

The Parameterized Complexity of Abduction*

Michael R. Fellows

michael.fellows@cdu.edu.au
Charles Darwin University (Australia)

Frances A. Rosamond

frances.rosamond@cdu.edu.au
Charles Darwin University (Australia)

Andreas Pfandler

pfandler@dbai.tuwien.ac.at
Vienna University of Technology (Austria)

Stefan Rümmele

ruemmele@dbai.tuwien.ac.at
Vienna University of Technology (Austria)

Abstract

Abduction belongs to the most fundamental reasoning methods. It is a method for reverse inference, this means one is interested in explaining observed behavior by finding appropriate causes. We study logic-based abduction, where knowledge is represented by propositional formulas. The computational complexity of this problem is highly intractable in many interesting settings. In this work we therefore present an extensive parameterized complexity analysis of abduction within various fragments of propositional logic together with (combinations of) natural parameters.

Introduction

The young PhD student Bob wakes up during the night and discovers that the light in his room is not working. Looking out of the window, he sees that in his neighbor's flat the light is on. He reasons that there is no blackout. Therefore he concludes that either the light bulb is broken or that he had forgotten to pay his bills.

This kind of reasoning is called abductive reasoning (Abduction for short) and belongs to the most fundamental reasoning methods. In contrast to deductive reasoning, it is a method for reverse inference. This means one is interested in explaining observed behavior by finding appropriate causes. It is widely believed that humans use abduction in their reasoning when searching for diagnostic explanations. In this paper we study logic-based abduction, where knowledge is represented by a (set of) propositional formula(s). This reasoning problem has many important applications such as system and medical diagnosis, planning, configuration and database updates.

In the propositional abduction problem we are given a propositional theory T , a set of hypotheses H and a set of manifestations M . The task is to find a solution $S \subseteq H$ such that $S \cup T$ is consistent and logically entails M . Thus, we require that the situation represented in S is possible in the system described by T and that S explains the observations.

The classical complexity of abduction has been extensively studied in the literature (Selman and Levesque 1990;

Eiter and Gottlob 1995; Creignou and Zanuttini 2006; Nordh and Zanuttini 2008; Creignou, Schmidt, and Thomas 2010). Unfortunately the computational complexity of this problem turned out to be highly intractable in many interesting settings, which imposes a severe obstacle to the broad applicability of this formalism. Although syntactical fragments of lower complexity were explored, there is still the need for improvement.

A successful way of dealing with intractability is the concept of parameterized complexity. There (structural) parameters and their influence on the complexity of the problem are studied. The aim is to find so called fixed-parameter tractable (FPT) algorithms with respect to some parameter k , i.e. algorithms with runtime $f(k) \cdot n^{O(1)}$, where f is some function depending only on k . Such algorithms are considered to be tractable when the parameter value is sufficiently small. For more details we refer to the next section.

Abduction has recently been shown to be fixed-parameter tractable when parameterized by treewidth (Gottlob, Pichler, and Wei 2010), but all other possible parameters remained unexplored. In this work we consider various fragments of propositional logic, namely Horn, definite Horn and Krom together with (combinations of) natural parameters.

Very related to the quest of searching for FPT algorithms is the search for efficient preprocessing techniques. More precisely the goal is to obtain in polynomial time an equivalent instance (called *kernel*) whose size is bounded by a function of the parameter. While it is trivial to construct a kernel of exponential size for an arbitrary FPT problem, obtaining in polynomial time a kernel of size polynomial in the parameter remains a central algorithmic challenge that may or may not be achievable. Our main contributions are the following:

- We perform a classical complexity analysis of a new abduction problem asking for solutions of certain size.
- We present several fixed-parameter tractability results and even a polynomial kernel in case of Krom formulas.
- For the remaining fixed-parameter tractable problems we prove that no polynomial kernel exists unless the Polynomial Hierarchy collapses to the third level.
- For parameterizations where we do not show fixed-parameter tractability, we present parameterized intractability results by either proving completeness in the W-hierarchy or hardness for para-NP.

An overview of the results can be found in Tables 1 and 2.

*The research of M.R. Fellows and F.A. Rosamond is supported by the Australian Research Council. The research of A. Pfandler and S. Rümmele is supported by the Austrian Science Fund (FWF): P20704-N18.

Copyright © 2012, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

	ABD[PROP]	ABD[HORN]	ABD[HORN] _{≤/=}	ABD[DEFHORN]	ABD[DEFHORN] _{≤/=}
M	para-NP-h**	para-NP-h**	para-NP-h (Thm 5)	P *	para-NP-h (Thm 5)
$H, M = 1$	para-NP-h (Prop 4)	npk (Thm 14)	npk (Thm 11)	P *	npk (Thm 11)
k	–	–	W[P]-c (Cor 9)	–	W[P]-c (Cor 9)
$k, M = 1$	–	–	W[P]-c (Thm 8)	–	W[P]-c (Thm 8)
tw	npk (Thm 14)	npk (Thm 14)	npk (Cor 15/16)	P *	npk (Cor 12)
(τ, H)	npk (Thm 14)	npk (Thm 14)	npk (Cor 15/16)	P *	npk (Cor 12)
V	npk (Thm 14)	npk (Thm 14)	npk (Cor 15/16)	P *	FPT (Prop 13)

* cf. (Eiter and Gottlob 1995)

** cf. (Selman and Levesque 1990)

Table 1: Results for PROP, HORN and DEFHORN.

	ABD[KROM]	ABD[KROM] _{<}	ABD[KROM] ₌
M	W[1]-c (Thm 26)	W[1]-c (Thm 25)	para-NP-h (Thm 7)
(H, M)	pk (Thm 29)	pk (Thm 29)	pk (Thm 29)
k	–	W[2]-c (Thm 21)	W[2]-c (Thm 21)
(k, M)	–	W[1]-c (Thm 22)	W[1]-c (Thm 23)
$k, M = 1$	–	P *	W[1]-c (Thm 23)
τ	npk (Thm 27)	npk (Thm 27)	npk (Thm 27)
V	pk (Thm 29)	pk (Thm 29)	pk (Thm 29)

* cf. (Creignou and Zanuttini 2006)

Table 2: Results for KROM.

Preliminaries

Let PROP be the class of all (*propositional*) formulas. The class of formulas in conjunctive normal form is denoted by CNF. It is convenient to view a formula in CNF also as a set of clauses and a clause as a set of literals. $KROM \subseteq CNF$ denotes the class of all formulas having clause size at most 2. *Horn* (*Definite Horn*) formulas are CNF formulas with at most (resp. exactly) one positive literal per clause.

We use standard notation and denote by $var(\varphi)$ the set of propositional variables occurring in a formula φ .

Let $Res(\varphi)$ be an operator extending $\varphi \in CNF$ by iteratively applying resolution and dropping tautological clauses until a fixed-point is reached. Applying resolution adds the clause $C \cup D$ to φ if $C \cup \{x\} \in \varphi$ and $D \cup \{\neg x\} \in \varphi$. Resolution on Krom formulas will always yield a Krom formula. In that case $Res(\varphi)$ can be computed in polynomial time. Let C be a non-tautological clause then $C \in Res(\varphi)$ if and only if $\varphi \models C$. For details, see e.g. (Leitsch 1997).

Let $\mathcal{C} \subseteq PROP$. A (*propositional*) *abduction instance* for \mathcal{C} -theories consists of a tuple $\langle V, H, M, T \rangle$, where V is the set of variables, $H \subseteq V$ is the set of hypotheses, $M \subseteq V$ is the set of manifestations, and $T \in \mathcal{C}$ is the theory, a formula over V . It is required that $M \cap H = \emptyset$.

Definition 1. Let $\mathcal{P} = \langle V, H, M, T \rangle$ be an abduction instance. $S \subseteq H$ is a *solution* (or *explanation*) to \mathcal{P} if $T \cup S$ is consistent and $T \cup S \models M$ (entailment). $Sol(\mathcal{P})$ denotes the set of all solutions to \mathcal{P} .

Let $\mathcal{C} \subseteq PROP$. The *solvability problem for propositional abduction* ABD[\mathcal{C}] for \mathcal{C} -theories is the following problem:

ABD[\mathcal{C}]
<i>Instance:</i> An abduction instance \mathcal{P} .
<i>Problem:</i> Decide $Sol(\mathcal{P}) \neq \emptyset$.

We introduce a version of the abduction problem where the size of the solutions is limited. Let $\sim \in \{=, \leq\}$.

ABD[\mathcal{C}] _{\sim}
<i>Instance:</i> An abduction instance \mathcal{P} and an integer k .
<i>Problem:</i> Is there a set $S \in Sol(\mathcal{P})$ s.t. $ S \sim k$.

Parameterized algorithmics (cf. (Downey and Fellows 1999; Flum and Grohe 2006; Niedermeier 2006)) is an approach to finding optimal solutions for NP-hard problems. The idea is to accept the seemingly inevitable combinatorial explosion, but to confine it to one aspect of the problem, the parameter. More precisely, a problem is fixed-parameter tractable (FPT) with respect to a parameter k if there is an algorithm solving any problem instance of size n in $f(k) \cdot n^{\mathcal{O}(1)}$ time for some computable function f . Analogously to classical complexity theory, (Downey and Fellows 1999) developed a framework providing reducibility and completeness notions. A *parameterized reduction* of a parameterized problem Π to a parameterized problem Π' is an FPT algorithm that transforms an instance (I, k) of Π to an instance (I', k') of Π' such that: (1) (I, k) is a yes-instance of Π if and only if (I', k') is a yes-instance of Π' , and (2) $k' = g(k)$, that is, k' depends only on k . This notion leads to a hierarchy of parameterized complexity classes, principally: $FPT \subseteq W[1] \subseteq W[2] \subseteq \dots \subseteq W[P] \subseteq XP$ where XP is the class of parameterized problems solvable in time $\mathcal{O}(n^{g(k)})$ for some function g . A parameterized problem Π is *para-NP-hard* if there is some fixed k for which Π restricted to instances (x, k) is NP-hard.

A common method in parameterized algorithmics is to provide polynomial-time executable data-reduction rules (Downey and Fellows 1999). It is easily shown that a parameterized problem Π is FPT if and only if there is a polynomial time data-reduction (or *kernelization*) algorithm that transforms a problem instance (I, k) of Π into an instance (I', k') of Π such that: (1) (I, k) is a yes-instance if and only if (I', k') is a yes-instance, (2) $k' \leq f(k)$, and (3)

$|I'| \leq g(k)$ for functions f and g depending only on k . Although in general the kernelization bound $g(k)$ may be exponential in k , it has been shown that many FPT problems admit *polynomial kernels*, that is, polynomial time kernelization algorithms where the kernelization bound is a polynomial function of k , $g(k) = k^{O(1)}$.

Recently, lower bound methods for polynomial kernelization have been developed, based on the following notion.

Definition 2. A parameterized problem $P \subseteq \Sigma^* \times \mathbb{N}$ is *compositional* if there exists an algorithm that computes, given a sequence $(x_1, k), \dots, (x_t, k) \in \Sigma^* \times \mathbb{N}$, a new instance $(x', k') \subseteq \Sigma^* \times \mathbb{N}$ s.t. the following properties hold: (1) The algorithm requires time polynomial in $\sum_{i=1}^t |x_i| + k$, (2) (x', k') is a yes-instance if and only if there is some $1 \leq i \leq t$ s.t. (x_i, k) is a yes-instance, and (3) $k' \leq k^{O(1)}$.

Theorem 3 ((Bodlaender et al. 2009)). *Let Π be a parameterized problem s.t. the unparameterized version of Π is NP-complete. If Π is compositional, then it does not admit a polynomial kernel unless the Polynomial Hierarchy collapses to the third level ($\text{PH} = \Sigma_3^P$).*

A polynomial parameter and time (PPT) reduction is a polynomial time reduction increasing the parameter only polynomially. (Bodlaender, Thomassé, and Yeo 2011) showed that a PPT reduction from Π to Π' preserves polynomial kernels if the unparameterized version of Π is NP-complete and the unparameterized version of Π' is in NP.

In the sequel, we will consider parameterizations by the *vertex cover number* and the *treewidth* of the *primal graphs* of abduction instances. For an instance $\mathcal{P} = \langle V, H, M, T \rangle$, such a graph has vertex set V and there is an edge between two vertices if they occur together in a clause of T .

The *vertex cover number* $\tau(G)$ of a graph G is the size of the smallest vertex cover of G . A vertex cover is a set of vertices containing at least one endpoint of each edge.

The treewidth $tw(G)$ is a measure for its “tree-likeness”, cf. (Robertson and Seymour 1986) for a definition. All our hardness results with respect to parameter τ carry over to parameterizing by tw , since bounded τ implies bounded tw .

An independent set is a subset of the vertices which does not contain both endpoints of any edge. INDEPENDENT SET asks for such a set of size k . The problem is W[1]-complete when parameterized by k .

A problem parameterized by k is in W[P] if there exists a nondeterministic algorithm running in time $f(k) \cdot n^{O(1)}$ using only $f'(k) \cdot \log n$ many nondeterministic steps, where f and f' are computable functions.

To show membership in the W-hierarchy we will use $\text{MC}[\Sigma_{t,u}]$, the model-checking problem over $\Sigma_{t,u}$ formulas. The class $\Sigma_{t,u}$ contains all first-order formulas of the form $\exists \bar{x}_1 \forall \bar{x}_2 \exists \bar{x}_3 \dots Q \bar{x}_t \varphi(\bar{x}_1, \dots, \bar{x}_t)$, where φ is quantifier free and Q is an \exists if t is odd and a \forall if t is even, and the quantifier blocks – with the exception of the first \exists block – are of length at most u . Given a finite structure \mathcal{A} and a formula $\varphi \in \Sigma_{t,u}$, $\text{MC}[\Sigma_{t,u}]$ asks whether \mathcal{A} is a model of φ . When parameterized by $|\varphi|$, $\text{MC}[\Sigma_{t,u}]$ is W[t]-complete for $t \geq 1$, $u \geq 1$ (Downey, Fellows, and Regan 1998).

For $m \in \mathbb{N}$, we use $[m]$ to denote the set $\{1, \dots, m\}$. Finally, $\mathcal{O}^*(\cdot)$ is defined in the same way as $\mathcal{O}(\cdot)$ but ignores polynomial factors.

Classical Complexity

Early work on the complexity of propositional abduction was done by (Selman and Levesque 1990). Among others they showed that $\text{ABD}[\text{HORN}]$ is NP-complete. A systematic complexity analysis was done by (Eiter and Gottlob 1995). Their results include that $\text{ABD}[\text{PROP}]$ is Σ_2^P -complete and that $\text{ABD}[\text{DEFHORN}]$ is in P. Note that the hardness results for PROP and HORN hold even for $|M| = 1$ since in those classes one can add a new clause to the theory where all existing manifestations imply a single new manifestation. The problem $\text{ABD}[\text{KROM}]$ was shown to be NP-complete (Nordh and Zanuttini 2008) while it is in P (Creignou and Zanuttini 2006) when restricted to $|M| = 1$. In the latter result, they use the fact that $\text{ABD}[\text{KROM}]$ restricted to a single manifestation has a solution if and only if it has a solution of size ≤ 1 . Therefore, $\text{ABD}[\text{KROM}]_{\leq}$ with $|M| = 1$ is in P by the same argument. The following proposition is a consequence from a remark in (Eiter and Gottlob 1995), which states that deciding $S \in \text{Sol}(\mathcal{P})$ for an instance \mathcal{P} is DP-complete. This also shows that for PROP parameters H and M are not sufficient.

Proposition 4. $\text{ABD}[\text{PROP}]$ and $\text{ABD}[\text{PROP}]_{\leq/ =}$ are DP-complete for $|H| = 0$, even if $|M| = 1$.

In order to motivate a parameterized complexity analysis, the remainder of this section is dedicated to showing that the problems $\text{ABD}[\text{HORN}]_{\leq/ =}$, $\text{ABD}[\text{DEFHORN}]_{\leq/ =}$, and $\text{ABD}[\text{KROM}]_{\leq/ =}$ are intractable in the classical setting. According to Definition 1, our reductions must ensure both consistency and entailment.

Theorem 5. $\text{ABD}[\text{HORN}]_{\leq/ =}$ and $\text{ABD}[\text{DEFHORN}]_{\leq/ =}$ are NP-complete, even if $|M| = 1$.

Proof. Membership is trivial. We show hardness by reduction from VERTEX COVER. Given a graph (N, E) and integer k . Does (N, E) have a vertex cover of size \leq (resp. $=$) k ? We construct an instance $(\langle V, H, M, T \rangle, k)$ of $\text{ABD}[\text{DEFHORN}]_{\leq/ =}$ as follows. Let $V := N \cup E \cup \{m\}$, where m is a new variable, $H := N$, $M := \{m\}$, and $T := (m \vee \bigvee_{e \in E} \neg e) \wedge \bigwedge_{\{x,y\} \in E} ((x \rightarrow e) \wedge (y \rightarrow e))$. Note that $T \cup S$ is satisfiable for every $S \subseteq H$. The first clause of T ensures that m is entailed if and only if each $e \in E$ is entailed. This in turn is the case if and only if S contains an endpoint of each edge and therefore is a vertex cover of (N, E) . \square

Corollary 6. $\text{ABD}[\text{KROM}]_{\leq}$ is NP-complete.

Proof. This can be shown similarly to the proof of Theorem 5. Thereby the first clause of theory T is removed, all edges are used as manifestations $M := E$, and we set $V := N \cup E$. \square

Theorem 7. $\text{ABD}[\text{KROM}]_{=}$ is NP-complete even if $|M| = 1$.

Proof sketch. Membership holds trivially and hardness can be shown by a reduction from INDEPENDENT SET. \square

Parameterized Complexity

In this section we study the parameterized complexity of abduction. The first part is mainly dedicated to the HORN and DEFHORN fragments, whereas the second part deals with the KROM fragment. Unless otherwise specified, V , H , M , and T refer to the components of an abduction instance (see Definition 1). Additionally, k denotes the bound on the solution size. When parameterizing by the cardinality of some set, we omit the vertical bars, e.g. “parameterized by M ” means parameterized by $|M|$. Two parameters together are denoted by a tuple, e.g. (k, M) instead of $k + |M|$.

We start by showing that the HORN fragments are intractable when parameterized by the solution size k .

Theorem 8. $\text{ABD}[\text{HORN}]_{\leq/} \text{ and } \text{ABD}[\text{DEFHORN}]_{\leq/} \text{ parameterized by } (k, M) \text{ are } \mathbf{W[P]} \text{-complete even if } |M|=1.$

Proof. For the $\mathbf{W[P]}$ -membership, note that the problem can be solved by nondeterministically guessing k times a (not necessarily distinct in case of \leq) hypothesis, each of which can be described by $\log n$ bits, and then deterministically verifying consistency and entailment. This checking part can be done in polynomial time for HORN as well as DEFHORN theories.

We show hardness by reduction from WEIGHTED MONOTONE CIRCUIT SAT, where an instance is given by a monotone circuit C (i.e. without NOT-gates) and an integer k . The question is whether there exists an assignment setting at most/exactly k many input gates to true s.t. the output gate is true as well. This problem is $\mathbf{W[P]}$ -complete, when parameterized by k , even when every AND-gate and every OR-gate is binary. We construct an instance $((V, H, M, T), k)$ for the abduction problem. First, we introduce a new variable for each gate of C and call the resulting set V . Let H be the set of input gates and let $M := \{m\}$, where m represents the output gate. Theory T is constructed as follows: For each AND-gate a with input i_1 and i_2 , we add $(i_1 \wedge i_2 \rightarrow a)$ to T . For each OR-gate o with input i_1 and i_2 , we add $(i_1 \rightarrow o) \wedge (i_2 \rightarrow o)$ to T . By construction, for each set $S \subseteq H$, $T \cup S$ is consistent. Furthermore, $T \cup S \models M$ if and only if activating only the input gates in S satisfies C . \square

Since the membership result above only uses parameter k , we immediately get the following corollary.

Corollary 9. $\text{ABD}[\text{HORN}]_{\leq/} \text{ and } \text{ABD}[\text{DEFHORN}]_{\leq/} \text{ parameterized by } k \text{ are } \mathbf{W[P]} \text{-complete.}$

On the other hand, the parameterization by the number of hypotheses is trivially FPT.

Proposition 10. $\text{ABD}[\text{HORN}] \text{ and } \text{ABD}[\text{HORN}]_{\leq/} \text{ parameterized by } H \text{ are FPT, solvable in time } \mathcal{O}^*(2^{|H|}).$

Proof. Let (V, H, M, T) be an abduction instance with theory $T \in \text{HORN}$. There exist $2^{|H|}$ many subsets of H . Given $S \subseteq H$, checking if $S \in \text{Sol}(\mathcal{P})$ can be done in polynomial time for HORN-theories (Eiter and Gottlob 1995). \square

Despite the problems being trivially FPT, it turns out that they do not admit a polynomial kernel, even when adding the solution size as a parameter. This follows from the more general result below.

Theorem 11. $\text{ABD}[\text{DEFHORN}]_{\leq} \text{ and } \text{ABD}[\text{DEFHORN}]_{=}$ parameterized by H do not admit a polynomial kernel unless the Polynomial Hierarchy collapses, even if $|M| = 1$.

Proof. We show the result for $\text{ABD}[\text{DEFHORN}]_{\leq}$ by a PPT reduction from SMALL UNIVERSE HITTING SET, where an instance is given by a family of sets $\mathcal{F} = \{F_1, \dots, F_l\}$ over an universe $U = \bigcup_{i=1}^l F_i$ with $|U| = d$, and an integer k . The question is to find a set $U' \subseteq U$ of cardinality $\leq k$ s.t. each set in the family has a non-empty intersection with U' . This problem, parameterized by k and d does not admit a polynomial kernel unless the Polynomial Hierarchy collapses (Dom, Lokshtanov, and Saurabh 2009). We construct an $\text{ABD}[\text{DEFHORN}]_{\leq}$ instance $((V, H, M, T), k)$ as follows. Let $X = \{x_1, \dots, x_l\}$ be a set of new variables representing elements of \mathcal{F} . Let $H := U$, $M := \{m\}$, where m is a new variable, $V := H \cup X \cup M$, and $T := (x_1 \wedge \dots \wedge x_l \rightarrow m) \wedge \bigwedge_{i \in [l]} \bigwedge_{e \in F_i} (e \rightarrow x_i)$. Manifestation m is entailed if and only if all variables in X are entailed. Variable $x_i \in X$ is entailed if and only if at least one of the elements in the set F_i is selected. Therefore, a solution $S \subseteq H$ corresponds to a hitting set of the same size. We can reduce $\text{ABD}[\text{DEFHORN}]_{\leq}$ to $\text{ABD}[\text{DEFHORN}]_{=}$, due to the monotonicity of DEFHORN formulas. To be more precise, if there is a solution $S \subseteq H$ of an $\text{ABD}[\text{DEFHORN}]_{\leq}$ instance, then also all $S' \supset S$ are solutions as well. \square

A similar reduction can be used to show that even the aggregate parameterization with both τ and H does not yield a polynomial kernel. Since $\text{tw}(G) \leq \tau(G)$ for every graph G , we immediately obtain the same result for parameter tw .

Corollary 12. $\text{ABD}[\text{DEFHORN}]_{\leq} \text{ and } \text{ABD}[\text{DEFHORN}]_{=}$ parameterized by (τ, H) do not admit a polynomial kernel unless the Polynomial Hierarchy collapses.

Proof. This can be shown by using the same reduction as in the proof of Theorem 11, but without the restriction to a single manifestation. That means, the conjunct $(x_1 \wedge \dots \wedge x_l \rightarrow m)$ is removed from T and $M := X$. Observe that $U = H$ is a vertex cover of the abduction instance. \square

For parameter V , even $\text{ABD}[\text{PROP}]$ is trivially FPT.

Proposition 13. $\text{ABD}[\text{PROP}]$ parameterized by V is FPT, solvable in time $\mathcal{O}^*(2^{2^{|V|}})$.

Proof. There are at most $2^{|H|} \leq 2^{|V|}$ possible solution candidates. For each of them we need to test consistency and entailment, which can be done in time $\mathcal{O}(2^{|H|}(n+2^{|V|}n))$. \square

Again we show that this parameterization is not sufficient for a polynomial kernel.

Theorem 14. $\text{ABD}[\text{HORN}]$ parameterized by V does not admit a polynomial kernel unless the Polynomial Hierarchy collapses.

Proof. We show that the problem parameterized by (V, H) is compositional. Parameter H does not change the problem since $H \subseteq V$, but allows us to assume in the composition that all instances have the same number of hypotheses. Let $\mathcal{P}_1, \dots, \mathcal{P}_t$ be a given sequence of instances of $\text{ABD}[\text{HORN}]$ where $\mathcal{P}_i = (V_i, H_i, M_i, T_i)$, $1 \leq i \leq t$, with $|V_i| = d$ and $|H_i| = e$. We assume without loss of generality that $V_i = V_j$

and $H_i = H_j$ for all $1 \leq i < j \leq t$ since otherwise we could rename the variables. We distinguish two cases.

Case 1: $t > 2^{2d}$. Let $n := \max_{i=1}^t \|\mathcal{P}_i\|$. Whether \mathcal{P}_i has a solution can be decided in time $\mathcal{O}(2^{2d}n)$ by the FPT algorithm from Proposition 13. We can check whether at least one of $\mathcal{P}_1, \dots, \mathcal{P}_t$ has a solution in time $\mathcal{O}(t2^{2d}n) \leq \mathcal{O}(t^2n)$ which is polynomial in $\sum_{i=1}^t \|\mathcal{P}_i\|$. If some \mathcal{P}_i has a solution, we output \mathcal{P}_i ; otherwise we output \mathcal{P}_1 , which has no solution. Hence, we have a composition algorithm in Case 1.

Case 2: $t \leq 2^{2d}$. We construct a new instance $\mathcal{P} := \langle V, H, M, T \rangle$ of ABD[HORN] as follows. Let $s := \lceil \log_2 t \rceil$. Let $V := V_1 \cup X \cup X' \cup Y \cup \{m\}$, where $X := \{x_1, \dots, x_s\}$, $X' := \{x'_1, \dots, x'_s\}$, $Y := \{y_1, \dots, y_s\}$, and m are $3s + 1$ new variables. Let $H := H_1 \cup X \cup X'$ and let $M := Y \cup \{m\}$. For each theory T_i we create a new one $T'_i := T_i \cup \{\neg m' \mid m' \in M_i\} \cup \{m\}$. Let C_1, \dots, C_{2^s} be a sequence of all 2^s possible clauses $\{l_1, \dots, l_s\}$ where l_j is either $\neg x_j$ or $\neg x'_j$, $1 \leq j \leq s$. For each theory T'_i we create a new one $T''_i := \{C \cup C_i \mid C \in T'_i\}$. Finally, let $T := \bigcup_{i=1}^t T''_i \cup \bigcup_{j=1}^s \{\{\neg x_j, y_j\}, \{\neg x'_j, y_j\}, \{\neg x_j, \neg x'_j\}\}$. Since the y_j 's are manifestations, the clauses $\{\neg x_j, y_j\}$, $\{\neg x'_j, y_j\}$, and $\{\neg x_j, \neg x'_j\}$, $1 \leq j \leq s$, ensure the equivalence $\neg x_j \equiv x'_j$ which is not directly expressible in HORN. Because of this equivalence, a solution S of \mathcal{P} has to contain exactly one of x_j or x'_j for each $1 \leq j \leq s$. Therefore, there is exactly one subclause (in the construction they were merged with other clauses) C_l , $1 \leq l \leq 2^s$, which is not satisfied by S . Hence, all theories T_i with $i \neq l$ are trivially satisfied and cannot entail the manifestation m . Thus, \mathcal{P} has a solution if and only if \mathcal{P}_l has one. Since $|V|$ and $|H|$ is polynomial in d respectively e , we have also a composition algorithm in Case 2.

Applying Theorem 3, the result follows. \square

Since ABD[HORN] has a solution if and only if there is a solution for $\text{ABD}[\text{HORN}]_{\leq}$ of size $\leq |H|$, we immediately get the following corollary.

Corollary 15. $\text{ABD}[\text{HORN}]_{\leq}$ parameterized by V does not admit a polynomial kernel unless the Polynomial Hierarchy collapses.

Corollary 16. $\text{ABD}[\text{HORN}]_{=}$ parameterized by V does not admit a polynomial kernel unless the Polynomial Hierarchy collapses.

Proof. We present a PPT reduction from ABD[HORN] parameterized by V . Let $\mathcal{P} = \langle V, H, M, T \rangle$ be an instance of ABD[HORN] with $H = \{h_1, \dots, h_e\}$. We construct a new instance $\mathcal{P}' := (\langle V \cup H' \cup M', H \cup H', M \cup M', T' \rangle, k)$ for $\text{ABD}[\text{HORN}]_{=}$ as follows. Let $H' := \{h'_1, \dots, h'_e\}$ and $M' := \{m_1, \dots, m_e\}$ be new variables. Let $k := e$ and let $T' := T \cup \bigcup_{i=1}^e \{\{\neg h_i, \neg h'_i\}, \{\neg h_i, m_i\}, \{\neg h'_i, m_i\}\}$. Then \mathcal{P} has a solution if and only if \mathcal{P}' has a solution of size k . The reason for this is that the clauses $\{\neg h_i, \neg h'_i\}, \{\neg h_i, m_i\}, \{\neg h'_i, m_i\}$ enforce that a solution S contains exactly one of the two hypotheses h_i, h'_i for each $1 \leq i \leq e$. Selecting h'_i in \mathcal{P}' is equivalent to not selecting h_i in \mathcal{P} , since the variables h'_i occur nowhere in T . \square

Next we study the KROM fragment. Thereby the following preprocessing function will be very useful.

Definition 17. Given an abduction instance for KROM theories $\langle V, H, M, T \rangle$. We define the function $\text{TrimRes}(T, H, M) := \{C \in \text{Res}(T) \mid C \subseteq X\}$, with $X = H \cup M \cup \{\neg x \mid x \in (H \cup M)\}$.

In other words, the function $\text{TrimRes}(T, H, M)$ first computes the closure under resolution $\text{Res}(T)$ and then keeps only those clauses which solely consist of hypotheses and manifestations.

First we show that while the complexity is lower than in the HORN fragment when parameterized by k , the problem remains intractable. Due to space limitations we omit the proofs of all lemmas.

Lemma 18. Let T be a satisfiable KROM theory and let S be a set of propositional variables. Then $T \wedge S$ is unsatisfiable if and only if there exist $x, y \in S$ such that $T \models \neg x \vee \neg y$.

Lemma 19. Let $\langle V, H, M, T \rangle$ be an abduction instance for KROM theories and let $S \subseteq H$. Then $T \wedge S$ is satisfiable if and only if $\text{TrimRes}(T, H, M) \wedge S$ is satisfiable.

Lemma 20. Let $\langle V, H, M, T \rangle$ be an abduction instance for KROM theories, $S \subseteq H$, $m \in M$, and $T \wedge S$ be satisfiable. Then $T \wedge S \models m$ implies that either $\{m\} \in \text{TrimRes}(T, H, M)$ or there exists some $h \in S$ with $\{\neg h, m\} \in \text{TrimRes}(T, H, M)$.

Theorem 21. $\text{ABD}[\text{KROM}]_{\leq}$ and $\text{ABD}[\text{KROM}]_{=}$ parameterized by k are W[2]-complete.

Proof. We show membership by reducing an abduction instance $(\langle V, H, M, T \rangle, k)$ to an $\text{MC}[\Sigma_{2,1}]$ instance (\mathcal{A}, φ) . First we check whether the empty set is already a solution. In that case we return a tautology. In the other case we first ensure that T is satisfiable and compute $\text{TrimRes}(T, H, M)$ as defined in Definition 17. We construct structure $\mathcal{A} := \langle A, \text{hyp}, \text{mani}, \text{fact}, \text{cl}, \text{pos}, \text{neg} \rangle$ as follows. Domain A contains an element for each hypothesis in H , each manifestation in M and two distinct elements denoted by *positive* and *negative*. Let the sets *hyp* (resp. *mani*) represent the hypotheses (resp. manifestations). We use the following notation. Let l be a literal, then $\text{pol}(l)$ denotes the element *positive* (resp. *negative*) if l is a positive (resp. negative) literal. Relation *fact* contains the pairs $\{(\text{pol}(l), l) \mid \{l\} \in \text{TrimRes}(T, H, M)\}$. Relation *cl* contains the tuples $\{(\text{pol}(l_1), l_1, \text{pol}(l_2), l_2) \mid \{l_1, l_2\} \in \text{TrimRes}(T, H, M), l_1 \neq l_2\}$. Finally, we have that $\text{pos} := \{\text{positive}\}$ and that $\text{neg} := \{\text{negative}\}$. We define

$$\begin{aligned} \psi &:= \text{pos}(p) \wedge \text{neg}(n) \wedge \bigwedge_{i \in [k]} \text{hyp}(h_i) \wedge \\ &\quad \bigwedge_{i \in [k]} \neg \text{fact}(n, h_i) \wedge \bigwedge_{i, j \in [k]} \neg \text{cl}(n, h_i, n, h_j), \\ \chi[x] &:= \text{fact}(p, x) \vee \bigvee_{j \in [k]} \text{cl}(n, h_j, p, x), \\ \varphi &:= \exists h_1 \dots \exists h_k \exists p \exists n \forall m \psi \wedge (\text{mani}(m) \rightarrow \chi[m]). \end{aligned}$$

Since T is satisfiable, we know by Lemma 18 that $T \wedge S$ is unsatisfiable if and only if there exist (not necessarily distinct) $x, y \in S$ such that $T \wedge x \wedge y$ is unsatisfiable. Remember that we used $\text{TrimRes}(T, H, M)$ to construct φ at the beginning of the reduction. It follows from Lemma 19 that $T \wedge S$ is satisfiable if and only if

for all $h_1, h_2 \in S$, $\{-h_1\} \notin \text{TrimRes}(T, H, M)$ and $\{-h_1, \neg h_2\} \notin \text{TrimRes}(T, H, M)$. This is encoded in φ by requiring $\neg \text{fact}(n, h_i)$ and $\neg \text{cl}(n, h_i, n, h_j)$ for all $i, j \in [k]$. Having ensured consistency it remains to check entailment. From Lemma 20 we know that in this setting it is sufficient to check whether each manifestation m is either contained as a fact in $\text{TrimRes}(T, H, M)$ or there is a single hypothesis $h \in S$ s.t. $\{-h, m\} \in \text{TrimRes}(T, H, M)$. In φ this is ensured by subformula χ . Therefore, $\text{ABD}[\text{KROM}]_{\leq}$ is in $\text{W}[2]$. Adding the conjuncts $\bigwedge_{1 \leq i < j \leq k} (h_i \neq h_j)$ to φ yields membership for $\text{ABD}[\text{KROM}]_{=}$.

We show hardness by reduction from RED-BLUE DOMINATING SET, where an instance is given by a bipartite graph $G = (N_{\text{red}} \cup N_{\text{blue}}, E)$ and an integer k . The question is whether there is a set $S \subseteq N_{\text{red}}$, $|S| \leq k$, s.t. each vertex in N_{blue} is adjacent to a vertex in S . This problem is $\text{W}[2]$ -complete when parameterized by k (Fernau 2008). We construct an instance $((V, H, M, T), k)$ of $\text{ABD}[\text{KROM}]_{\leq}$. Let $V := N_{\text{red}} \cup N_{\text{blue}}$, $H := N_{\text{red}}$, $M := N_{\text{blue}}$, and $T := \bigwedge_{n \in N_{\text{red}}, b \in N_{\text{blue}}} (n \rightarrow b)$, where $N[n]$ contains n and all its adjacent vertices. A set $S \subseteq H$ of size $\leq k$ is a solution for $\text{ABD}[\text{KROM}]_{\leq}$ if and only if it is a solution for the dominating set problem. Since each superset of a dominating set is a dominating set as well, there is a solution of size k if and only if there is a solution of size $\leq k$ (assuming $|N_{\text{red}}|$ is big enough). Thus, hardness also holds for $\text{ABD}[\text{KROM}]_{=}$. \square

The next two theorems show that in KROM, adding M reduces the complexity by one level in the W-hierarchy.

Theorem 22. $\text{ABD}[\text{KROM}]_{\leq}$ parameterized by (k, M) is $\text{W}[1]$ -complete.

Proof. We show $\text{W}[1]$ -membership by reducing an instance $((V, H, M, T), k)$ to an $\text{MC}[\Sigma_1]$ instance (\mathcal{A}, φ) . The reduction is similar to the one used in the proof of Theorem 21. Let $\mathcal{A} := \langle A, \text{hyp}, \text{mani}, \text{fact}, \text{cl}, \text{pos}, \text{neg} \rangle$, ψ , and χ be defined as before, let $|M| = d$, and

$$\varphi := \exists h_1 \dots \exists h_k \exists p \exists n \exists m_1 \dots \exists m_d \psi \wedge \bigwedge_{1 \leq i < j \leq d} (m_i \neq m_j) \wedge \bigwedge_{i \in [d]} (\text{mani}(m_i) \wedge \chi[m_i]).$$

Next, we show hardness by reduction from MULTICOLORED INDEPENDENT SET, where an instance is given by a graph $G = (N, E)$, a size bound k and a k -coloring of the vertices $c : N \rightarrow \{c_1, \dots, c_k\}$. The task is to find a subset $N' \subseteq N$, $|N'| = k$, s.t. for all $x, y \subseteq N'$: $\{x, y\} \notin E$ and $c(x) \neq c(y)$. This problem is $\text{W}[1]$ -complete when parameterized by k (Fellows et al. 2009). We construct an abduction instance $((V, H, M, T), k)$ as follows. Let $V := N \cup \{c_1, \dots, c_k\}$, $H := N$, $M := \{c_1, \dots, c_k\}$, and $T := \bigwedge_{\{x, y\} \in E} (\neg x \vee \neg y) \wedge \bigwedge_{n \in N} (n \rightarrow c(n))$. The k different manifestations (colors) imply that a solution contains at least k hypotheses. The independent set property is ensured by the consistency check. \square

Theorem 23. $\text{ABD}[\text{KROM}]_{=}$ is $\text{W}[1]$ -complete, when parameterized by (k, M) , even when $|M| = 1$.

Proof. Membership can be shown analogously to the proof of Theorem 22 by adding $\bigwedge_{1 \leq i < j \leq k} (h_i \neq h_j)$ to formula φ .

Hardness is shown by reduction from INDEPENDENT SET. We reduce a graph $G = (N, E)$ and integer k to an abduction instance $((V, H, M, T), k)$. Let $V := N \cup \{m\}$, $H := N$, $M := \{m\}$, and $T := m \wedge \bigwedge_{\{x, y\} \in E} (\neg x \vee \neg y)$. By construction, entailment is always fulfilled. \square

Recall that by Theorem 7, $\text{ABD}[\text{KROM}]_{=}$ parameterized by M is para-NP-hard. We show now that interestingly $\text{ABD}[\text{KROM}]_{\leq}$ parameterized by M is $\text{W}[1]$ -complete.

Lemma 24. If an instance of $\text{ABD}[\text{KROM}]$ has a solution, then it has a solution S such that $|S| \leq |M|$.

Theorem 25. $\text{ABD}[\text{KROM}]_{\leq}$ is $\text{W}[1]$ -complete, when parameterized by M .

Proof. Hardness follows immediately from Theorem 22.

For the membership consider the reduction to $\text{MC}[\Sigma_1]$ from the proof of Theorem 22. By Lemma 24 we know that $\text{ABD}[\text{KROM}]$ has a solution if and only if there is a solution of size $\leq |M|$. Thus, we can replace in formula φ from the proof of Theorem 22 any occurrence of k by $b := \min(k, |M|)$. The length of this formula can be bounded in terms of M . \square

The following theorem generalizes the classical results of KROM-abduction. It is NP-complete in general and in P when $|M| = 1$. In fact, the P-membership for every fixed number of manifestations follows from the $\text{W}[1]$ -completeness.

Theorem 26. $\text{ABD}[\text{KROM}]$ is $\text{W}[1]$ -complete, when parameterized by M .

Proof. Membership follows from Theorem 25.

We show hardness by reduction from INDEPENDENT SET. Let (G, k) with graph $G = (N, E)$ and vertices $N = \{v_1, \dots, v_l\}$ be an instance of INDEPENDENT SET. The construction is inspired by (Lackner and Pfandler 2012). We construct an instance $((V, H, M, T), k)$ for the abduction problem. Let $H := \{h_i^j \mid i \in [l], j \in [k]\}$, $M := \{m_i \mid i \in [k]\}$, and $V := N \cup H \cup M$. Next we create $T := T_{\text{IS}} \wedge \bigwedge_{i \in [4]} T_i$, where $T_{\text{IS}} := \bigwedge_{\{x, y\} \in E} (\neg x \vee \neg y)$ and

$$\begin{aligned} T_1 &:= \bigwedge_{i \in [l], j \in [k]} (h_i^j \rightarrow m_j), & T_2 &:= \bigwedge_{i \in [l], j, j' \in [k], j \neq j'} (h_i^j \rightarrow \neg h_i^{j'}), \\ T_3 &:= \bigwedge_{i, i' \in [l], i \neq i', j \in [k]} (h_i^j \rightarrow \neg h_{i'}^j), & T_4 &:= \bigwedge_{i \in [l], j \in [k]} (h_i^j \rightarrow v_i). \end{aligned}$$

Formula T_{IS} encodes the independent set property. In order to ensure that k many vertices are picked, we introduce k hypotheses $h_i^1 \dots h_i^k$ for each vertex v_i . Selecting hypothesis h_i^j corresponds to selecting vertex i as the j -th pick in the independent set. Formula T_1 ensures that the k manifestations are only entailed if for each $j \in [k]$ (the j -th pick) at least one hypothesis of h_1^j, \dots, h_l^j is selected. Subformula T_2 ensures that for each vertex at most one hypothesis is picked; T_3 ensures that we select at most one hypothesis as the same pick. Finally, T_4 forces a vertex to true if one of the corresponding hypotheses was chosen. \square

The fixed-parameter tractability of KROM-abduction parameterized by vertex cover number follows from the FPT result for parameter treewidth (Gottlob, Pichler, and Wei

2010) and the fact that $tw(G) \leq \tau(G)$ for every graph G . We show that this parameterization does not lead to a polynomial kernel.

Theorem 27. *ABD[KROM] and $ABD[KROM]_{\leq/}$ parameterized by τ do not admit a polynomial kernel unless the Polynomial Hierarchy collapses.*

Proof. We show this by PPT-reduction from SAT of $\varphi \in \text{CNF}$ parameterized by $var(\varphi)$. This problem does not admit a polynomial kernel unless the Polynomial Hierarchy collapses (Chen, Flum, and Müller 2011). We create an abduction instance $\langle V, H, M, T \rangle$ as follows. Let $V := X \cup X' \cup M$, where $X := var(\varphi)$, $X' := \{x' \mid x \in X\}$, and M contains a manifestation for each clause in φ . Let $H := X \cup X'$, and

$$T := \bigwedge_{x \in X} ((x \vee x') \wedge (\neg x \vee \neg x')) \wedge \bigwedge_{c \in \varphi} \left(\bigwedge_{x \in c} (x \rightarrow c) \wedge \bigwedge_{\neg x \in c} (x' \rightarrow c) \right).$$

For instances of $ABD[KROM]_{\leq/}$, we additionally set $k := |X|$. Observe that the primal graph of T can be covered by the set $X \cup X'$. Thus, τ can be bounded by $2 \cdot |var(\varphi)|$. \square

Next we will show that $ABD[KROM]$ and $ABD[KROM]_{\leq/}$ have a polynomial kernel when parameterized by (H, M) . Thereby *TrimRes* from Definition 17 is used as a kernelization function.

Lemma 28. *Let $\langle V, H, M, T \rangle$ be an abduction instance for KROM theories, let $S \subseteq H$, let $m \in M$, and let $T \wedge S$ be satisfiable. Then $T \wedge S \wedge \neg m$ is unsatisfiable if and only if $TrimRes(T, H, M) \wedge S \wedge \neg m$ is unsatisfiable.*

Theorem 29. *$ABD[KROM]$ and $ABD[KROM]_{\leq/}$ have a polynomial kernel when parameterized by (H, M) .*

Proof. Given an instance $\langle V, H, M, T \rangle$ with $T \in \text{KROM}$. We can test T for unsatisfiability in polynomial time and output a trivial no-instance in case the answer is yes. Otherwise we compute in polynomial time $TrimRes(T, H, M)$ which has size $O((|H| + |M|)^2)$. Indeed $\langle H \cup M, H, M, TrimRes(T, H, M) \rangle$ gives a kernel for our instance. Given a set $S \subseteq H$, by Lemma 19 testing the satisfiability of $T \wedge S$ is equivalent to testing $TrimRes(T, H, M) \wedge S$. Testing whether $T \wedge S \models M$ is equivalent to testing if $T \wedge S \models m$ for all $m \in M$. By Lemma 28 each of those tests can be done on $TrimRes(T, H, M)$. \square

It follows immediately that $ABD[KROM]$ and $ABD[KROM]_{\leq/}$ have a polynomial kernel when parameterized by V . Note that these problems are already FPT when parameterized by H alone.

Conclusion

We have drawn a detailed picture of the parameterized complexity of abduction as depicted in Tables 1 and 2. Although there are many cases where $ABD[\text{HORN}]$ is FPT, it never admits a polynomial kernel. For $ABD[\text{KROM}]$ we were able to show the existence of a polynomial kernel when parameterized by (H, M) . The question whether $ABD[\text{DEFHORN}]$

parameterized by V admits a polynomial kernel remains open. Future work includes the search for further parameters yielding fixed-parameter tractability and the analysis of other syntactical fragments.

References

- Bodlaender, H. L.; Downey, R. G.; Fellows, M. R.; and Hermelin, D. 2009. On problems without polynomial kernels. *J. Comput. Syst. Sci.* 75(8):423–434.
- Bodlaender, H. L.; Thomassé, S.; and Yeo, A. 2011. Kernel bounds for disjoint cycles and disjoint paths. *Theor. Comput. Sci.* 412(35):4570–4578.
- Chen, Y.; Flum, J.; and Müller, M. 2011. Lower bounds for kernelizations and other preprocessing procedures. *Theory Comput. Syst.* 48(4):803–839.
- Creignou, N., and Zanuttini, B. 2006. A complete classification of the complexity of propositional abduction. *SIAM J. Comput.* 36(1):207–229.
- Creignou, N.; Schmidt, J.; and Thomas, M. 2010. Complexity of propositional abduction for restricted sets of boolean functions. In *Proc. KR'10*. AAAI Press.
- Dom, M.; Lokshtanov, D.; and Saurabh, S. 2009. Incompressibility through colors and ids. In *Proc. ICALP'09, Part I*, volume 5555 of *LNCS*, 378–389. Springer.
- Downey, R. G., and Fellows, M. R. 1999. *Parameterized Complexity*. Springer.
- Downey, R. G.; Fellows, M. R.; and Regan, K. W. 1998. Descriptive complexity and the W hierarchy. In *Proof Complexity and Feasible Arithmetic*, volume 39 of *AMS-DIMACS Volume Series*, 119–134. AMS.
- Eiter, T., and Gottlob, G. 1995. The complexity of logic-based abduction. *J. ACM* 42(1):3–42.
- Fellows, M. R.; Hermelin, D.; Rosamond, F. A.; and Viallette, S. 2009. On the parameterized complexity of multiple-interval graph problems. *Theor. Comput. Sci.* 410(1):53–61.
- Fernau, H. 2008. Roman domination: a parameterized perspective. *Int. J. Comput. Math.* 85(1):25–38.
- Flum, J., and Grohe, M. 2006. *Parameterized Complexity Theory*. Springer.
- Gottlob, G.; Pichler, R.; and Wei, F. 2010. Bounded treewidth as a key to tractability of knowledge representation and reasoning. *Artif. Intell.* 174(1):105–132.
- Lackner, M., and Pfandler, A. 2012. Fixed-parameter algorithms for finding minimal models. In *Proc. KR'12*. AAAI Press. To appear.
- Leitsch, A. 1997. *The resolution calculus*. Texts in theoretical computer science. Springer.
- Niedermeier, R. 2006. *Invitation to Fixed-Parameter Algorithms*. Oxford University Press.
- Nordh, G., and Zanuttini, B. 2008. What makes propositional abduction tractable. *Artif. Intell.* 172(10):1245–1284.
- Robertson, N., and Seymour, P. D. 1986. Graph minors. II. algorithmic aspects of tree-width. *J. Algorithms* 7(3):309–322.
- Selman, B., and Levesque, H. J. 1990. Abductive and default reasoning: A computational core. In *Proc. AAAI'90*, 343–348.