Chapter 1

Introduction

1.1 Complexity Theory

Suppose that you are given a set of villages connected via dirt roads and are asked to tarmac a set of streets such that each pair of villages is connected via asphalted roads. Your budget is limited, hence you want to know whether this task can be solved with a given amount of money. This problem is fairly easy to solve: starting from an arbitrary village v_1, tarmac the shortest dirt road connecting v_1 to some not yet accessible village v_2. Now again choose the shortest dirt road as above starting from either v_1 or v_2, and so forth until all villages are connected.

This algorithm, also known as Prim's algorithm [Jar30, Pri57], will provide you with a minimal cost solution which you can compare to your budget. Moreover, the resources required to solve the problem are quite limited: one only needs to keep track of the set of villages already connected to each other and find the shortest dirt road leading from these to some not yet accessible village. But what if you are instead asked to tarmac a round trip that visits each village exactly once rather than an arbitrary set of streets. For this modification, the above strategy will no longer work. Indeed, no one has yet found an algorithm running in subexponential-time that answers the question whether you can tarmac a round trip. But can we be sure that no such algorithm exists? And in which way does the additional restriction make the problem computationally more involved?

These questions are typically studied in an area of theoretical computer science called *(computational) complexity theory*. This area analyzes the resources required to solve a computational problem and classifies these according to their inherent difficulty. One of the main goals of this area is to understand which problems are easy to solve, which problems are computationally hard, and of course, why. The class of easy decision problems is denoted by P and comprises those problems that are efficiently (that is, polynomial-time) solvable. The first of the above problems belongs to this class. For the second problem no polynomial-time algorithm is known; however, given a solution we can easily verify its correctness. Such problems are called *efficiently verifiable*, and the class of all such problems is denoted NP. As any efficiently solvable problem is also efficiently verifiable, we have $\mathrm{P} \subseteq \mathrm{NP}$. And while the question whether $\mathrm{P} = \mathrm{NP}$ or $\mathrm{P} \subsetneq \mathrm{NP}$ is one of the most important open problems in computer science, the inability to

prove or refute P = NP led to the development of a rich theory of computational complexity.

An important role in this context play the hardest problems in NP in the sense that an algorithm for any such problem can be transformed into one for any problem in NP. These problems are called NP-*complete*. An efficient algorithm for an NP-complete problem would thus allow for the efficient solution of all problems in NP, that is, P = NP. The first problem shown to be NP-complete was SAT, the satisfiability problem for propositional formulae [Coo71].

In this thesis, we will encounter problems that do not fall into the classes P or NP, for example, problems whose complement lies in NP. This class of problems for which the absence of solutions can be verified in polynomial-time, is known as coNP. We also require classes for problems that are harder to solve than SAT in the sense that they are only known to be efficiently verifiable if provided with an oracle that is able to instantaneously answer queries to a language in NP. These problems are called efficient verifiable relative to an NP-oracle. For example, the problem to determine whether the lexicographic smallest assignment of a formula sets to *true* a certain proposition is known to be efficiently verifiable relative to an NP-oracle but not known to be in NP or coNP. One can now consider problems that are efficiently verifiable relative to such problems, and so on. The concept of efficient verification relative to an oracle thus naturally leads to a hierarchy of complexity classes known as the *polynomial hierarchy*. The $(i+1)$th level of this hierarchy comprises the class $\Sigma^{\mathrm{P}}_{i+1}$ of problems known to be efficiently verifiable given an oracle for the ith level and the class Π^{P}_{i+1} of their complements, where the Σ^{P}_0 and Π^{P}_0 are defined as P.

We will use this rich framework of complexity theory to classify the complexity of computational problems connected to logics for knowledge representation and commonsense reasoning.

1.2 Nonmonotonic Logic

One of the most intriguing aspects of human reasoning is its flexibility and speed. Despite the fact that in most situations we do not have all relevant knowledge at hand, commonsense enables one to draw conclusions by virtue of plausible assumptions. These assumptions might be invalidated by new information about the world; therefore human reasoning is said to be *nonmonotonic*.

For example, suppose that you need some advice from a colleague. As his office is empty and it is noon, you conclude that he is gone for lunch; a conclusion derived from an assumption about his usual behaviour. Learning that he is on a business trip now invalidates your old conclusion.

From the very beginning of knowledge representation and reasoning, it has been argued that classical logic is not suited to formalize the process of human reasoning, mainly for its inherent monotonicity: once a statement is derivable it may never be invalidated regardless of whatever knowledge one might gain.

To overcome this deficiency, nonmonotonic logics have been introduced around 1980 [McC80, MD80, Rei80]. These logic can be distinguished by the way they facilitate nonmonotonic behaviour:

1. by extension with new inference rules,
2. by extension with modal operators,
3. by modification of the semantics.

In this thesis, we will examine one logic from each of the above approaches and study the complexity of natural problems arising in these. In particular, we focus on the following well-known logics.

Default logic has been introduced by Reiter [Rei80] and extends classical (first-order or propositional) logic with inference rules of the form $\frac{\alpha:\beta}{\gamma}$, called *default rules*. The default rule $\frac{\alpha:\beta}{\gamma}$ allows to conclude γ if the premise α is derivable and the justification β can consistently be assumed.

Autoepistemic logic has been introduced by Moore [Moo85] and extends classical logic with a unary "introspective" operator L expressing belief. For a formula φ, $L\varphi$ states that an ideally rational agent can derive φ.

Circumscription has been introduced by McCarthy [McC80]. Rather than extending classical logic, it restricts the notion of satisfiability and inference to consider the minimal model of a formula only. It has been shown that circumscription as defined by Lifschitz [Lif85] is equivalent to reasoning under the *extended closed world assumption*, which for a designated set P allows to assume $\neg p$ whenever $p \in P$ is not derivable [GPP89].

The extensions introduced by default or autoepistemic logic condition the derivable knowledge on a set of beliefs. Therefore maximal stable sets of knowledge supersede the traditional deductive closure. For default logic these are called *stable extensions*; for autoepistemic logic, *stable expansions*. A default or an autoepistemic theory may possess multiple or no such maximal stable sets of knowledge. Thus the following questions naturally arise: Does a given set of formulae admit a maximal stable set of knowledge? A lack thereof would correspond to the case that for all possible sets of beliefs one eventually arrives at contradictory information. The problem hence asks whether one can obtain consistent knowledge of the world. This problem is a rough analogue of the satisfiability problem in propositional logics and will henceforth be referred to as the *extension* (respectively *expansion*) *existence* problem.

Beyond, the potential presence of multiple maximal stable sets of knowledge leads to two different interpretations for the question whether a certain information is derivable: the first, *credulous reasoning* (also referred to as *brave reasoning*), asks whether a formula is contained in at least one stable extension (respectively expansion) of the knowledge base; the second, *skeptical reasoning* (also referred

to as *cautious reasoning*), asks whether the formula is contained in all stable extensions (respectively expansions). On an intuitive level, credulously entailed knowledge can be considered "possible", while skeptically entailed knowledge is "certain" in the sense that any possible interpretation of the world entails it. The associated decision problems are natural generalizations of the propositional implication problem and will henceforth be referred to as the *credulous reasoning* problem and the *skeptical reasoning* problem.

In the restricted semantics of minimal models no corresponding notion of maximal stable sets of knowledge exists. The corresponding notion in circumscription are minimal (or *circumscriptive*) models, which exist if and only if the given knowledge base is satisfiable. Therefore the problem of determining their existence of is equal to SAT. For circumscription we are hence restricted to the study of the skeptical reasoning problem, that is, to decide whether for a given set of formulae Γ and a formula φ, whether φ is true in all minimal models of Γ.

1.3 RESULTS

While for extensions of first-order logic all of the above decision problems are undecidable, they are decidable for extensions of propositional logic—but presumably harder than the traditional satisfiability or implication problem: they are known to be complete for the second level of the polynomial hierarchy [Nie90, CL90, Got92, EG93]. For this reason, several semantic restrictions and parameterizations of these problems have been studied in the literature (see [CL90, KS91, NR94, KK03, Nor04, CHS07], amongst others).

In this thesis, we take a different approach and perform a systematic study of the complexity of the above extension (respectively expansion) existence and reasoning problems obtained by restricting the set of allowed Boolean connectives. To this end, we generalize the underlying problems to allow for arbitrary Boolean connectives rather than the Boolean standard base $\{\wedge, \vee, \neg\}$ and classify the complexity of these problems parameterized by the set of allowed Boolean connectives for all possible finite sets of Boolean connectives.

This approach has first been taken by Lewis [Lew79], who showed that the satisfiability problem is NP-complete if and only if the negation of the implication ($x \nrightarrow y$) can be composed from the given Boolean connectives. Such a dichotomous behaviour cannot be taken for granted due to Ladner's theorem: if P $\neq$ NP then there exists infinitely many degrees of complexity between P and NP-completeness [Lad75]. Since then, Lewis' approach has been applied to a wide range of problems including equivalence and implication problems [Rei03, BMTV09a], satisfiability and model checking in modal and temporal logics [BHSS06, BSS$^+$08, BMS$^+$09, MMTV09, MMS$^+$09], and abduction [CST10].

Herein we study whether a similarly polytomous complexity classification is possible for the extension (respectively expansion) existence and reasoning problems mentioned above. Our goal is to exhibit fragments of lower complexity

which might lead to better algorithms for cases in which the set of Boolean connectives can be restricted. Furthermore we aim to understand the sources of hardness and to provide an understanding which connectives take the role of $x \nrightarrow y$ in the context of the nonmonotonic logics mentioned above, that is, which connectives account for jumps in the complexity of the problems. These connectives may help to identify candidates for parameters in the study of parameterized complexity of nonmonotonic logics.

To be more precise, let B denote the finite set of available Boolean connectives. Although at first sight, an infinite number of sets B of allowed Boolean connectives has to be examined, we prove, making use of results from universal algebra, that for all considered problems the complexity does not depend on the particular set but rather on the *clone* $[B]$ of B, that is, the set of functions which can be implemented from B using projections and arbitrary composition.

Decision Problems

We show that both the complexity of the extension existence problem in default logic and the complexity of the expansion existence problem in autoepistemic logic are polytomous (see Theorems 4.1.1 and 4.2.1):

the extension existence problem remains Σ_2^{p}-complete for all sets B such that $[B \cup \{1\}] = \mathsf{BF}$; becomes Δ_2^{p}-complete for monotone sets B that contain conjunctions, disjunctions and the constant 0; is NP-complete if $[B \cup \{1\}]$ contains $\neg$ and comprises affine functions only; and becomes tractable in all other cases (with this case splitting into P-complete, NL-complete, and trivial sub-cases). The expansion existence problem for autoepistemic logic, on the other hand, remains Σ_2^{p}-complete for all B such that $[B \cup \{0,1\}]$ includes the Boolean functions $\wedge$ and $\vee$, is NP-complete if $[B]$ contains $\vee$ and the Boolean constants only, and becomes polynomial-time decidable in all other cases (with this case splitting into three different complexity degrees inside P).

For the credulous and skeptical reasoning problems in default logic and autoepistemic logic, the situation is more diverse as there are two sources for the complexity: On the one hand, we need to determine a finite characterization of a candidate for a stable extension (respectively expansion). And, on the other hand, we have to verify that this candidate is indeed a finite characterization as desired—a task that requires to test for formula implication. Depending on the Boolean connectives allowed, one or both tasks can be performed in polynomial time or even become trivial. In principle, this yields five possible cases for the complexity of the problems, and we will see that all five cases actually occur. In principle, this yields five possible cases for the complexity of the problems, and we will see that all four cases actually occur.

We obtain Σ_2^{p}-completeness for the skeptical reasoning problems and Π_2^{p}-completeness for the credulous reasoning problems for all clones where both the stable extension and the implication problem attain their highest complexity.

For default logic, the complexity of the reasoning problems reduces to Δ_2^{p} for clones that allow for an efficient computation of stable extensions but whose implication problem remains coNP-complete. More precisely, these problems are Δ_2^{p}-complete if a stable extension may not exist and becomes coNP-complete otherwise. Conversely, if the implication problem becomes easy but determining an extension candidate is hard, then the credulous reasoning problem is NP-complete, while the skeptical reasoning problems is coNP-complete. Similarly for autoepistemic logic, the credulous and skeptical reasoning problems become complete for respectively NP and coNP if the implication problem is tractable but determining an expansion candidate is hard. Finally, for clones that allow for solving both tasks in polynomial time all reasoning problems become tractable (with these cases splitting up into different complexity degrees ranging from membership in AC^0 to completeness for P). We hence obtain polytomous classifications of the computational complexity of the problems, where for the credulous reasoning problem in default logic, notably, complete fragments for all classes of the polynomial hierarchy below Σ_2^{p} occur. In contrast to this, the complexity of credulous and skeptical reasoning in autoepistemic logic decreases in coarser steps. These results are presented in Theorems 5.1.1, 5.1.5, 5.2.1 and 5.2.4.

As for circumscription, the complexity of the skeptical reasoning problem is Π_2^{p}-complete for all clones such that the implication problem and the problem to determine the minimality of models are intractable. If all available functions are affine or monotone, then the complexity of the problem is contained in coNP, where it is coNP-complete in the former case as long as $\vee$ remains expressible using the available functions and the constant 1. This decrease in the complexity comes from different sources: for monotone functions the test for minimality of models becomes tractable, while for affine functions the implication problem becomes tractable. Finally, if the set of available functions is further restricted to contain either only negations or only conjunctions, then the problem becomes polynomial-time solvable (its complexity drops to respectively $\mathrm{AC}^0[2]$-completeness or membership in AC^0). This is summarized in Theorem 5.3.1. We point out that the implication problem and the problem to determine the minimality of models do not completely determine the complexity of the skeptical reasoning problem: for all sets B such that $[B \cup \{0,1\}]$ contains $\vee$ and the Boolean constants only, the latter problem remains coNP-complete whereas the implication problem and minimality of models can be decided in polynomial time.

Counting Problems

Besides the decision variants, another natural question is concerned with the number of stable extensions (respectively expansions) or the number of minimal models. This question refers to *counting problems*. Recently, counting problems have gained quite a lot of attention in nonmonotonic logics. For circumscrip-

tion, the counting problem (that is, determining the number of minimal models of a propositional formula) has been studied in [DHK05, DH08]. For propositional abduction, a nonmonotonic formalism for computing explanations, some complexity results on the problem of counting the number of "solutions" to a propositional abduction problem were presented in [HP07, CST10]. Algorithms based on bounded treewidth have been proposed in [JPRW08] for the counting problems in abduction and circumscription. Here, we consider the complexity of the problem to count the number of stable extensions, stable expansions and minimal models of a given knowledge base. To the best of our knowledge, the first problem is addressed here for the first time.

In particular, we show in Theorem 6.1.1 that for sets B of Boolean connectives such that $[B \cup \{1\}]$ is functional complete counting the number of stable extensions is complete for the second level of the counting hierarchy; becomes Δ_2^{p}-complete for all monotone sets B such that $[B \cup \{1\}] = \mathsf{M}$; is #P-complete for affine sets B such that $\neg$ can be implemented from $B \cup \{1\}$; and becomes efficiently computable in all other cases. In autoepistemic logic, the complexity of counting the number of stable expansions is trichotomous and decreases analogously to the complexity of the stable expansion problem, see Theorem 6.2.1.

We think it is important to note that for our classification of the two counting problems above the conceptually simple parsimonious reductions are sufficient, while for related classifications in the literature less restrictive (and more complicated) reductions such as subtractive or complementive reductions had to be used (see, for example, [DHK05, DH08, BBC$^+$09] and some of the results of [HP07]). Parsimonious reductions are not only the conceptually simplest ones since they are direct analogues of the usual many-one reductions among languages. They also form the strongest (or strictest) type of reduction with a number of good properties, for example, all relevant counting classes are closed under parsimonious reductions.

Lastly, the complexity of counting the number of minimal models is classified in Theorem 6.3.1. Unlike the preceding counting problems, here we have sets of Boolean functions for which the problem to decide whether a given assignment is a circumscriptive model is tractable while the corresponding counting problem is #P-complete (namely affine sets of Boolean functions that implement the ternary exclusive-or). In all remaining cases, its complexity can be derived from the complexity of the skeptical reasoning problem in circumscription in the way that completeness for the second level of the polynomial hierarchy translates to #·coNP-completeness, completeness for the first level translates to #P-completeness, and membership in P translates to membership in FP. However, mind that the decision problem underlying the circumscriptive model counting problem is the question whether there exists a minimal model for the given formula—a problem equivalent to the satisfiability problem for propositional formulae. It thus represents a counting problem whose underlying decision problem is, though intractable, supposedly easier to solve than the decision problems underlying the generic complete problem for #·coNP.