Minimal Inference Problem Over Finite Domains: The Landscape of Complexity

Michał Wrona(✉)

Theoretical Computer Science Department,
Faculty of Mathematics and Computer Science,
Jagiellonian University, Kraków, Poland
wrona@tcs.uj.edu.pl

Abstract. The complexity of the general inference problem for propositional circumscription in Boolean logic (or equivalently over the two-element domain) has been recently classified. This paper generalizes the problem to arbitrary finite domains. The problem we study here is parameterized by a set of relations (a constraint language), from which we are allowed to build a knowledge base, and a linear order on the domain, which is used to compare models.

We use the algebraic approach provided originally in order to understand the complexity of the constraint satisfaction problem to give first non-trivial dichotomies and tractability results for the minimal inference problem over finite domains.

1 Introduction

The need for logics that could capture human way of thinking triggered the development of an area of Artificial Intelligence called nonmonotonic reasoning. A number of formalisms emerged. One of the most important and best studied is circumscription introduced by McCarthy [17]. The circumscription of a formula is the set of its minimal models that are supposed to represent possible situations that are consistent with common sense.

It is often the case [6,9,11,18] that models are compared according to the preorder ($\leq_{(P,Z)}$) induced by a partition of variables V into three subsets P, Z, Q (possibly empty) where P — variables that are subject to minimizing, Q — variables that maintain the fixed value, and Z — variables whose value can vary. Now, for two assignments $\alpha, \beta : V \to D$ we will have ($\alpha \leq_{(P,Z)} \beta$) if $\alpha[Q] = \beta[Q]$ (α is equal to β on variables in Q) and $\alpha[P] \leq \beta[P]$ (α is less than or equal to β on variables in P) where $\leq$ is the coordinatewise extension of the natural order on $\{0, 1\}$ with $0 < 1$.

As every logical formalism does, circumscription gives rise to two main computational problems: the model-checking problem and the inference problem. In this paper we concentrate on the inference problem for propositional circumscription, called also *the minimal inference problem*. In the most general formulation

The research was partially supported by NCN grant number 2014/14/A/ST6/00138.

M. Balduccini and T. Janhunen (Eds.): LPNMR 2017, LNAI 10377, pp. 101–113, 2017.
DOI: 10.1007/978-3-319-61660-5_11

an instance of this problem in the Boolean case consists of two CNF formulas: *a knowledge base* φ and a *query* ψ over the same set of variables V partitioned into P, Z, Q. In the minimal inference problem we ask if every $(\leq_{(P,Z)})$-min model (minimal model wrt. $(\leq_{(P,Z)})$) of φ is a model of ψ. Since this task can be performed for every clause of ψ independently and we are interested in the complexity of the problem up to polynomial time reduction, we can assume that ψ is a disjunction of propositional literals (a clause). This problem is in general Π_2^P-complete [10]. Thus, one considers the general minimal inference problem GMININF(Γ) parameterized by a set of relations Γ over $\{0, 1\}$.

- *Instance of GMININF*(Γ)*:* a conjunction of atomic formulas φ of the form: $R_1(x_1^1, \ldots, x_{k_1}^1) \wedge \cdots \wedge R_l(x_1^l, \ldots, x_{k_l}^l)$, where every R_i with $i \in [l]$ is a relation symbol in a signature of Γ, over variables V partitioned into P, Z, Q, and a propositional clause ψ over V.
- *Question:* is every $(\leq_{(P,Z)})$-min model of φ a model of ψ?

The complete complexity classification of GMININF(Γ) with respect to Γ has been obtained after a series of papers, e.g. [6,8,10,15] in [9]. Under usual complexity theoretical assumptions, in this case that: P $\subsetneq$ coNP $\subsetneq \Pi_2^P$, it is shown that GMININF(Γ) is either Π_2^P-complete, or coNP-complete, or in P. This raises a question about a similar classification in many-valued logics, or, more generally, over arbitrary finite domains. Especially that circumscription over larger finite domains has been studied in the literature, e.g. in [7,20].

The same course of events took place in the case of the problem CSP(Γ) where a question is whether a given conjunction of atomic formulas over the signature of Γ is satisfiable. In [22], Schaefer established the dichotomy between NP-complete and in P in the case where Γ is over the two-element domain. Then researchers turned to a so-called Feder-Vardi conjecture that states that a similar dichotomy holds for arbitrary finite domains. The understanding and many advanced partial results[1], see [16] for a recent survey, were possible thanks to the development of the so-called algebraic approach [5,12]. This approach has been also already applied to the model checking and the inference problem in propositional circumscription over arbitrary finite domains in [19]. Here we use algebra for the inference problem in a bit different formulation.

Certainly every relation $R \subseteq D^n$ over any finite domain D can be defined by a conjunction of disjunctions of disequalities (a CNF of disequalities) of the form $(x \neq d)$ where x is a variable and $d \in D$ simply by the formula: $\bigwedge_{(d_1,\ldots,d_n) \notin R} (x_1 \neq d_1 \vee \cdots \vee x_n \neq d_n)$. This implies that in the most general version of the minimal inference problem the input may consist of a CNF of disequalities that states for a knowledge base, a CNF of disequalities that states for a query and a *linear order* $\mathcal{O} = (D; \leq^{\mathcal{O}})$. The preorder $\leq_{(P,Z)}^{\mathcal{O}}$ is defined as in

[1] Recently, three different groups of researchers announced a proof of the dichotomy.

the two-element case with a difference that we use $\leq^{\mathcal{O}}$ on coordinates instead of $0 < 1$. Since CNFs of disequalities coincide with clauses if $|D| = 2$, we have that the minimal inference problem in this formulation is Π_2^P-hard. It is straightforward to show that the problem is in fact Π_2^P-complete. Thus, it is natural to ask about the complexity of the parametrized version GMININF$(\Gamma, \mathcal{O})$ defined below. As in the Boolean case we can assume that a query consists of a single disjunction of disequalities.

- *Instance of GMININF*$(\Gamma, \mathcal{O})$*:* a conjunction of atomic formulas φ of the form: $R_1(x_1^1, \ldots, x_{k_1}^1) \wedge \cdots \wedge R_l(x_1^l, \ldots, x_{k_l}^l)$, where every R_i with $i \in [l]$ is a relation symbol in a signature of Γ, over variables V partitioned into P, Z, Q, and a disjunction of disequalities ψ of the form $(x_1 \neq d_1 \vee \cdots \vee x_k \neq d_k)$ over V.
- *Question:* is every $(\leq^{\mathcal{O}}_{(P,Z)})$-min model of φ a model of ψ?

For an example, consider Γ over $D = \{1, 2, 3\}$ containing the relation $R_{\neq} := \{(d_1, d_2) \in D^2 \mid d_1 \neq d_2\}$ and the order $1 <^{\mathcal{O}} 2 <^{\mathcal{O}} 3$. A formula $R_{\neq}(x_1, x_2) \wedge R_{\neq}(x_2, x_3) \wedge R_{\neq}(x_3, x_4) \wedge R_{\neq}(x_2, x_3)$ and a disjunction $(x_1 \neq 2 \vee x_4 \neq 1)$ form an instance of GMININF$(\Gamma, \mathcal{O})$. This problem is Π_2^P-complete. Throughout the paper we give parametrizations of GMININF$(\Gamma, \mathcal{O})$ of lower complexity.

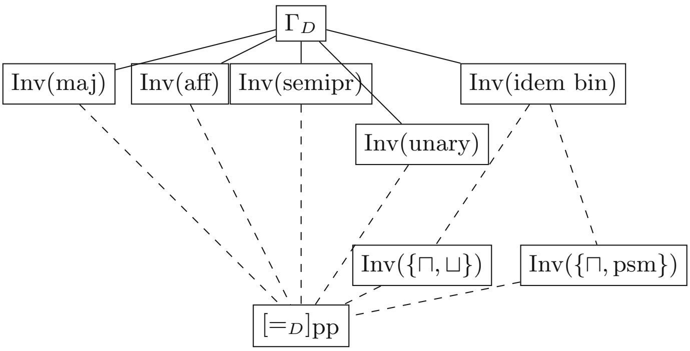

Fig. 1. An illustrative presentation of the lattice of relational clones over a domain D.

Contribution. Our attack on the complexity classification of GMININF$(\Gamma, \mathcal{O})$ is based on the algebraic approach. All the notions we use in this paper are defined carefully in Sect. 2. As explained there in order to complete the classification task, it is enough to establish the complexity only for relational clones that are constraint languages closed under primitive positive definitions. Equivalently, a relational clone is the set, denoted by Inv(F), of all relations invariant under (preserved by) all operations in some set F. The relational clones over D are organized in the lattice ordered by inclusion. The larger is the relational clone, the complexity of the problem is harder. An illustrative presentation of the lattice over some D containing only kinds of relational clones we look at in this paper is presented in Fig. 1. For a pair of relational clones linked by a line the one placed higher contains the one which is below. The solid lines indicate that there are no relational clones inbetween, the dashed ones that there might be some. The top element of the lattice is the set of all relations Γ_D over D. The only operations that preserve all relations in Γ_D are projections. The bottom of the lattice is the set of relations primitively positively definable by means of equality. Such are preserved by all operations over D.

The problem GMININF$(\Gamma_D, \mathcal{O})$ is Π_2^P-complete for all $|D| \geq 2$. To find subproblems of lower complexity we climb down the lattice of relational clones. The natural choice for the languages to be studied first are *maximal constraint languages* that in the lattice lay directly below Γ_D. According to Rosenberg's Five Types Theorem [21], it is enough to consider languages of the form Inv$(\{f\})$ where f is a unary operation, a semiprojection, a binary idempotent operation, a majority operation or an affine operation, see Fig. 1. Their complexity was analyzed for many problems parametrized by constraint languages: for CSPs the result can be found in [1,4] We show that GMININF$(\Gamma, \mathcal{O})$ is coNP-hard for all such languages. Furthermore for maximal conservative (i.e., containing all unary relations) languages the problem is either Π_2^P-complete or coNP-complete. See [3] for the CSP classification over conservative languages.

In order to find tractable (polynomial-time decidable) classes we climb down the lattice of relational clones even further. In particular we give dichotomies between coNP-complete and P for:

- GMININF(Inv$(\{\sqcap, \sqcup\}), \mathcal{O})$ where $\sqcap, \sqcup$ are the join and the meet of some lattice, and
- GMININF(Inv$(\{\sqcap, f\}), \mathcal{O})$ where $\sqcap$ is the meet of some semilattice and f a newly defined *pms*-operation.

This gives us new large tractability classes and new coNP-hardness results. Futhermore, our algorithms are based on polymorphisms which is crucial for further generalizations, consult again [16] to see the importance of polymorphisms in providing algorithms for the CSP.

Related Work. Classifications such as a trichotomy for GMININF(Γ) for Γ over the two-element domain as one presented in [9] are much easier to be obtained than analogical results over arbitrary finite domains. The reason is that in the two-element case, the lattice of relational clones (so-called Post's lattice) has countably many and very well described elements. Thus, in order to obtain a classification it is enough to consider the problem for each of them. The situation is very different already over the three-element domain where there are uncountably many relational clones and the lattice is not comprehensible.

The results on the minimal inference problem over arbitrary finite domains included in [19] concern a version of GMININF$(\Gamma, \mathcal{O})$ where the query ψ is a single atomic formula $R(x_1, \dots, x_k)$ where R corresponds to some k-ary relation over the domain of Γ, and the order $\mathcal{O}$ is a part of the input. The authors of [19] provide mainly preliminary results that may be seen as tools for classifying the problem. Some of these results such as Theorem 1 we reprove here for our version of the problem and use heavily in our paper. The only complexity result in [19], however, is a dichotomy between Π_2^P-complete and in coNP for conservative languages Γ over the three-element domain[2] and even this specific result follows in a rather straightforward way from the dichotomy for CSP(Γ) over the three-element domain [2].

In this paper, for the first time, we provide complexity results characteristic for the minimal inference problem over larger domains. We expect the classification for the problem GMININF to be completely different than and not easily obtainable from the one for CSP. This is already true over two elements. For the first time, we provide here polynomial algorithms. The tractability results for GMININF$(\mathrm{Inv}(\{\sqcap, \sqcup\}), \mathcal{O})$ and GMININF$(\mathrm{Inv}(\{\sqcap, f\}), \mathcal{O})$ substantially generalize these for GMININF$(\mathrm{Inv}(\wedge, \vee))$ and GMININF$(\mathrm{Inv}(\{\wedge, (x \wedge (y \vee \neg z))\}))$ from the two-element world. Our dichotomy: between Π_2^P-complete and in coNP for conservative maximal languages generalize a dichotomy between these two classes for the two-element domain from [9] and for the three-element conservative case (up to a small difference in the definition of the problems) from [19].

2 Preliminaries

We write $t = (t[1], \dots, t[n])$ for a tuple of elements and $[n]$ to denote the set $\{1, \dots, n\}$. The reverse of the order $\mathcal{O}_1 = (D; \leq^{\mathcal{O}})$ is $\mathcal{O}_2 = (D; (\leq^{\mathcal{O}})^{-1})$ where $(\leq^{\mathcal{O}})^{-1}$ is the relation $\{(a, b) \in D \mid (b, a) \in \leq^{\mathcal{O}}\}$.

Constraint Languages. In this paper a *(constraint) language* over (always finite) domain D, denoted by capital Greek letters such as Γ, is a set of relations over D. A signature of Γ denoted usually by τ is a set of relation symbols associated to relations in Γ. For the sake of simplicity we usually use the same symbols to denote both a relation symbol and the corresponding relation. We also assume that the domain of a relation or a language under consideration is the set D.

A *primitive positive formula (pp-formula)* over a signature τ is a first-order formula built exclusively from conjunction, existential quantifiers and atomic formulas over τ and equalities, that is atomic formulas of the form $(x = y)$. We say that a relation R has a pp-definition over a set of relations Γ if there exists a pp-formula over the signature of Γ that holds exactly on those tuples that are contained in R. We say that a set of relations Δ has a pp-definition over Γ if every relation in Δ has a pp-definition over Γ. A set of relations with a pp-definition over Γ is denoted by $[\Gamma]_{\mathrm{pp}}$ and called a *relational clone*.

[2] The paper claims that the dichotomy is for all languages over the three-element domain. However, this is not true since the proof of Theorem 3.6 is flawed for domains with more than two elements [14].

We define a Γ-formula to be a conjunction of atomic formulas over a signature of Γ. Observe that a Γ-formula is a special form of a pp-formula where quantifiers and equality are not in use. Let $\psi := R(x_1, \ldots, x_k)$ be an atomic Γ-formula with $x_1, \ldots, x_k$ not necessarily different and $W \subseteq \{x_1, \ldots, x_k\}$. We write $\psi|_W$ for an atomic formula $R|_W(y_1, \ldots, y_l)$ where $R|_W$ is a projection of R to coordinates corresponding to variables in W and $y_1, \ldots, y_l$ is a subsequence of $x_1, \ldots, x_k$ containing only variables in W. For a Γ-formula $\varphi = \psi_1 \wedge \ldots \wedge \psi_n$ over V and $W \subseteq V$ we write $\varphi|_W$ to indicate $\psi_1|_W \wedge \ldots \wedge \psi_n|_W$.

Operations, Polymorphisms. Let Γ be a language over domain D. An operation $f : D^n \to D$ is a *polymorphism* of an m-ary relation R if for all m-tuples $t_1, \ldots, t_n \in R$, it holds that the tuple $(f(t_1[1], \ldots, t_n[1]), \ldots, f(t_1[m], \ldots, t_n[m]))$ is also in R. An operation f is a polymorphism of a language Γ if it is a polymorphism of every relation in Γ. If $f\colon D^n \to D$ is a polymorphism of Γ, R, we say that f *preserves* Γ, R. The set of relations preserved by a set of operations F is denoted by $\mathrm{Inv}(F)$. The following Galois correspondence links sets of polymorphisms and relational clones, see e.g. [12].

Lemma 1. *Let Γ be a constraint language. Then $\Gamma_1 \subseteq [\Gamma_2]_{pp}$ if and only if $Pol(\Gamma_2) \subseteq Pol(\Gamma_1)$.*

Here we list some kinds of operations that are of use for this paper. We say that an operation $f : D^n \to D$ is *idempotent* if for all $x \in D$ we have $f(x, \ldots, x) = x$. An operation $f : D \to D$ that is bijective is said to be a *permutation*. An operation $f : D^n \to D$ is a *projection* if there exists $i \in [n]$ such that $f(x_1, \ldots, x_i, \ldots, x_n) = x_i$ for all $x_1, \ldots, x_n \in D$. We say that a ternary operation $f : D^3 \to D$ is a *majority operation* if for all $x, y \in D$ we have $f(x, x, y) = f(x, y, x) = f(y, x, x) = x$, and that f is *affine* if for all $x_1, x_2, x_3 \in D$ we have $f(x_1, x_2, x_3) = x_1 - x_2 + x_3$, where $+$ and $-$ are the operations of an Abelian (commutative) group $(D; +, -)$. An operation $f : D^n \to D$ with $D \geq 3$ is said to be a *semiprojection* if there exists $i \in [n]$ such that for all $x_1, \ldots, x_n \in D$ we have $f(x_1, \ldots, x_n) = x_i$ whenever $|\{x_1, \ldots, x_n\}| < n$ and f is not a projection. A *semilattice operation* s on a set D is an idempotent operation satisfying universally the identities $s(x, y) = s(y, x)$ and $s(s(x, y), z) = s(x, s(y, z))$. The first of these identities implies that s is commutative and the other that s is associative. We define a ternary operation $f : D^3 \to D$ to be a *pms*-operation compliant with a semilattice operation s if for all $x, y \in D$ it satisfies: $f(x, y, y) = x$, $f(y, x, y) = y$ and $f(y, y, x) = s(x, y)$.

The General Minimal Inference Problem. We now give a careful definition of the GMININF problem and provide basic results that help classifying the complexity of the problem. Some of them for a variant of the problem we study here are already available in the literature [8,9,19].

Let V be a set of variables and D a finite set. We use small Greek letters: α, β, γ to denote assignments of the type $V \to D$. We say that $\alpha : V_1 \to D$ is a restriction of $\beta : V_2 \to D$ to (variables in) V_1 if $V_1 \subseteq V_2$ and for all $v \in V_1$, we have $\alpha(v) = \beta(v)$. In this case we also say that β is an extension of α to (variables in) V_2. For $\alpha : V \to D$ and $V_1 \subseteq V$, we write $\alpha[V_1]$ to indicate the restriction of α to V_1.

Let $\mathcal{O} = (D; \leq^{\mathcal{O}})$ be a linear order. We extend the order to assignments in the natural way. For $\alpha, \beta : V \to D$ we have $\alpha \leq^{\mathcal{O}} \beta$ if for all $v \in V$ it holds $\alpha(v) \leq^{\mathcal{O}} \beta(v)$. We write $\alpha <^{\mathcal{O}} \beta$ if $\alpha \leq^{\mathcal{O}} \beta$ and for at least one $v \in V$ it holds $\alpha(v) <^{\mathcal{O}} \beta(v)$. For the purposes of this paper we need also a special preorder on assignments denoted by $(\leq^{\mathcal{O}}_{(P,Z)})$ and defined as follows. Let $\alpha, \beta : V \to D$ and P, Z, Q be a partition of V. We have $(\alpha \leq^{\mathcal{O}}_{(P,Z)} \beta)$ if $\alpha[Q] = \beta[Q]$ and $\alpha[P] \leq^{\mathcal{O}} \beta[P]$ and $(\alpha <^{\mathcal{O}}_{(P,Z)} \beta)$ if $\alpha[Q] = \beta[Q]$ and $\alpha[P] <^{\mathcal{O}} \beta[P]$. Let Γ be a constraint language over D and φ be a Γ-formula over variables V. An assignment $\alpha : V \to D$ is a model of φ if for every conjunct $R(x_1, \ldots, x_k)$ of φ we have $(\alpha(x_1), \ldots, \alpha(x_k)) \in R$. We say that a model α of φ is a $(\leq^{\mathcal{O}}_{(P,Z)})$-minimal $((\leq^{\mathcal{O}}_{(P,Z)})$-min) model of φ if there is no model β of φ such that $\beta <^{\mathcal{O}}_{(P,Z)} \alpha$.

We now rephrase the definition of GMININF$(\Gamma, \mathcal{O})$ from the introduction using new notions introduced in this section. The definitions are equivalent.

Definition 1. *[GMININF$(\Gamma, \mathcal{O})$]*

- *INSTANCE: A Γ-formula φ over variables V partitioned into three sets P, Z, Q and a disjunction of disequalities ψ over V.*
- *QUESTION: Is every $(\leq^{\mathcal{O}}_{(P,Z)})$-min model of φ a model of ψ?*

For finite Γ we measure the complexity of GMININF$(\Gamma, \mathcal{O})$ as a function of the length of a Γ-formula. In this paper we consider Γ and $\mathcal{O}$ that make the problem Π_2^P-complete, coNP-complete or in P. For an infinite set of relations Γ we use the usual convention. We say that GMININF$(\Gamma, \mathcal{O})$ is Π_2^P-hard (coNP-hard) if there is a finite $\Gamma' \subseteq \Gamma$ such that GMININF$(\Gamma', \mathcal{O})$ is Π_2^P-hard (coNP-hard) and that GMININF$(\Gamma, \mathcal{O})$ is in Π_2^P, coNP or P if for every finite $\Gamma' \subseteq \Gamma$ the problem GMININF$(\Gamma', \mathcal{O})$ is in Π_2^P, coNP or P, respectively.

The computational complexity of GMININF$(\Gamma, \mathcal{O})$ is fully captured by the relational clone $[\Gamma]_{\text{pp}}$, or equivalently the set of polymorphisms of Γ.

Theorem 1. *Let Γ_1, Γ_2 be constraint languages such that $\Gamma_1 \subseteq [\Gamma_2]_{pp}$ (or equivalently $Pol(\Gamma_2) \subseteq Pol(\Gamma_1)$), then there is a polynomial-time reduction from GMININF$(\Gamma_1, \mathcal{O})$ to GMININF$(\Gamma_2, \mathcal{O})$.*

This is usually easier to look at GMINEXT than GMININF.

Definition 2. *[GMINEXT$(\Gamma, \mathcal{O})$]*

- *INSTANCE: A Γ-formula φ over variables V partitioned into three sets P, Z, Q and a partial assignment $\alpha : V_1 \to D$ with $V_1 \subseteq V$.*
- *QUESTION: Is there an extension $\beta : V \to D$ of α such that β is a $(\leq^{\mathcal{O}}_{(P,Z)})$-min model of φ?*

The following proposition reveals the connection between the problems. We have that the complement of GMININF$(\Gamma, \mathcal{O})$ and GMINEXT$(\Gamma, \mathcal{O})$ are polynomially equivalent. The reduction from the complement of GMININF$(\Gamma, \mathcal{O})$ to GMINEXT$(\Gamma, \mathcal{O})$ comes to replacing a disjunction $\psi := (x_1 \neq d_1 \vee \cdots \vee x_k \neq d_k)$ with $\alpha : \{x_1, \ldots, x_k\} \to D$ such that $\alpha(x_i) = d_i$ for $i \in [k]$. Now, if there is a $(\leq^{\mathcal{O}}_{(P,Z)})$-min model β of φ (the same for both instances) that is not a model of ψ it satisfies $\beta(x_i) = d_i$ for $i \in [k]$, and hence β extends α. On the other hand, if there is a $(\leq^{\mathcal{O}}_{(P,Z)})$-min model β of φ extending α, then it certainly does not satisfy ψ. For the reduction from GMINEXT$(\Gamma, \mathcal{O})$ to the complement of GMININF$(\Gamma, \mathcal{O})$, we replace α with $\psi := (x_1 \neq \alpha(x_1) \vee \cdots \vee x_k \neq \alpha(x_k))$, where $\{x_1, \ldots, x_k\}$ is the domain of α. The proof is analogous.

Proposition 1. *Let $\mathcal{O} = (D; \leq^{\mathcal{O}})$ be a linear order and Γ a constraint language over D. The problem GMINEXT$(\Gamma, \mathcal{O})$ is Σ_2^P-hard, NP-hard, in NP, in P if and only if GMININF$(\Gamma, \mathcal{O})$, is Π_2^P-hard, coNP-hard, in coNP, in P, respectively.*

3 Maximal Constraint Languages

In this section we give a lower bound for GMININF$(\Gamma, \mathcal{O})$ over maximal constraint languages Γ and a dichotomy for conservative maximal languages.

Definition 3. *Let Γ_D be the set of all relations over domain D. A constraint language $\Gamma \subseteq \Gamma_D$ is* maximal *if $[\Gamma]_{pp} \subsetneq \Gamma_D$ and for every $R \notin \Gamma$, we have that $[\Gamma \cup R]_{pp} = \Gamma_D$. A constraint language Γ over D is* conservative *if Γ contains all subsets of D as unary relations.*

To build the classification we use some methods [4] and some results [6,9] known from the literature. We start with Rosenberg's theorem.

Theorem 2 (Rosenberg Theorem). *Every maximal constraint language has the form Inv$(\{f\})$ where the operation f is of one of the following types:*

1. *a unary operation which is either a permutation or else acts identically on its range;*
2. *a binary operation which is not a projection;*
3. *a majority operation;*
4. *an affine operation;*
5. *a semiprojection.*

We need to know the complexity of GMININF$(\mathrm{Inv}(f), \mathcal{O})$ for every type of operations from Theorem 2. In the case where $f : D^3 \to D$ is an affine operation we focus on two particular relations R_+ and R_{++} defined as follows. Let $\perp, \Box \in D$ be the two least elements in $\mathcal{O}$, i.e., $\perp <^{\mathcal{O}} \Box <^{\mathcal{O}} x$ for all $x \in D \setminus \{\perp, \Box\}$. We will have:

- $R_+ = \{(x, y) \mid x + y = \perp + \Box\}$,
- $R_{++} = \{(x, y, z) \mid x + y + z = \perp + \perp + \Box\}$,

where $+$ comes from an Abelian group $(D; +, -)$. Since R_+ and R_{++} are defined by a single equation, it is straightforward to show that they are both in $\mathrm{Inv}(f)$.

Lemma 2. *Let $(D; +, -)$ be an Abelian group and $\mathcal{O} = (D; \leq^{\mathcal{O}})$ a linear order. Then GMININF$(\Gamma, \mathcal{O})$ with $\Gamma = \{R_+, R_{++}\}$ is coNP-hard.*

We first prove that the problem for maximal constraint languages is coNP-hard.

Theorem 3. *Let $\mathcal{O} = (D; \leq^{\mathcal{O}})$ be a linear order and Γ a maximal constraint language over Γ, then GMININF$(\Gamma, \mathcal{O})$ is coNP-hard.*

We are able to prove the full dichotomy only under an additional assumption that languages are conservative.

Theorem 4. *Let $\mathcal{O} = (D; \leq^{\mathcal{O}})$ be a linear order and Inv$(\{f\})$, for some operation f, a conservative maximal constraint language over D. Then we have exactly one of the following:*

1. *f is a unary operation, or there is a two-element $D' \subseteq D$ such that $f|_{D'}$ is a projection and then GMININF(Inv$(\{f\}), \mathcal{O})$ is Π_2^P-complete.*
2. *f is a commutative binary operation, a majority operation, or an affine operation and then GMININF(Inv$(\{f\}), \mathcal{O})$ is coNP-complete.*

4 The Minimal Inference Problem and Semilattice Operations

Let $\mathcal{O} = (D; \leq^{\mathcal{O}})$ be a linear order, Γ a constraint language over D and φ a Γ-formula. We say that a model α of φ is the least (the greatest) model of φ wrt. $\leq^{\mathcal{O}}$ if for every model β of φ, it holds $\alpha \leq^{\mathcal{O}} \beta$ ($\beta \leq^{\mathcal{O}} \alpha$). The least (the greatest) model does not have to exist. However, if Γ is preserved by some well-behaved semilattice operation, then we have the following.

Observation 1. *Let $\mathcal{O} = (D; \leq^{\mathcal{O}})$ be a linear order and Γ a constraint language over D preserved by the meet $\sqcap$ (the join $\sqcup$) of some meet-semilattice (join-semilattice) $\mathcal{L} = (D; \leq^{\mathcal{L}})$ such that $\leq^{\mathcal{L}} \subseteq \leq^{\mathcal{O}}$. Let φ be a Γ-formula. Then there exists a model α of φ such that α is the least (the greatest) model of φ wrt the order $\leq^{\mathcal{L}}$ and at the same time α is the least (the greatest) model of φ wrt the order $\leq^{\mathcal{O}}$.*

In the case described by the previous observation we can quickly compute the least (the greatest) wrt both $\leq^{\mathcal{L}}$ and $\leq^{\mathcal{O}}$ model of a Γ-formula extending a given assignment α. The procedure is very well known [12] and may be performed by enforcing *generalized arc consistency.* We refer to this procedure by $\mathrm{leastext}(\mathcal{O}, \varphi, \alpha, V_2)$ ($\mathrm{greatext}(\mathcal{O}, \varphi, \alpha, V_2)$).

Proposition 2. *Let $\mathcal{O} = (D; \leq^{\mathcal{O}})$ be a linear order and Γ a constraint language preserved by the meet $\sqcap$ (the join $\sqcup$) of some meet-semilattice (join-semilattice) $\mathcal{L} = (D; \leq^{\mathcal{L}})$ such that $\leq^{\mathcal{L}} \subseteq \leq^{\mathcal{O}}$. Then the procedure leastext$(O, \varphi, \alpha, V_2)$* (the procedure greatext $(O, \varphi, \alpha, V_2)$) *for a Γ-formula φ and a partial assignment $\alpha : V_1 \to D$ with $V_1 \subseteq V_2 \subseteq V$ returns:*

– *false if there is no model of* $\varphi|_{V_2}$ *extending* α*;*
– *the least (the greatest), wrt* $\leq^{\mathcal{O}}$*, model of* $\varphi|_{V_2}$ *extending* α*.*

The procedures work in polynomial time.

In the case we consider here $\alpha[P]$ of a $(\leq^{\mathcal{O}}_{(P,Z)})$-min model α is uniquely determined by $\alpha[Q]$.

Observation 2. *Let* $\mathcal{O} = (D; \leq)$ *be a linear order and* Γ *a constraint language preserved by the meet* $\sqcap : D^2 \to D$ *of some meet-semilattice* $\mathcal{L} = (D; \leq^{\mathcal{L}})$ *such that* $\leq^{\mathcal{L}} \subseteq \leq^{\mathcal{O}}$*. Let* φ *be a* Γ*-formula over variables* V *partitioned into* P, Z, Q*. Then* α *is a* $(\leq^{\mathcal{O}}_{(P,Z)})$*-min model of* φ *if and only if* $\alpha[P] = \beta[P]$ *where* $\beta =$ *leastext*$(\mathcal{O}, \varphi, \alpha[Q], V)$*.*

Theorem 5.2 in [13] states that a relation R is preserved by the meet $\sqcap$ of a linear order $\mathcal{L} = (D; \leq^{\mathcal{L}})$ iff it can be defined by a conjunction of clauses of the form $(x_1 \geq^{\mathcal{L}} a_1 \vee \cdots \vee x_k \geq^{\mathcal{L}} a_k) \to (x_i >^{\mathcal{L}} b_i)$ where $x_1, \ldots, x_k$ are variables; $a_1, \ldots, a_k, b_i$ with $i \leq k$ are in D. We note that this result can be extended to all *semilattices*.

5 Lattice

In this section we consider relations in $\mathrm{Inv}(\{\sqcap, \sqcup\})$ where $\sqcap, \sqcup$ are the meet and the join of some lattice $\mathcal{L} = (D; \leq^{\mathcal{L}})$. For an example of such a relation consider $D = \{\bot, a_1, a_2, a_3, b_2, \top\}$ and the order $\leq^{\mathcal{L}}$ such that $\bot \leq^{\mathcal{L}} a_1 \leq^{\mathcal{L}} a_2 \leq^{\mathcal{L}} b_2 \leq^{\mathcal{L}} \top$ and $a_1 \leq^{\mathcal{L}} a_3 \leq^{\mathcal{L}} \top$. Thus, in particular a_2, b_2 are not comparable with a_3. It is now straightforward to show that $R := ((x_1 \geq^{\mathcal{L}} a_1 \wedge x_2 \geq^{\mathcal{L}} a_2) \to x_2 \geq^{\mathcal{L}} b_2) \wedge (x_1 \geq a_1 \to x_2 \geq a_3)$ is preserved by $\sqcup$ and $\sqcap$.

In Fig. 2, we present an algorithm for $\mathrm{GMINEXT}(\mathrm{Inv}(\{\sqcap, \sqcup\}), \mathcal{O})$ for the case where $\sqcap, \sqcup$ are the meet and the join of some lattice $\mathcal{L} = (D; \leq^{\mathcal{L}})$ that can be extended to $\mathcal{O}$. The algorithm works in polynomial time. By Proposition 1, it gives us a polynomial algorithm for the problem GMININF.

Lemma 3. *Let* $\mathcal{O} = (D; \leq^{\mathcal{O}})$ *be a linear order and* Γ *a constraint language preserved by* $\sqcap, \sqcup$ *that are the meet and the join of some lattice* $\mathcal{L} = (D; \leq^{\mathcal{L}})$ *such that* $\leq^{\mathcal{L}} \subseteq \leq^{\mathcal{O}}$*. For a given* Γ*-formula* φ *over variables* V *partitioned into* P, Z, Q *and a partial assignment* $\alpha : V_1 \to D$ *with* $V_1 \subseteq V$ *we have that* α *can be extended to a* $(\leq^{\mathcal{O}}_{(P,Z)})$*-min model of* φ *iff the algorithm Lattice returns* **true***. The algorithm Lattice works in polynomial time.*

We now turn to the hardness result. We use the notation $Z^{c,d}_{a,b}$, where $a \neq c$ and $b \neq d$, for the relation $\{(a, b), (a, d), (c, d)\}$.

Lemma 4. *Let* $\mathcal{O} = (D; \leq^{\mathcal{O}})$ *and* $a, b, c, d, e, f \in D$ *such that* $a \neq b$*,* $c <^{\mathcal{O}} d$ *and* $e <^{\mathcal{O}} f$*. Then GMININF*$(\Gamma, \mathcal{O})$ *where* $\Gamma = \{Z^{a,d}_{b,c}, Z^{b,f}_{a,e}\}$ *is coNP-hard.*

Finally we give the dichotomy.

Parameters: a linear order $\mathcal{O} = (D; \leq^{\mathcal{O}})$, a constraint language Γ over D such that Γ is preserved by both the meet $\sqcap$ and the join $\sqcup$ of the lattice $\mathcal{L} = (D; \leq^{\mathcal{L}})$ with $\leq^{\mathcal{L}} \subseteq \leq^{\mathcal{O}}$.
Data: A Γ-formula φ over variables V partitioned into P, Z, Q and a partial assignment $\alpha : V_1 \to D$ for some $V_1 \subseteq V$.
Result: If there exists a $(\leq^{\mathcal{O}}_{(P,Z)})$-min model of φ extending α, then `true`, else `false`.

1. $Q_2 := Q \setminus V_1$
2. $\alpha_G :=$ greatext$(\mathcal{O}, \varphi, \alpha, V_1 \cup Q_2)$
3. If ($\alpha_G ==$ `false`) return `false`
4. $\beta_G :=$ leastext$(\mathcal{O}, \varphi, \alpha_G, V)$
5. If ($\beta_G ==$ `false`) return `false`
6. $\beta_{CG} :=$ leastext$(\mathcal{O}, \varphi, \beta_G[Q], V)$
7. If ($\beta_G[P] == \beta_{CG}[P]$) return `true`
8. Return `false`

Fig. 2. Algorithm lattice

Theorem 5. *Let $\mathcal{O} = (D; \leq^{\mathcal{O}})$ be a linear order and $\sqcup, \sqcap$ the meet and the join of some lattice $\mathcal{L} = (D, \leq^{\mathcal{L}})$. Then we have one of the following.*

- *If $\leq^{\mathcal{L}} \subseteq \leq^{\mathcal{O}}$ or $(\leq^{\mathcal{L}})^{-1} \subseteq \leq^{\mathcal{O}}$, then GMININF(Inv($\{\sqcap, \sqcup\}$), $\mathcal{O}$) is in P.*
- *If neither $\leq^{\mathcal{L}} \subseteq \leq^{\mathcal{O}}$ nor $(\leq^{\mathcal{L}})^{-1} \subseteq \leq^{\mathcal{O}}$, then GMININF(Inv($\{\sqcap, \sqcup\}$), $\mathcal{O}$) is coNP-hard.*

6 Semilattice and a *pms*-operation

In this section we consider GMININF(Inv($\{\sqcap, f\}$), $\mathcal{O}$) in the case where $\sqcap$ is the meet of some meet-semilattice $\mathcal{L} = (D, \leq^{\mathcal{L}})$ and f is a *pms*-operation compliant with $\sqcap$. When it comes to examples of relations preserved by both operations, it is straightforward to prove that all relations definable by conjunctions of clauses of the form $(x \geq^{\mathcal{L}} d)$ and $(\neg x_1 \geq^{\mathcal{L}} d_1 \vee \cdots \vee \neg x_k \geq^{\mathcal{L}} d_k)$ are in Inv($\{\sqcap, f\}$) where f is a pms-operation that for three pairwise different $d_1, d_2, d_3 \in D$ returns $(d_1 \sqcap d_2 \sqcap d_3)$.

We now turn to the complexity of the problem. Again we work rather with GMINEXT(Inv, ($\{\sqcap, f\}$), $\mathcal{O}$) than GMININF(Inv($\{\sqcap, f\}$), $\mathcal{O}$). The procedure in Fig. 3 solves the problem in P on the condition that $\leq^{\mathcal{L}} \subseteq \leq^{\mathcal{O}}$.

Lemma 5. *Let $\mathcal{O} = (D; \leq^{\mathcal{O}})$ be a linear order and Γ a constraint language preserved by both the meet $\sqcap$ of some meet-semilattice $\mathcal{L} = (D; \leq^{\mathcal{L}})$ such that $\leq^{\mathcal{L}} \subseteq \leq^{\mathcal{O}}$ and a pms-operation f compliant with $\sqcap$. For a given Γ-formula φ over variables V partitioned into P, Z, Q and a partial assignment $\alpha : V_1 \to D$ with $V_1 \subseteq V$ we have that α can be extended to a $(\leq^{\mathcal{O}}_{(P,Z)})$-min model of φ iff the algorithm MeetPMS returns* `true`*. The algorithm MeetPMS works in polynomial time.*

Parameters: a linear order $\mathcal{O} = (D; \leq^{\mathcal{O}})$, a constraint language Γ over D preserved by both the meet $\sqcap$ of some meet-semilattice $\mathcal{L} = (D; \leq^{\mathcal{L}})$ such that $\leq^{\mathcal{L}} \subseteq \leq^{\mathcal{O}}$ and a *pms*-operation compliant with $\sqcap$.
Data: A Γ-formula φ over variables V partitioned into P, Z, Q and a partial assignment $\alpha : V_1 \to D$ for some $V_1 \subseteq V$.
Result: If there exists a $(\leq^{\mathcal{O}}_{(P,Z)})$-min model of φ extending α, then `true`, else `false`.

1. $\beta_L :=$ leastext$(\mathcal{O}, \varphi, \alpha, V)$
2. If ($\beta_L ==$ `false`) return `false`
3. $\beta_{CL} :=$ leastext$(\mathcal{O}, \varphi, \beta_L[Q], V)$
4. If ($\beta_L[P] == \beta_{CL}[P]$) return `true`
5. Return `false`

Fig. 3. Algorithm MeetPMS

We close this section by complementing Lemma 5 with a hardness result.

Theorem 6. *Let $\mathcal{O} = (D; \leq^{\mathcal{O}})$ be a linear order, $\sqcap$ the meet of some meet-semilattice $\mathcal{L} = (D; \leq^{\mathcal{L}})$ and f a pms-operation compliant with $\sqcap$. Then we have one of the following.*

- *If $\leq^{\mathcal{L}} \subseteq \leq^{\mathcal{O}}$, then GMININF(Inv($\{\sqcap, f\}$), $\mathcal{O}$) is in P.*
- *If $\leq^{\mathcal{L}} \not\subseteq \leq^{\mathcal{O}}$, then GMININF(Inv($\{\sqcap, f\}$), $\mathcal{O}$) is coNP-hard.*

7 Summary and Future Work

In this article we have systematically studied the complexity of the minimal inference problem over arbitrary finite domains. We considered a version of the problem parameterized by a constraint language and a linear order. We provided a dichotomy for maximal conservative languages: between Π_2^P-complete and coNP-complete and two tractability results complemented by coNP-hardness results. This gives two dichotomies: between coNP-complete and in P. Furthermore, one of the tractability results is based on a newly discovered operation: a pms-operation. Identifying tractable classes and appropriate polymorphisms is crucial when one works in algebraic approach.

We believe that our research will soon result in more advanced classifications, e.g., for all conservative languages or all languages over the three-element domain. Both classifications were provided for the CSP, see [2,3].

References

1. Bulatov, A.A.: A graph of a relational structure and constraint satisfaction problems. In: Proceedings of the Symposium on Logic in Computer Science (LICS), Turku, Finland (2004)

2. Bulatov, A.A.: A dichotomy theorem for constraint satisfaction problems on a 3-element set. J. ACM **53**(1), 66–120 (2006)
3. Bulatov, A.A.: Complexity of conservative constraint satisfaction problems. ACM Trans. Comput. Log. **12**(4), 24 (2011)
4. Bulatov, A.A., Krokhin, A.A., Jeavons, P.: The complexity of maximal constraint languages. In: Proceedings of the Symposium on Theory of Computing (STOC), pp. 667–674 (2001)
5. Bulatov, A.A., Krokhin, A.A., Jeavons, P.G.: Classifying the complexity of constraints using finite algebras. SIAM J. Comput. **34**, 720–742 (2005)
6. Cadoli, M., Lenzerini, M.: The complexity of propositional closed world reasoning and circumscription. J. Comput. Syst. Sci. **48**(2), 255–310 (1994)
7. Doherty, P., Kachniarz, J., Szałas, A.: Using contextually closed queries for local closed-world reasoning in rough knowledge databases. In: Skowron, A., Polkowski, L., Pal, S.K. (eds.) Rough-Neural Computing: Techniques for Computing with Words. Cognitive Technologies, pp. 219–250. Springer, Heidelberg (2004)
8. Durand, A., Hermann, M.: The inference problem for propositional circumscription of affine formulas is coNP-complete. In: STACS, pp. 451–462 (2003)
9. Durand, A., Hermann, M., Nordh, G.: Trichotomies in the complexity of minimal inference. Theor. Comput. Syst. **50**(3), 446–491 (2012)
10. Eiter, T., Gottlob, G.: Propositional circumscription and extended closed-world reasoning are IIp2-complete. Theor. Comput. Sci. **114**(2), 231–245 (1993)
11. Gelfond, M., Przymusinska, H., Przymusinski, T.C.: On the relationship between circumscription and negation as failure. Artif. Intell. **38**(1), 75–94 (1989)
12. Jeavons, P., Cohen, D., Gyssens, M.: Closure properties of constraints. J. ACM **44**(4), 527–548 (1997)
13. Jeavons, P., Cooper, M.C.: Tractable constraints on ordered domains. Artif. Intell. **79**(2), 327–339 (1995)
14. Jonsson, P.: The complexity of mincsp and csp, personal communication
15. Kirousis, L.M., Kolaitis, P.G.: A dichotomy in the complexity of propositional circumscription. Theor. Comput. Syst. **37**(6), 695–715 (2004)
16. Krokhin, A.A., Zivny, S. (eds.): The Constraint Satisfaction Problem: Complexity and Approximability, Dagstuhl Follow-Ups, vol. 7. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2017)
17. McCarthy, J.: Circumscription - a form of non-monotonic reasoning. Artif. Intell. **13**(1–2), 27–39 (1980)
18. McCarthy, J.: Applications of circumscription to formalizing common-sense knowledge. Artif. Intell. **28**(1), 89–116 (1986)
19. Nordh, G., Jonsson, P.: An algebraic approach to the complexity of propositional circumscription. In: Proceedings of the Symposium on Logic in Computer Science (LICS), pp. 367–376 (2004)
20. Przymusinski, T.C.: Three-valued nonmonotonic formalisms and semantics of logic programs. Artif. Intell. **49**(1–3), 309–343 (1991)
21. Rosenberg, I.G.: Minimal Clones I: The Five Types. Lectures in Universal Algebra (Proc. Conf. Szeged, 1983). Colloq. Math. Soc. J. Bolyai, vol. 43, pp. 405–427 (1986)
22. Schaefer, T.J.: The complexity of satisfiability problems. In: Proceedings of the Symposium on Theory of Computing (STOC), pp. 216–226 (1978)