitions, one for the determiner THE and one for the determiner MY. The second definition makes use of the first one. A definition may generally refer to other defined symbols as long as the definition expansion terminates, i.e., there are no cyclic definitions.

7.4 The Semantics

Logics are, in general, defined by the semantics of their logical constants. In KIMBA, the semantics of a logical constant is given operationally by the definition of a concurrent procedure that manipulates the values of finite-domain integer variables. A user can program these procedures, and combine them into structures that define which logical constants are available within a logic and in which way they behave.

7.4.1 Logic Definition Structures

A logic based on the *MQC* specification language can be defined in KIMBA by building a Logic Definition Structure (LDS). LDSs are unique to KIMBA. Unlike other finite model generators restricted to some fixed (classical) logic, LDSs provide KIMBA with a mechanism for implementing different logics in a modular way. The logics in KIMBA can vary in the logical constants that are available, their semantics, and in the number of truth values that can occur as the denotation of a formula.

An LDS is a mapping from constant symbols into propagator procedures. The constant symbols mentioned are the logical constants of the logic, for which the propagator procedures define an operational semantic. Figure 7.7 shows a partial LDS for classical 2-valued logic. The constants for the various connectives are mapped directly into predicates that are provided by OZ, for instance `'FD.conj'` as the propagator procedure that treats 2-valued conjunction. LDSs may define the semantics of truth constants such as `'true'` and `'false'` and the semantics of arithmetic predicates on integers such as `'>'`. We also have a slot for defining the equality predicate `'='`. This allows us to implement special forms of equality if necessary.

Additionally, LDSs also determines the type of all logical constants and the number of truth values and a translation from the truth values into a human-readable form. For instance, we usually want to interpret the highest truth value as “true” and the lowest one as “false” while other truth values might have different purposes. As mentioned earlier, the symbol `'x'` indicates a polymorphic type. A typical example is the definition of the quantifier `'exists'` that has a type $(\alpha \rightarrow o) \rightarrow o$, or, in KIMBA's syntax, a type `[o [o x]]`.

An LDS does not necessarily have to contain all logical constants that are used in other logics. A user may give only a basic set of connectives, quantifiers and determiners (if any) and implement other logical constants by the definition mechanism presented in the previous section.

7.4.2 Propagator Procedures

Implementing a logic in KIMBA means to implement the **propagator procedures** for the logical constants. A propagator procedure in OZ is a concurrent procedure that observes and propagates the changes in the values of a set of variables. In our case, these variables are finite-domain integer vari-

```

ClassicLDS =
lds(truths: are('false' 'true')
  constants: are('true': (o # 1)
    'false': (o # 0)
    '=': ([o x x] # proc {$ FA SA Root}
      if FA == SA then
        Root = 1 else
        Root = 0 end
      end)
    '>': ([o n n] # proc {$ [FA SA] Root}
      if FA > SA then
        Root = 1 else
        Root = 0 end
      end)
    'succ': ([o n n] # proc {$ [FA SA] Root}
      if FA == SA+1 then
        Root = 1 else
        Root = 0 end
      end)
    'not': ([o o] # FD.nega)
    'and': ([o o o] # FD.conj)
    'implies': ([o o o] # FD.impl)
    'exists': ([o [o x]] #
      proc {$ [Insts] Var}
        {FD.reified.sum Insts '>:' 0 Var}
      end)
    'atLeast': ([o n [o x]] #
      proc {$ [Insts Card] Var}
        {FD.reified.sum Insts '>=:'
          Card Var}
      end)
    'some': ([o [o x] [o x]] #
      proc {$ [InstsA InstsB] Var}
        {FD.reified.sum
          {List.zip InstsA InstsB
            fun {$ IA IB} {FD.conj IA IB} end}
          '>:' 0 Var}
      end)))
end)))

```

Figure 7.7: A partial LDS for a classical two-valued logic

```

proc {N1 V1 V2}
  [V1 V2] ::: 0#2
  thread
    if V1 = 0 then V2 = 2
    [] V1 = 1 then V2 = V1
    [] V1 = 2 then V2 = 0
    [] V2 = 0 then V1 = 2
    [] V2 = 1 then V1 = V2
    [] V2 = 2 then V1 = 0
  end
end
end
end

proc {N2 V1 V2}
  [V1 V2] ::: 0#2
  {FD.minus 2 V1 V2}
end

```

Figure 7.8: Two variants of a 3-valued negation

ables that represent the interpretations of formulas. A propagator procedure operationally defines a relation between the variables it observes. In KIMBA, each propagator procedure implements the relation defined by its associated logical constant, and the mapping of logical constants to their associated propagator procedures is defined by the logics' LDS.

Connectives

A propagator procedure for a binary connective will observe three variables, one for the whole formula and two for the two components of the formula. If any of the values of these variables changes, the values of the other two variables will be modified accordingly, if possible. Otherwise, propagation fails and the variable assignment is rejected.

KIMBA's generic translation of deduction into OZ constraints allows us to design dedicated, optimised propagator procedures for each logic. OZ itself already has an efficient set of propagators for the standard connectives in classical two-valued logic.

Figure 7.8 shows two operationally equivalent propagator procedures for negation, N1 and N2, in a 3-valued logic. The parameters of the procedures are two finite-domain integer variables. Variable V1 corresponds to the truth value of the formula in the scope of the negation, while V2 is used for the truth value of the whole formula.

Procedure **N1** first restricts its two parameters to values between 0 (“false”) and 2 (“true”), where 1 is used for the truth value “undefined”. Then it starts a concurrent process which waits for one of these parameters to be determined and determines the other one accordingly. Propagation works in both directions, so whenever either the truth value of the formula in the scope of the negation or the the truth value of the whole formula becomes determinate, then so becomes the other value.

Procedure **N2** implements the same behaviour, simply by constraining the values of V_1 and V_2 by the equation $V_2 = 2 - V_1$. Experiments show that the “symbolic” propagation used in **N1** is far more effective than the “arithmetic” propagation in **N2** as long as the number of truth values in the logic considered is relatively small.

Monadic Quantifiers

Monadic quantifiers Q such as \exists, \forall etc., are used for formulating the properties of single sets. In **KIMBA**, quantified formulas $F = Q(P_{\alpha \rightarrow o})$ are translated into constraints over the instantiations $P(\mathcal{C}_\alpha)$ that are possible using a finite domain \mathcal{C}_α of constant symbols of type α . For classical logics, **OZ** provides some built-in propagators for **constraint reification** that come handy for implementing these quantifiers. Figure 7.9 shows two procedures that implement existential and uniqueness quantification. The truth conditions implemented by the propagator procedures can be formalised as follows.

$$(4.1) \quad \llbracket \mathit{exists}(P_{\alpha \rightarrow o}) \rrbracket_{\mathcal{I}} = 1 \text{ iff } \sum_{i=1}^n \llbracket P(X_i) \rrbracket_{\mathcal{I}} \geq 1 \text{ for } X_i \in \mathcal{C}_\alpha$$

$$(4.2) \quad \llbracket \mathit{unique}(P_{\alpha \rightarrow o}) \rrbracket_{\mathcal{I}} = 1 \text{ iff } \sum_{i=1}^n \llbracket P(X_i) \rrbracket_{\mathcal{I}} = 1 \text{ for } X_i \in \mathcal{C}_\alpha$$

The propagator procedures have two parameters, **Var** for the truth value of the whole quantified formula, and a list of truth variables **Insts** for the instantiations $P(X_i)$ over the current domain. **OZ** only supports reification for 2-valued classical logic, so all variables considered are 0/1-integer variables.

The procedures each create a reified sum constraint over the list of variables in **Insts**. If the sum constraint becomes determined, i.e., when it is known whether it is satisfied or violated, the reification constraint determines the truth value **Var** as well. If, on the other hand, **Var** becomes determined, then the sum over **Insts** is constrained accordingly.

OZ does not provide constraint reification into many-valued truth variables and in the case of many-valued logics, the semantics of the reified

```

proc {Exists Var [Insts]}
  {FD.reified.sum Insts '>:' 0 Var}
end

proc {Unique Var [Insts]}
  {FD.reified.sum Insts '=: ' 1 Var}
end

```

Figure 7.9: Implementing quantification as constraint reification

constraints must be programmed in the form of a more complex case analysis. Figure 7.10 shows the implementation of the existential quantification in some 3-valued logic.

Diadic Quantifiers

Diadic quantifiers (also known as generalised determiners) are the logical constants that formulate relations between two sets. The difference between the implementation of a monadic quantifier and that of a diadic quantifier is an additional parameter for the propagator procedure that consists of a list of truth variables for the instantiations of a second set. Figure 7.11 exemplifies how diadic quantifiers can be defined in general. The procedure `More` defines the operational semantics of the diadic quantifier `MORE`. The truth values for two instantiations over the domain are given as lists `InstsA` and `InstsB`. The propagator procedure reifies the constraint that the sum of the truth values in `InstsA` must be greater than the sum of the values in `InstsB`. Hence, the truth value of a formula `MORE(P)(Q)` is determined by the two sums of the instantiations $P(\Sigma_\alpha)$ and $Q(\Sigma_\alpha)$.

The diadic quantifier `NO` is defined as follows: the formula `NO(P)(Q)` is true iff $P(x) \wedge Q(x)$ is false for all x of the appropriate type. The implementation uses the constraint that the sum of the truth values of all conjuncts $P(x) \wedge Q(x)$ must be 0 iff `NO(P)(Q)` is true. Again, we use constraint reification as an economical way to express the relationship between the truth value of a formula and the instantiations of its parts. The diadic quantifier `SOME` is defined almost exactly as `NO`. Here, this sum over the conjuncts must be bigger than 0 if the top-level formula is to be evaluated to true.

```
proc {Exist3 Var [Insts]}
  thread
    cond Var = 2 then {FD.atLeast 1 Insts 2}
    [] Var = 0 then Insts ::: 0
    [] Var = 1
    then
      {FD.exactly 0 Insts 2}
      {FD.atLeast 1 Insts 1}
    [] Insts ::: 0
    then Var = 0
    [] {FD.exactly 0 Insts 2}
      {FD.atLeast 1 Insts 1}
    then Var = 1
    [] {FD.atLeast 1 Insts 2}
    then Var = 2
    end
  end
end
```

Figure 7.10: Implementing a 3-valued existential quantification

```
proc {More Var [InstsA InstsB]}
  {FD.reified.sum InstsA '>:'
    {FoldL InstsB fun {$ I Z} {FD.plus I Z} end 0}
  Var}
end

proc {No Var [InstsA InstsB]}
  {FD.reified.sum
    {List.zip InstsA InstsB
      fun {$ IA IB} {FD.conj IA IB} end}
    '=: ' 0 Var}
end

proc {Some Var [InstsA InstsB]}
  {FD.reified.sum
    {List.zip InstsA InstsB
      fun {$ IA IB} {FD.conj IA IB} end}
    '>:' 0 Var}
end
```

Figure 7.11: Some implementations of diadic quantifiers

```

proc {Exactly Var [Insts Card]}
  {FD.reified.sum Insts '=' Card Var}
end

proc {AtLeast Var [Insts Card]}
  {FD.reified.sum Insts '>=' Card Var}
end

proc {AtMost Var [Insts Card]}
  {FD.reified.sum Insts '<=' Card Var}
end

```

Figure 7.12: Propagator procedures for cardinality quantifiers

KIMBA has a special group of diadic quantifiers that are called **cardinality quantifiers** (see Chapter 3). Instead of defining the relation of two sets, a cardinality quantifier imposes a constraint on one set and its cardinality. For classical logics, KIMBA implements the three cardinality quantifiers `ATLEAST`, `ATMOST` and `EXACTLY` using constraint reification. The implementation is shown in Figure 7.12.

7.4.3 The Translation

The procedure `expand` shown in Figure 7.13 implements the translation from formulas `F` into constraints. The procedure is defined as a method that is inherited to all of KIMBA's proof engines (cf. Section 7.5). The main procedure of KIMBA simply follows the syntactic structure of an input formula. The input parameters are a formula `F` given in KIMBA's syntax, and a root variable `Root` that is the finite-domain integer variable that is associated with the interpretation of `F`. The procedure then splits over the cases that are possible: `F` could be an equality atom, a truth constant, a formula with an unary connective, a formula with a binary connective, a formula dominated by some quantifier, or an atom. In the case of complex formulas, the method `expand` must initiate a further translation of the components and instantiations. In any case, the semantics of the logical constants are all look up in the currently used logic definition structure `self.logic`.

The translation recursively generates a number of propagators by calling

the propagator procedures of the logical constants. Some input problems create several thousand concurrent processes that implement the relations between the interpretations of the formulas and their components or instantiations. Propagation is concurrent, so the process of determining variables can partly take place during the translation. An unsatisfiable set of constraints can sometimes be detected before any search has been done. After the translation has been finished and the initial propagation has been completed, KIMBA starts the actual search for models where undetermined variables are provisionally restricted further, and the next iteration of propagation starts.

7.5 Proof Engines and Controlling Search

KIMBA is a modular system that has been built in an object oriented way. It consists mainly of a uniform method `expand` for translating the specification using the logic definition structures, and a variety of classes of **proof engines** that implement different algorithms for searching solutions within the search space defined by the constraint satisfaction problems.

7.5.1 Proof Engines

The generic proof engine intuitively prepares the enumeration of models for a given input problem as follows. First, it constructs an initial constraint frame from the problem specification which is a minimal frame of domains for the constant symbols that occur in the problem. This implies for instance that the domain of first-order constants is not empty and contains at least one constant `C1`.

Then, the proof engine applies the the translation method `expand` to all input formulas and constrains all their root variables to the maximum truth value defined by the input's logic. The result is a set of finite-domain integer variables that correspond to the interpretation of the atoms defined by the input and the current constant frame.

In the so-called *First-Order* proof engines, the search space is potentially infinite. If it is not possible to assign to each formula a unique integer value such that the result defines a model, then the input is unsatisfiable over the current constant frame and the first-order proof engine restarts model search by extending the current domain of first-order individuals with a newly

```

meth expand(F Root)
  case F
  of ['=' A B]                                     %% F is equality
  then {self.logic.constants.'='.2 A B Root}
  elseif H|Rs                                       %% F is application
  then
    if {HasFeature self.logic.constants H} %% F is complex
    then
      HType = self.logic.constants.H.1
      Propagator = self.logic.constants.H.2
    in
      case HType of o then                          %% H is truth const
        Root = Propagator
      elseif [o o]                                  %% H is unary
      then C in
        {Propagator C Root}
        {self expand(Rs.1 C)}
      elseif [o o o]                                %% H is binary
      then Ca Cb in
        {Propagator Ca Cb Root}
        {self expand({Nth F 2} Ca)}
        {self expand({Nth F 3} Cb)}
      else                                          %% H is quantifier
        {Propagator                                %% or predicate in LDS
        {Map Rs
        fun {$ I} {self termToDenotation(I $)} end} Root}
        end
      else                                          %% F is complex atom
        {self storeAtomic(F Root)}
      end
    else                                          %% F is symbolic
      {self storeAtomic(F Root)}
    end
  end
end

```

Figure 7.13: KIMBA's translation from formulas to constraints

generated constant and by building up a new constraint tableau.

The following are the proof engines that are available in the current implementation.

Propositional: An engine that leaves out the iterative extension of the first-order part of the constant frame. This engine's main application is the search for models within a given domain of individuals, as is for instance usual in puzzle applications.

Minimal: A variant of the **Propositional** engine that rejects models which are not subset-minimal. The engine first searches for propositional models in the usual way and eliminates non-minimal models by proving subset-minimality (see Section 3.5) with the standard **Propositional** engine.

First-Order: The **Propositional** engine with a mechanism for iterative deepening over the universe of discourse. While the engine's results can be controlled by branch-and-bound search, its main application is to enumerate the finite models of an input over an increasing first-order domain.

Local Minimal: This first-order engine implements local minimality, i.e., combines the **First-Order** engine and the subset-minimality constraint as implemented by the **Minimal** engine. The main application of this engine have been the analyses of definites as presented in Chapter 4.

Predicate Minimal: The minimisation of a certain predicate *ab*, in KIMBA usually written as $\$R$, is an essential part of conservative minimality. The present engine is a sub-engine that only produces *ab*-minimal solutions for the current translation.

Conservative: This is a first-order engine for conservative minimality as presented in Section 5.3.3. It is based on a two stage computation that first determines a *ab*-minimal model using the **Predicate Minimal** engine. The models are enumerated as in the engine **First-Order**, but every model produced must (a) have a minimal *ab*-index and (b) be subset-minimal with respect to all other *ab*-minimal models in the current constant frame.

```

class Minimal from Propositional
  meth prove
    self.population =: {Length @universe}
    {self.expandAll}
    choice
      % propagare weights; we ignore ab-index
      {FD.sum @bools '=' self.weight}
      choice
        {FD.distribute generic(order: nbSusps value: min)
         {List.toTuple vars @bools}}
      end
    end
    % refuting non-minimal solution
    {RejectNonMinimal self}
  end
end
end

```

Figure 7.14: The Minimal engine

As a first example, we show the implementation of the class for `Minimal` engines in Figure 7.14. The class inherits from the standard `Propositional` class of engines. It only replaces the top-level proof method `prove`. The `prove` method first constraints the first-order domain size to the size of the current universe of individual constants. By this, it prevents any iterative deepening that is inherited from the generic engine. Next, the method `expandAll` expands all formulas and initiates the first stage of propagation. Next, the **atom weight**, i.e., the number of atoms evaluated to true, is constrained to the sum of the truth values of all atoms in the Herbrand base `bool`. The `FD.distribute` procedure then starts distributing and propagating, controlled by an external search engine. The search process stops when either no model could be found or the set of truth values in `bool` could be determined consistently.

The difference between the `Minimal` engine and the standard propositional engine is the final call to the procedure `RejectNonMinimal`. This procedure verifies that the currently computed model is a locally minimal one. `RejectNonMinimal` proves that there is no other model within the same domain size that validates the input specification by using a real subset of the atoms validated by the current model. This proof problem is a proposi-

tional one, and `RejectNonMinimal` simply uses the standard `Propositional` engine for deciding this problem. If the current model is not minimal, then the procedure signals an error and the current model is rejected.

Figure 7.14 shows the first-order engine `Conservative`. It includes the method `computeLowestBound` that implements the first stage of the computation, namely the identification of the minimum of *ab*-predicates that are possible in a models of the current constant frame. The main method `prove` then basically does the same as the `Minimal` engine, except that the search space is opened up to the addition of new individual constants into the current constant frame.

7.5.2 Search

The solutions for a constraint problem can be computed by applying one of OZ's built-in encapsulated search procedures to the constraint problem generated by the proof engine. The search procedures available are for instance depth-first, branch-and-bound, and visual search using the `EXPLORER` tool. The `EXPLORER` visualises the search spaces as trees; we have used such search trees to visualise the differences in methods for minimal model computation in Section 5.2.4.

The search for models is influenced by the way in which the distribution strategy selects variables for distribution. As a general heuristic guide, `KIMBA` restricts first those variables whose value affects as many other variables and constraints as possible. It also minimises the number of positive literals that are validated in the model at the same time. The process of variable distribution and propagation is repeated until all integer variables have a unique value.

If we use a depth-first left-to-right search strategy, then each first-order engine will always find those models first whose domain of individuals is minimal. Additionally, search can be bounded by the number of literals that are validated. This implements a strong heuristic for minimising the assumptions that are made in a model. By combining this restriction with domain minimality, `KIMBA` can be used for efficiently computing one domain minimal models of the input whose presentation as a finite set of positive literals is as short as possible.

`KIMBA`'s standard search is based on a branch-and-bound approach where properties of earlier computed solutions are used as upper bounds for the rest of the search space. For instance, the search can be bounded such that each

```

class Conservative from PredicateMinimal
  meth computeLowestBound
    proc {SearchAux S}
      S = {New PredicateMinimal init}
      {CopyEngine S self}
      {S prove}
    end
    %% search for ab-minimal, domain-minimal model
    Best = {Search.base.best SearchAux
      proc {$ Ea Eb}
        Ea.cost >: Eb.cost
        Ea.population >=: Eb.population
      end}
  in
    %% constrain current ab-index to the best possible
    if Best == nil then fail
    else self.cost =: Best.1.cost end
  end

  meth prove
    choice
      self.population =: {Length @universe}
      {self computeLowestBound} %% ab-minimal model
      {self expandAll}
      choice %% the index are the 'costly' atoms
        {FD.sum @bools '=' self.weight}
        {FD.sum @costly '=' self.cost}
      choice
        {FD.distribute generic(order: nbSusps value: min)
          {List.toTuple vars @bools}}
      end
      {RejectNonMinimal self PredicateMinimal}
    end
    [] %% first-order iterative deepening
    {self addIndividual} {self prove}
  end
end
end
end

```

Figure 7.15: The Conservative engine

model computed must not use more individuals than a previously computed model. This bound effectively excludes models that are not domain minimal. The search can be restricted by a combination of the following values:

- the number of atoms that are validated within a model
- the number of individual constants in the universe of discourse
- the *ab*-index of the model

Additionally, the bound can be defined as total, i.e., each model must have a smaller bound than all previous ones, or as partial, where a sequence of models can have the same bounds. Totally bounded search restricts the search more quickly and produces the model with the lowest bound faster. However, the number of models that are produced by the enumeration is also reduced greatly, and we might miss models with interesting properties.

To be more specific, Figure 7.16 shows `OneModel`, the generic best-solution search used in KIMBA. The procedure has three arguments, namely `Problem` that is the input problem's specification, `Engine` which names the proof engine to be used, and `WeightP`, the constraint procedure which specifies the branch-and-bound search. The one-model search simply applies the encapsulated branch-and-bound best-solution search that is built into OZ to a procedure `SearchAux`. This auxiliary procedure simply creates a proof engine using both the problem specification and the weight constraint, and then applies the `prove` method to this new engine. `WeightP` may be for instance the procedure `MinimalAtomWeight` which makes sure that the best model found by `OneModel` is one that has a minimal atom weight—the procedure constrains the weight of each successor solution `Eb` to be smaller than that of an earlier solution `Ea`. While `Search.base.best` explores the search space in a depth-first, left-to-right manner, the bound given by `MinimalAtomWeight` makes sure that each model found has a lower atom weight than all earlier ones. The last model found therefore must be the best according to the given bound.

7.6 System Performance

KIMBA is a research prototype whose main design goal has been simplicity rather than speed. Nevertheless, it is interesting to compare the system's

```

fun {OneModel Problem Engine WeightP}
  proc {SearchAux S}
    S = {New Engine create(Problem WeightP)}
    {S prove} %% that's it folks!
  end
in
  {Search.base.best SearchAux WeightP}
end

proc {MinimalAtomWeight Ea Eb}
  Ea.weight >: Eb.weight
end

```

Figure 7.16: Search for one model, and a branch-and-bound procedure.

performance with that of automated reasoning systems that can perform similar tasks. KIMBA's intended application in linguistics has been to determine preferred interpretations with specific minimal model properties, a task where it has no real competitors because there are no systems around that implement the forms of minimality that we found is needed. Nevertheless, KIMBA is also a general finite model generator that can for instance be used as a refutation proof procedure for propositional logics as well. In the following, we will discuss KIMBA's performance and weak points.

7.6.1 Identifying Single Solutions

The table below shows KIMBA's performance on some selected problems from the TPTP's puzzle domain [86]. Times were taken on a Sparc Ultra 1.

Problem	PUZ001-1	PUZ005-1	PUZ017-1	PUZ031-1 (8)
Time	0.3s	4.6s	0.7s	1.6s

The original TPTP specifications are first-order clause sets that have been reformalised for KIMBA using first- and higher-order specifications in classical 2-valued logic. The TPTP formalisations of the logical puzzles are unsatisfiable and can be solved by applying first-order refutation procedures. In

contrast to this, the KIMBA formalisation for each puzzle is satisfiable, and each model produced corresponds to a solution.

Logical puzzles, such as the example PUZ017-1 given in full in Appendix A.1, can often be solved easily by model generators that are based on constraint solving such as FINDER [82] or KIMBA. The rules given in a puzzle’s logic specification are translated in very effective constraints on the combinatorial possibilities. PUZ017-1 is hard for most theorem provers because its first-order clause representations leads to a large search space. KIMBA’s exceptionally good performance can be explained by an elegant higher-order formalisation. For instance, the specification *every job is held by at least one person* in PUZ017-1 has a natural formalisation in KIMBA as follows:

$$\text{EVERY}(\lambda J_{(t \rightarrow \circ)} \text{job}(J))(\lambda J_{(t \rightarrow \circ)} \text{ATLEAST}(J)(1))$$

With this form of quantification, KIMBA keeps entities of different type separately. In a standard first-order form, all quantified entities must be individuals. This usually leads to a larger than necessary domain for which model generation and theorem proving becomes exponentially more complex. A similar kind of improvement is possible in conventional deduction systems when using sorted logics. The problem mentioned above presents no problem for other constraint-based model generators which can make use of a sorted formulation (i.e., FINDER and SEM). FINDER can solve PUZ017-1 in about 0.05s.

Puzzles are typical examples of **single-solution problems** for model generators, i.e., combinatorial problems where one arbitrary model suffices as a result. A class of single-solution problems where finite model generators have traditionally been very successful are quasi-group existence problems. A **quasi-group** is a set with a binary operation \circ which satisfies unique solvability of equations. That is, for all a, b there exist unique c, d such that $c \circ a = b$ and $a \circ d = b$. A quasi-group existence problem [83] is finding an n by n multiplication table over the elements $\{1, \dots, n\}$ such that the table satisfies the properties of a quasi-group and some additional constraints.

For instance, the “QG5” quasi-group problem is finding an table that satisfies the equation $((b \circ a) \circ b) \circ b = a$ for all elements a and b . Furthermore, the table must also be idempotent, that is, that $a \circ a = a$ for all a . The tables, if they exist, are Latin Squares in that there are no repeated entries in any row or column. Quasi-group existence problems are characterised by very large search spaces with very few solutions.

As interesting open questions about quasi-groups can be answered by finite model generators, many systems support the solution of quasi-groups by providing efficient data-structures that have been tailored especially for this task. For instance, `FINDER` internally represents a quasi-group as a multiplication table of integers with a built-in Latin-square property. The current implementation of `KIMBA` has not been optimised for such applications and does not employ any low-level optimisations. As result, it is heavily outperformed by other systems, and is not able to construct models for any but the most trivial finite (quasi)group problems.

7.6.2 `KIMBA` as a Propositional Theorem Prover

Cook's pigeon-hole problem, `MSC007-2` from the `TPTP`, is a classical benchmark proof problem for propositional theorem provers. The task is to prove that there can be no way of putting $n + 1$ pigeons into n pigeon holes such that each hole contains at most one pigeon. Given a set of $n + 1$ pigeons which all are individuals, and a set of n holes which are first-order properties denoting the set of pigeons they contain, the formalisation is as follows.

$$(6.1) \quad \forall x_i \exists P_{i \rightarrow o} \text{hole}(P) \wedge P(x)$$

$$(6.2) \quad \forall P_{i \rightarrow o} \text{hole}(P) \Rightarrow \text{ATMOST}(P)(1)$$

Formula (6.1) states that every individual, i.e., every pigeon, must have a property P that is a hole, i.e., must be in a hole. Formula (6.2) then formalises that every hole can have at most one member, i.e., there can be at most one individual per hole. Of course, with a fixed set of pigeons and holes, the two formulas actually are only a compact higher-order specification of an unsatisfiable set of propositional formulas. For `KIMBA`, the task is to prove that the two formulas interpreted over a given set n of holes and $n + 1$ pigeons are unsatisfiable. As the `Propositional` engine completely enumerates the finite models over the initial constant frame, it can be used to perform this task.

The following table shows the performance of `KIMBA` on problem instantiations with four to eight holes. It also gives the number of atoms whose interpretations must be determined, and the number of branching nodes in the search space defined by the problem.

Holes	4	5	6	7	8
Time	< 0.05s	0.18s	1.30s	9.80s	85.00s
Atoms	24	35	48	63	80
Nodes	23	119	719	5039	40319

The size of the search space, as indicated by the branching nodes, grows exponentially with the size of atoms and results in an exponentially larger proof time. The naive way in which KIMBA explores the whole search space is not optimal, which implies that KIMBA in general is not well suited for the task of proving the non-existence of models. In particular, the finite model generator SEM is more than two orders of magnitude faster than KIMBA for the pigeon-hole problems. The reason for this gap in efficiency lies in the fact that KIMBA has no mechanism for suppressing **isomorphic models**. The actual order of pigeons in the holes does not matter, as any permutation of $n + 1$ pigeons cannot be put into n holes. Yet, a different permutation results in a different interpretation, and all of these are investigated and refuted by KIMBA. This is relevant for many practical applications, as the same complexity problem frequently occurs whenever a problem specifications allows the permutation of elements with respect to certain properties.

The suppression of isomorphic models as proposed by Zhang and Zhang [93] probably is the most promising single optimisation that is possible in finite model generation. A model is isomorphic if it is equal to another one modulo a renaming of the constant symbols. The SEM system uses an effective heuristic that avoids the generation of isomorphic models without losing completeness. However, it is not clear how the approach can be implemented in KIMBA because it relies on properties of clause sets that cannot be tested within a system that does not use clause normalisation. The pigeon-hole problems require an exhaustive search in a very large search space that almost entirely consists of isomorphic interpretations. Because of isomorphic models, KIMBA also is not efficient enough for proving the non-existence of certain quasi-groups. Our translation of deduction into constraint solving does not effectively restrict large search spaces for refutation proof problems.

7.6.3 Generating Minimal Models

The property of being a subset-minimal model can only be decided by a reasoning process that must consider a substantial number of interpretations regardless of what is already known about the current one. Hence, the generation of subset-minimal models is in general one step up the complexity

hierarchy in comparison to the generation of arbitrary models [73]. Local minimality and conservative minimality both suffer from this problem. In order to verify the minimality constraint, we must initiate a proof process that traverses a potentially very large search space. In the examples for definites given in Chapter 4, the space of interpretations usually was small and thus the effect hardly noticeable—in all examples, the preferred models can be determined in under a second on a conventional Pentium II 266Mhz PC. The interpretation of reciprocal sentences shows more clearly the relations between the search space of interpretations, the number of conservative minimal models, and the cost of determining a single arbitrary model.

Example	Models	Conservative	First	All
Pirates	26/51	23	0.12s	0.45s
Pitchers	24/121	24	1.66s	12.48s
Measles	148/2305	64	5.93s	41.47s
Like	432/863	1	0.12s	0.36s

The two numbers in the row “Models” give the number of valid interpretations that the logical form has and the number of nodes the search space for the problem has. Thus, for the “Pirates” example from page 109, we have 26 models in a search space with 51 nodes, i.e., branchings. The row “Conservative” gives the number of conservative models, which is 23 for the “Pirates” example. As we can see, conservative minimality has eliminated $26 - 23 = 3$ models which are not conservative minimal. The time for determining the first conservative minimal model is 0.12s, while the time for enumerating all conservative models is 0.45s. In the Boston “Pitchers” example (cf. page 129), we have 24 valid interpretations and 121 nodes. All models of the logical form also are conservative minimal. The larger search space results in an increased time for finding the first conservative minimal model—the proof procedure must verify for each model whether it is also locally minimal with respect to all other *ab*-minimal models. While the search space in comparison to the “Pirates” example has not even tripled, the time for determining the first conservative model has increased by an order of magnitude. This is the worst case which may not occur in all cases where the search space of interpretations grows larger. In the “Measles” examples (cf. page 130), the search space as a whole is 19 times larger than in the “Pitchers” example with a similar logical encoding. Yet, the time complexity both for determining a first conservative model and all conservative models

is below one order of magnitude. In some cases, a larger space of interpretations may even play no rôle at all. In the “Like” example (cf. page 125) that exemplifies SR reciprocity, we have 432 models in a search space with 863 nodes. However, the logical encoding has only one conservative model which is also very easy to determine, as it is also the first one that the model generator can find in the whole search space. The *ab*-minimality of this model effectively eliminates the major part of the search space. As it takes only 0.12 seconds to find the model and to identify it as the *ab*-minimal one, the remaining search space can be traversed in about 0.36 seconds, which is even less than in the very simple “Pirates” example. Thus, we can conclude that the minimal model constraints on the interpretations may be able to restrict the search space effectively, even though it has been shown that (subset-)minimal model generation may in cases lead to an exponential worst-case complexity. At the same time, the worst case can be testified for some examples. As KIMBA completely enumerates the conservative minimal models, the effect of the combinatorial explosion caused by a larger universe of discourse is plainly visible. The following tables shows the “Pirates” example with three, four, and five pirates who stare at each other.

Example	Models	Conservative	First	All
Pirates (3)	26/51	23	0.12s	0.45s
Pirates (4)	281/561	81	0.27s	4.65s
Pirates (5)	3494/6993	1024	1.99s	87.00s

As we can see, the number of interpretations, the search space, and the number of conservative models explodes with the size of the reciprocal set. This is a direct result of the way in which we interpret reciprocity, as the linguistically valid interpretations must involve permutations of elements with respect to the reciprocal relation. For practical applications, the increase in time for determining a first model may be especially worrisome, as this can be taken as the instance of a preferred linguistic interpretation. A possible counter-measure for the effect would be the elimination of isomorphic models as mentioned in Section 7.6.2, a task that cannot be performed by KIMBA in its current form. On the other hand, the theoretical properties of subset-minimality make it impossible to avoid this effect in general.

Chapter 8

Conclusion

You have always been a cunning linguist, Bond.
(Money Penny, *Tomorrow never dies*)

The primary hypothesis of computational logic-based semantics is that logic can be used to capture the meaning of natural language. A logic consist of a formal language, i.e., a syntax, and a semantics that maps the expressions of this language into some domain and finally to truth conditions. To define a suitable logical representation for natural language is a worthwhile pursuit on its own right, as such a representation permits us to investigate natural language on a formal, unambiguous level. Works in theoretical logic-based semantics, such as Dalrymple et al.'s research on reciprocity, use a logical device to formally describe some properties of the semantics of natural language. Without inference, however, there is no *computational* logic-based semantics.

8.1 Why Inference is Worth the Effort

If logic can in fact capture meaning, then there is all reason to suspect that inferences on the logical form can be used to model the process of natural language understanding. It is conceivable that the way in which humans interpret language is determined to a very large degree by some system of rules, that is, is systematic. Otherwise, we would not even be able to detect when the process of communication completely fails.

The research documented in the present work focuses on inference rather

than on representation. The starting point was the observation that certain models of a semantic representation and a context consisting of situational and world knowledge can be taken as the meaning of the represented natural-language sentence. Our early analyses were not based on some of the state-of-the-art semantic representation languages that have been designed especially for such phenomena as pronoun resolution or multi-level ambiguities. Instead, we started out with plain first-order predicate logics because these come with a plethora of automatic reasoning systems for model generation and deduction.

Finite model generation is an interesting research tool in computational semantics as it gives us, for the first time, the ability to enumerate a large class of models of the input. This enumeration often makes explicit some subtle errors in semantic representations. It also make clear that semantic representation without contextual knowledge are near to worthless for the purpose of interpretation. There is almost no example where the models of a logical form alone correspond to interpretations that are acceptable. Logic and inference require world knowledge reasoning as the third aspect of a computational logic-based form of natural-language interpretation. Fortunately, common-sense reasoning can often be performed as part of the process of computing a model, and many faulty interpretations can be turned into correct ones by adding the missing knowledge.

Automated reasoning systems are now at a level of performance that allows their application in the real world. Even a simple system like KIMBA, which is based on a lean approach to inference and that is not primarily built with efficiency in mind, profits greatly from the improvements in constraint-solving technology in recent years. Our higher-order logical language *MQC* evolved from experimentation with what is still tractable within the limits of a translation of logic into finite-domain integer constraints. The language *MQC* as well as the KIMBA system owes much of its appearance to the fact that we started out with a standard approach of inference based on a classical, well-understood form of logic. One of the essential results of the efforts presented here is that the use of efficient inference techniques for classical logics is at least as much an advantage as the logic's restricted expressivity may appear an annoyance to semanticists. The sacrifices that one must bring when using conventional logics or restricted knowledge representation languages are often paid for by the ease of access and the ability for experimentation that the existing well-designed systems provide.

Nevertheless, it is not surprising that our experiments sometimes reveal

that a more expressive logical language is desirable not only for theoretical reasons. Some problems on the computational level can be solved without question only by a better representation. As an example, we refer to the improved predictions that a formalisation of unicity gives for singular definite descriptions. The present work in any case avoids topics such as dynamic anaphoric binding or scope ambiguities, as we believe that the questions raised by these phenomena cannot be answered adequately in the formal framework of classical logics that we have chosen. Still, it can be argued that one of the main purposes of a formal representation should be to allow for a simple and possibly tractable form of inference. While a formal language that, for instance, can represent structural ambiguities of quantifier scoping within the logical form may be theoretically attractive, it will be of little use in practice if the representation itself makes the desired forms of inference impracticable. We therefore conclude that both levels of semantics, representation and inference, should no longer be discussed separately.

8.2 Contributions

The present work argues that model generation is a form of reasoning that can be used to conceptualise some inference processes in natural-language understanding. That is, we may use model generation as a formal model of inference processes whose actual forms remain inaccessible to us.

The international conference *Inference in Computational Semantics*, first held in 1999, showed that there is a growing interest in the computational linguistics community to discuss such forms of inference and the computational tools that approximate them. Inference is one of the keys to natural-language understanding, and a formal method that models even small aspects of the interpretation of language therefore may give us new insights in the way in which language works.

The guinea pigs that we have investigated in detail in our linguistic laboratory are singular definite descriptions and reciprocal sentences. Definites were the starting point for a great deal of research in modern computational semantics. Our approach of interpretation as model generation concentrates on the inference tasks that originate in the problem of resolving singular definites in a logically encoded context. Our experiments show that model generation is able to deal with some of the stumbling stones of other approaches, as model generation embeds reasoning about linguistic knowledge

into reasoning about the situational context and the world. Unlike as in deductive approaches, we do not have to give some set of entities that may be candidate referents. Domain-minimal model generation implements, in a very natural way, the preference of anaphoric resolution over bridging over accommodation.

Not all questions in connection with definites can be answered by semantics and model generation alone. Definite descriptions show a surprisingly complex behaviour, and the actual truth conditions of a definite can not easily be determined. With the introduction of the identifying property, our theory is no longer a theory of *computational* semantics as we cannot give a formal theory of how identifying properties can be determined.

In contrast to this, the meaning of reciprocal sentences seems to be governed by one principle of interpretation. This principle, the MMH, can be modelled by a certain form of minimal model generation. Our analysis is able to identify the correct interpretations for simple reciprocal sentences within large sets of model candidates. As far as we know, there are no other approaches that can deal equally well with reciprocals. The semantic analysis on the level of interpretations gives us a “continuum of meaning” in contrast to some idiosyncratic class system of reciprocal semantics. In practice, the model generation approach is a convenient means to experiment with the linguistic theory. Theoretically, our treatment of reciprocity is appealing as it requires only one semantic representation for reciprocal expressions and does not treat them as n -way ambiguous constituents. However, as in the case of definites, the present work is cannot be seen as a thorough analysis of the linguistic phenomenon. For instance, a compositional treatment of the interaction of reciprocal expressions and quantifiers requires a different logical representation than the one that we have given. In any way, the maximisation of logical contribution on the level of the models of the logical form currently has no computational alternative, and we believe that model generation is an adequate formal model for the process of reciprocal interpretation.

The reasoning that humans perform when interpreting everyday language is that of reasoning on small and finite sets. Standard first-order predicate logic has its problems with this kind of reasoning, as many properties in connection with finiteness are not first-order expressible, at least not without further technical machinery in the encoding. We have found finite model generation to be appealing because many first-order “inexpressible” truth conditions of natural-language quantifiers can be expressed as computable

constraints over finite models. If we interpret the semantic representations of Montague-style higher-order logical semantics only over finite domains, we obtain an interesting fragment of higher-order logics where standard first-order model generation techniques can be used. This fragment inherits the elegant compositional way in which expressions can be composed and defined in the λ -calculus. The language \mathcal{MQL} can be taken as a first model of a formal language that combines desirable properties from higher-order and first-order logics. KIMBA exemplifies that the introduction of concepts like β -reduction, higher-order definitions, or restricted higher-order quantification does not make inference inherently more complex in the way that full higher-order reasoning often does.

8.3 Models as Meaning

Even with all necessary background knowledge present, finite model generation computes in general many models whose content does not reflect the meaning of the logical form. Human conversation implies several constraints for what is an acceptable interpretation and what is not. Model generation must be extended by methods that implement these conventions. Some handy tools for this are available as results of non-linguistic applications of model generation. First and foremost, the notion of a subset-minimal model is central for eliminating those interpretations which give redundant information that cannot be justified by any evidence at all. Then, for such phenomena as definite descriptions, the minimisation of the universe of discourse is convenient for implementing the preference of anaphor resolution over accommodation. Finally, we have used the minimisation of a certain index predicate to model the maximisation of the logical contribution of the scope relation in reciprocal sentences. The forms of minimisation that we have found useful in natural-language interpretation and analysis highly interact, and we must make sure that the intersection of different minimality constraints does not become too strict to be of any use. As a compromise, we have defined a combination of three minimality constraints called conservative minimality that implements a simultaneous minimisation of individuals, logical assumptions, and index value. The minimisation of the universe can be justified on linguistic grounds, but it is also necessary in order to stay in a computable fragment of model generation. There is in general no way of deciding whether some arbitrary first-order model is subset-minimal, but

local minimality is decidable because of the domain closure of local minimal models. Likewise, the minimal index can be computed for the set of domain minimal finite models, but not for the set of all models.

The simple techniques that we have used to characterise preferred readings do obviously not exactly model the complex inference processes of natural-language understanding. There can be no question that what any form of minimality will give us is only a heuristic to identify semantic interpretations. As all heuristics, it may sometimes fail. Minimal model generation frequently gives us models which are not preferred readings in any way, or even models whose content has no correspondence to the meaning of the represented utterances at all. However, the term “minimality” is open to any reasoning process that empirically yields better results, and we hope that further research will discover new and better methods that characterise which models of the logical form are linguistically interesting.

Appendix A

Some Example Problems

A.1 The Job Puzzle

We consider the logical puzzle PUZ017-1 from the TPTP problem library for automated reasoning systems [86]. Its natural-language formalisation looks like this:

There are four people: Roberta, Thelma, Steve, and Pete. Among them they hold eight different jobs. Each holds exactly two jobs. The jobs are: chef, guard, nurse, telephone operator, police officer (either gender), teacher, actor, and boxer. The job of a nurse is held by a male. The husband of the chef is the telephone operator. Roberta is not a boxer. Pete has no education past the ninth grade. Roberta, the chef, and the police officer went golfing together. Question : Who holds which job?

The problem's formalisation in KIMBA's syntax is given in Figure A.1. Figure A.2 shows the search space of the problem when using the propositional minimal model engine `Minimal` (see Section 7.5.2). The sixteen diamonds represent the minimal models of the problem, while the 53 squares represent failures caused either by some inconsistency or by non-minimality of interpretations. The minimal models prove that the puzzle does not have only one solution, but in fact sixteen. On a Pentium-II/266 workstation, the exhaustive search for all minimal models of the problem takes about 6 seconds.

```

p(remark: 'Who gets the job.'
  logic: classical
  types: are(o:o i:i n:n)
  constants: are(
    roberta:i thelma:i steve:i pete:i
    chef:[o i] guard:[o i] nurse:[o i] operator:[o i]
    officer:[o i] teacher:[o i] actor:[o i]
    boxer:[o i] male:[o i] educated:[o i]
    husband:[o i i] job:[o [o i]])
  definitions: are(female: lam(x#i ['not' ['male' x]])
    the: lam(p#[o i]
      lam(q#[o i]
        ['and' [exactly p 1]
          forall(x#i ['implies' [p x] [q x]]))))))
  formulas:
  [
    [job chef] [job guard] [job nurse] [job operator]
    [job officer] [job teacher] [job actor] [job boxer]
    ['not' [job male]] ['not' [job educated]]
    %% gender
    [male steve] [male pete] [female roberta] [female thelma]
    %% every job is held by someone
    [every lam(j#[o i] [job j]) lam(j#[o i] [atLeast j 1])]
    %% everyone has exactly 2 jobs
    forall(x#i [exactly lam(j#[o i] ['and' [job j] [j x]] 2)])
    %% the nurse is male, the chef is female
    [the nurse male] [the chef female]
    %% the chef has a husband who is the operator
    [the chef lam(x#i [unique lam(y#i ['and' [husband y x]
      [operator y]])))]
    %% husbands are male
    [every lam(x#i exists(y#i [husband x y])) male]
    %% wives are female
    [every lam(x#i exists(y#i [husband y x])) female]
    %% there is atmost one husband per wife
    forall(x#i [atMost lam(y#i [husband y x]) 1])
    %% all actors are educated and roberta is not a boxer
    [every actor educated] ['not' [boxer roberta]]
    %% Pete has no education past the ninth grade.
    ['not' [educated pete]]
    %% nurses, teachers, and officers are educated
    [every nurse educated] [every teacher educated]
    [every officer educated]
    %% the chef is neither a police officer nor roberta
    [no chef officer] ['not' [chef roberta]]])

```

Figure A.1: The puzzle PUZ017-1 from the TPTP library [86]




Figure A.2: The search tree for PUZ017-1 with 16 minimal models.

A.2 Reciprocals: The Boston Pitchers

Figure A.3 shows a logical encoding for the reciprocal sentence *the Boston pitchers sat alongside each other* for a discourse situation with four pitchers: Tom Bolton (tb), Jeff Reardon (jr), Larry Anderson (la), and Jeff Gray (jg).

In the encoding, we use some of KIMBA's built-in predicates for formalising the position relation, for instance the predicate `succ` that corresponds to the successor function on natural numbers. The encoding also includes the definitions `iao`, `price`, and `rcp` for the semantic of reciprocals. The predicate `$R` denotes our index predicate.

A model generation process yields 24 different conservative minimal models, each of which corresponds to one of the $4! = 24$ different interpretations of *the Boston pitchers sit alongside each other* in a discourse situation where we have four Boston pitchers. In our example, we individual constants for the pitchers *Larry Anderson* (la), *Tom Bolton* (tb), *Jeff Reardon* (jr), and *Jeff Gray* (jg). A typical model appears as follows.

$$(2.1) \quad \{ ab(jg, jr), ab(jg, tb), ab(jr, jg), ab(jr, tb), ab(la, tb), \\ ab(tb, jg), ab(tb, jr), ab(tb, la), bospt(jg), bospt(jr), \\ bospt(la), bospt(tb), pos(sital, 1, tb), pos(sital, 2, jr), \\ pos(sital, 3, la), pos(sital, 4, jg), sital(tb, jr), \\ sital(jr, tb), sital(jg, la), sital(jr, la), sital(la, jg), \\ sital(la, jr) \}$$

The situation described by the model is illustrated by the following table.

position	1	2	3	4
pitcher	Tom Bolton	Jeff Reardon	Larry Anderson	Jeff Gray

```

p(logic: 'classical' engine: 'conservative'
  constants: are(
    la:i tb:i jr:i jg:i %% the guys
    sital:[o i i] %% sitting alongside
    bospt:[o i] %% Boston Pitchers
    pos:[o i n [o i i]]
    1:n 2:n 3:n 4:n 5:n 6:n 7:n %% some numbers that we might need
    '$R':[o i i] %% our index relation
  definitions: are(%% Intermediate Alternative Ordering
    iao: lam(p#[o i]
      lam(r#[o i i]
        ['and' [atLeast p 2]
          [every p
            lam(x#i
              exists(y#i
                ['and' ['and' [p y]
                  ['not' ['=' y x]]]
                ['or'
                  [r x y]
                  [r y x]]])])])])
          %% The price of not being in the reciprocal
        price: lam(p#[o i]
          lam(r#[o i i]
            [every p
              lam(x#i
                ['and' ['not' ['$R' x x]]
                  [every lam(y#i ['and' [p y]
                    ['not' ['=' x y]])]
                  lam(y#i ['equiv'
                    ['not' [r x y]
                    ['$R' x y]]])])])])
          %% the semantic od reciprocals
        rcp: lam(p#[o i]
          lam(r#[o i i] ['and' [iao p r] [price p r]]))
    formulas: [
      %% siting alongside
      forall(x#i forall(y#i ['equiv'
        [sital x y]
        exists(p#n
          ['and' [pos sital p x]
            exists(l#n
              ['and'
                ['or' ['succ' p l] ['succ' l p]]
                [pos sital l y]]])])])
      forall(p#n [atMost lam(x#i [pos sital p x]) 1])
      forall(x#i [atMost lam(p#n [pos sital p x]) 1])
      forall(p#n ['implies' exists(y#i [pos sital p y])
        forall(l#n ['implies' ['<' l p]
          exists(x#i [pos sital l x]]))])
      %% 4 boston pitchers
      [bospt la] [bospt tb] [bospt jr] [bospt jg]
      %% the Boston pitchers sit alongside each other
      [rcp bospt sital] ]
    ]
  )

```

Figure A.3: The Boston pitchers sat alongside each other

All conservative minimal models of the input carry an index of 6, and the conservative minimal models are also identical to the *ab*-minimal models of the input.

Bibliography

- [1] S. Abdennadher and H. Schütz. Model generation with existentially quantified variables and constraints. In M. Hanus, J. Heering, and K. Meinke, editors, *Proc., Algebraic and Logic Programming. 6th International Joint Conference, ALP '97 – HOA '97*, number 1298 in LNCS, pages 256–272. Springer, 1997.
- [2] P. B. Andrews. General models descriptions and choice in type theory. *Journal of Symbolic Logic*, 37(2):385–394, 1972.
- [3] P. B. Andrews. *An Introduction to Mathematical Logic and Type Theory: To Truth Through Proof*. Academic Press, 1986.
- [4] N. Asher and H. Wada. A computational account of syntactic, semantic and discourse principles for anaphora resolution. *Journal of Semantics*, 6:309–344, 1988.
- [5] H. P. Barendregt. *The Lambda Calculus*. North Holland, 1984.
- [6] P. Baumgartner. Hyper tableaux — the next generation. In H. de Swart, editor, *Proc., Automated Reasoning with Analytic Tableaux and Related Methods, TABLEAUX 98*, volume 1397 of *LNAI*. Springer, 1998.
- [7] P. Baumgartner, P. Fröhlich, U. Furbach, and W. Nejdl. Tableaux for diagnosis applications. In D. Galmiche, editor, *Proc. TABLEAUX 97*, number 1071 in LNCS, pages 76–90, Pont-a-Mousson, France, May 1997. Springer.
- [8] P. Baumgartner, U. Furbach, and I. Niemelä. Hyper tableaux. In *Logics in AI: European Workshop, JELIA '96*, number 1126 in *LNAI*, 1996.

- [9] P. Baumgartner and M. Kühn. Abductive coreference by model construction. In *ICoS-1 Inference in Computational Semantics*, Institute for Logic, Language and Computation, University of Amsterdam, August 1999.
- [10] B. Beckert and R. Hähnle. Analytic tableaux. In W. Bibel and P. Schmitt, editors, *Automated Deduction: A Basis for Applications*, volume I, chapter 1, pages 11–41. Kluwer, 1998.
- [11] B. Beckert, R. Hähnle, and P. H. Schmitt. The *even more* liberalized δ -rule in free variable semantic tableaux. In G. Gottlob, A. Leitsch, and D. Mundici, editors, *Proceedings of the third Kurt Gödel Colloquium KGC'93, Brno, Czech Republic*, volume 713 of *LNCS*, pages 108–119. Springer, August 1993.
- [12] B. Beckert and J. Posegga. lean^{AP} : Lean tableau-based deduction. *Journal of Automated Reasoning*, 15(3):339–358, 1995.
- [13] C. Benz Müller and M. Kohlhase. Extensional higher order resolution. In C. Kirchner and H. Kirchner, editors, *Proc., 15th International Conference on Automated Deduction (CADE), Lindau, Germany*, number 1421 in Springer LNAI, pages 56–72, 1998.
- [14] C. Benz Müller and M. Kohlhase. LEO, a higher order theorem prover. In C. Kirchner and H. Kirchner, editors, *Proc., 15th Conference on Automated Deduction, CADE 98*, number 1421 in LNAI, pages 139–144, Lindau, Germany, 1998. Springer.
- [15] P. Blackburn, J. Bos, M. Kohlhase, and H. de Nivelle. Theorem proving for natural language understanding. In *Proc., Workshop Problem Solving Methods in Automated Deduction at CADE 98*, 1998.
- [16] R. Boyer, E. Lusk, W. McCune, R. Overbeek, M. Stickel, and L. Wos. Set theory for first-order logic: Clauses for Gödel's axioms. *Journal of Automated Reasoning*, 2:287–327, 1986.
- [17] T. Brüggemann, F. Bry, N. Eisinger, T. Geisler, S. Panne, H. Schütz, S. Torge, and A. Yahya. Satchmo: Minimal model generation and compilation (system description). In J.-L. Imbert, editor, *Proc., Cinquièmes Journées Francophones de Programmation en Logique et Programmation*

- par Contraintes. Prototypes, JFPLC 96.*, pages 9–14, Clermont-Ferrand, June 1996.
- [18] F. Bry, N. Eisinger, H. Schütz, and S. Torge. SIC: Satisfiability checking for integrity constraints. In P. Fraternali, U. Geske, C. Ruiz, and D. Seipel, editors, *Proc. 6th International Workshop on Deductive Databases and Logic Programming, DDLP '98*, GMD Report 22, Manchester, UK, June 1998.
- [19] F. Bry and S. Torge. A deduction method complete for refutation and finite satisfiability. In *Proc. 6th European Workshop on Logics in AI (JELIA)*, LNAI, 1998.
- [20] F. Bry and A. Yahya. Minimal model generation with positive unit hyper-resolution tableaux. In P. Miglioli, U. Moscato, D. Mundici, and M. Ornaghi, editors, *Proc. TABLEAUX'96*, number 1071 in LNCS, pages 143–159, Terrasini, Palermo, Italy, May 1996. Springer.
- [21] B. Carpenter. *Type-Logical Semantics*. The MIT Press, Cambridge, Massachusetts, 1993.
- [22] T. Castell and H. Fargier. Between SAT and CSP: Propositional satisfaction problems and clausal CSPs. In H. Prade, editor, *Proc. 13th European Conference on Artificial Intelligence, Brighton*, pages 214–218. John Wiley & Sons, 1998.
- [23] A. Church. A formulation of the simple theory of types. *Journal of Symbolic Logic*, 5:56–68, 1940.
- [24] A. Collins, J. Brown, and K. Larkin. Inference in text understanding. In R. Spiro, B. Bruce, and W. Brewer, editors, *Theoretical Issues in Language Comprehension*. Lawrence Erlbaum Associates, New Jersey, 1978.
- [25] R. Cooper. The interpretation of pronouns. In F. Heny and H. Schnelle, editors, *Syntax and semantics 10*, pages 61–92. Academic Press, New York, 1979.
- [26] J. Cunningham. Comprehension by model-building as a basis for an expert system. In *Proc., 5th Technical Conference of the British Society*

- Specialist Group on Expert Systems*. University of Warwick, Dec.17–19 1985.
- [27] M. Dalrymple, M. Kanazawa, Y. Kim, S. Mchombo, and S. Peters. Reciprocal expressions and the concept of reciprocity. In P. Barbosa, D. Fox, P. Hagstrom, and M. McGinnis, editors, *Is the Best Good Enough?: Proc. Workshop on Optimality in Syntax*. The MIT Press, 1996.
- [28] M. Davis, G. Logeman, and D. Loveland. A machine program for theorem-proving. *Communications of the Association for Computing Theory*, 5:394–397, 1962.
- [29] M. Davis and H. Putnam. A computing procedure for quantification theory. *Journal of the ACM*, 7(3):201–215, July 1960.
- [30] H. de Nivelle. A resolution decision procedure for the guarded fragment. In C. Kirchner and H. Kirchner, editors, *Proc. CADE-15*, number 1421 in LNAI, 1998.
- [31] M. D. Ernst, T. D. Millstein, and D. S. Weld. Automatic SAT-compilation of planning problems. In *Proc., of the 15th International Joint Conference on Artificial Intelligence, IJCAI 97*, pages 1169–1176, Nagoya, Aichi, Japan, Aug. 1997.
- [32] W. M. Farmer. A partial-function version of Church’s simple theory of types. *Journal of Symbolic Logic*, 55:1269–1291, 1990.
- [33] C. Fermüller, A. Leitsch, T. Tammet, and N. Zamov. *Resolution Methods for the Decision Problem*. LNAI 679. Springer, 1993.
- [34] M. Fitting. *First-Order Logic and Automated Theorem Proving*. Springer, 1990.
- [35] G. Friedrich and W. Nejdl. MOMO—model-based diagnosis for everybody. In *Proceedings 6th IEEE Conference on AI Applications*, pages 206–213, Santa Barbara, CA, USA, 1990.
- [36] M. Fujita, R. Hasegawa, Y. Shirai, and F. Kumeno. Program synthesis by a model generation theorem prover. Technical report, Institute for New Generation Computer Technology, Tokyo, 1991.

- [37] G. Gallo and G. Urbani. Algorithms for testing the satisfiability of propositional formulae. *Journal of Logic Programming*, 7(1):45–62, July 1989.
- [38] C. Gardent, M. Kohlhase, and K. Konrad. Higher-order colored unification: A linguistic application. *Technique es science informatiques*, 18(2), 1999.
- [39] M. Giese and W. Ahrendt. Hilbert’s epsilon-terms in automated theorem proving. In N. Murray, editor, *Proc. TABLEAUX’99*, number 1617 in LNCS, pages 171–185, Saratoga Springs, NY, USA, June 1999. Springer.
- [40] M. L. Ginsberg. A circumscriptive theorem prover. *Artificial Intelligence*, 39:209–230, 1989.
- [41] J. Groenendijk, M. Stokhof, and F. Veltman. Coreference and modality. In *Handbook of Contemporary semantic theory*. Blackwell, Oxford, 1995.
- [42] N. Haddock. *Incremental Semantics and Interactive Syntactic Processing*. PhD thesis, University of Edinburgh, 1989. advisor: M. Steedman.
- [43] F. Harche, J. N. Hooker, and G. L. Thompson. A computational study of satisfiability algorithms for propositional logic. *ORSA Journal on Computing*, 6(4):423–435, 1994.
- [44] R. Hasegawa. Model generation theorem provers and their applications. In L. Sterling, editor, *Proceedings of the 12th International Conference on Logic Programming*, pages 7–8, Cambridge, MA, USA, June 13–18 1995. MIT Press.
- [45] R. Hasegawa, H. Fujita, and M. Koshimura. MGTP: A model generation theorem prover – its advanced features and applications. In P. Baumgartner, R. Hähnle, and J. Posegga, editors, *Proc., 4th International Workshop on Theorem Proving with Analytic Tableaux and Related Methods, TABLEAUX 95*, number 918 in LNCS, pages 1–15. Springer, 1995.
- [46] I. Heim. *The semantics of definites and indefinite noun phrases in English*. PhD thesis, University of Massachusetts, Amherst, 1982.

- [47] L. Henkin. Completeness in the theory of types. *Journal of Symbolic Logic*, 15(2):81–91, 1950.
- [48] A. J. W. Hilton. The reconstruction of latin squares with applications to school timetabling and to experimental design. *Mathematical Programming Study*, 13:68–77, 1980.
- [49] J. Hintikka. Model minimization – an alternative to circumscription. *Journal of Automated Reasoning*, 4(1):1–13, 1988.
- [50] J. Hobbs, M. Stickel, D. Appelt, and P. Martin. Interpretation as abduction. *Artificial Intelligence*, 63:69–142, 1993.
- [51] H. Horacek and K. Konrad. Presenting herband models with linguistically motivated techniques. In *Proc., CIMCA 99*, LNCS, forthcoming, Vienna, Austria, 1999. *fehlt noch!*
- [52] D. Jackson and C. Damon. Elements of style: Analyzing a software design feature with a counterexample detector. In *International Symposium on Software Testing and Analysis*, 1996.
- [53] R. G. Jeroslow and J. Wang. Solving propositional satisfiability problems. *Annals of Mathematics and Artificial Intelligence*, 1:167–187, 1990.
- [54] N. Kadmon. Uniqueness. *Linguistics and Philosophy*, 13:273–324, 1990.
- [55] H. Kamp. A theory of truth and semantic representation. In J. Groenendijk, T. Janssen, and M. Stokhof, editors, *Formal Methods in the Study of Language*, pages 277 – 322. Mathematisch Centrum Tracts, Amsterdam, 1981.
- [56] H. Kamp and U. Reyle. *From Discourse to Logic*. Kluwer, Dordrecht, 1993.
- [57] L. Karttunen. Discourse referents. In J. McCawley, editor, *Syntax and Semantics 7*, pages 363–385. Academic Press, New York, 1976.
- [58] H. Kautz and B. Selman. Planning as satisfiability. In *Proc. ECAI 92*, pages 359–363, Vienna, Austria, 1992.

- [59] H. Kautz and B. Selman. Pushing the envelope: Planning, propositional logic and stochastic search. In *Proc., 13th National Conference on Artificial Intelligence, AAAI 96*, Portland, OR, USA, 1996.
- [60] S. Klingenbeck. *Counter Examples in Semantic Tableaux*. Number 156 in DISKI. infix, 1996. PhD Thesis.
- [61] D. T. Langendoen. The logic of reciprocity. *Linguistic Inquiry*, 9(2):177–197, 1978.
- [62] D. Lewis. Scorekeeping in a language game. In R. Bauerle, U. Egli, and A. van Stechow, editors, *Semantics from different points of view*, pages 172–185. Springer, Berlin, 1979.
- [63] S. Lorenz. A tableau prover for domain minimization. *Journal of Automated Reasoning*, 13:375–390, 1994.
- [64] D. W. Loveland, D. W. Reed, and D. S. Wilson. SATCHMORE: SATCHMO with RElevancy. *Journal of Automated Reasoning*, 14:325–351, 1995.
- [65] R. Manthey and F. Bry. SATCHMO: A theorem prover implemented in Prolog. In *Proc. CADE'88*, pages 415–434, 1988.
- [66] D. McAllester and R. Givan. Natural language syntax and first-order inference. *Artificial Intelligence*, 56, 1992.
- [67] J. McCarthy. Circumscription – a form of non-monotonic reasoning. *Artificial Intelligence*, 13(1-2):27–39, 1980.
- [68] W. McCune. A Davis-Putnam program and its application to finite first-order model search: Quasigroup existence problems. Technical Memorandum ANL/MCS-TM-194, Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL, USA, 1994.
- [69] W. McCune. Automatic proofs and counterexamples for some ortholattice examples. *Information Processing Letters*, 65:285–291, 1998.
- [70] R. Montague. The proper treatment of quantification in ordinary English. In R. Montague, editor, *Formal Philosophy. Selected Papers*. Yale University Press, New Haven, 1974.

- [71] W. Nejdl and P. Fröhlich. Minimal model semantics for diagnosis – techniques and first benchmarks. In *Proc., 7th International Workshop on Principles of Diagnosis*, Val Morin, 1996.
- [72] I. Niemelä. Implementing circumscription using a tableau method. In W. Wahlster, editor, *Proc., European Conference on Artificial Intelligence, ECAI 96*, pages 80–84, Budapest, Hungary, 1996.
- [73] I. Niemelä. A tableau calculus for minimal model reasoning. In P. Miglioli, U. Moscato, D. Mundici, and M. Ornaghi, editors, *Proc. TABLEAUX'96*, number 1071 in LNCS, pages 278–294, Terrasini, Palermo, Italy, 1996. Springer.
- [74] P. F. Patel-Schneider. A decidable first-order logic for knowledge representation. *Journal of Automated Reasoning*, 6:361–388, 1990.
- [75] N. Peltier. *Nouvelles Techniques pour la Construction de Modèles finis ou infinis en Dédution Automatique*. PhD thesis, Institut National Polytechnique de Grenoble, 1997. <ftp://ftp.imag.fr/pub/Mediatheque.IMAG/theses/97-Peltier.Nicolas/>.
- [76] W. Philip. *Children who know each other*. ms., OTS, Utrech University, 1996.
- [77] C. S. Pierce. Abduction and induction. In J. Buchler, editor, *Philosophical Writings of Pierce*, pages 150–156. Dover Books, 1955.
- [78] Programming Systems Lab Saarbrücken, 1998. Oz Webpage: <http://www.ps.uni-sb.de/oz/>.
- [79] B. Russell. On denoting. *Mind*, 14:479–493, 1905.
- [80] H. Salkin and K. Mathur. *Foundations of Integer Programming*. North-Holland, 1989.
- [81] H. Schütz. Generating minimal herbrand models step by step. In N. Murray, editor, *Proc. TABLEAUX'99*, number 1617 in LNCS, pages 171–185, Saratoga Springs, NY, USA, June 1999. Springer.
- [82] J. Slaney. FINDER (Finite Domain Enumerator): Notes and guide. Technical Report TR-ARP-1/92, Australian National University Automated Reasoning Project, Canberra, 1992.

- [83] J. Slaney, M. Fujita, and M. Stickel. Automated reasoning and exhaustive search: Quasigroup existence problems. *Computers and Mathematics with Applications*, 29:115–132, 1995.
- [84] R. M. Smullyan. *First-Order Logic*. Springer, 1968.
- [85] P. F. Strawson. On referring. *Mind*, 59:320–344, 1950.
- [86] C. Suttner and G. Sutcliffe. The TPTP problem library (TPTP v2.2.0). Technical Report 97-04, Department of Computer Science, James Cook University, Townsville, Australia, 1998.
- [87] B. Webber. So what can we talk about now? In M. Brady and R. Berwick, editors, *Computational Models of Discourse*, pages 331–371. MIT Press, Cambridge MA, 1982.
- [88] D. Westerstahl. Determiners and context sets. In J. van Benthem and A. T. Meulen, editors, *Generalised quantifiers in natural language*. Foris, Dordrecht, 1991.
- [89] Y. Winter. What does the strongest meaning hypothesis mean? In *Proc., Sixth Semantics and Linguistic Theory Conference*, 1996.
- [90] H. Zeevat. Presupposition and accomodation in update semantics. *Journal of Semantics*, 9:379–412, 1992.
- [91] H. Zhang and M. E. Stickel. Implementing the Davis-Putnam algorithm by tries. Technical report, Computer Science Department, The University of Iowa, Iowa City, Iowa, USA, Aug. 1994.
- [92] J. Zhang. *The generation and applications of finite models*. Dissertation, Institute of Software, Academia Sinica, Bejiing, 1994.
- [93] J. Zhang and H. Zhang. SEM: A system for enumerating models. In *Proc., International Joint Conference on Artificial Intelligence, IJ-CAI 95*, 1995.

Index

- abduction, 136
- abductive explanation, 137
- accommodation, 83
- allowable hypothesis, 137
- anaphor, 82
- answer, 146
- antecedent group, 106
- atom, 22, 53
- atom weight, 178

- background theory, 137
- base types, 48
- binary logical connectives, 51
- Bliksem, 27
- bound variable, 49

- cardinality quantifiers, 174
- closed λ -term, 49
- combinatorial, 23
- complement, 22
- compositionality, 41
- comprehension axioms, 59
- Conservative Minimality, 127
- constant frame, 60
- constant frame \mathcal{C}_Φ , 60
- constraint reification, 170
- constraint solver, 68
- context, 137

- database, 146
- deductive databases, 21

- definite descriptions, 81
- definites, 81
- denotation, 50
- Denotatpflicht, 59
- diadic quantifier constants, 51
- diagnosis, 21, 27
- discourse model, 85
- distribution, 160
- domain minimal, 77
- domain-minimal, 26

- equality sign, 51
- expert systems, 21

- Falcon, 38
- Finder, 28
- FinFiMo, 28
- finite constant frame, 60
- finite mathematics, 21
- finite model, 24
- finite model generators, 28
- formula, 53
- frame, 50
- function application, 49
- function domain, 50
- function-free quantified, 58

- generalised determiners, 42
- generalised frames, 53
- generalised interpretation, 53
- generalised quantifiers, 42

- ground tableaux, 29
- guarded fragment, 27
- Herbrand equality, 66
- Herbrand models, 23
- higher-order definitions, 52
- identifying property, 90
- Inclusive Alternative Ordering, 110
- inference costs, 147
- initial constant frame, 60
- Intermediate Alternative Reciprocity, 111
- Intermediate Reciprocity, 110
- interpretation of the simply typed λ -calculus, 50
- interpretations, 23
- isomorphic models, 185
- languages with exponents, 39
- LDPP, 38
- lean deduction, 162
- Lisp, 41
- literal, 22, 53
- literal sets, 35
- locally minimal model, 77
- logic, 51
- Logic Definition Structure, 167
- Logic Programming, 146
- logical constants, 51
- logically equivalent, 55
- Mace, 38
- Maximise Meaning Hypothesis, 117
- MGTP, 38
- MM-Satchmo, 26
- model, 23
- model of a set of formulas, 54
- monadic quantifier constants, 51
- MoreThan, 45
- observation, 137
- One-Way Weak Reciprocity, 109
- overgeneration problem of abduction, 136
- parameter constants, 51
- parameters, 51
- planning, 21
- positive model, 25
- predicate-specific minimal models, 27
- problem specification, 164
- program synthesis, 21
- proof engines, 175
- propagation, 160
- propagator procedures, 167
- propositional decision procedures, 35
- PUHR, 32
- quasi-group, 183
- query, 146
- range-restricted rules, 32
- SAT, 28
- Satchmo, 28, 32
- Sato, 38
- saturated, 30
- scheduling, 21
- scope relation, 106
- Sem, 38
- set exclusion operator, 106
- signed formulas, 70
- simply typed λ -calculus, 41
- single-solution problems, 183
- specification, 22

- standard frames, 53
- Strong Reciprocity, 108
- Strongest Meaning Hypothesis (SMH), 105
- Strongest Meaning Hypothesis for Reciprocals, 115
- subset condition, 88
- subset-minimal, 25

- theory, 22
- truth conditions, 14
- type ι of individuals, 48
- type o of truth values, 48
- type constructor, 48
- types, 48

- unary logical connectives, 51
- unicity condition, 88
- unification, 31
- unique-name assumption, 61
- universe of discourse, 23

- variable assignment, 50

- Weighted Abduction, 146
- weighted abduction, 138, 146
- weights, 147