Available online at www.sciencedirect.com
ScienceDirect
Theoretical Computer Science 369 (2006) 82–115
Theoretical Computer Science
www.elsevier.com/locate/tcs

Reasoning under minimal upper bounds in propositional logic☆

Thomas Eiter[a,*], Georg Gottlob[b]

[a] *Institute of Information Systems, Vienna University of Technology, Favoritenstraße 9-11, A-1040 Vienna, Austria*
[b] *Computing Laboratory, Oxford University, Parks Road, Oxford, OX1 3QD, UK*

Received 14 October 2005; received in revised form 19 July 2006; accepted 19 July 2006

Communicated by G. Ausiello

Abstract

Reasoning from the minimal models of a theory, as fostered by circumscription, is in the area of Artificial Intelligence an important method to formalize common sense reasoning. However, as it appears, minimal models may not always be suitable to capture the intuitive semantics of a knowledge base, aiming intuitively at an exclusive interpretation of disjunctions of atoms, i.e., if possible then assign at most one of the disjuncts the value true in a model. In this paper, we consider an approach which is more lenient and also admits non-minimal models, such that inclusive interpretation of disjunction also may be possible in cases where minimal model reasoning adopts an exclusive interpretation. Nonetheless, in the spirit of minimization, the approach aims at including only positive information that is necessary. This is achieved by closing the set of admissible models of a theory under minimal upper bounds in the set of models of the theory, which we refer to as curbing. We demonstrate this method on some examples, and investigate its semantical and computational properties. We establish that curbing is an expressive reasoning method, since the main reasoning tasks are shown to be PSPACE-complete. On the other hand, we also present cases of lower complexity, and in particular cases in which the complexity is located, just as for ordinary minimal model reasoning, at the second level of the Polynomial Hierarchy, or even below.

Keywords: Artificial intelligence; Circumscription; Computational complexity; Curbing; Knowledge representation; Minimal models; Non-monotonic reasoning; Propositional logic

1. Introduction

In Artificial Intelligence, Occam's Razor ("entia non sunt multiplicanda praeter necessitatem") is widely used to model human problem solving. In particular, in the area of knowledge representation and reasoning, this principle is underlying circumscription [29,34,35], which is one of the most well-known approaches to non-monotonic reasoning. Circumscription aims at minimizing the positive information contained in the models of a logical theory T that constitutes a knowledge base. To this end, it selects the minimal models of T under a preference relation $\leqslant$ on the set of all models of T, according to which a model M is better than (or as good as) a model M' if each elementary fact which

☆ Some results in this paper appeared in preliminary form in the Proceedings of LPAR 2000 [18] and the Proceedings of IJCAI-93 [19].

* Corresponding author. Tel.: +43 15880118460; fax: +43 15880118493.
E-mail addresses: eiter@kr.tuwien.ac.at (T. Eiter), Georg.Gottlob@comlab.ox.ac.uk (G. Gottlob).

0304-3975/$ - see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.tcs.2006.07.054

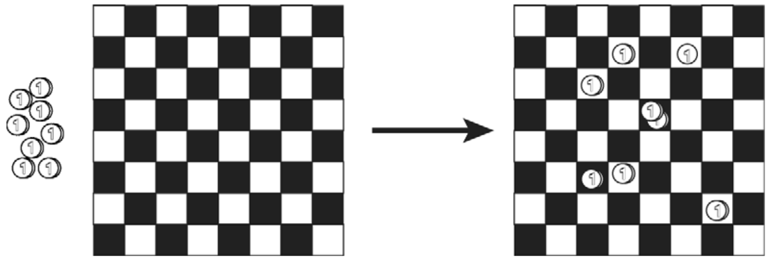

Fig. 1. Possible outcome of throwing coins on a chessboard, circumscribing the coin locations.

is true in M is also true in M'. For example, if T consists of the three formulas $p \to q, r \to q$, and p, then there are two models, of T on the atoms p, q, and r: one in which all three atoms are true, and another one in which p and q are true and r is false; circumscription selects from these two models the second (in which r is false), since it is better than the first. The theory of circumscription is far developed, and several variants exist; we refer to [21,30] for an introduction.

As noted by various authors [11,13,40–43], reasoning under minimal models runs into problems in connection with disjunctive information. The minimality principle of circumscription often enforces the *exclusive* interpretation of a disjunction $a \vee b$ by adopting the models in which either a or b is true but not both. There are many situations in which an *inclusive* interpretation is desired and seems more natural. For example, consider the following scenario due to Raymond Reiter. [1]

Example 1.1. Suppose a coin is thrown into an area which is divided into a black and a white field. Circumscription applied to the propositional theory

$$T_0 = \{\, lies\_on\_black \;\vee\; lies\_on\_white \,\}$$

excludes the model in which *lies_on_black* and *lies_on_white*1 hold, i.e., that the coin falls into both fields, and yields that the coin is either in the black or in the white field, according to the two minimal models in which *lies_on_black* is true and *lies_on_white* is false, respectively *lies_on_black* is false and *lies_on_white* is true. This is certainly not satisfactory. An extension of this example is even more impressive. Imagine a handful of coins thrown onto a chessboard; given that the propositions c_{ij} represent for coin c that it lies on field f, where $f \in \{a1, a2, \ldots, a8, b1, \ldots, h8\}$, the location of c is described by the disjunction

$$loc\_c = c\_lies\_on\_a1 \vee c\_lies\_on\_a2 \vee \cdots \vee c\_lies\_on\_h8.$$

Applying circumscription to the conjunction $T_0 = loc\_c_1 \wedge \cdots \wedge loc\_c_k$ of the coin locations, we obtain that each coin c_i lies on exactly one field, and thus does not touch both a black and a white field (cf. Fig. 1).

None of the well-known variants of circumscription [21,30] seems to handle inclusive disjunction of positive information suitably in general. In order to redress this problem and to provide a way for handling inclusive disjunctions of positive information properly, we proposed the *curbing* method in [19], which is a generalization of circumscription for handling inclusive disjunctions of positive information at the semantic level. This method fosters the notion of a *curb model* of a theory, which is roughly speaking either a minimal model of the theory T, or recursively defined as a model of T which is a *minimal upper bound* (mub) of a set of smaller curb models in the partially ordered set of models of T under the ordering $\leqslant$. The use of mubs is guided by the minimization principle; note that least upper bounds (LUB), which would be more natural, do not always exist. In the example from above, curbing selects besides the two minimal models also the model where both *lies_on_black* and *lies_on_white* are true. While in this and the chessboard example, curbing selects all models of the theory, this is not the case in general (see e.g. the further examples in Section 2.1). [2]

The contributions of this paper are briefly summarized as follows.

[1] Personal communication.

[2] Note that in the chessboard example, no information about the size of the coins resp. the board is given, and axioms for contiguous occupation of fields are missing. However, it is arguably reasonable that such domain-specific knowledge must be provided for a domain-independent method to work well.

(1) We formalize curbing for (possibly infinite) propositional theories on a given set of atoms $\mathcal{A}$. Note that the propositional setting plays an important role in many knowledge representation applications, since knowledge base formalizations used in practice are often propositional in nature; indeed, in AI applications, frequently first-order predicate logic is used for knowledge representation which considers a fixed, finite domain of discourse. To this end, domain closure axioms are used, often in combination with unique name axioms for individuals. This is essentially a propositional setting, to which the predicate formalization can be easily reduced. This also holds for certain first-order settings with infinite domains, e.g., in logic programming (under customary Herbrand semantics) cf. [1,32,14,36], which by definition are reduced to the propositional case. Thus, the propositional setting is important for applications.
(2) We provide purely semantic and constructive characterizations of curb models, in a customary minimization setting where only a part P of the atoms $\mathcal{A}$ is minimized and a further part Z is floating (i.e., projected from the ordering $\leqslant$), while all other atoms are considered to be fixed. Furthermore, we analyze the relationship to circumscription and provide for finite theories a syntactic characterization of the theories for which curbing and circumscription coincide.
(3) We consider restricted notions of curbing, by limiting the iterations for taking mubs and/or the number of models for which a mub is formed. These restrictions can be viewed as semantic and computational approximations of an idealized reasoner, who closes her set of accepted models under mubs. As we show, under these restrictions for certain classes of theories still the full set of curb models is obtained, which proves useful for reasoning algorithms.
(4) We study the class of theories for which collections of curb models always have a least (unique minimal) upper bound, which we call *LUB theories*, as well as the more general class of *Weak LUB* theories, in which each curb model is the LUB of some set of curb models. Both classes have interesting properties; for example, they require only a bounded number of steps to obtain all curb models by iteratively taking mubs.
(5) Finally, we investigate computational aspects of curbing. We first give a simple algorithm for computing the curb models of a given finite theory. This algorithm uses polynomial space in the size of the output, which might have exponential size. By using this algorithm, curb inference (i.e., truth in all curb models) can be decided in exponential space and time. However, we give an algorithm for curb model checking which runs in quadratic space, which implies that also logical inference is decidable in quadratic space. On the other hand, we show that both problems are PSPACE-complete, and that they are PSPACE-hard even in case of *global* curbing, i.e., when curbing is applied to *all* atoms. Furthermore, we single out cases of lower complexity (under common hypotheses in complexity theory), in particular cases where curbing is feasible with complexity in Π_2^P, which is the complexity of circumscription [16], or even with lower complexity. We also discuss how curbing can be polynomially reduced to circumscription in these cases.

The PSPACE-completeness results, whose proofs are non-trivial, strongly indicate that curbing is from the point of computational expressiveness a much more powerful reasoning method than circumscription in general. In particular, they imply that curb model checking and curb inference are not solvable in polynomial time using a circumscriptive theorem prover in general (which, however, by our results is feasible under restrictions). On the other hand, a curb reasoning engine may be based on a quantified Boolean formula (QBF) solver (see `http://www.qbflib.org/solvers.html` and [25] for a recent evaluation report), to which model checking and inference can be reduced in polynomial time.

Since its introduction in [19], the curbing method has been used and studied by other authors. For instance, Scarcello et al. [45] provide a fixpoint semantics for propositional (Boolean) curbing and derive complexity results for curbing the class of dual Horn-quadratic theories, i.e., clausal theories where each clause contains at most two literals and at least one positive literal. In the present paper, we sharpen some of these results. Liberatore [27,28] considers belief update operators from [49,50], which are related to a restricted version of curbing.

The rest of this paper is organized as follows. In the next section, we give some preliminaries and consider further examples that motivate reasoning under mubs. After that, we give in Section 3 a formal definition of curbing and consider some properties, viz. different characterizations, behavior with respect to disjunctions, and relationship to circumscription. The subsequent Section 4 is devoted to restricted notions of curbing. After that, we consider classes of theories in Section 5 for which the set of curb models has a particular structure. Section 6 then considers computation and complexity of curbing. Related work is discussed in Sections 7, and 8 presents some conclusions.

2. Preliminaries and examples

We assume a (possibly infinite) set of propositional atoms $\mathcal{A}$. In this paper, $\mathcal{A}$ will often be denumerable, i.e., finite or countably infinite ($\mathcal{A} = \{p_i | i = 0, 1, \ldots\}$); in general, we assume that $\mathcal{A}$ is indexed by some ordinal η such that $\mathcal{A} = \{p_\alpha | \alpha \in \eta\}$ (i.e., $\mathcal{A}$ is enumerated by some well-ordering). A *theory* is any finite or infinite set of propositional formulas on $\mathcal{A}$, which are built using the usual Boolean connectives, including constants $\top$ (truth) and $\bot$ (falsity).

An *interpretation* (or *model*) M is an assignment $M : \mathcal{A} \to \{0, 1\}$ of truth values 0 (false) or 1 (true) to all atoms, which extends to formulas as usual. We identify M also with the set of atoms $p \in \mathcal{A}$ such that $M(p) = 1$. The restriction of a model M to a set of atoms $\mathcal{A}' \subseteq \mathcal{A}$ is denoted by $M[\mathcal{A}']$. Furthermore, for any propositional formula φ and model M, we denote by $\varphi[M]$ resp. $\varphi[M[\mathcal{A}']]$ the result of replacing in φ every $p \in \mathcal{A}$ resp. $p \in \mathcal{A}'$ by $\top$ if $M(p) = 1$ and by $\bot$ if $M(p) = 0$.

An interpretation M satisfies a formula φ (is a model of φ), denoted $M \models \varphi$, if $M(\varphi) = 1$, and M satisfies a theory T (is a model of T), denoted $M \models T$, if $M \models \varphi$ for each $\varphi \in T$; we denote by $mod(\varphi)$ (resp., $mod(T)$) the set of all models of φ (resp. T).

Definition 2.1. We call a pair $(P; Z)$ of subsets $P, Z \subseteq \mathcal{A}$ such that $P \cap Z = \emptyset$ a *minimization setting*; informally, the atoms in P are to be minimized in parallel while those Z in may vary and all other atoms are fixed. Furthermore, we denote by $\leqslant_{P;Z}$ the respective preference relation on the set models for $\mathcal{A}$. That is, $M \leqslant_{P;Z} M'$ (M is more or equally preferable to M') iff $M[P] \subseteq M'[P]$ and $M[Q] = M'[Q]$, where $Q = \mathcal{A} \setminus P \cup Z$. As usual, $M <_{P;Z} M'$ stands for $M \leqslant_{P;Z} M' \wedge M \not\leqslant_{P;Z} M$, and we write $M =_{P;Z} M'$ for $M \leqslant_{P;Z} M' \wedge M \leqslant_{P;Z} M'$. We refer to the special case where $P = \mathcal{A}$ and $Z = \emptyset$ as *global minimization* and omit P; Z in subscripts, etc.

We use $\max_{P;Z}(\mathcal{M})$ to denote the maximal elements of $\mathcal{M}$ under $\leqslant_{P;Z}$, that is, $\max_{P;Z}(\mathcal{M}) = \{M \in \mathcal{M} | \nexists M' \in \mathcal{M} : M <_{P;Z} M'\}$; for global minimization, P; Z is omitted.

The *circumscription* of a theory T with respect to a minimization setting $(P; Z)$ is semantically given by the set of its P; Z-minimal models, denoted $mmod_{P;Z}(T)$, where a model M in any set of models $\mathcal{M} \subseteq 2^{\mathcal{A}}$ is P; Z-minimal, if there is no model $M' \in \mathcal{M}$ such that $M' <_{P;Z} M$. We remind that $(P; Z)$ minimization has been introduced in order to facilitate the conclusion of new positive evidence, which cannot be drawn under global minimization, as shown in the following example.

Example 2.1. The theory $T = \{bird,\ bird \wedge \neg ab \to flies\}$ on $\mathcal{A} = \{b, ab, flies\}$ has in the minimization setting $P = \{ab\}$ and $Z = \{flies\}$ the single P; Z-minimal model $M_1 = \{bird, flies\}$. Therefore, we can conclude *flies* from the circumscription of T w.r.t. $(P; Z)$. On the other hand, the global minimal models of T are M_1 and $M_2 = \{bird, ab\}$, and this conclusion is not possible; note that $M_1 <_{P;Z} M_2$.

For any finite T, its circumscription under $(P; Z)$ is expressed by the set of formulas

$$\{\neg p | p \in P \setminus P_T\} \cup \{Circ(\varphi; P_T; Z_T)\},$$

where P_T and Z_T are the atoms from P, respectively, Z occurring in T, and Q_T are all other atoms in T; $\varphi = \varphi(P_T; Q_T; Z_T)$ is the conjunction of all formulas in T; and $Circ(\varphi; P_T; Z_T)$ is the QBF

$$Circ(\varphi; P_T; Z_T) = \varphi(P_T; Q_T; Z_T) \wedge \forall P'_T Z'_T((\varphi[P'_T; Z'_T] \wedge P'_T \leqslant P_T) \to P_T \leqslant P'_T).$$

Here P'_T and Z'_T are sets of fresh atoms (not occurring in φ) in 1–1 correspondence to P and Z, respectively, and $P_T \leqslant P'_T$ (resp., $P'_T \leqslant P_T$) denotes the formula $\bigwedge_{p \in P_T}(p \to p')$ resp. $\bigwedge_{p \in P_T}(p' \to p)$. By eliminating quantifiers $\exists, \forall$ as usual, $Circ(\varphi; P_T; Z_T)$ can be rewritten to an equivalent propositional formula. That is, replace subformulas $\forall x \varphi(x, y_1, \ldots, y_k)$ with $\varphi(\bot, y_1, \ldots, y_k) \wedge \varphi(\top, y_1, \ldots, y_k)$ and subformulas $\exists x \varphi(x, y_1, \ldots, y_k)$ with $\varphi(\bot, y_1, \ldots, y_k) \vee \varphi(\top, y_1, \ldots, y_k)$, respectively.

2.1. Curb models

Let us first describe two further scenarios in which non-minimal models are desirable.

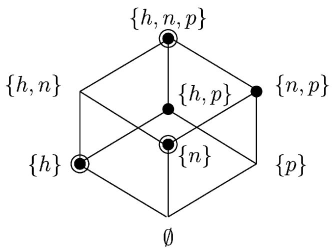

$$T_1 = \{h \vee n,\ (h \wedge n) \rightarrow p\}$$

Fig. 2. The hammer–nail-painting example.

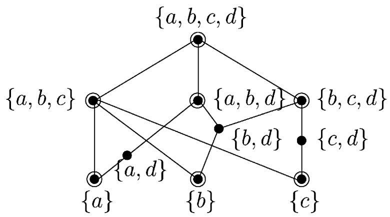

$$\begin{aligned} T_2 = \{\, & a \vee b \vee c, \\ & (a \wedge b) \rightarrow (c \vee d), \\ & (b \wedge c) \rightarrow (a \vee d), \\ & (a \wedge c) \rightarrow b \,\} \end{aligned}$$

Fig. 3. The party example.

Example 2.2. Suppose Alice is in a room with a painting, which she hangs on the wall if she has a hammer and a nail. It is known that Alice has a hammer or a nail, and possibly both. This scenario is represented by the theory T_1 in Fig. 2. The desired models are $\{h\}$, $\{n\}$, and $\{h, n, p\}$, which are encircled. Circumscribing T_1 under global minimization yields the two minimal models $\{h\}$ and $\{n\}$ (see Fig. 2). Since p is false in all minimal models, circumscription tells us that Alice does not hang the painting up. Furthermore, it tells us that Alice has either a hammer or a nail, but not both. On the other hand, the model $\{h, n, p\}$, which intuitively corresponds to the inclusive interpretation of the disjunction $h \vee n$, seems plausible. Therefore, one may expect that the conclusion of Alice not hanging the paper up is not drawn, and similarly not that Alice holds only one item.

Example 2.3. Suppose you have invited some friends to a party. You know for certain that one of Alice, Bob, and Chris will come, but you do not know whether Doug will come. You know in addition the following habits of your friends. If Alice and Bob go to a party, then Chris or Doug will also come; if Bob and Chris go, then Alice or Doug will go. Furthermore, if Alice and Chris go, then Bob will also go. This is represented by the theory T_2 in Fig. 3. Now what can you say about who will come to the party? Consider the models of T_2 in Fig. 3. Circumscription under global minimization yields the minimal models $\{a\}$, $\{b\}$, and $\{c\}$, which interpret the clause $a \vee b \vee c$ exclusively in the sense that it is minimally satisfied. However, there are other plausible models. For example, $\{a, b, c\}$. This model embodies an inclusive interpretation of a and b within $a \vee b \vee c$; it is also minimal in this respect. The model $\{a, b, d\}$ is another model with this property. Similarly, $\{b, c, d\}$ is a minimal model for an inclusive interpretation of b and c. The models $\{a, d\}$, $\{b, d\}$, and $\{c, d\}$ seem not plausible, however, since a scenario in which Doug and only one of Alice, Bob or Chris are present does not seem well-supported.

In the light of these examples, the question arises how minimization can be extended to work more satisfactorily. Obviously, also non-minimal models must be allowed, even if such models may contain positive information that is not contained in any minimal model, as shown by Example 2.2. On the other hand, the fruitful principle of minimality should not be abandoned by adopting models that are intuitively not concise. Such a situation is given when the theory at hand has a single minimal model. Adhering to minimality, there seems to be no reason for adopting non-minimal models in this case, since there is no ambiguity in which values the atoms to be minimized should be assigned. The situation is different if there is more than one minimal model. Each of these models represents a possible way to minimize positive information coherently. Here, an *inclusive account* of opposite positive information in different minimal models M_1

and M_2, which is given by their symmetric difference, is not considered, since any model M which includes both M_1 and M_2 is rejected under minimization.

Extending the circumscription approach, the *Curbing* method aims at giving an inclusive account of positive information at the semantic level, based on well-understood mathematical concepts. The basic idea is to adopt for minimal models M_1, M_2 any model M which includes both M_1 and M_2, but is a minimal such model; in other words, M is a *minimal upper bound* (*mub*) for M_1 and M_2. Intuitively, adopting M semantically accounts for disjunction, which must be present in the theory (otherwise, a single minimal model would exist); we defer a discussion of this issue to Section 3.1.

For illustration, in Example 2.2 $\{h, n, p\}$ is a mub for $\{h\}$ and $\{n\}$ in the models of T_1 (notice that $\{h, n\}$ is not a model of T_1), and in Example 2.3, $\{a, b, c\}$ is a mub for $\{a\}$ and $\{b\}$ in the models of T_2; the model $\{a, b, d\}$ is another one, so several mubs can exist. In order to capture general inclusive interpretations, mubs of arbitrary families of minimal models of the theory T at hand are adopted.

It appears that in general not all "intuitive" models are obtainable as mubs of families of minimal models. The intuitive model $\{a, b, c, d\}$ in Example 2 shows this. It is, however, a mub of the "good" models $\{a\}$ and $\{b, c, d\}$ (as well as of $\{a, b, c\}$ and $\{a, b, d\}$). This suggests that not only mubs of families of minimal models, but also mubs of any family of intuitive models should be selected.

The curbing approach to generalize model minimization is thus the following: adopt the least set of models of a theory which contains all circumscriptive (i.e., minimal) models and which is closed under including mubs. This approach yields in Examples 2.2 and 2.3 the sets of "intuitive" models, which are encircled in Figs. 2 and 3. The closedness property is a natural idealization in lack of further information about which minimal upper bounds should be adopted in a particular situation, and all adopted models are considered to be en par. In Example 2.3, the models $\{a, b, c\}$ and $\{a, b, d\}$ are mubs of $\{a\}$ and $\{b\}$, and they are considered to be en par with $\{c\}$; thus, mubs among $\{a, b, c\}$, $\{a, b, d\}$, and $\{c\}$ should be adopted as well.

As for inferencing, closedness is sound in that no conclusions are accepted which would not be drawn in the actual situation. Different rationales for discriminating between models when taking mubs can be imagined, and then closedness under mubs might be abandoned (cf. Section 4). To consider the least closed set of models which is closed under mubs is guided by the principle to discard any model of the theory which is not "foundedly" (in terms of mubs) included, and thus to obtain at a smallest set of models and arrive in this way at a largest set of conclusions.

3. Formal definition of curbing

We formally define the concept of a curb model as follows. First we define the notion of a mub of a set of models, and then property that a set of models is closed under mubs.

Definition 3.1. Let $\mathcal{M}_0 \subseteq 2^{\mathcal{A}}$ be a set of models on atoms $\mathcal{A}$ and let $(P; Z)$ be a minimization setting. A model $M \in \mathcal{M}_0$ is a $P; Z$-upper bound of a set of models $\mathcal{M} \subseteq 2^{\mathcal{A}}$ with respect to $\mathcal{M}_0$, if $N \leqslant_{P;Z} M$ for every $N \in \mathcal{M}$. Moreover, such an M is a $P; Z$-mub of $\mathcal{M}$ w.r.t. $\mathcal{M}_0$, if there exists no $P; Z$-upper bound N of $\mathcal{M}$ w.r.t. $\mathcal{M}_0$ such that $N <_{P;Z} M$. The set of all $P; Z$-upper bounds (resp., $P; Z$-mubs) of $\mathcal{M}$ w.r.t. $\mathcal{M}_0$ is denoted by $ub^{\mathcal{M}_0}_{P;Z}(\mathcal{M})$ (resp., $mub^{\mathcal{M}_0}_{P;Z}(\mathcal{M})$).

Definition 3.2. A set of models $\mathcal{M}$ is closed under $P; Z$-mubs with respect to $\mathcal{M}_0$, if $mub^{\mathcal{M}_0}_{P;Z}(\mathcal{N}) \subseteq \mathcal{M}$ holds for every $\mathcal{N} \subseteq \mathcal{M}$.

If $\mathcal{M}_0 = mod(T)$ for some theory T, we simply write T in place of $mod(T)$. Clearly, $mod(T)$ is closed under $P; Z$-mubs with respect to T for any $(P; Z)$. Furthermore, every set of models of T which is closed under $P; Z$-mubs with respect to T must contain $mmod_{P;Z}(T)$ (recall that $mmod_{P;Z}(T)$ denotes the set of all $P; Z$-minimal models of T). Indeed, the $P; Z$-mubs of the empty set $\mathcal{M} = \emptyset$ w.r.t. T are just the $\leqslant_{P;Z}$-minimal models of T.

We thus define the curb models of a theory T as follows.

Definition 3.3. For any theory T and minimization setting $(P; Z)$, the set of $P; Z$-curb models of T, denoted by $cmod_{P;Z}(T)$, is the least set of models of T which is closed under $P; Z$-mubs w.r.t. T.

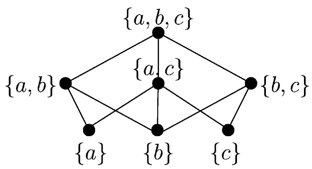

Fig. 4. Set of curb models whose ranks are not strictly decreasing along chains.

Notice that curb models are well-defined only if a unique smallest closed set exists. The latter is immediate from the fact that the family of all sets $\mathcal{M} \subseteq mod(T)$ which are closed under $P; Z$-mubs w.r.t. T is intersection-closed (i.e., if $\mathcal{M}$ and $\mathcal{N}$ belong to it, then also $\mathcal{M} \cap \mathcal{N}$ belongs to it).

Revisiting Example 1.1, the set $cmod(T_0)$ coincides with $mod(T_0)$. (In the case of global minimization ($P = \mathcal{A}$, $Z = \emptyset$), we omit subscripts $P; Z$ as usual.) In Examples 2.2 and 2.3, the sets $cmod(T_1)$ and $cmod(T_2)$ for T_1, respectively, T_2 are given by the models encircled in Figs. 2 and 3, respectively.

Curb models may be alternatively described as follows.

Proposition 3.1. *A model M is a $P; Z$-*curb *model of T iff M belongs to every set $\mathcal{M} \subseteq mod(T)$ which is closed under $P; Z$-mubs w.r.t. T.*

Proof. Immediate from the fact that the family of all sets $\mathcal{M} \subseteq mod(T)$ as in the statement is closed under intersection (which implies the existence of the least such $\mathcal{M}$ w.r.t. set inclusion). □

An easy recursive characterization of curb models is the following. For any T and model M, let us denote by $cmod^{<M}_{P;Z}(T) = \{M' \in cmod_{P;Z}(T) | M' <_{P;Z} M\}$ the set of all curb models of T which are smaller than M.

Proposition 3.2. *A model M is a $P; Z$-*curb *model of T iff M is a member of $mub^T_{P;Z}(cmod^{<M}_{P;Z}(T))$.*

Proof. The if direction obviously holds since $cmod^{<M}_{P;Z}(T) \subseteq cmod_{P;Z}(T)$. For the only if direction, if $M \in cmod_{P;Z}(T)$ then by the minimality of $cmod_{P;Z}(T)$ we have that $M \in mub^T_{P;Z}(\mathcal{M})$ for some $\mathcal{M} \subseteq cmod_{P;Z}(T)$ such that $M' <_{P;Z} M$, for all $M' \in \mathcal{M}$; otherwise, the set $\{M' \in cmod_{P;Z}(T) | M' \neq_{P;Z} M\}$ is closed under $P; Z$-mubs w.r.t. T, which contradicts the minimality of $cmod_{P;Z}(T)$. Since $\mathcal{M} \subseteq cmod^{<M}_{P;Z}(T)$, it follows that $M \in mub^T_{P;Z}(cmod^{<M}_{P;Z}(T))$. □

Alternatively, we can characterize curb models in terms of inductive definability. For this purpose, we define the notion of α-curb model for ordinals $\alpha \geqslant 0$.

Definition 3.4. Let T be a theory and let $(P; Z)$ be a minimization setting. Then, for every $M \in mod(T)$ and ordinal $\alpha \geqslant 0$, M is an α-curb model of T w.r.t. $P; Z$ (for short, an α-curb$_{P;Z}$ model of T), if

- $M \in mmod_{P;Z}(T)$, if $\alpha = 0$,
- $M \in mub^T_{P;Z}(\mathcal{M})$ for some $\mathcal{M} \subseteq mod(T)$, such that each $M \in \mathcal{M}$ is a β-curb$_{P;Z}$ model of T for some $\beta < \alpha$.

The rank of M w.r.t. $(P; Z)$, denoted $rank^T_{P;Z}(M)$, is the least such α.

As shown in [19], the concept of α-curb model provides a constructive means to obtain the set of curb models, since each curb model is constructible from curb models of lower rank. Notice that an inductive construction based on Proposition 3.2 does not work, since $M' <_{P;Z} M$ does not imply $rank^T_{P;Z}(M') < rank^T_{P;Z}(M)$, as shown by the following example.

Example 3.1. Consider the theory $T = \{a \vee b \vee c\}$ whose models are shown in Fig. 4.

Here, all models are global curb models, and the model $\{a, b, c\}$ has rank 1 (as it is the mub of the minimal models $\{a\}$, $\{b\}$, and $\{c\}$) as well as any of the smaller models $\{a, b\}$, $\{a, c\}$, and $\{b, c\}$.

Therefore, the equation $rank^T_{P;Z}(M) = \sup\{rank^T_{P;Z}(M')+1 | M' \in cmod^{<M}_{P;Z}(T)\}$ does not hold in general. However, we obtain M also as a mub of all smaller models $M' <_{P;Z}$ such that $rank^T_{P;Z}(M') < rank^T_{P;Z}(M)$, which gives the following result.

Theorem 3.3. *Let T be a theory on $\mathcal{A}$ and let $(P; Z)$ be a minimization setting. Then, $M \in cmod_{P;Z}(T)$ if and only if M is an α-curb$_{P;Z}$ model of T for some ordinal α. Furthermore,*

$$rank^T_{P;Z}(M) = \sup\{rank^T_{P;Z}(M') + 1 |\ M' \in cmod^{<M}_{P;Z}(T), rank^T_{P;Z}(M') < rank^T_{P;Z}(M)\}, \tag{1}$$

and $|rank^T_{P;Z}(M)| \leqslant |\mathcal{A}|$ (in particular, $|rank^T_{P;Z}(M)| \leqslant \aleph_0$ if $\mathcal{A}$ is denumerable).

Proof. We show this result, differently as in [19], exploiting well-known results in fixpoint theory.

Define the operator $\Lambda_T : 2^{2^{\mathcal{A}}} \to 2^{2^{\mathcal{A}}}$ on the domain $2^{2^{\mathcal{A}}}$ of sets of interpretations for $\mathcal{A}$ by

$$\Lambda_T(\mathcal{M}) = \bigcup_{\mathcal{N}\subseteq\mathcal{M}} mub^T_{P;Z}(\mathcal{N}).$$

Clearly, Λ_T is a monotone operator, and thus by the well-known Knaster–Tarski Lemma, Λ_T has a least fixpoint, $lfp(\Lambda_T)$. It is given by the first element S_γ in the sequence S_α, $\alpha \geqslant 0$, defined by $S_0 = \emptyset$, $S_{\alpha+1} = \Lambda_T(S_\alpha)$ for successor ordinals $\alpha + 1$ and $S_\alpha = \bigcup_{\beta<\alpha} S_\beta$ for limit ordinals $\alpha > 0$, such that S_γ is a fixpoint of Λ_T. Since the fixpoints of Λ_T are those sets of models of T closed under P; Z-mubs w.r.t. T, it follows that $S_\gamma = cmod_{P;Z}(T)$. By construction, each model $M \in S_\alpha$ is a α-curb$_{P;Z}$ model of T, and if α is the smallest ordinal with this property, then $rank^T_{P;Z}(M) = \alpha - 1$ if α is a successor ordinal and $rank^T_{P;Z}(M) = \alpha$ otherwise. Therefore, $rank^T_{P;Z}(M)$ must satisfy Eq. (1). Finally, an easy induction yields that $|rank^T_{P;Z}(M)| \leqslant |M|$. □

3.1. Inclusive interpretation of disjunctive clauses

Curbing as introduced above aims at selecting more models than circumscription in case of multiple minimal models. Syntactically, this means that the theory T at hand must entail some disjunctive clause γ of the form $a \vee b \vee \delta$, where a and b are distinct atoms and δ is a (possibly empty) clause, such that none of a and b can be removed from it. Furthermore, there exist minimal models M_a and M_b of T in which γ is satisfied only by virtue of a, respectively, b. Unless enforced by other formulas in T, no minimal model M of T will satisfy both a and b, that is, give an inclusive account of the disjunction of a and b, and $\neg a \vee \neg b$ holds in all minimal models. This is the case in Examples 1.1, 2.2, and 2.3, for instance. While circumscription fosters there exclusive interpretations of the clauses loc_c, $h \vee n$, and $a \vee b \vee c$, respectively, curbing lends itself also for an inclusive interpretation.

On the other hand, an important note is that neither circumscription will effect always an exclusive interpretation of positive disjunctions, nor that curbing always supports also an inclusive interpretation, even if this would be possible. This is shown by the following two examples.

Example 3.2. Consider the theory $T = \{p \vee q,\ p \vee r,\ q \vee r\}$ on $\mathcal{A} = \{p, q, r\}$. Then T has the three minimal models $\{p, q\}$, $\{p, r\}$, and $\{q, r\}$, and none of the three clauses in T is interpreted exclusively in all of them. Notice, however, that inclusive interpretation is made in each minimal model as little as needed, and two out of three clauses are interpreted exclusively. Curbing also selects the non-minimal model $\{p, q, r\}$ which reflects the simultaneous inclusive interpretation of all clauses in T.

Example 3.3. Consider the theory $T = \{p \vee q,\ q \vee r,\ \neg(p \wedge q \wedge r)\}$ on $\mathcal{A} = \{p, q, r\}$. Its minimal models are $M_1 = \{q\}$ and $M_2 = \{p, r\}$, which are also its curb models. Thus like circumscription, curbing interprets the positive clauses in T exclusively. However, since T has also the models $\{p, q\}$ and $\{q, r\}$, both clauses might be interpreted inclusively. Informally, the behavior of curbing is explained by the fact that a joint account of the positive information opposite in M_1 and M_2 is not possible, and that, by the principle of minimality, models $\{p, q\}$ and $\{q, r\}$ are arguably not concise.

The above examples, to which further ones might be added, show that at the syntactic level, the behavior of circumscription and curbing with respect to exclusive and inclusive interpretation of clauses has to be considered with care. To this end, the notions of "exclusive" and "inclusive" interpretation would have to be made formally precise. In order to deal with representation matters, a canonical form of syntactic representation, such as in terms of prime implicates of a theory, may be in order here to surpass syntax matters. Note that in all examples given above, the clauses in the theories are prime implicates. We do not pursue this issue in further depth here, however, as we do not attempt to give a formal definition of inclusive interpretation of disjunction. Rather, the intent is to present a method which deals with disjunctive information on a model-theoretic basis, based on well-established mathematical principles, and *as a feature* admits intuitively inclusive interpretations of disjunctive clauses sometimes (but does not strictly enforce them). We defer further discussion of the interpretation of disjunction and the role of syntax to Section 7.

3.2. Reasoning under curbing

Like under circumscription, inference under curbing may be defined by truth in all curb models. That is, a formula F is a consequence of a theory T under curbing with respect to a minimization setting $(P;\,Z)$, if $cmod_{P;Z}(T) \models F$, i.e., $M \models F$ for each $M \in cmod_{P;Z}(T)$.

Since all minimal models are curb models, the following proposition is immediate.

Proposition 3.4. *For every $(P;\,Z)$, $T \models F$ only if $cmod_{P;Z}(T) \models F$ only if $mmod_{P;Z}(T) \models F$.*

That is, curbing is complete with respect to inference in classical logic, but does not yield more consequences than circumscription. This, however, is coherent with the view that circumscription should desirably be "softened" in some cases, and drawing conclusions that rule out inclusive alternatives is prevented. Looking back at our examples, in Example 2.2, this would be the conclusion that Alice does not hang the painting up, since $mmod_{P;Z}(T_1) \models \neg p$ while $cmod_{P;Z}(T_1) \not\models \neg p$; in Example 2.3, circumscription infers that Doug stays at home, since $mmod_{P;Z}(T_2) \models \neg d$ while $cmod_{P;Z}(T_2) \not\models \neg d$.

Drawing these conclusions under circumscription can also be prevented by choosing appropriate minimization settings. Indeed, we obtain the set of "intuitive" models in Example 2.2 for $P = \{p\}$ and $Z = \emptyset$ (i.e., fix h and n), and in Example 2.3 for $P = \{d\}$ and $Z = \emptyset$ (i.e., fix a, b, and c). While in these cases, one might argue that such a minimization setting is intuitive for circumscription, this seems not to be the case for preventing the circumscriptive conclusion in Example 2.3 that exactly one person is going to the party, for instance; all persons being equal, any minimization setting which treats the persons differently seems to be questionable. Therefore, either all propositions should be minimized in parallel (i.e., global minimization is applied), or all fixed, or all floating. This means that applying circumscription is only effective for global minimization, since in the other settings it (knowingly) coincides with classical logic. Similar considerations apply in Example 1.1 to the circumscriptive conclusion that each coin is located in exactly one field.

As concerns the inference of positive information, curbing behaves like circumscription, which as well-known in the folklore does not infer new positive information on minimized atoms. More precisely,

Proposition 3.5. *Let T be a theory and let $(P;\,Z)$ be a minimization setting. Then for any formula F built with connectives $\vee$, $\wedge$, and $\neg$ such that each atom from P occurs only under an even number of negations and no atom from Z occurs in F, it holds that $T \models F$ iff $cmod_{P;Z}(T) \models F$ iff $mmod_{P;Z}(T) \models F$.*

Proof. The formula F can be rewritten to a disjunctive normal form DNF $D_1 \vee \cdots D_n$ in which each disjunct D_i has only positive literals over P. Thus, $M \models F$ and $M \leqslant_{P;Z} M'$ implies $M' \models F$. Since for every $M' \in mod(T)$ there exists some $M \in mmod_{P;Z}(T)$ such that $M \leqslant_{P;Z} M'$, we have that $mmod_{P;Z}(T) \models F$ implies $T \models F$. The result thus follows by Proposition 3.4. $\square$

As the above examples show, capturing the (global) curb models under a suitable minimization setting in circumscription is more a technical engineering task than an intuitive process. Furthermore, we point here to Theorem 6.7 in Section 6, which shows that the set of curb models of a theory T cannot be represented by polynomial-size formulas under arbitrary circumscription in general (and thus in particular, not by just picking the right minimization setting for T) under standard complexity hypotheses.

On the other hand, a projection setting such as $(P; Z)$ minimization for curbing is needed for capturing the set of models under circumscription. This is shown by Example 2.1, for instance, where the theory T has the two (global) minimal models $M_1 = \{b, f\}$ and $M_2 = \{b, ab\}$ of which only M_1 should be selected.

Selecting a minimization setting $(P; Z)$ for curbing is an issue similar as for circumscription, which has been studied in the literature (see [30] and references therein). We do not aim at providing here an engineering method by which, more syntactically, a set of models can be captured by the clever choice of a minimization setting. Rather, we take here the view that curbing should be applied as a complementary method to circumscription, if a (natural) minimization setting is chosen, to account for possible disjunction at the semantic level. As argued above, this can be fruitfully applied to the party example, where curbing was used under global minimization. For an example with non-empty fixed and floating variables, let us consider a variant of the bird-flies example.

Example 3.4. Consider the theory from Example 2.1 augmented with the knowledge that the bird is a duck, that ducks normally can swim, and that the bird is abnormal. That is, we have

$$\begin{aligned} T = \{&\mathit{bird} \wedge \neg \mathit{ab} \rightarrow \mathit{flies},\\ &\mathit{duck} \wedge \neg \mathit{ab}' \rightarrow \mathit{swims},\\ &\mathit{bird},\ \ \mathit{duck},\ \ \mathit{ab} \vee \mathit{ab}'\ \ \},\end{aligned}$$

where ab' represents abnormality with respect to the ability to swim. Then, in the suggestive minimization setting where $P = \{ab, ab'\}$ and $Z = \{\mathit{flies}, \mathit{swims}\}$, we can conclude $\mathit{flies} \vee \mathit{swims}$ from the circumscription of T w.r.t. $(P; Z)$. Indeed, in each $P; Z$-minimal model of T, either ab is false or ab' is false (but not both), and thus either *flies* or *swims* is true in it. This is because circumscription enforces that exactly one of the two abnormalities holds.

Under curbing of T w.r.t. $P; Z$, which reflects the view that both abnormalities might occur simultaneously, we cannot conclude $\mathit{flies} \vee \mathit{swims}$, since $\{\mathit{bird}, \mathit{duck}, ab, ab'\}$ is a $P; Z$-curb model of T. Of course, under circumscription the same effect is obtained by moving ab and ab' to the fixed variables. However, this conceptually means to abandon the minimization of abnormality. In more complex scenarios, with several notions of abnormality relating to different properties and constraints among them, the result may be different, since certain combinations of abnormalities are accounted by mubs.

For example, suppose that T contains additional axioms $\mathit{bird} \wedge \neg ab'' \rightarrow \mathit{animal}$ and $ab'' \rightarrow (ab \wedge ab')$, stating that a bird is normally an animal and that abnormality with respect to this property implies abnormality with respect to flying and swimming. Under the policy to minimize the abnormalities, while *bird* and *duck* are fixed and the other variables float, we can conclude under circumscription both *animal* and $\mathit{flies} \vee \mathit{swims}$. Under curbing with this setting, we can conclude the former but not the latter (since there is a $P; Z$-curb model $\{\mathit{bird}, \mathit{duck}, \mathit{animal}, ab, ab'\}$). If as previously, we would move the abnormalities to the fixed variables, then neither *animal* nor $\mathit{flies} \vee \mathit{swims}$ can be concluded. However, concluding *animal* seems reasonable. Note that the result of curbing can also be obtained by circumscription when ab and ab' are moved to the fixed variables (and only ab is minimized); however, similar as in the party example, the abnormality properties are then not treated uniformly by the policy. This is less intuitive, in particular if T would be rewritten such that the disjunction $ab \vee ab'$ is no longer explicitly apparent.

3.3. Coincidence of curbing and circumscription

A natural question is when circumscription and curbing coincide. Clearly, this is the case when taking mubs is infeasible. This might be due to the following basic reasons: (1) the theory might have a single minimal model (if it has any model); (2) taking upper bounds is infeasible. For example, $T = \{p \vee q, \neg(p \wedge q)\}$ has the minimal models $M_1 = \{p\}$ and $M_2 = \{q\}$, for which no upper bound exists. By standard arguments, (2) is tantamount to the non-existence of mubs, which follows from an easy proposition.

Proposition 3.6. *Let T be a theory and let $(P; Z)$ be a minimization setting. For every $\mathcal{M} \subseteq mod(T)$, $ub^T_{P;Z}(\mathcal{M}) \neq \emptyset$ implies $mub^T_{P;Z}(\mathcal{M}) \neq \emptyset$.*

Proof. Indeed, if $ub^T_{P;Z}(\mathcal{M})$ is non-empty, construct a (global) minimal model of the theory

$$T' = T \cup \bigcup_{M \in \mathcal{M}} M[P] \quad \cup Q_{\mathcal{M}} \cup \{\neg q | q \in \mathcal{A} \setminus (P \cup Z \cup Q_{\mathcal{M}})\},$$

where $Q_{\mathcal{M}} = \bigcap_{M \in M} M[Q]$ and $Q = \mathcal{A} \setminus (P \cup Z)$, using a (well-ordered) enumeration $\eta = x_0, x_1, \ldots, x_\alpha, \ldots, \alpha \geqslant 0$, of $\mathcal{A}$ in which all atoms of P occur before those of Z. Starting from $\alpha = 0$, we assign each atom x_α the value 0, if the partial truth assignment to all x_β, $\beta < \alpha$, is extendible to some model of T' in which $x_\alpha = 0$, and the value 1 otherwise; in this way, a minimal model N of T' is obtained, which satisfies $N \in mub^T_{P;Z}(\mathcal{M})$, and thus $mub^T_{P;Z}(\mathcal{M}) \neq \emptyset$. □

Condition (1) yields the following easy result; recall that every *Horn* theory T, i.e., set of clauses $L_1 \vee \cdots \vee L_k$, $k \geqslant 1$, such that at most one of the literals L_i is a positive atom, has a least model.

Proposition 3.7. *Suppose T is a Horn theory. Then, $cmod_{P;Z}(T) = mmod_{P;Z}(T)$ for any minimization setting $(P; Z)$.*

In particular, this means that curbing of Horn logic programs is tantamount to circumscribing them.

Condition (2) is enforced, for instance, by the syntactic condition of blocking in a disjunctive normal form (DNF), which is as follows. As usual, a *term* is any satisfiable conjunction of literals $D = L_1 \wedge \cdots \wedge L_k$, $k \geqslant 0$, which we also identify with the set $\{L_1, \ldots, L_k\}$ of its literals. For any set of atoms A, we denote by $Pos_A(D)$ (resp., $Neg_A(D)$) the set of atoms $p \in A$ which occur positively (resp., negatively) in D. Two terms D and D' are *blocking* with respect to A, if both $Pos_A(D) \cap Neg_A(D') \neq \emptyset$ and $Neg_A(D) \cap Pos_A(D') \neq \emptyset$ holds. A DNF $D_1 \vee D_2 \vee \cdots \vee D_n$ is *blocking* with respect to A, if each pair of distinct terms D_i and D_j in it is blocking w.r.t. A. For example, the DNF $F = (a \wedge \neg b \wedge c) \vee (\neg a \wedge b \wedge \neg d)$ is blocking with respect to $P = \{a, b\}$, but not with respect to $P = \{a, c\}$.

Proposition 3.8. *For any theory T and minimization setting $(P; Z)$, $cmod_{P;Z}(T) = mmod_{P;Z}(T)$ if T has some blocking DNF representation w.r.t. P.*

Proof. Towards a contradiction, assume that $M \notin mmod_{P;Z}(T)$ is a $P; Z$-mub of $\mathcal{M} \subseteq mmod_{P;Z}(T)$ such that $|\mathcal{M}| \geqslant 2$. Then M must satisfy some disjunct D_j in any DNF $F = D_1 \vee \cdots \vee D_n$ for T. Consider any $M' \in \mathcal{M}$. We note the following easy lemma.

Lemma 3.9. *Suppose $F = D_1 \vee \cdots \vee D_n$ is a DNF for a theory T, and let $(P; Z)$ be any minimization setting. Then, for every $M \in mmod_{P;Z}(T)$, there exists some D_i* in D *such that $M \vDash D_i$ and $M[P] = Pos_P(D_i)$.*

Let for M' be D_i as in the lemma. Since $|\mathcal{M}| \geqslant 2$, we have $M' <_{P;Z} M$, and thus $M[P] \setminus M[P'] \neq \emptyset$. Moreover, since $M'[P] = Pos_P(D_i)$, also $M'[P] \setminus M[P] \neq \emptyset$ holds if F is blocking w.r.t. P. However, this contradicts $M' <_{P;Z} M$. □

Thus, for the DNF F from above, circumscription and curbing in any minimization setting $(P; Z)$ such that $P = \{a, b\}$ coincide. However, while sufficient, the condition of Proposition 3.8 is not necessary. This is shown by the following example.

Example 3.5. Consider the theory $T = \{D_1 \vee D_2 \vee D_3\}$ where $D_1 = a \wedge b \wedge \neg d$, $D_2 = b \wedge \neg c$, and $D_3 = \neg a \wedge \neg b \wedge c \wedge d$. Its global minimal models are $M_1 = \{b\}$ and $M_2 = \{c, d\}$, for which no upper bound w.r.t. T exists; hence, $cmod(T) = mmod(T)$. On the other hand, there is no blocking DNF for T w.r.t. $P = \{a, b, c, d\}$.

A weakened form of blocking is suitable for a syntactical characterization of finitely representable theories for which curbing and circumscription coincides. Denote for any term D and set of atoms A by $D[A]$ its projection to A, i.e., the maximal subterm of D with literals on A. We call terms $D_1, \ldots, D_n$ *A-satisfiable*, if $D_1[A] \cup \cdots \cup D_n[A]$ contains no pair of opposite literals. Recall that a hitting set of a family $\mathcal{S} = \{S_1, \ldots, S_m\}$ of sets S_i is any subset $H \subseteq \bigcup_i S_i$ which has non-empty intersection with each S_i, $i = 1, \ldots, m$.

Theorem 3.10. *Let T be any theory on $\mathcal{A}$ which is representable by some DNF $F = D_1 \vee \cdots \vee D_n$. Let $(P; Z)$ be any minimization setting, and let $Q = \mathcal{A} \setminus (P \cup Z)$. Then, $cmod_{P;Z}(T) = mmod_{P;Z}(T)$ if and only if for every pair D_{i_1}, D_{i_2} of Q-satisfiable terms in F such that $Pos_P(D_{i_1}) \nsubseteq Pos_P(D_{i_2}) \nsubseteq Pos_P(D_{i_1})$ (i.e., $Pos_P(D_{i_1})$ and $Pos_P(D_{i_2})$ are incomparable), the following two conditions hold*:

(i) *Either D_{i_1} and D_{i_2} are blocking w.r.t. P, or the family $\mathcal{S}_{i_1,i_2} = \{D_{i_3}[Q] \setminus (D_{i_1} \cup D_{i_2}) | D_{i_3}$ in F, D_{i_1}, D_{i_2}, D_{i_3} are Q-satisfiable and $Pos_P(D_{i_3}) \subseteq Pos_P(D_{i_1}) \cap Pos_P(D_{i_2})\}$ has no satisfiable hitting set.*

(ii) *For every term D_{i_3} in F such that $Pos_P(D_{i_1}) \cup Pos_P(D_{i_2}) \subseteq Pos_P(D_{i_3})$ and D_{i_1}, D_{i_2}, D_{i_3} are Q-satisfiable, the family $\mathcal{S}_{i_1,i_2,i_3} = \{D_{i_4}[Q] \setminus (D_{i_1} \cup D_{i_2} \cup D_{i_3}) | D_{i_4}$ in F, $D_{i_1}, \ldots, D_{i_4}$ are Q-satisfiable and $Pos_P(D_{i_4}) \subseteq Pos_P(D_{i_1}) \cap Pos_P(D_{i_2})\}$ has no satisfiable hitting set.*

Informally, Condition (i) enforces blocking for terms D_{i_1} and D_{i_2} which represent minimal models, such that no upper bound of them is a model of any of the two terms. Condition (ii) takes explicit upper bounds by taking other terms into account.

Proof. For the only-if part, suppose towards a contradiction that $cmod_{P;Z}(T) = mmod_{P;Z}(T)$ but either condition (i) or (ii) fails for some terms D_{i_1} and D_{i_2} in F.

In case (i), this means that D_{i_1} and D_{i_2} are not blocking w.r.t. P and there is a satisfiable hitting set H of S_{i_1,i_2}. This means that there is no term D_{i_3} in F such that $Pos_P(D_{i_3}) \subseteq Pos_P(D_{i_1}) \cap Pos_P(D_{i_2})$ and $D_{i_3} \models D_{i_1}[Q] \cup D_{i_2}[Q] \cup \neg.H$, where $\neg.H = \{\neg p | p \in H\} \cup \{p | \neg p \in H\}$ denotes the set of literals opposite to those in H. This implies that there exist incomparable $M_1, M_2 \in mmod_{P;Z}(T)$ such that $M_1[Q] = M_2[Q]$, and for $j \in \{1, 2\}$, we have $M_j \models D_{i_1}[Q] \cup D_{i_2}[Q] \cup \neg.H$ and $M_j[P] \subseteq Pos_P(D_{i_j})$. Without loss of generality, suppose that $Pos_P(D_{i_1}) \cap Neg_P(D_{i_2}) = \emptyset$; this is feasible since D_{i_1} and D_{i_2} are not blocking w.r.t. P. Then, the model $M = Pos_P(D_2) \cup M_1[P] \cup M_2[Q \cup Z]$ satisfies D_{i_2}, and thus $M \in ub^T_{P;Z}(\{M_1, M_2\})$. By Proposition 3.6, it follows that $mub^T_{P;Z}(\{M_1, M_2\}) \setminus mmod_{P;Z}(T) \neq \emptyset$, and thus $cmod_{P;Z}(T) \neq mmod_{P;Z}(T)$. This is a contradiction, which proves case (i).

In case (ii), we have some term D_{i_3} in F such that (a) $Pos_P(D_{i_1}) \cup Pos_P(D_{i_2}) \subseteq Pos_P(D_{i_3})$ and D_{i_1}, D_{i_2}, D_{i_3} are Q-satisfiable, and (b) S_{i_1,i_2,i_3} has a satisfiable hitting set H ($\subseteq (Q \cup \neg.Q) \setminus (D_{i_1} \cup D_{i_2} \cup D_{i_3})$). By condition (a), there exist models M_1, M_2, and M_3 of T, such that for $j \in \{1, 2, 3\}$, $M_j \models D_{i_j}$ and $M_j[P] = Pos_P(D_{i_j})$, and $M_3 \in ub^T_{P;Z}(\{M_1, M_2\})$. Condition (b) implies that there is no term D_{i_4} in F such that $D_{i_1}, \ldots, D_{i_4}$ are Q-satisfiable, $Pos_P(D_{i_4}) \subseteq Pos_P(D_{i_1}) \cap Pos_P(D_{i_2})$, and $D_{i_4} \cup D_{i_1}[Q] \cup D_{i_2}[Q] \cup D_{i_3}[Q] \cup \neg.H$ is satisfiable. Thus we may assume that $M_j \models D_{i_1}[Q] \cup D_{i_2}[Q] \cup D_{i_3}[Q] \cup \neg.H$ for $j \in \{1, 2, 3\}$ and that no $M \in mod(T)$ satisfies $M \leqslant_{P;Z} M_1$ and $M \leqslant_{P;Z} M_2$. Hence there exist incomparable $N_1, N_2 \in mmod_{P;Z}(T)$ such that $N_1 \leqslant_{P;Z} M_1$, and $N_2 \leqslant_{P;Z} M_2$. Since $M_3 \in ub^T_{P;Z}(\{N_1, N_2\})$, similar as above by applying Proposition 3.6 we arrive at a contradiction.

For the if part, consider any DNF for T which verifies the above conditions, and suppose towards a contradiction that $cmod_{P;Z}(T) \neq mmod_{P;Z}(T)$. Then, there exists some $M \notin mmod_{P;Z}(T)$ and $\mathcal{M} \subseteq mmod_{P;Z}(T)$ such that $M \in mub^T_{P;Z}(\mathcal{M})$. Hence, there are incomparable models $M_1, M_2 \in \mathcal{M}$ such that $M \in ub^T_{P;Z}(\{M_1, M_2\})$ (and thus $M[Q] = M_1[Q] = M_2[Q]$). By Lemma 3.9, for every $j \in \{1, 2\}$ there is a term D_{i_j} in F such that $M_j \models D_{i_j}$ and $M_j[P] = Pos_P(D_{i_j})$. Therefore, D_{i_1}, D_{i_2} are Q-satisfiable and $Pos_P(D_{i_1})$, $Pos_P(D_{i_2})$ are incomparable. Since $M \in mod(T)$, there must exist some term D_k in F such that $M \models D_k$. Condition (ii) for D_{i_1}, D_{i_2} implies that there is no term D_k in F such that $Pos_P(D_k) \supseteq Pos_P(D_{i_1}) \cup Pos_P(D_{i_2})$ and $M \models D_k$. Indeed, suppose such a term D_k would exist. By $P; Z$-minimality of M_1 and M_2, there is no term D_{i_4} in F such that $M[Q] \models D_{i_4}[Q]$ and $Pos_P(D_{i_4}) \subseteq Pos_P(D_{i_1}) \cap Pos_P(D_{i_2})$. Hence, $\mathcal{S}_{i_1,i_2,k}$ would have the satisfiable hitting set $H = (Q \setminus M) \cup \neg.M[Q]$; this would contradict that Condition (ii) is satisfied by D_{i_1} and D_{i_2}.

From $P; Z$-minimality of M_1 and M_2, it thus follows that $M \models D_k$ for some term D_k in F such that, for some $j \in \{1, 2\}$, $Pos_P(D_k)$ and $Pos_P(D_{i_j})$ are incomparable and D_k, D_{i_j} are not blocking w.r.t. P. By Condition (i) for $D_{i'_1} = D_k$ and $D_{i'_2} = D_{i_j}$, the family $\mathcal{S}_{i'_1,i'_2}$ has no satisfiable hitting set. Thus in particular, $H = (Q \setminus M) \cup \neg.M[Q]$ is not a hitting set of $\mathcal{S}_{i'_1,i'_2}$, i.e., there exists some set $S_l \in \mathcal{S}_{i'_1,i'_2}$ such that $H \cap S_l = \emptyset$. This implies that there exists some term D_l in F such that $M[Q] \models D_l[Q]$ and $Pos_P(D_l) \subset Pos_P(D_{i_j})$. Hence, there exists some model $N <_{P;Z} M_j$ such that $N \models D_l$ and thus $N \models T$. However, this means that $M_j \notin mmod_{P;Z}(T)$, which is a contradiction. □

We remark that testing the satisfiable hitting set condition on S_{i_1,i_2} and S_{i_1,i_2,i_3} in Theorem 3.10 is coNP-complete in general; however, such a complex test is justified by the fact that, as easily seen, deciding whether $cmod_{P;Z}(F) = mmod_{P;Z}(F)$ for a given DNF F is coNP-complete in general.

If $Q = \emptyset$ (no atoms are fixed), conditions (i) and (ii) amount to the following ones, which can be checked in polynomial time:

(i′) Either D_{i_1} and D_{i_2} are blocking w.r.t. P, or there is some term D_{i_3} in F such that $D_{i_1}, D_{i_2}, D_{i_3}$ are Q-satisfiable and $Pos_P(D_{i_3}) \subseteq Pos_P(D_{i_1}) \cap Pos_P(D_{i_2})$.

(ii′) For every term D_{i_3} in F such that $Pos_P(D_{i_1}) \cup Pos_P(D_{i_2}) \subseteq Pos_P(D_{i_3})$ and $D_{i_1}, D_{i_2}, D_{i_3}$ are Q-satisfiable, there exists a term D_{i_4} in F such that $D_{i_1}, \ldots, D_{i_4}$ are Q-satisfiable and $Pos_P(D_{i_4}) \subseteq Pos_P(D_{i_1}) \cap Pos_P(D_{i_2})$.

Note that in Example 3.5 the (unordered) pairs D_1, D_3 and D_2, D_3 are those incomparable under global minimization. In both cases, the terms are blocking, and thus condition (i) holds. Condition (ii) is vacuously satisfied since the remaining term in the DNF does not match the hypothesis. Thus, global curbing and circumscription for T coincide.

The following example illustrates that restricting the test to terms D_{i_1}, D_{i_2} in F with minimal positive part on P (as one might think) is not possible, and that the hitting set part of Condition (i) is needed.

Example 3.6. Consider the theory $T = \{D_1 \vee D_2 \vee D_3 \vee D_4\}$ where $D_1 = a \wedge \neg b \wedge c \wedge \neg d$, $D_2 = \neg a \wedge b \wedge c \wedge \neg d$, $D_3 = d \wedge e$, and $D_4 = \neg c \wedge d$ under global minimization. Then, D_1, D_2, and D_4 are the terms with minimal positive part, and the pairs (D_1, D_2), (D_1, D_4), and (D_2, D_4) are blocking. Hence, Condition (i) is satisfied and Condition (ii) is vacuously true. However, $cmod(T) \neq mmod(T)$, since the minimal models $M_1 = \{a, c\}$ and $M_2 = \{b, c\}$ have the mub $M = \{a, b, c, d, e\}$ in T. Note that D_1 and D_3 violate Condition (i), since they are not blocking and F has no negative term.

4. Restricted curbing

In the previous section, we have defined the curb models as those models which an idealized reasoner should adopt under the principle of closedness under mubs, without further discrimination among these models. In this section, we consider two restricted notions of curbing. The first notion puts a bound on the number of iterations which can be made on including, starting from the minimal models, mubs of models. This may be interpreted as reasoning with limited resources. In another view, it may be seen as reasoning with a bounded horizon of plausibility, i.e., models which are obtained as mubs of other models are considered to be less plausible then these models. In iteration of mubs, the plausibility thus decreases, and may be considered to be below a threshold after a certain number of iterations.

The second restriction constrains taking mubs of sets which contain a bounded number of models. Intuitively, this corresponds to limiting the basic reasoning capability of the agent itself (but not the resources).

As shown below, these two restrictions have different effects on the expressiveness of the formalism. Both can be seen as sound *approximations* of the curb models, which in the limit (under a growing number of iterations resp. cardinality of sets of models for taking mubs) coincide with the curb models.

4.1. Bounded iteration depth

The first restriction limits iterated inclusion of mubs. Informally, we choose only the models that are α-good for some ordinal α such that $|\alpha| \leqslant |\delta|$, where the ordinal δ is a limit on the depth in building mubs.

Definition 4.1. For any theory T and minimization setting $(P; Z)$, the set of *δ-depth curb models of T w.r.t.* $(P; Z)$ is given by $\delta\text{-}cmod_{P;Z}(T) = \{M \in cmod_{P;Z}(T) | rank^T_{P;Z}(M) \leqslant \delta\}$.

Example 4.1. For example, $\{h\}$ and $\{n\}$ in Example 2.2 are both 0-curb models of T_1, while $\{h, n, p\}$ is 1-curb but not 0-curb. Thus, the curb models of T_1 coincide with its 1-curb models. In Example 2.3, $\{a, b, c, d\}$ is a 2-curb model of T_2 but not a 1-curb-model; in fact, $1\text{-}cmod(T_2)$ clips from the curb models of T_2 just this model.

In certain cases, bounding the depth for building curb models does not lead to a loss of curb models. We shall encounter this in Section 5.

4.2. Bounded mub width

Another possibility is to limit the cardinality of model sets from which mubs are formed. Intuitively, this corresponds to limiting the number of inclusively interpreted disjuncts.

Definition 4.2. Let T be any theory and $(P; Z)$ be any minimization setting. For any cardinal $\kappa \geqslant 0$, a set of models $\mathcal{M} \subseteq mod(T)$ of a theory T is closed under κ-width $P; Z$-mubs w.r.t. T, if $mub^T_{P;Z}(\mathcal{N}) \subseteq \mathcal{M}$ for every $\mathcal{N} \subseteq \mathcal{M}$ such that $|\mathcal{N}| \leqslant \kappa$. The set of κ-width $P; Z$-curb models of T, denoted by $cmod^{\kappa}_{P;Z}(T)$, is the least set of models of T which is closed under κ-width $P; Z$-mubs w.r.t. T.

Clearly, $cmod^0_{P;Z}(T)$ and $cmod^1_{P;Z}(T)$ coincide with the circumscription of T w.r.t. $(P; Z)$. For $\kappa \geqslant 2$, we obtain the following result.

Theorem 4.1. *Let T and $(P; Z)$ such that no infinite increasing chain $M_0 <_{P;Z} M_1 <_{P;Z} \ldots$ of models $M_i \in cmod_{P;Z}(T)$, $i \geqslant 0$, exists. Then, for every $\kappa \geqslant 2$, $cmod_{P;Z}(T) = cmod^{\kappa}_{P;Z}(T)$.*

Proof. We show this by induction on the level $\ell(M) \geqslant 0$ of $M \in cmod_{P;Z}(T)$, which is defined by

$$\ell(M) = \begin{cases} 0 & \text{if } M \in mmod_{P;Z}(T), \\ 1 + \min\{\ell(M') | M' \in \max_{P;Z}(cmod^{<M}_{P;Z}(T))\} & \text{otherwise.} \end{cases}$$

Notice that in absence of increasing infinite chains, $\ell(M)$ is well-defined and finite.

The base case $\ell(M) = 0$ is trivial. Consider then the case $\ell(M) > 0$, and let $\mathcal{M} = cmod^{<M}_{P;Z}(T)$. By Proposition 3.2, $M \in mub^T_{P;Z}(\mathcal{M})$. If $\mathcal{M} = \emptyset$, the statement trivially holds. Otherwise, let $M' \in \mathcal{M}$ be maximal w.r.t. $\leqslant_{P;Z}$; since all chains in $cmod_{P;Z}(T)$ are finite, such an M' must exist. Let $M'' \in \mathcal{M}$ such that $M'' \nleqslant_{P;Z} M'$. Also such an M'' must exist; otherwise, we would have $M \notin mub^T_{P;Z}(\mathcal{M})$. By the maximality of M', $M \in mub^T_{P;Z}(\{M', M''\})$. Since $\ell(M'), \ell(M'') < \ell(M)$, by the induction hypothesis $M', M'' \in cmod^2_{P;Z}(T)$. Hence, $M \in cmod^2_{P;Z}(T) \subseteq cmod^{\kappa}_{P;Z}(T)$. Conversely, $M \in cmod^{\kappa}_{P;Z}(T)$ clearly implies $M \in cmod_{P;Z}(T)$. This proves the result. □

Theorem 4.1 implies a dichotomy result on the expressivity of κ-bounded disjuncts: either we get only the minimal models (for $\kappa = 0, 1$) or all curb models. Note that the hypothesis of the theorem trivially holds if P is finite, and it also holds if T is finite.

Note that even in the case where P is finite, Theorem 4.1 is not an immediate consequence of Theorem 3.3. A simple inductive argument on the rank of curb models does not work, as can be seen from Example 3.1: for $\kappa = 2$, the curb model $\{a, b, c\}$, whose rank is 1, is obtained as a 2-mub only of curb models M and M' of which at least one has also rank 1.

Theorem 4.1 fails if infinite chains of curb models occur, as shown by the following example.

Example 4.2. Let $\mathcal{A} = \{x_i | i = 0, 1, 2, \ldots\}$ *and consider the following theory* T:

$$\begin{aligned} T = \{&x_0 \vee x_1, \\ &(x_{2i} \wedge x_{2i+1}) \rightarrow (x_{2i+2} \vee x_{2i+3}), \\ &(x_{2i} \leftrightarrow \neg x_{2i+1}) \rightarrow \neg x_j, \quad i \geqslant 0, \ j > 2i+1 \ \}. \end{aligned}$$

The models of T are $\mathcal{A}$ and all models $\{x_0, x_1, \ldots, x_{2i}\}$ and $\{x_0, x_1, \ldots, x_{2i-1}, x_{2i+1}\}$, for $i \geqslant 0$, which under global minimization are ordered as shown in Fig. 5. Each model of T is a curb model; however, the model $\mathcal{A}$ is not a mub of any finite family of smaller (curb) models of T.

As for iterating mubs, the power of finite disjunctions of models is limited to what can be obtained in a finite (but unbounded) number of steps. Let, for any ordinal δ and cardinal κ, denote δ-$cmod^{\kappa}_{P;Z}(T)$ the set of all δ-curb models of T w.r.t. $(P; Z)$, if $P; Z$-mubs are replaced by κ-width $P; Z$-mubs in the definition.

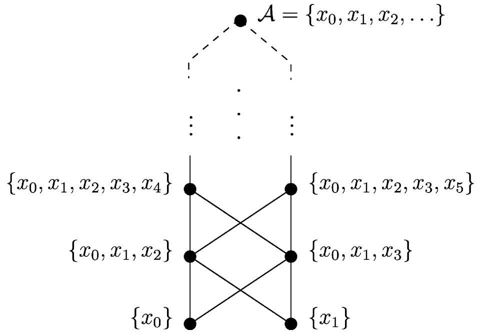

Fig. 5. The infinite tower.

Proposition 4.2. *Let κ be finite. Then, $cmod^{\kappa}_{P;Z}(T) = \omega\text{-}cmod^{\kappa}_{P;Z}(T)$.*

Proof. To show this result, we consider the variant $\Lambda_{T,\kappa}$ of the operator Λ_T in the proof of Theorem 3.3 which results by adding the condition "$|\mathcal{N}| \leqslant \kappa$" on $\mathcal{N}$. Also $\Lambda_{T,\kappa}$ is a monotone operator and, moreover, finitizable, i.e., $\Lambda_{T,\kappa}(\mathcal{M}) = \bigcup_{\mathcal{M}' \subseteq \mathcal{M}: |\mathcal{M}'| < \aleph_0} \Lambda_{T,\kappa}(\mathcal{M}')$. Hence the least fixpoint of $\Lambda_{T,\kappa}$ is given by an element S_γ of the sequence S_α, $\alpha \geqslant 0$, for $\Lambda_{T,\kappa}$ analogous to the one for $\Lambda_{T,\kappa}$ such that $\gamma \leqslant \omega$. Consequently, $M \in cmod^{\kappa}_{P;Z}(T)$ implies $M \in \omega\text{-}cmod^{\kappa}_{P;Z}(T)$. The converse is trivial. $\square$

Thus, in case of a denumerable set of atoms $\mathcal{A}$, each $^{\kappa}$-width curb model of T can be finitely constructed by taking mubs of κ-width, and all curb models of T are obtainable in this way if there is no infinitely increasing chain of curb models. In order to effectively prune the set of curb models, some further constraint is needed, for example, an additional bound on the iteration depth. In Example 4, under the simultaneous bounds $\kappa = 2$ and $\delta = 1$ on width and depth, respectively, the model $\{a, b, c\}$ is no longer obtainable.

5. Least upper bounds

In general, a set of curb models does not have a unique mub in the models of a theory. In this section, we first investigate the case where this property holds. We then consider a generalization in which not all mubs are unique, but all curb models can be reconstructed from lubs.

5.1. LUB property

We start with the following definitions.

Definition 5.1. Given a minimization setting $(P; Z)$, a model M is a $P; Z$-least upper bound (lub) of a set of models $\mathcal{M}$ w.r.t. a theory T, if $M =_{P;Z} N$ holds for each $N \in mub^T_{P;Z}(\mathcal{M})$.

Definition 5.2. A theory T has the $P; Z$-LUB property, if every non-empty $\mathcal{M} \subseteq cmod_{P;Z}(T)$ such that $mub^T_{P;Z}(\mathcal{M}) \neq \emptyset$ has a $P; Z$-lub w.r.t. T.

For example, the theory in the hammer–nail-painting scenario (Example 2.2) has the LUB property.

Our first aim is to provide a characterization of LUB theories. In fact, a simple criterion exists provided that only finite disjunctions are needed to construct all curb models of a theory T, i.e., every $P; Z$-curb model of T is a $P; Z$-mub of some *finite* set $\mathcal{M} \subseteq cmod_{P;Z}(T)$ with respect to T. Let us call any such theory with this property $P; Z$-*mub-compact*. We note the following simple lemma.

Lemma 5.1. *Let, for $i \in \{1, 2\}$ be $M_i \in mub^T_{P;Z}(\mathcal{M}_i)$, where $\emptyset \subset \mathcal{M}_1 \subseteq \mathcal{M}_2 \subseteq cmod_{P;Z}(T)$. If M_1 is a $P; Z$-lub of $\mathcal{M}_1$ w.r.t. T, then $M_1 \leqslant_{P;Z} M_2$.*

Theorem 5.2. *Let T be a $P; Z$-mub-compact theory. Then T has the $P; Z$-LUB property iff every $\{M_1, M_2\} \subseteq cmod_{P;Z}(T)$ such that $mub^T_{P;Z}(\{M_1, M_2\}) \neq \emptyset$ has a $P; Z$-lub w.r.t. T.*

Proof. For the *if* direction, we show by induction on finite cardinality $\kappa \geqslant 0$ that every set $\mathcal{M} \subseteq cmod_{P;Z}(T)$ such that $|\mathcal{M}| \leqslant \kappa$ and $mub^T_{P;Z}(\mathcal{M}) \neq \emptyset$ has a $P; Z$-lub w.r.t. T. For $\kappa \leqslant 2$, this holds by the hypothesis. For $\kappa > 2$, let $M \in \mathcal{M}$ be maximal under $\leqslant_{P;Z}$. By the induction hypothesis, $\mathcal{M} \setminus \{M\}$ has a $P; Z$-lub M' w.r.t. T. Let $N \in mub^T_{P;Z}(\mathcal{M})$ be arbitrary. By Lemma 5.1 $M' \leqslant_{P;Z} N$. Again by the hypothesis, the set $\{M, M'\}$ has a $P; Z$-lub M'' w.r.t. T. Since $M, M' \leqslant_{P;Z} N$, it follows $M'' \leqslant_{P;Z} N$. Since $M'' \in ub^T_{P;Z}(\mathcal{M})$, minimality of N implies $M'' =_{P;Z} N$. This means that N is a $P; Z$-lub of $\mathcal{M}$ w.r.t. T, which concludes the induction. From the $P; Z$-mub-compactness of T and Lemma 5.1, we now easily obtain that T has the $P; Z$-LUB-property. The *only if* direction is trivial. □

The theory T in Example 3.1 illustrates this theorem (recall that all models are global curb models): Each pair of models has a lub, and indeed each nonempty set of models has a lub, thus P satisfies the LUB property. This remains true if fresh atoms were added both to the top model and arbitrarily to the models in the middle layer in Fig. 4.

An important consequence of this theorem is that in case of finite theories, the LUB property is tantamount to the property that each pair of different curb models has a lub. This means that the LUB property is invariant to bounding the width of mubs.

The following result is an easy corollary of Theorems 4.1 and 5.2.

Corollary 5.3. *Let T be a theory. Suppose that $mod(T)$ forms an upper semi-lattice with respect to $\leqslant_{P;Z}$ and that no infinite chain $M_0 <_{P;Z} M_1 <_{P;Z} \ldots$ of models $M_i \in cmod_{P;Z}(T)$, $i \geqslant 0$, exists. Then T has the $P; Z$-LUB property, and $cmod_{P;Z}(T) = 1\text{-}cmod_{P;Z}(T)$.*

We now consider the number of steps which are needed to construct all curb models of a theory which satisfies the LUB property. Surprisingly, no iteration is needed in this case.

Theorem 5.4. *Suppose the theory T has the $P;Z$-LUB property. Then, $cmod_{P;Z}(T) = 1 - cmod_{P;Z}(T)$.*

Proof. We show by induction on $rank^T_{P;Z}(M) \geqslant 0$ that every $M \in cmod_{P;Z}(T)$ is a 1-curb model of T. For $rank^T_{P;Z}(M) \leqslant 1$, this is obvious. Consider thus $rank^T_{P;Z}(M) > 1$. Then, by Theorem 3.3 $M \in mub^T_{P;Z}(\mathcal{M})$ for $\mathcal{M} = \{M' \in cmod^{<M}_{P;Z}(T) | rank^T_{P;Z}(M') < rank^T_{P;Z}(M)\}$. Let $\mathcal{M}_m = \mathcal{M} \cap mmod_{P;Z}(T)$ be the set of all $P; Z$-minimal models of T in $\mathcal{M}$.

If $\mathcal{M}_m = \emptyset$, then $M \in mmod_{P;Z}(T)$ and the statement holds. Otherwise, $\mathcal{M}_m$ has the $P; Z$-lub M_m w.r.t. T. By the induction hypothesis, for each $M' \in \mathcal{M}$ we have $M' \in mub^T_{P;Z}(\mathcal{M}_{M'})$ for some $\mathcal{M}_{M'} \subseteq mmod_{P;Z}(T)$ ($\mathcal{M}_{M'} = \emptyset$ if $M' \in mmod_{P;Z}(T)$). Since T has the LUB property and clearly $\mathcal{M}_{M'} \subseteq \mathcal{M}_m$, we have by Lemma 5.1 $M' \leqslant_{P;Z} M_m$ for each $M' \in \mathcal{M}$. Thus $M_m \in ub^T_{P;Z}(\mathcal{M})$; since M is a $P; Z$-lub of $\mathcal{M}$ w.r.t. T, we have $M \leqslant_{P;Z} M_m$. On the other hand, since $\mathcal{M}_m \subseteq \mathcal{M}$ and M_m is a lug of $\mathcal{M}_m$ w.r.t. T, Lemma 5.1 implies that $M_m \leqslant_{P;Z} M$. It follows $M_m =_{P;Z} M$. Thus, $M \in 1\text{-}cmod_{P;Z}(T)$ and the statement holds. □

An immediate consequence is that LUB theories do not require iteration on taking mubs. Therefore, reasoning does not become more difficult than under circumscription in the worst case (see Section 6).

5.1.1. Examples of LUB theories

Examples of theories which enjoy the LUB property are all *dual-Horn* theories, which are clausal theories in which each clause contains at most one negative literal, or positive theories (which contain only *positive* formulas, i.e., formulas built with connectives $\vee$, $\wedge$, and $\neg$ such that each atom occurs only under an even number of negations). In both cases, the set of models of a theory T is closed under union, and hence, trivially T satisfies the LUB property in any minimization setting $(P; Z)$.

Notice that closedness of $mod(T)$ under unions was shown in [45] for quadratic dual-Horn theories, where a theory is *quadratic* (or *Krom*), if it contains only clauses of size at most 2. The authors used this lemma to establish that Curb Model Checking for quadratic dual-Horn theories is in Σ^P_2.

The LUB property holds in fact for more general classes of theories, even in presence of fixed and floating atoms. One such class are P; Z-NegOrDual-Horn theories, which are defined as follows.

Definition 5.3. Let T by a clausal theory and let $(P; Z)$ be a minimization setting. Then T is P; *Z-NegOrDual-Horn*, if for every clause $\gamma \in T$ it holds that, if γ contains some positive literal on P, then γ contains at most one negative literal on P and no literal on Z.

For example, $T = \{p_1 \vee p_2 \vee \neg p_3 \vee q, \neg p_1 \vee \neg p_2 \vee z, \neg p_3 \vee \neg q \vee \neg z\}$ is P; Z-NegOrDual-Horn for $P = \{p_1, p_2, p_3\}$ and $Z = \{z\}$. Then, we have the following result.

Theorem 5.5. *Let T be a P; Z-NegOrDual-Horn theory. Then T has the P; Z-LUB property.*

Proof. We show the following: for every $N \in ub^T_{P;Z}(\mathcal{M})$ such that $\mathcal{M} \subseteq mod(T)$ and $|\mathcal{M}| \geqslant 2$,

$$M := \left(\bigcup_{M' \in \mathcal{M}} M'[\mathcal{A} \setminus Z]\right) \cup N[Z]$$

is a model of T. Consider any clause $\gamma \in T$. If $N \models L$ for some literal $L \in \gamma$ such that either L is from $\mathcal{A} \setminus P$ or L is from P and negative, then by definition of M clearly $M \models L$ and hence $M \models \gamma$. Suppose that none of the two cases applies and, hence, $N \models L$ holds for some positive literal L from P such that $L \in \gamma$. The P; Z-NegOrDual-Horn condition in Definition 5.3 then implies that γ must not contain a literal on Z. Consider any model $M' \in \mathcal{M}$. If $M' \models L'$, $L' \in \gamma$ for some positive literal L' from P, then $M \models L'$ and thus $M \models \gamma$. If no such model M' exists, then γ must contain a negative literal L'' from P such that $M' \models L''$ for every $M' \in \mathcal{M}$. It follows $M \models L''$ and hence $M \models \gamma$. Since γ was arbitrary, this shows that $M \models T$, which proves the result. □

We observe that the P; Z-NegOrDual-Horn property applies e.g., under global minimization to clausal theories in which each clause is either positive or negative, and thus describes a class of formulas for which the satisfiability problem is NP-complete.

An easy corollary of Theorem 5.5 is the following:

Corollary 5.6. *Any quadratic theory T in which no clause contains a positive literal on P and a literal on Z fulfills the P; Z-LUB property.*

The above condition trivially holds if $Z = \emptyset$. We remark that it can be similarly be shown that each quadratic theory fulfills the LUB property w.r.t. $(P; Z)$ if all clauses containing literals on Z are Horn.

5.2. Weak LUB property

The LUB property defined above requires that *every* non-empty family of curb models of a theory has a lub if some mub exists. This property is quite strong. One possibility to relax it is to abandon the requirement that each set of curb models has a lub, but to retain that each curb model can be obtained as a lub of some family of smaller curb models. This motivates the following definition.

Definition 5.4. A theory T has the weak P; Z-least upper bound (Weak P; Z-LUB) property, if every model $M \in cmod_{P;Z}(T) \setminus mmod_{P;Z}(T)$ is the P; Z-lub of some set $\mathcal{M} \subseteq cmod_{P;Z}(T)$ w.r.t. T.

Notice that the LUB property implies the Weak LUB property, but not vice versa. This is shown by the following example.

Example 5.1. Suppose a theory T on $\mathcal{A} = \{a, b, c, d\}$ has the models shown in Fig. 6. All models are curb models under global minimization, and $M_1 = \{a, b, c\}$, $M_2 = \{b, c, d\}$ are the lubs of the sets $\{\{a\}, \{b\}, \{c\}\}$ and $\{\{b\}, \{c\}, \{d\}\}$, respectively. However, the curb models $\{b\}$ and $\{c\}$ do not have a lub; thus, the theory satisfies the Weak LUB property but not the LUB property.

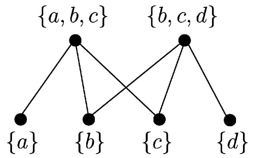

Fig. 6. The Weak LUB property does not imply the LUB property.

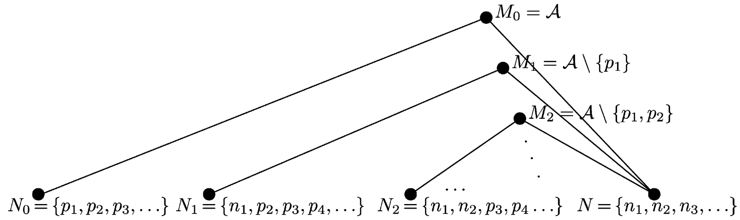

Fig. 7. A family of curb models which is not well-founded ($\mathcal{A} = \{p_i, n_i | i = 1, 2, \ldots\}$).

Intuitively, if a theory satisfies the Weak LUB property, then any model M in a set of curb models $\mathcal{M}$ can be replaced by a set of curb models $\mathcal{M}'$ whose lub is M, without affecting the mubs of the family, i.e., $\mathcal{M}$ has the same mubs as $\mathcal{M} \setminus \{M\} \cup \mathcal{M}'$. By repeating this replacement, $\mathcal{M}$ can be turned into a set $\mathcal{M}^*$ of minimal models which has the same mubs as $\mathcal{M}$. This is actually the case, provided that the set of curb models has the following property.

Definition 5.5. The set $cmod_{P;Z}(T)$ is well-founded, if every decreasing chain $M_0 \geqslant_{P;Z} M_1 \geqslant_{P;Z} \cdots \geqslant_{P;Z} M_\alpha \geqslant_{P;Z} \cdots$ of P; Z-curb models has a smallest element, i.e., there exists some index α such that $M_\beta =_{P,Z} M_\alpha$ for each $\beta \geqslant \alpha$.

Notice that in the context of minimal model reasoning, theories were sometimes called well-founded if every model M of a theory T includes a minimal model of T [30], which is different.

As shown by the following example, the set of curb models of a theory is not necessarily well-founded,

Example 5.2. Consider the following theory T on $\mathcal{A} = \{p_i, n_i | i = 1, 2, \ldots\}$:

$$\begin{aligned} T = \{&n_1 \vee p_1, \\ &\quad p_i \to p_{i+1},\ \neg n_i \to \neg n_{i+1}, \\ &\quad (n_i \wedge \neg n_{i+1}) \to (\neg p_i \wedge p_{i+1}), \quad i \geqslant 1\}. \end{aligned}$$

Informally, $\mathcal{A}$ contains an atom for each non-zero integer, where p_i stands for the positive integer i and n_i for the negative integer $-i$. Every model of T describes a set of integers, where an integer belongs to the set if its corresponding atom has value 1 in the model. The sets captured by T are as follows (see Fig. 7): the integers except 0 (given by model M_0); all negative integers (model N); all integers except some interval $[0, 1, 2, \ldots, i]$, $i \geqslant 1$ (model M_i); and all positive integers but where the interval $[1, k]$, $k \geqslant 0$, is replaced by the interval $[-k, -1]$ (model N_k; in particular, $N_0 = \{p_1, p_2, p_3, \ldots\}$).

As easily checked, the minimal models of T under global minimization are N and all N_k, $k \geqslant 0$. Furthermore, all models of T are curb models: indeed, every other model M_i, $i \geqslant 0$, is a mub of the models N and N_i (in fact, their union; see Fig. 7). Clearly, $M_0 > M_1 > \cdots$ forms an infinite decreasing chain of curb models. This chain has no smallest element, and hence $cmod_{P;Z}(T)$ is not well-founded.

Theorem 5.7. *Let T be a theory which fulfills the Weak P; Z-LUB property. If $cmod_{P;Z}(T)$ is well-founded, then every $M \in cmod_{P;Z}(T) \setminus mmod_{P;Z}(T)$ is the P; Z-lub of some $\mathcal{M} \subseteq mmod_{P;Z}(T)$.*

Proof. Let $\mathcal{B} \subseteq cmod_{P;Z}(T) \setminus mmod_{P;Z}(T)$ be the set of non-minimal P; Z-curb models of T which are not the P; Z-lub of some $\mathcal{M} \subseteq mmod_{P;Z}(T)$. Towards a contradiction, assume that $\mathcal{B} \neq \emptyset$. Since $cmod_{P;Z}(T)$ is well-founded,

$\mathcal{B}$ must have a minimal element M. (To obtain such an M, construct in $\mathcal{B}$ a maximal decreasing chain $M_0 >_{P;Z} M_1 >_{P;Z} \cdots$ and take the smallest element from this chain, which must exist.) Since T has the Weak $P; Z$-LUB property, M is the $P; Z$-lub of some set $\mathcal{M} \subseteq cmod_{P;Z}(T)$. The minimality of M and the Weak $P; Z$-LUB property of T imply that every $N \in \mathcal{M}$ is the $P; Z$-lub of some set $\mathcal{M}_N \subseteq mmod_{P;Z}(T)$. We show that M is the $P; Z$-lub of $\mathcal{M}' = \bigcup_{N \in \mathcal{M}} \mathcal{M}_N$ w.r.t. T. Clearly, $M \in ub^T_{P;Z}(\mathcal{M}')$. Assume then that M is not minimal, i.e., there exists some $M' \in ub^T_{P;Z}(\mathcal{M}')$ such that $M' <_{P,Z} M$. Since each $N \in M$ is a $P; Z$-lub of $\mathcal{M}_N$ w.r.t. T, $M' \in ub^T_{P;Z}(\mathcal{M})$. However, this contradicts $M \in mub^T_{P;Z}(\mathcal{M})$. It follows that M is minimal, i.e., $M \in mub^T_{P;Z}(\mathcal{M}')$.

On the other hand, every $M' \in ub^T_{P;Z}(\mathcal{M}')$ must satisfy $M \leqslant_{P;Z} M'$. Therefore, M is a $P; Z$-lub of $\mathcal{M}'$ w.r.t. T. Since $\mathcal{M}' \subseteq mmod_{P;Z}(T)$, by definition $M \notin \mathcal{B}$. This is a (global) contradiction, which proves the result. □

Like LUB theories, also Weak LUB theories do not require iteration on taking mub's.

Corollary 5.8. *Suppose that a theory T satisfies the Weak P; Z-LUB property. Then* $cmod_{P;Z}(T) = 1 - cmod_{P;Z}(T)$.

The converse of Theorem 5.7 (if a theory has the Weak LUB property and every non-minimal curb model is the lub of some set of minimal models, then the set of curb models of T is well-founded) is not true. This is shown by Example 5.2, in which T has even the LUB property. Furthermore, Theorem 5.7 fails in general if the set of curb models is not well-founded. This is shown by the following example.

Example 5.3. Reconsider Example 5.2, and add to $\mathcal{A}$ two fresh atoms a and b. Define the theory T' by

$$T' = T \cup \{\neg(a \wedge b),\ a \vee b \vee (p_i \leftrightarrow \neg n_i),\ \neg n_i \rightarrow (\neg a \wedge \neg b) | i \geqslant 0\}.$$

Then T' has the models N, $N \cup \{a\}$, $N \cup \{b\}$, N_i, $M_i \cup \{a\}$, and $M_i \cup \{b\}$, for all $i \geqslant 0$, where N, N_i and M_i are as in Example 5.2. All models of T' except $N \cup \{a\}$ and $N \cup \{b\}$ are curb models under global minimization, and each of them is the lub of some set of curb models. Indeed, N and all N_i, $i \geqslant 0$ are minimal models; each $M_i \cup \{a\}$, $i \geqslant 0$, is the lub of $\{N_i, M_{i+1} \cup \{a\}\}$, while $M_i \cup \{b\}$ is the lub of $\{N_i, M_{i+1} \cup \{b\}\}$. Thus, each model of T is a 1-curb model, even under 2-bounded mubs. However, none of the models $M_i \cup \{a\}$ and $M_i \cup \{b\}$ is the lub of a set of minimal models.

A suggestive attempt to strengthen the Weak LUB property is to use ordinals. We say that the curb models of a theory T have the *inductive Weak LUB property*, if every non-minimal model M of T is the $P; Z$-lub of a family of models $\mathcal{M} \subseteq cmod_{P;Z}(T)$ such that $rank^T_{P;Z}(N) < rank^T_{P;Z}(M)$, for every $N \in \mathcal{M}$. Notice that the set of curb models in Example 5.2 has the inductive Weak LUB property (which, as a consequence, does not imply well-foundedness). However, the following result is an easy consequence of our results from above.

Corollary 5.9. *Let T be theory such that $cmod_{P;Z}(T)$ is well-founded. Then, T has the inductive Weak LUB property if and only if T has the Weak LUB property.*

Thus, inductive constructibility is implicit in presence of a well-founded set of curb models.

6. Computation and complexity

In this section, we consider computational issues for curbing applied to finite theories. We first present a simple algorithm for computing the curb models and then address the computational complexity of curbing.

Algorithm CURB_MODELS in Table 1 computes the set of all curb models of a given theory T with respect to a minimization setting $(P; Z)$. They are computed bottom up, exploiting Proposition 3.2. In a naive implementation, the models $M \in mod(T)$ that have cardinality i on P can be enumerated on line (3) by cycling through all models M of this cardinality and testing $M \models T$; more efficient generation is possible using an algorithm which generates all models of T of size i relative to P; note that arbitrary models may be constrained to this size by using auxiliary atoms and clauses "counting" the model size on P, and thus a general algorithm for generating all models of a theory may be employed. For an efficient implementation of the test on line (4), on line (5) links from M to the models in

Table 1
Algorithm for computing the curb models of a theory

Algorithm. CURB_MODELS(T, P, Z) : *set_of_models;*
Input: Finite propositional theory T on a finite set of atoms $\mathcal{A}$,
sets $P, Z \subseteq \mathcal{A}$ such that $P \cap Z = \emptyset$.
Output: Set $cmod_{P;Z}(T)$ of all $P; Z$-curb models of T.
$\mathcal{M} := \emptyset$;
for $i := 0$ **to** $|P|$ **do**
  for each $M \in mod(T)$ **such that** $|M \cap P| = i$ do
    if $|\max_{P;Z}\{M' \in \mathcal{M} | M' <_{P;Z} M\}| \neq 1$ **then**
      $\mathcal{M} := \mathcal{M} \cup \{M\}$
return $\mathcal{M}$;

$\max_{P;Z}\{M' \in \mathcal{M}|M' <_{P;Z} M\}$ (which coincides with $\max_{P;Z}(cmod^{<M}_{P;Z}(T))$) may be installed; then, proceeding downwards from $i = |M \cap P| - 1$, these links may be employed to recursively eliminate with any model $N <_{P;Z} M$ also all models $N' \leqslant_{P;Z} N$ from the search space.

It is easy to see that even a naive implementation of the algorithm in Table 1 is feasible to run within space polynomial in the size of T and $\mathcal{M}$. Therefore, we note the following result.

Proposition 6.1. *Algorithm* CURB_MODELS$(T; P; Z)$ *correctly computes the set of models* $cmod_{P;Z}(T)$ *in space polynomial in the size of* T, $\mathcal{A}$, *and* $|cmod_{P;Z}(T)|$.

Proof. The proof by induction on $j \geqslant 0$ that for each $M \in mod(T)$ such that $|M \cap P| = j$, it holds that $M \in cmod_{P;Z}(T)$ iff $M \in \mathcal{M}$ after executing the loop for $i = j$, is straightforward from Proposition 3.2 and the facts that, in the finite case, for each $M' \in cmod^{<M}_{P;Z}(T)$ it holds that $|M' \cap P| < |M \cap P|$, and that for every M and $\mathcal{M}'$ such that $M \neq_{P;Z} M'$ for every $M' \in \mathcal{M}'$, it holds that $M \in mub^{T}_{P;Z}(\mathcal{M}')$ iff $M \in mub^{T}_{P;Z}(\max_{P;Z}\{M' \in \mathcal{M}'|M' <_{P;Z} M\})$. □

We now turn to the complexity of curbing. We consider the following decision problems, which are major issues for reasoning procedures:

- *Curb Model Checking*: Given a finite theory T on a finite set $\mathcal{A}$ of atoms, a model $M \subseteq \mathcal{A}$, and disjoint sets $P, Z \subseteq \mathcal{A}$, decide whether $M \in cmod_{P;Z}(T)$, i.e., whether M is a $P; Z$-curb model of T.
- *Curb Inference*: Given a finite theory T on a finite set $\mathcal{A}$ of atoms, disjoint sets $P, Z \subseteq \mathcal{A}$, and a Boolean formula F, decide whether $cmod_{P,Z}(T) \models F$, i.e., $M \models F$ for each $M \in cmod_{P;Z}(T)$.

We recall that model checking for circumscription is coNP complete [7] and inferencing under circumscription is Π^P_2-complete [16]. The class Π^P_2 is the class co-Σ^P_2, where $\Sigma^P_2 = \mathrm{NP}^{\mathrm{NP}}$ contains all decision problems solvable in polynomial time by a non-deterministic Turing machine with an NP oracle. It holds that NP $\cup$ coNP $\subseteq \Sigma^P_2 \subseteq$ PH $\subseteq$ PSPACE, where PH $= \bigcup_{k \geqslant 0} \Sigma^P_k$ is the Polynomial Hierarchy, and each of the inclusions is widely believed to be strict. For a background on complexity theory, we refer to [39].

The two problems Curb Model Checking and Curb Inference can be easily solved using Algorithm CURB_MODELS, by first generating the curb models and then deciding the problem at hand. However, even with some obvious optimizations (compute not all models but stop as soon as the problem can be decided) this method is not efficient, since it uses exponential space in the size of the problem input in general. Also algorithms based on minimal model reasoning, such as in terms of reduction to circumscription as discussed in Section 6.3, require exponential space under standard complexity hypotheses (cf. Theorem 6.7 and the subsequent discussion). A more careful analysis of the problems in Section 6.1 reveals that they are actually solving in polynomial space.

In this section, we shall prove that Curb Model Checking and Curb Inference are PSPACE-complete. Note that it was conjectured in [45,27] that curbing has higher complexity than circumscription. This is intuitively supported by a result of Bodenstorfer [5] stating that in an explicitly given set of models, witnessing that some particular model is a curb model may involve an exponential number of smaller curb models. That is, any "proof" that a model has the curb property, given by a proper family of curb models, may have non-polynomial size in general. On the other hand, this

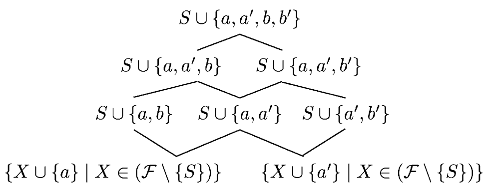

Fig. 8. Cloning a family $\mathcal{F}$ with unique maximal model S.

"proof" can be recursively generated in polynomial space, and thus Curb Model Checking is feasible in polynomial space.

Despite the comparatively high complexity of curbing in the general case, we shall single out several cases in which the complexity is lower, including cases in which the complexity is in NP, respectively, coNP, or even tractable. In the course of this, we improve previous complexity results in [45].

6.1. PSPACE membership

In order to make the intuition more precise that the curb property of a model might refer to a large number of smaller curb models, let us introduce the following concept.

Definition 6.1. Given a set of models $\mathcal{M} \subseteq 2^{\mathcal{A}}$ and a minimization setting $(P; Z)$, a $P; Z$-*support* of a model $M \in \mathcal{M}$ with respect to $\mathcal{M}$ is any subset $\mathcal{S} \subseteq \mathcal{M}$ such that (i) $M \in \mathcal{S}$ and (ii) for every $N \in \mathcal{S}$, $N \in mub^{\mathcal{M}}_{P;Z}(\mathcal{S}')$ for some $\mathcal{S}' \subseteq \mathcal{S} \setminus \{N\}$.

Note that every global minimal model M of a theory T ($\mathcal{M} = mod(T)$) has the support $\{M\}$ w.r.t. $\mathcal{M}$, and that all models in any support are curb models of T. Furthermore, every curb model of T has some support w.r.t. T.

Bodenstorfer [5] has defined a family $\mathcal{F}_n$, $n \geqslant 0$, of sets of models on an set of $O(n)$ atoms, such that $\mathcal{F}_n$ contains exponentially many models (in n), and $\mathcal{F}_n$ itself is the only support of the unique maximal model M_n of $\mathcal{F}_n$. Informally, $\mathcal{F}_0 = \{\{a_0\}\}$, and the family $\mathcal{F}_n$ is constructed inductively by cloning $\mathcal{F}_{n-1}$ and adding some sets which ensure that the maximal model needs all models for a proof of being a curb model (see Fig. 8). It may thus seem that Curb Model Checking and Curb Inference require exponential space. A straightforward algorithm for Curb Model Checking is a variant of Algorithm CURB_MODELS which tries to generate a support for the model M bottom up starting from the smallest models. This is clearly exponential in both time and space.

However, Algorithm CURB_CHECK in Table 2 shows that Curb Model Checking is feasible in polynomial space.

Theorem 6.2. *Curb Model Checking can be solved in quadratic space, more precisely, in space* $O(|\mathcal{A}|(|M|+|T|))$.

Proof. The correctness of Algorithm CURB_CHECK follows from Theorem 4.1, which implies that we need only to consider mubs of *pairs* of $P; Z$-curb models. Concerning the running time, CURB_CHECK is straightforward to implement such that its body uses only space $O(|M|+|\mathcal{A}|)$, which is linear in the size of the input. Furthermore, it is easily shown by induction that the recursion depth is bounded by $|P|$. Hence, the algorithm runs in space $O(|P|(|M| + |\mathcal{A}|))$. $\square$

Algorithm INFER in Table 3 then solves Curb Inference, and we obtain the following corollary to Theorem 6.2.

Corollary 6.3. *Curb Inference is in* PSPACE*(in particular, it is feasible in space* $O(|P|(|M|+|\mathcal{A}|)+|F|)$*).*

6.2. PSPACE hardness

It turns out that the PSPACE upper bound for Curb Model Checking and Curb Inference is in fact tight, since as we show in this section the problem is also PSPACE-hard. Since the proof is technically involved, we give here a sketch of the main idea.

Table 2
Algorithm for testing the curb model property

Algorithm. CURB_CHECK(M, T, P, Z) : *Boolean;*
Input: Truth assignment M, finite propositional theory T on a finite set of atoms $\mathcal{A}$, sets $P, Z \subseteq \mathcal{A}$ such that $P \cap Z = \emptyset$.
Output: "*true*" iff $M \in cmod_{P;Z}(T)$.
if $M \not\models T$ **then return** *false*;
minimal := *true*;
for each $M_1 <_{P;Z} M$ **do**
  if $M_1 \models T$ **then** *minimal* := *false*;
if *minimal* **then return** *true*;
for each models M_1, M_2 such that $M_1, M_2 <_{P;Z} M$,
  $M_1 \not\leqslant_{P;Z} M_2$, and $M_2 \not\leqslant_{P;Z} M_1$ **do**
  if CURB_CHECK$(M_1, T, P, Z) \wedge$ CURB_CHECK(M_2, T, P, Z) **then**
  begin *mub* := *true*;
    for each M_3 such that $M_1, M_2 <_{P;Z} M_3 <_{P;Z} M$ **do**
      if CURB_CHECK(M_3, T, P, Z) **then** *mub* := *false*;
    if *mub* **then return** *true*;
  end;
return *false*;

Table 3
Algorithm for curb Inference

Algorithm. INFER(T, F, P, Z) : *Boolean;*
Input: Finite propositional theory T and propositional formula F on finite set of atoms $\mathcal{A}$, sets $P, Z \subseteq \mathcal{A}$ such that $P \cap Z = \emptyset$.
Output: "*true*" iff $cmod_{P,Z}(T) \models F$.
for each $M \subseteq \mathcal{A}$ **do**
  if CURB_CHECK$(M, T, P, Z) \wedge (M \not\models F)$ **then return** *false*;
return *true*;

We take Bodenstorfer's construction as a starting point, since it gives us a method to construct instances where curb models provably have supports of non-polynomial size—any PSPACE-hard instances of Curbing Model Checking must have this property, since otherwise the instance can be solved by a guess and check algorithm with complexity in Σ_2^P. Extending this construction, we show then how the canonical PSPACE-complete problem of evaluating a given QBF $F = Q_n a_n Q_{n-1} a_{n-1} \cdots Q_1 a_1 E$ can be reduced to Curb Model Checking in polynomial time. Roughly, Bodenstorfer's construction is extended in a way such that the possible value assignments to an atom a_i are modeled by cloning, and the evaluation of the quantifier Q_i is encoded by further auxiliary atoms. The unique largest model of the theory constructed will then be a curb model precisely if F evaluates to true. A slight extension of the construction shows PSPACE-hardness of Curb Inference.

6.2.1. *Describing the exponential support family $\mathcal{F}_n$*

We describe Bodenstorfer's family $\mathcal{F}_n$ by a Boolean formula φ_n, such that $\mathcal{F}_n = mod(\varphi_n)$. The letters we use are $\mathcal{A}_n = \{a_i, a_i', b_i, b_i' \mid 1 \leqslant i \leqslant n\} \cup \{a_0\}$. We define the formula φ_n inductively, where we set $\varphi_0 = a_0$, $M_0 = \{a_0\}$, $\mu_0 = a_0$ and for $n > 0$:

$$\varphi_n = (\mu_{n-1} \wedge \gamma_n) \vee (\neg\mu_{n-1} \wedge \varphi_{n-1} \wedge (a_n \leftrightarrow \neg a_n') \wedge \neg b_n \wedge \neg b_n'),$$

where

$$\gamma_n = (a_n \wedge b_n \wedge \neg a_n' \wedge \neg b_n') \vee (a_n \wedge a_n' \wedge \neg b_n \wedge \neg b_n') \vee (a_n' \wedge b_n' \wedge \neg a_n \wedge \neg b_n) \vee (a_n \wedge b_n \wedge a_n' \wedge \neg b_n')$$
$$\vee (a_n' \wedge b_n' \wedge a_n \wedge \neg b_n) \vee (a_n \wedge b_n \wedge a_n' \wedge b_n'),$$

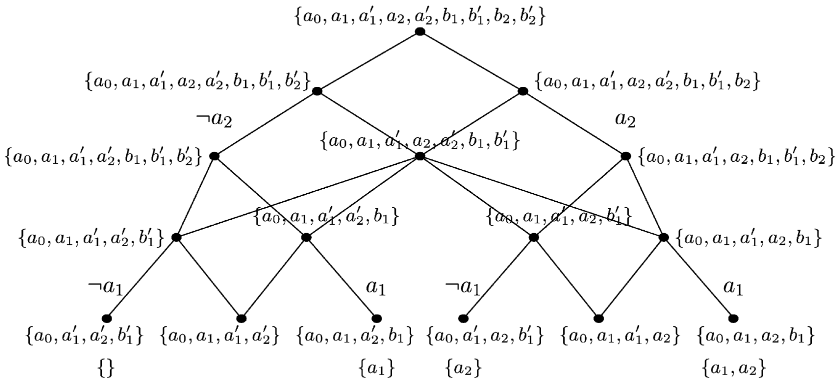

Fig. 9. The set of models $mod(\varphi_2)$.

$$M_n = M_{n-1} \cup \{a_n, a_n', b_n, b_n'\};$$

$$\mu_n = \bigwedge_{p \in M_n} p.$$

Note that the left disjunct of φ_n gives rise to six models, which extend M_{n-1} by the following sets of atoms: $A_{n,1} = \{a_n, b_n\}, A_{n,0} = \{a_n', b_n'\}, B_n = \{a_n, a_n'\}, C_{n,1} = \{a_n, a_n', b_n\}, C_{n,0} = \{a_n, a_n', b_n'\}$ and $D_n = \{a_n, a_n', b_n, b_n'\}$.

Informally, $A_{n,1}$ (resp., $A_{n,0}$) represents the assignment of 1 (resp., 0) to a_n. The right disjunct of φ_n generates recursively assignments to the other atoms $a_{n-1}, \ldots, a_1$, such that certain minimal models of φ_n under global minimization represent assignments to the atoms $a_1, \ldots, a_n$ (see Fig. 9). Note that $M_n = M_{n-1} \cup D_n$ (i.e., all atoms are 1) is, as easily seen, the unique maximal model of the formula φ_n. The set $mod(\varphi_n)$ of models of φ_n over $\mathcal{A}_n$, defines the family $\mathcal{F}_n$ as described in [5]. Thus, each model $M \in mod(\varphi_n)$ is a curb model under global minimization, and each support of M_n must have exponential size.

6.2.2. *Evaluating a QBF on mod*(φ_n)

We now show that a QBF

$$F = Q_n a_n Q_{n-1} a_{n-1} \cdots Q_1 a_1 E,$$

where each $Q_i \in \{\forall, \exists\}$ and E is a Boolean formula on atoms $a_1, \ldots, a_n$, can be "evaluated" on the family $mod(\varphi_n)$ of curb models under global minimization.

Roughly, the idea is as follows: $mod(\varphi_n)$ can be layered into n overlapping layers of models, where each layer i contains the models recursively generated by the left disjunct of the formula φ_i. In each layer we have three levels of models. Neighbored layers i and $i-1$ overlap such that the bottom level of i is the top level of $i-1$ (see Fig. 10). The minimal models in $mod(\varphi_n)$ are the bottom models of layer 1, and can be considered as the top model of an artificial layer 0. Similarly, the maximal model M_n in $mod(\varphi_n)$ can be viewed as a bottom model of an artificial layer $n+1$.

In order to "evaluate" the QBF F, we will obtain a formula $\psi(F)$ from F by adding conjunctively a set of formulas $\Gamma(F)$ to φ_n. Thus, $\psi(F) = \varphi_n \wedge \Gamma(F)$. The formulas in Γ will be chosen such that the overall structure of the set of curb models of $\psi(F)$ does not differ from the one of the set of models of φ_n. In particular, each model M of φ_n will correspond to some curb model $f(M)$ of $\psi(F)$ which augments M by certain atoms that describe the truth status of subformulas of F.

By adjoining $\Gamma(F)$ to φ_n, we "adorn" the models in $mod(\varphi_n)$ with additional atoms which help us in evaluating the formula F along the layers. At a layer i in $mod(\varphi_n)$, we have fixed an assignment to the atoms $a_{i+1}, \ldots, a_n$ already, where a_j is 1 if a_j occurs in the model, and a_j is 0 if a_j' occurs in the model, for all $j \geqslant i+1$ (there are some ill-defined assignments in top elements of layer i, in which both a_{i+1} and a_{i+1}' occur; these assignments will be ignored). Then, at two sets at the bottom of the layer i which correspond to the possible extensions of the assignment to $a_{i+1}, \ldots, a_n$ by setting a_i either 1 (effected by the set $A_{i,1}$) or to 0 (by $A_{i,0}$), we "evaluate" the formula $Q_{i-1} a_{i-1} \cdots Q_1 a_1 E(a_i, a_{i+1}, \ldots, a_n)$ where the atoms $a_i, \ldots, a_n$ are fixed to the assignment. If that formula evaluates

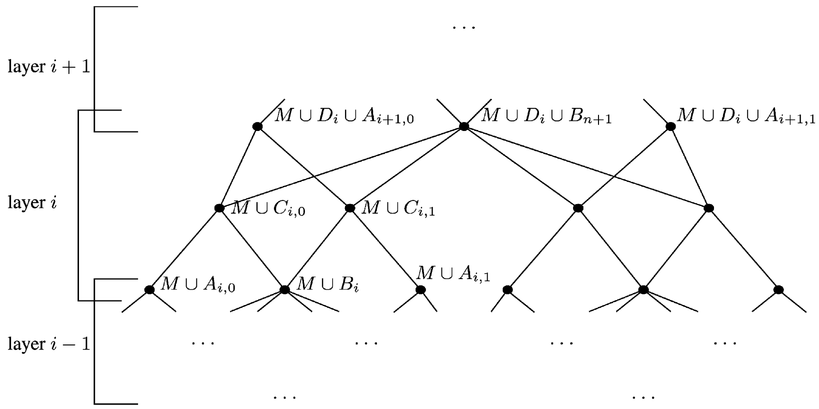

Fig. 10. Layers in $mod(\varphi_n)$.

to 1, then if a_i is 1, an atom v_i is included (resp., if a_i is 0 an atom v_i') at this bottom element. The quantifier Q_i is then evaluated by including in the top element "above" the two bottom sets an atom t_i if, in case of $Q_i = \exists$, either v_i or v_i' occurs in one of the two bottom elements, and in case of $Q_i = \forall$, v_i and v_i' occur in the bottom elements. The top element is itself a bottom element at the next layer $i+1$, and the atom t_i is used there to see whether the formula $Q_i a_i \cdots Q_1 a_1 E(a_{i+1}, \ldots, a_n)$ evaluates to 1.

In what follows, we formalize this intuition. We introduce a set of new atoms $\mathcal{A}_n' = \{v_i, v_i', t_i \mid 1 \leqslant i \leqslant n\} \cup \{t_0\}$.

The following formulas are convenient for our purpose:

$$\begin{aligned} ass_i &= a_i \leftrightarrow \neg a_i', \\ \lambda_i &= (\neg b_{i+1} \vee \neg b_{i+1}') \wedge (a_{i+1} \wedge a_{i+1}' \rightarrow \neg b_{i+1} \wedge \neg b_{i+1}'), \quad 1 \leqslant i \leqslant n; \\ \Lambda_1 &= \lambda_1, \\ \Lambda_j &= \lambda_j \wedge \neg\lambda_{j-1}, \quad 2 \leqslant j \leqslant n. \end{aligned}$$

Informally, ass_i tells whether the model considered assigns a_i legally a truth value. The formula λ_i says that the model is at layer i or below. The formula Λ_i says that the model is at layer i. The models at the bottom of layer i which are of interest to us are those in which ass_i is 1; all other models of the entire layer violate ass_i.

At layer $i \geqslant 1$, we evaluate the formula using the following formulas:

$$\Lambda_i \wedge ass_i \wedge t_{i-1} \wedge a_i \rightarrow v_i, \quad \Lambda_i \wedge ass_i \wedge t_{i-1} \wedge a_i' \rightarrow v_i'.$$

For $i = 1$, we add

$$E \rightarrow t_0,$$

which under curbing evaluates the quantifier-free part after assigning all atoms. Depending on the quantifier Q_i, we add a clause as follows:

- If $Q_i = \exists$, then we add

$$\Lambda_i \wedge (v_i \vee v_i') \rightarrow t_i,$$

- if $Q_i = \forall$, then we add

$$\Lambda_i \wedge v_i \wedge v_i' \rightarrow t_i.$$

For "garbage collection" of the new atoms used at lower layers, we use a formula $trap_i$ which adds all v_j, v_j', t_j' of lower layers to all elements of layer i which correspond to an illegal assignment to a_i:

$$trap_i = \Lambda_i \wedge \neg ass_i \rightarrow t_0 \wedge \bigwedge_{j=1}^{i-1} (v_j \wedge v_j' \wedge t_j).$$

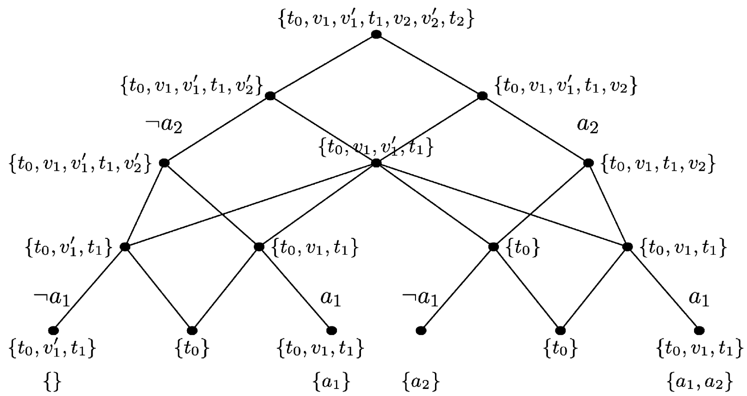

Fig. 11. Evaluating $F = \forall a_2 \exists a_1 (a_2 \to a_1)$: extending M to $f(M) = M \cup X$ (X shown).

Informally, models corresponding to different extensions of an assignment will always have a mub which is upper bounded by the bottom model at layer i which is an illegal assignment.

Let Γ_i be the conjunction of all formulas introduced for layer i, where $1 \leqslant i \leqslant n$, and let $\Gamma(F) = \bigwedge_{i=1}^{n} \Gamma_i$. Then we define

$$\psi(F) = \varphi_n \wedge \Gamma(F).$$

Note that $\psi(F)$ has a single maximal model M_F, given by $M_F = M_n \cup \{v_i, v_i', t_i \mid 1 \leqslant i \leqslant n\}$ (i.e., all atoms have value 1).

Let us call a model $M \in mod(\psi(F))$ an *assignment model*, if either $M[\mathcal{A}_n] = M_n$, or (b) $M \models \Lambda_i \wedge ass_i$, i.e., either M extends the maximal model of φ_n or M is at the bottom of layer i and assigns a_i a unique truth value. In case (a), we view M at the bottom of an artificial layer $n+1$. M represents a (partial) assignment σ_M to $a_i, \ldots, a_n$ defined by $\sigma_M(a_j) = 1$ if $a_j \in M$ and $\sigma_M(a_j) = 0$ if $a_j' \in M$, for all $j = i, \ldots, n$.

We establish the following lemma.

Lemma 6.4. *For each model $M \in mod(\varphi_n)$, there exists a curb model $f(M)$ of $\psi(F)$, such that*:

(1) $f(M)[\mathcal{A}_n] = M$ *(i.e., $f(M)$ coincides with M on the atoms of φ_n)*,
(2) *If M is an assignment model at layer $i \in \{1, \ldots, n+1\}$, then $f(M)$ contains t_{i-1} iff the formula*

$$F_i = Q_{i-1} a_{i-1} Q_{i-2} a_2 \cdots Q_1 a_1 E(a_1, \ldots, a_{i-1}, \sigma_M(a_i), \ldots, \sigma_M(a_n))$$

is true,
(3) *If M is at layer $i \in \{1, \ldots, n\}$ but not an assignment model*, *then*

$$f(M) = \begin{cases} M \cup \mathcal{A}_{i-1}' & \text{if } M = M_{n-1} \cup B_n, \\ f(M_{n-1} \cup A_{n,k}) \cup f(M_{n-1} \cup B_n) & \text{if } M = M_{n-1} \cup C_{n,k},\ k \in \{0, 1\}. \end{cases}$$

(4) *$f(M_n)$ is the single maximal curb model of $\psi(F)$, and if $Q_n = \forall$, then $t_n \in f(M_n)$ iff $f(M_n) = \mathcal{A}_n \cup \mathcal{A}_n'$.*

An example of the construction of $f(\cdot)$ for the formula $F = \forall a_2 \exists a_1 (a_2 \to a_1)$ is shown in Fig. 11.

Proof. We first note that each model M' of $\psi(F)$ is of the form $M \cup S$, where $M \in mod(\varphi_n)$ and $S \subseteq \mathcal{A}_n'$, and each $M \in mod(\varphi_n)$ gives rise to at least one such M' (just add $\mathcal{A}_n'$ to M).

We prove the lemma showing by induction on $n \geqslant 0$ how to construct such a correspondence $f(M)$.

The base case $n = 0$ (in which F contains no atoms and is either $\bot$ or $\top$) is easy: $mod(\varphi_0) = \{\{a_0\}\}$ and, if $F = \top$, then $mod(\psi(F)) = \{\{a_0, t_0\}\}$ and $f(\{a_0\}) = \{a_0, t_0\}$, and if $F = \bot$, then $mod(\psi(F)) = \{\{a_0\}, \{a_0, t_0\}\}$ and $f(\{a_0\}) = \{a_0\}$.

Consider the case $n > 1$ and suppose the statement holds for $n-1$. Let $M \in mod(\varphi_n)$. We consider two cases.

(1) $M \models \lambda_{n-1}$ *and* $M \not\models a_n \wedge a'_n$. Then, $M \models a_n \leftrightarrow \neg a'_n$, and either M is an assignment model at the bottom of layer n (in this case, M satisfies the left disjunct of φ_n) or some model not at layer n (in this case, M satisfies the right disjunct of M). In any case, $N = M \setminus \{a_n, a'_n, b_n, b'_n\}$ is a model of φ_{n-1}. By the induction hypothesis, it follows that for N we have a curb model $\hat{f}(N)$ of $\psi(F')$, where $F' = Q_{n-1}a_{n-1} \cdots Q_1 a_1 E'$ and $E' = E[a_n/\top]$ if $a_n \in M$ and $E' = E[a_n/\bot]$ if $a'_n \in M$ (i.e., $a_n \notin M$), such that $\hat{f}(N)$ fulfills the items in the lemma. We define $f(M)$ as follows. If $N \subset M_{n-1}$, then $f(M) := M \cup \hat{f}(N)$; otherwise, if $N = M_{n-1}$, then $f(M) = M \cup f(N) \cup S_M$, where

$$S_M = \begin{cases} \emptyset & \text{if } t_{i-1} \notin \hat{f}(N), \\ \{v_n, t_n\} & \text{if } t_{i-1} \in \hat{f}(N),\ Q_n = \exists \ \text{ and } \ a_i \in M, \\ \{v'_n, t_n\} & \text{if } t_{i-1} \in \hat{f}(N),\ Q_n = \exists \ \text{ and } \ a'_i \in M, \\ \{v_n\} & \text{if } t_{i-1} \in \hat{f}(N),\ Q_n = \forall \ \text{ and } \ a_i \in M, \\ \{v'_n\} & \text{if } t_{i-1} \in \hat{f}(N),\ Q_n = \forall \ \text{ and } \ a'_i \in M. \end{cases}$$

As easily checked, $f(M)$ is a model of $\psi(F)$. Furthermore, $f(M)$ is either a minimal model of $\psi(F)$ (if $n = 1$), or the mub of curb models $f(M_1)$ and $f(M_2)$ such that $M_1, M_2 \in mod(\varphi_{n-1})$, $M_1, M_2 \subset M$, and M is a mub of M_1, M_2 w.r.t. φ_{n-1}. (If not, then $\hat{f}(N)$ were not a mub of $\hat{f}(N_1)$, $\hat{f}(N_2)$ w.r.t. $\psi(F')$, which is a contradiction.) We can see that $f(M)$ fulfills items 1–3 in the lemma.

(2) $M \not\models \lambda_{n-1}$ or $M \models a_n a'_n$, i.e., M is at layer n but not an assignment model at its bottom. We consider the following possible cases for M:

(2.1) $M = M_{n-1} \cup B_n$: If $n = 1$, then M is a minimal model of φ_n, and $f(M) = M \cup \{t_0\}$ is a minimal model of $\psi(F)$, thus $f(M)$ is a curb model of $\psi(F)$; otherwise (i.e., $n > 2$), M is a mub of any arbitrary models $M_1, M_2 \in mod(\varphi_n)$ such that M_1 contains a_n and M_2 contains a'_n, respectively, and $M_i \setminus \{a_n, a'_n, b_n, b'_n\} \subset M_{n-1}$, for $i \in \{1, 2\}$. Since, by construction, $\hat{f}(M_i) \subseteq M_{n-1} \cup \mathcal{A}'_{n-1} =: f(M)$, this set is an upper bound of $f(M_1)$ and $f(M_2)$ w.r.t. $\psi(F)$; from formula $trap_{n-1}$ it follows that $f(M)$ is a mub of $f(M_1)$,$f(M_2)$. Thus, $f(M)$ is a curb model of $\psi(F)$.

(2.2) $M = M_{n-1} \cup C_{n,k}$, $k \in \{0, 1\}$: As easily checked, $f(M) = f(M_{n-1} \cup A_{n,k}) \cup f(M_{n-1} \cup B_n)$ ($=M_{n-1} \cup B_n \cup S_{M_{n-1} \cup A_{n,k}}$) is a model of $\psi(F)$. Since, as already shown, both $f(M_{n-1} \cup A_{n,k})$ and $f(M_{n-1} \cup B_n)$ are curb models of $\psi(F)$, clearly $f(M)$ is a mub of them and thus a curb model of $\psi(F)$.

(2.3) $M = M_n$: We define

$$f(M) = f(M_{n-1} \cup C_{n,0}) \cup f(M_{n-1} \cup C_{n,1}) \cup \begin{cases} \{t_n\} & \text{if } Q_n = \forall \ \text{ and } \ v_n, v'_n \in X; \\ \emptyset & \text{otherwise.} \end{cases}$$

Observe that $f(M) = M_n \cup \mathcal{A}'_{n-1} \cup X$, where $X \subseteq \{v_n, v'_n, t_n\}$. Then, as easily checked, $f(M)$ is a model of $\psi(F)$. Clearly, $f(M)$ is a mub of $f(M_{n-1} \cup C_{n,0})$ and $f(M_{n-1} \cup C_{n,1})$, and thus, $f(M)$ is a curb model of $\psi(F)$.

We now show that $f(M)$ in (2.1)–(2.3) satisfies items 1–3 in the lemma. Obviously, this is true for (2.1) and (2.2). For the case (2.3), from the definitions of $f(\cdot)$ in (1) and (2.1)–(2.2) it follows that $t_n \in f(M)$ if and only if $t_{n-1} \in f(M_{n-1} \cup A_{n,k})$ holds for some $k \in \{0, 1\}$ if $Q_n = \exists$ and for both $k \in \{0, 1\}$ if $Q_n = \forall$. By the induction hypothesis, $t_{n-1} \in f(M_{n-1} \cup A_{n,k})$ is true iff the QBF $Q_{n-1}a_{n-1} \ldots Q_1 a_1 E'$, where $E' = E[a_n/\top]$ if $k = 1$ and $E' = E[a_n/\bot]$ if $k = 0$, evaluates to 1. Thus, $t_n \in f(M)$ iff the QBF F evaluates to 1. Hence, $f(M)$ satisfies items 1–3 of the lemma.

As for item 4, furthermore, in the case where $Q_n = \forall$, we have by definition of $f(M)$ that $t_n \in f(M)$ iff $f(M) = M_n \cup \mathcal{A}'_{n-1} \cup \{v_n, v'_n, t_n\} = \mathcal{A}_n \cup \mathcal{A}'_n$.

Finally, it remains to show that $f(M_n)$ is the unique maximal curb model of $\psi(F)$. As easily seen, every finite Boolean theory which has a single maximal model has a single maximal curb model; thus $\psi(F)$ has a single maximal curb model M'. From the induction hypothesis, it follows that $M_k = f(M_{n-1} \cup A_{n,k})$ is the unique maximal curb model M'_k of $\psi(F)$ such that $M'[\mathcal{A}_n] \subseteq M_{n-1} \cup A_{n,k}$, for $k \in \{0, 1\}$. Since $M_2 = f(M_{n-1} B_n)$ is the unique maximal curb model N of $\psi(F)$ such that $N[\mathcal{A}_n] \subseteq M_{n-1} \cup B_n$, we conclude from the structure of layer n (where mubs are in fact lubs), that M' is a mub of M_0, M_1, M_2. Since, by construction, $f(M)$ is an upper bound of M_1, M_2, M_3, it follows $M' = f(M)$.

This proves that the claimed statement holds for n, and completes the induction. □

We thus obtain the following result.

Theorem 6.5. (1) *Given a finite theory T and a model M of T, deciding whether $M \in cmod(T)$ is* PSPACE-*hard.*
(2) *Given a finite theory T and an atom p, deciding whether $cmod(T) \models \neg p$ is* PSPACE-*hard.*

Proof. By items 2 and 4 in Lemma 6.4, $M = \mathcal{A}_n \cup \mathcal{A}'_n$ is a curb model of $\psi(F)$ for a QBF $F = \forall a_n Q_{n-1} a_{n-1} \cdots Q_1 a_1 E$ iff F evaluates to 1. Furthermore, F evaluates to 0 iff no curb model of $\psi(F)$ contains t_n, i.e., $cmod(\psi(F)) \models \neg t_n$. Deciding whether a given QBF of this form evaluates to 1 (resp. 0) is a well-known PSPACE-complete problem, and $\psi(F)$ is easily constructed in polynomial time from F. This proves the result. □

Combined with the PSPACE-membership results of Theorem 6.2 and Corollary 6.3, we arrive at the main result of this section.

Theorem 6.6. (1) *Curb Model Checking is* PSPACE-*complete, and the problem remains* PSPACE-*hard even under global minimization.*
(2) *Curb Inference is* PSPACE-*complete, and remains* PSPACE-*hard even for single literal inference under global minimization.*

This result contrasts with the already mentioned complexity of circumscription, for which model checking is coNP complete [7] and inference is Π_2^P complete [16]. There similarly Π_2^P-hardness holds for inference of a single negative literal under global minimization. We remark that in this setting, for both curb and circumscription inference of a positive literal L from a theory T is equivalent to classical consequence of L from T, and thus is a coNP-complete problem. This is an immediate consequence of Proposition 3.5. From Theorem 6.6, we can further conclude that we cannot represent the curb models of a theory T by the classical models or the minimal models of any other theory T', possibly under a different minimization setting, which is not exponentially larger. In fact, this holds for any representation of the curb models which has polynomial overhead. More precisely,

Theorem 6.7. *There exists no representation $R(cmod_{P;Z}(T))$ of $cmod_{P;Z}(T)$ of size polynomial in the size of T and $|\mathcal{A}|$ in any formalism which permits model checking within* PH, *i.e., given $R(cmod_{P;Z}(T))$ and an interpretation M, decide whether $M \in cmod_{P;Z}(T)$, unless* PSPACE = PH.

Proof. Indeed, if such a representation $R(cmod_{P;Z}(T))$ would exist, then Curb Model Checking would be in PH, since we could guess $R(cmod_{P;Z}(T))$ in polynomial time and verify the guess and that M is represented by $R(cmod_{P;Z}(T))$ using an oracle in PH as follows. By hypothesis, deciding whether M is represented by the guess R' is feasible with a Σ_k^P oracle, for some $k \geqslant 2$. To verify the guess R' for $R(cmod_{P;Z}(T))$, it is sufficient to check that the set of models represented by R' (i) includes only models of T, (ii) includes all $P; Z$-minimal models of T, (iii) is closed under 2-mubs w.r.t. T, and (iv) violates (ii), or (iii) if any model represented by R' is discarded. Each of (i), (ii), and (iii) is decidable with a Σ_{k+1}^P oracle, and hence condition (iv) is decidable with a Σ_{k+2}^P oracle. Therefore, deciding $M \in cmod_{P;Z}(T)$ is feasible in $\Sigma_{k+3}^P \subseteq$ PH. From Theorem 6.6, PSPACE = PH follows. □

Particular examples of knowledge representation formalisms which permit model checking in PH are classical logic, circumscription, and nested circumscriptive theories with nesting depth bounded by a constant [31,9]. A similar result can be shown for representing the curb inferences of a theory T in terms of the inferences from a "theory" T' in another formalism. Indeed, Curb Model Checking under global minimization can be easily expressed in terms of Curb Inference. Therefore, any theory T' whose theorems on $\mathcal{A}$ represent precisely the curb inferences of T under global minimization must in essence represent the curb models of T. Thus, informally curbing can be exponentially more concise than these formalisms, and in particular any reduction to circumscription as the one described in the next subsection, is expected to be exponential. For further issues concerning succinct representation, we refer the reader to [8,12] and references therein.

6.3. Bounded depth

We now consider the complexity of the variant of curbing in which mubs are only taken up to bounded depth. The following lemma about recognizing mubs is useful.

Lemma 6.8. *Given a theory T, models M and $M_1, \ldots, M_n$, $n \geqslant 0$, and a minimization setting $(P; Z)$, deciding whether $M \in mub^T_{P;Z}(\{M_1, \ldots, M_n\})$ is in* coNP.

Proof. Indeed, just test whether $M \in mod(T)$ and no $M' \in mod(T)$ exists such that $M'[Q] = M[Q]$, where $Q = \mathcal{A} \setminus (P \cup Z)$, and $(M[P] \setminus \bigcup_{i=1}^n M_i[P]) \subseteq M'[P] \subset M[P]$ holds. $\square$

Theorem 6.9. *The δ-bounded version of Curb Model Checking (resp., Curb Inference), i.e., deciding whether $M \in \delta\text{-}cmod_{P;Z}(T)$ (resp., $\delta\text{-}cmod_{P;Z}(T) \vDash F$), where δ is a finite constant, is Σ_2^P-complete (resp., Π_2^P-complete). Hardness holds even for global minimization and $\delta = 1$ (resp., $\delta = 0$).*

Proof. If $M \in \delta\text{-}cmod_{P;Z}(T)$ holds, then M has a polynomial size $P; Z$-support $\mathcal{S}_M$ w.r.t. T. Indeed, $M \in \delta\text{-}cmod_{P;Z}(T)$ implies that $M \in mub^T_{P;Z}(\mathcal{M})$ for some $\mathcal{M} = \{M_1, \ldots, M_k\} \subseteq (\delta{-}1)\text{-}cmod_{P;Z}(T)$ such that $k \leqslant |M[P]|$. By a recursive argument and the fact that δ is constant, it follows that M has some $P; Z$-support $\mathcal{S}_M$ w.r.t. T containing $O(|M[P]|^\delta)$ models.

To show $M \in \delta\text{-}cmod_{P;Z}(T)$, we can exploit Lemma 6.8 and guess a polynomial size $P; Z$-support $\mathcal{S}_M$ for M w.r.t. T such that chains under $\leqslant_{P;Z}$ in $\mathcal{S}_M$ have length bounded by $\delta + 1$, and verify its correctness in polynomial time with an NP oracle (check for every $M' \in \mathcal{S}_M$ that $M' \vDash T$ and that M' is a $P; Z$-mub of $\{M'' \in \mathcal{S}_M | M'' <_{P;Z} M'\}$ w.r.t. T; testing boundedness of chains is easy). Therefore, Curb Model Checking is in Σ_2^P.

The hardness under the claimed restriction is immediately obtained from the proof of Theorem 5 in [17]: there, in essence a Boolean formula φ on atoms $\mathcal{A} = X \cup X' \cup Y \cup \{w\}$ is constructed, where $X' = \{x_i' | x_i \in X\}$ is a copy of X, for which it is Σ_2^P-complete to decide whether φ has some global minimal model M such that $M \vDash w$. This φ has the properties that (1) for each $S \subseteq X$, φ has some global minimal model M such that $M[X] = S$, (2) each global minimal model M satisfies $M[X'] = \{x_i' | x_i \in X \setminus M\}$, (3) for each model $M \in mod(\varphi)$, $M \vDash w$ implies $M \vDash Y$, and (4) the maximal interpretation $M = \mathcal{A}$ is a model of T.

Therefore, $M = \mathcal{A}$ is a 1-curb model for $T = \{\varphi\}$ iff φ has some (globally) minimal model containing w; this proves Σ_2^P-hardness of Curb Model Checking, even if $\delta = 1$. Note that for $\delta = 0$, the problem amounts to model checking for propositional circumscription, which is coNP complete [7].

From the result on Curb Model Checking, it follows that deciding whether some $M \in \delta\text{-}cmod_{P;Z}(T)$ exists such that $M \nvDash F$ is in $\mathrm{NP}^{\mathrm{NP}} = \Sigma_2^P$. Hence, Curb Inference is in Π_2^P. The Π_2^P-hardness for $\delta = 0$, which coincides with circumscription, was shown in [16, Lemma 3.1], (cf. also [17, Theorem 5]). We remark that the construction used in the proof of [16, Lemma 3.1] can be easily generalized to show Π_2^P-hardness for arbitrary finite constant $\delta \geqslant 0$, e.g., by incorporating the structure of a finite version of the theory in Example 4.2. $\square$

Thus, the inference problem for curbing under a number of iterations bounded by a constant has the same degree of complexity as circumscription in general. In fact, from the proof of Theorem 6.9, we can extract an easy reduction of computing the δ-$\mathrm{curb}^k_{P;Z}$ models of a finite theory T to minimal model computation. Let, for $i = 1, \ldots, k$, be $P^{(i)}$ and $Z^{(i)}$ a copy of P, respectively, Z, and let

$$T_k' = T \cup \bigcup_{i=1}^k \{\varphi(P^{(i)}; Q; Z^{(i)}) | \varphi(P; Q; Z) \in T\} \cup \{p^{(i)} \to p | p \in P\}.$$

Then, for the minimization setting $(P'; Z')$ where $P' = P \cup P^{(1)} \cup \cdots \cup P^{(k)}$ and $Z' = Z \cup Z^{(1)} \cup \cdots \cup Z^{(k)}$, the following is easy to see.

Proposition 6.10. *For any theory T and minimization setting $(P; Z)$, $\{M[\mathcal{A}] | M \in mmod_{P';Z'}(T_k')\} = 1\text{-}curb^k_{P;Z}(T)$, i.e., the 1-depth k-width $P; Z$-curb models of T are the projection of the $P'; Z'$-minimal models of T_k' to $\mathcal{A}$.*

In particular, for $k = |P|$, this yields the set of all 1-depth curb models of T.

Example 6.1. Consider Example 2.2 again, and take $T = \{h \vee n, (h \wedge n) \to p\}$ and global minimization (i.e., $P = \{h, n, p\}$). Then, $|P| = 3$, and for $k = 3$, we have $P' = \{h, n, p\} \cup \{h^{(i)}, n^{(i)}, p^{(i)} | 1 \leqslant i \leqslant 3\}$, $Z' = \emptyset$, and

$$T_3' = T \cup \{h^{(i)} \vee n^{(i)}, (h^{(i)} \wedge n^{(i)}) \to p^{(i)} | 1 \leqslant i \leqslant 3\} \cup \{h^{(i)} \to h, n^{(i)} \to n, p^{(i)} \to p | 1 \leqslant i \leqslant 3\}.$$

The $P'; Z'$-minimal models of T'_1 (i.e., its global minimal models) are $M_1 = \{h, h^{(1)}, h^{(2)}, h^{(3)}\}$, $M_2 = \{n, n^{(1)}, n^{(2)}, n^{(3)}\}$, and all models $M_S = \{h, n, p\} \cup \{h^{(i)} | i \in S\} \cup \{n^{(i)} | i \in \{1, 2, 3\} \setminus S\}$, where $\emptyset \subset S \subset \{1, 2, 3\}$. Projected to $\mathcal{A} = \{h, n, p\}$, we obtain the three models $\{h\}$, $\{n\}$, and $\{h, n, p\}$. These are the 1-curb models of T.

By applying the construction T'_k iteratively to T, i.e., setting $T_{1,k} = T'_k$ and $T_{i+1,k} = T_{i,k}{}'_k$, we obtain the δ-curb$^k_{P;Z}$ models of T as the projection of the $(P_{\delta,k}; Z_{\delta,k})$-minimal models of the theory $T_{\delta,k}$ to $\mathcal{A}$. In particular, for $\delta = |P|$ and $k = 2$, we obtain all $P; Z$-curb models of T. The construction of $T_{\delta,k}$ is exponential in general (as implied by Theorem 6.7), but it is polynomial if δ is bounded by a constant. Thus, computing all curb models of a theory within bounded depth can be polynomially mapped to computing the minimal models of a propositional theory. For computing minimal models, algorithms such as [33,37,6,38,2–4] or recent tools based on mappings to disjunctive logic programming [20,24,44,26,47] might be exploited. Similarly, Curb Inference can be polynomially reduced to a circumscriptive theorem prover, like [22,23,15]. However, for larger depth and width this approach may lack efficiency in practice, since the number of auxiliary atoms increases quickly, and alternative algorithms might be more efficient. This remains for further investigation.

6.4. LUB and Weak LUB theories

From Theorems 5.7 and 6.9, we immediately get the following complexity results.

Theorem 6.11. *For theories T which have the Weak $P; Z$-LUB property, Curb Model Checking is in Σ_2^P and Curb Inference is in Π_2^P.*

Since the LUB Property implies the Weak LUB Property, we thus obtain:

Corollary 6.12. *For theories T which satisfy the $P; Z$-LUB property, Curb Model Checking is in Σ_2^P and Curb Inference is in Π_2^P.*

For $P; Z$-NegOrDual-Horn theories, we obtain by exploiting Theorems 5.5 and 5.4 the following result.

Theorem 6.13. Curb Model Checking *is* NP-*complete and* Curb Inference *is* coNP-*complete if the input theory T is $P; Z$-NegOrDual-Horn. The* NP-*hardness resp.* coNP-*hardness holds even under global minimization.*

Proof. The proof of Theorem 5.5 implies that M is a $P; Z$-curb model of T iff (1) $M \models T$ and (2) for each atom $p \in M[P]$, there exists some $M_p \in mmod_{P;Z}(T)$ such that $M_p \leqslant_{P;Z} M$ and $M_p \models p$. Since no clause in T which contains a positive literal on P contains a literal on Z, this is equivalent to (1′) $M \models T$ and, assuming this holds, (2′) for each atom $p \in M[P]$, there exists some $M_p \in mmod(T')$ such that $M_p \models p$, where T' on atoms P is obtained from T as follows: (i) remove all negative clauses and all clauses which contain some literal on Z or some literal L on Q such that $M \models L$, (ii) remove all literals on Q from the remaining clauses, and (iii) add $\neg p$ for each $p \in P \setminus M[P]$. Informally, these formulas effect that every atom $p \in P$ such that $M(p) = 0$ is fixed to $\bot$. As easily seen, $mod(T')$ corresponds 1-1 with the models $\{M'[P] | M' \in mod(T), M' \leqslant_{P;Z} M\}$.

The theory T' is dual Horn, i.e., each clause contains at most one negative literal. Hence, the satisfiability problem for T' and minimal model checking, i.e., testing $M \in mmod(T')$, is possible in polynomial time, cf. [7]. Consequently, for every $p \in M[p]$ some $M_p \in mmod(T')$ such that $M_p \models p$ can be guessed and verified in polynomial time. It follows that deciding $M \in cmod_{P;Z}(T)$ is in NP.

The NP-hardness follows from the construction in the proof of coNP-hardness of circumscriptive literal inference from a dual Horn theory in [10, Theorem 11]. There, for any CNF φ on atoms $X = \{x_1, \ldots, x_n\}$ a dual Horn theory T on atoms $\mathcal{A} = \{x_i, y_i, z_i | 1 \leqslant i \leqslant n\}$ is constructed such that T has some (globally) minimal model containing x_n iff φ is satisfiable. These minimal models correspond 1-1 to the models of φ, such that $M \in mod(\varphi)$ corresponds to the minimal model $M \cup \{y_i | 1 \leqslant i \leqslant n, M \not\models x_i\} \cup \{z_1, \ldots, z_n\}$ of T. Without loss of generality, φ has the property that $M \models \varphi$ implies that $X \setminus M \models \varphi$. Since T has by Theorems 5.5 and 5.4 the LUB property, $\mathcal{A} \in cmod(T)$ iff $\mathcal{A} \in$ 1-$cmod(T)$, i.e., the mub of some set $\mathcal{M} \subseteq mmod(T)$. Since by the proof of Theorem 5.5, $mod(T)$ is closed under unions, by the property of φ equivalently some $M \in mmod(T)$ exists such that $M \models x_n$. This proves NP-hardness, even under global minimization.

The NP membership of Curb Model Checking implies that Curb Inference is in coNP. The coNP-hardness, even under global minimization, follows from Theorem 6.14. □

In particular, we can apply this result to quadratic theories T such that no clause in T contains a positive literal on P and a literal on Z; this improves the respective result in [45]. In fact, for these theories, we have the following result.

Theorem 6.14. *For quadratic theories T in which no clause has a positive literal on P and a literal on Z, Curb Model Checking is polynomial and Curb Inference is* coNP-*complete. Hardness for* coNP *holds even under global minimization.*

Proof. For quadratic T of this form, the theory T' in the proof of Theorem 6.13 is quadratic. The guess for $M_p \in mmod(T')$ such that $M_p \models p$ can be eliminated, since as a consequence of [10, Theorem 20], deciding whether such an M_p exists is feasible in polynomial time. Hence, Curb Model Checking is polynomial.

The coNP-membership of Curb Inference for quadratic T of this form follows from Theorem 6.13. The coNP-hardness is shown by a reduction from the unsatisfiability problem. Given a CNF $\varphi = \bigwedge_{i=1}^{m} \gamma_i$ on atoms $X = \{x_1, \ldots, x_n\}$, construct a theory $T = \{x_i \vee x'_i, \neg x_i \vee \neg x'_i, \ |1 \leqslant i \leqslant n\} \cup \{\neg.L \vee y_j | L \in \gamma_j, 1 \leqslant j \leqslant m\}$, where $x'_1, \ldots, x'_n$ and $y_1, \ldots, y_m$ are fresh atoms and $\neg.L$ is the opposite of literal L (viewing γ_j as set of literals). The minimal models of T naturally correspond 1-1 to the assignments for X, where for $M \in mmod(T)$ we have $M(y_j) = 0$ iff $M[X] \models \neg\gamma_j$. Thus, $mmod(T) \models F$ for $F = \neg y_1 \vee \cdots \vee \neg y_m$ iff φ is unsatisfiable. Since $ub^T(\mathcal{M}) = \emptyset$ for each $\mathcal{M} \subseteq mmod(T)$ such that $|\mathcal{M}| \geqslant 2$, it follows that $cmod(T) \models F$ iff φ is unsatisfiable. This proves the result. □

Notice that by the virtue of Theorem 5.4 and Corollary 5.8, by Proposition 6.10 we have a polynomial time transformation of computing the $P; Z$-curb models of a LUB or Weak LUB theory T into computing the projection of the $P'; Z'$-minimal models of a theory T', and furthermore a polynomial time transformation of Curb Inference into circumscriptive inference. For example, since the hammer–nail-painting theory in Example 2.2 has the LUB-property, in Example 6.1 all its curb models are obtained from the $P'; Z'$-minimal models of T'_3. In particular, since the theory T'_k is quadratic whenever T has this property, circumscriptive inference algorithms for quadratic theories (cf. [10]) might be employed for quadratic theories as in Theorem 6.14.

7. Related work

We have considered an approach to relax the minimal model approach of circumscription [29,30], by adding mub's of models in the set of models of a theory. At the syntactic level, this amounts to admitting, but not enforcing, inclusive interpretation of positive disjunctions in certain situations. In this way, conclusions about negative information as drawn under circumscription may be prevented.

The issue of inclusive interpretation of disjunction has been investigated in the areas of deductive databases and logic programming by a number of different authors, cf. [11,13,40–43,49,50]. There, formulas are often restricted to clauses of particular form which are interpreted as rules. The semantics of theories ("programs"), which are collections of rules over a first-order predicate language, is commonly defined in terms of Herbrand models and reduces to the propositional case studied in this paper.

A general observation on the works quoted above, and on many others in the fields, is that the semantics of theories depends (as well-known), on the syntactical form of representation. Two theories T and T' that are equivalent in classical logic might have different semantics, even for theories as simple as a single clause. We emphasize that this is not an flaw, but a desired feature: the form of clauses is associated with a particular intuitive meaning which the semantics aim to capture. Therefore, also the notion of inclusive interpretation of disjunction (at the syntax level) should somehow depend on this meaning. Curbing, instead, adopts like circumscription a purely model-theoretic view (in terms of classical logic) and is thus independent of syntactical representation. Any two logically equivalent theories have the same curb models. Therefore, unsurprisingly curbing and the semantics for inclusive disjunction in [11,13,40–43] behave differently. Thus, if a particular meaning is associated with the syntax of statements, then the curbing approach presented here may like circumscription be inappropriate for respective knowledge bases. Both curbing and circumscription are geared towards reducing the set of models of a classical theory under some reasoning rationale, rather than capturing the semantics of a particular language. In the sequel, we briefly discuss the above formalisms and their relation to curbing.

The Disjunctive Database Rule (DDR) [41] has been proposed to allow cautious derivation of negative literals from a disjunctive deductive database, which is given by a set of formulas similar to the theories T_1 and T_2 in Examples 2.2 and 2.3. In spirit, DDR strives for a maximally inclusive interpretation of disjunction that is founded by rule application, by replacing disjunctive rules $a \to b \vee c$ by $a \to b$ and $a \to c$, and considering the unique minimal model of the resulting theory. Accordingly, in T_1 and T_2 DDR does not infer negative literals, and thus coincides with curbing in this case. However, DDR is basically different from curbing. It is not aimed at models in between the minimal models and the maximal model. Furthermore, DDR is syntax dependent. E.g., for $T = \{a \vee b,\ a\}$, DDR infers no negative information, while for the theory $T' = \{a,\ b \to b\}$ which is classically equivalent to T, it infers $\neg b$.

A more sophisticated approach that allows also to deal with negative clauses was introduced in [42,11] by the concepts of "Possible Models Semantics" (PMS) and "Possible Worlds Semantics" (PWS), which turned out to be equivalent. Like DDR, it replaces disjunctive rules $a \to b \vee c$ by non-disjunctive rules, but considers, in separate cases, the addition of any non-empty subset of the split rules $\{a \to b,\ a \to c\}$ and selects the minimal models. The emerging "possible models" coincide with the sustained models of a theory defined in [13], which uses a level mapping on the atoms to foster derivability from facts. The PMS has also been generalized to cover Negation by Failure [43]. As opposed to DDR, PMS, PWS, and sustained models are interested in possible positive information, and thus select also models in between minimal and maximal inclusive interpretation of disjunction. In Example 2.2, PMS coincides with curbing, i.e., the possible models (equivalently, sustained models) and the curb models are the same. However, if the clause $h \vee n \vee p$ is added to T_1, which is subsumed by the clause $h \vee n$ and thus has no effect on the models, PMS selects all models of T_1, while curbing selects the same models as before. For Example 2.3, we have a similar picture. Hence, PWS, PMS, and sustained models are all syntax-dependent and basically semantically different from our method. Also computationally, PMS and PWS are different from curbing since they reduce computational complexity, while curbing increases it (for more information on the complexity of PMS and PWS, see [17]).

Another approach to treat disjunction inclusively is the extension of the well-founded semantics [46] to the weak well-founded semantics (WWF) for disjunctive logic programs in [40]. In case of negation-free programs, this semantics coincides with the DDR [40], which implies syntax-dependency of WWF. Hence, also WWF is quite different from curbing. For many further semantics of deductive databases, similar observations can be made.

In [49,50], the authors considered the issue of updates with disjunctive information, based in Winslett's Possible Model Approach (PMA) [48]. They argued that updates should, despite minimal change, also allow inclusive interpretation of disjunction, and developed two update operators to this aim, called *Minimal Change with Exception* (*MCE*) and *Minimal Change with Maximal Inclusive Disjunctive* (*MCD*), respectively. Although the formulation is different, they are related to curbing with iteration depth bounded by $\delta = 1$; we refer to [27,28,50] for more details.

In the circumscription literature, several variants of circumscription have been introduced. A form of nested circumscription has been introduced in [31,9], which informally allows to build hierarchical theories whose minimal models are defined in an inductive manner. In the propositional case, the formalism has like curbing PSPACE-complexity [9], but its expressiveness appears to be different; we refer to [9] for a more detailed comparison.

8. Conclusion

In this paper, we have considered an approach to reasoning from propositional theories which is more lenient to inclusive interpretation of disjunctive information than minimal model reasoning as fostered by circumscription [29,34,35,30], which has important applications in Artificial Intelligence. The curbing approach "softens" circumscription by adding minimal upper bounds of models to the accepted models, and thus inhibits (sometimes) unintuitive conclusions while keeping the minimization principle. Different from related approaches in deductive databases and non-monotonic logic programming, where syntax of theories plays a prominent role, curbing operates like circumscription at the semantical level of classical logic. It is thus a mathematical tool for reasoning under minimal models together with minimal upper bounds of models, which may be exploited to reduce the set of models of classical theories under some reasoning rationale.

Furthermore, we have presented restricted notions of curbing, by bounding the number of iterative disjunction steps and the number of models participating in a disjunction, respectively, and we studied structural properties of the set of models which effect bounded iteration. Finally, we studied the computational complexity of curbing, showing that the main reasoning tasks are PSPACE-complete in general but have lower complexity in restricted cases.

In the present paper, we have focused on reasoning under minimal upper bounds in propositional logic. In the seminal paper [19], curbing has been considered in the setting of predicate logic, where a theory consists of a finite set of first-order sentences (tantamount, a single first-order sentence φ). In this setting, curbing is naturally formalized by a sentence of third-order predicate logic, given that the definition of the set of curb models of a first-order sentence φ involves sets of sets of models. However, in [19] it was also shown that curbing can be formalized in second-order logic (assuming the standard ZFC framework, i.e., Zermelo-Fraenkel Set Theory with the Axiom of Choice). More precisely, Eiter et al. [19] show how to construct, given a first-order sentence $\varphi(P; Z)$ and lists P and Z of predicate constants, where the predicates in P are minimized while those in Z are floating and all other predicates are fixed, a sentence $Curb(\varphi(P, Z); P, Z)$ of second-order predicate logic, such that its models are precisely the $P; Z$-curb models of $\varphi(P, Z)$. Most of the semantic results discussed in the present paper can be reformulated for the first-order case, while the treatment gets more complicated; we refer to [19] for more details.

Several issues remain for future work. On the semantical side, one issue is studying the structure of the set of curb models more in detail, and to single out interesting conditions under which curbing can be simplified. The results on the collapse of curbing and circumscription, as well as the LUB and Weak LUB properties are steps in this direction. Another issue is a generalization of curbing which allows to combine minimization of certain predicates while taking minimal upper bounds on others.

Another issue is to investigate the relationship of curbing to other non-monotonic reasoning formalisms in more detail. Here, extending the results on the relationship between curbing and variants of circumscription such as nested circumscription [9,31], which are also independent of syntax, would be interesting. Furthermore, exploring the relationship between curbing and syntax-based approaches which foster inclusive interpretation of disjunction, e.g. [11,13,40–43] more deeply remains for further work. Here, it would be interesting to determine meaningful classes of theories on which curbing and such approaches behave equivalently, or to provide embeddings with possibly extended alphabets.

On the computational side, a natural issue is a comprehensive study of the computational properties of curbing, in which similar as for circumscription [7,10] the computational complexity of important syntactic classes of Boolean theories is determined. Closely related to this are efficient algorithms and implementation. In particular, an efficient implementation of curbing on top of QBF solvers would be interesting, as well as the efficient usage of circumscription algorithms.

Acknowledgments

We thank our colleagues and readers of preliminary versions of this paper, as well as the referees for their useful comments. We are especially grateful to Hendrik Decker for extensive comments on a preliminary version of this paper, and Chiaki Sakama for clarifying remarks. The first author gives a special thanks to the West End Hotel in Edinburgh, UK, for its stimulating atmosphere in which some of the results in this paper emerged. This work has been partially supported by the European Commission through the Network of Excellence on Computational Logic (IST-2001-33570 COLOGNET) and by the Austrian Science Fund (FWF) under project Z29-N04.

References

[1] K Apt, Logic programming, in: J. van Leeuwen (Ed.), Handbook of Theoretical Computer Science, Vol. B, Elsevier, North-Holland, Amsterdam, 1990, pp. 493–574, (Chapter 10).

[2] C. Avin, R. Ben-Eliyahu-Zohary, Algorithms for computing X-minimal models, in: T. Eiter, W. Faber, M. Truszczyński (Eds.), Proc. Sixth Internat. Conf. on Logic Programming and Nonmonotonic Reasoning (LPNMR-01), Lecture Notes in Computer Science, Vol. 2173, Springer, Berlin, 2001, pp. 322–335.

[3] R. Ben-Eliyahu, L. Palopoli, Reasoning with minimal models: efficient algorithms and applications, Artif. Intell. 96 (2) (1997) 421–449.

[4] R. Ben-Eliyahu-Zohary, An incremental algorithm for generating all minimal models, Artif. Intell. 169 (1) (2005) 1–22.

[5] B. Bodenstorfer, How many minimal upper bounds of minimal upper bounds, Computing 56 (1996) 171–178.

[6] F. Bry, A.H. Yahya, Minimal model generation with positive unit hyper-resolution tableaux, in: P. Miglioli, U. Moscato, D. Mundici, M. Ornaghi (Eds.), Proc. Fifth Internat. Workshop on Theorem Proving with Analytic Tableaux and Related Methods (TABLEAUX '96), Lecture Notes in Computer Science, Vol. 1071, Springer, Berlin, 1996, pp. 143–159.

[7] M. Cadoli, The complexity of model checking for circumscriptive formulae, Inform. Process. Lett. 44 (1992) 113–118.

[8] M. Cadoli, F. Donini, P. Liberatore, M. Schaerf, Space efficiency of propositional knowledge representation formalisms, J. Artif. Intell. Res. 13 (2000) 1–31.
[9] M. Cadoli, T. Eiter, G. Gottlob, Complexity of nested circumscription and nested abnormality theories, ACM Trans. Comput. Logic 6 (2) (2005) 232–272.
[10] M. Cadoli, M. Lenzerini, The complexity of propositional closed world reasoning and circumscription, J. Comput. System Sci. 43 (1994) 165–211.
[11] E. Chan, A possible worlds semantics for disjunctive databases, IEEE Trans. Knowledge Data Eng. 5 (2) (1993) 282–292.
[12] A. Darwiche, P. Marquis, A knowledge compilation map, J. Artif. Intell. Res. 17 (2002) 229–264.
[13] H. Decker, J.C. Casamayor, Sustained models and sustained answers in first-order databases, in: M. Alpuente et al. (Eds.), Proc. 1994 Joint Conf. on Declarative Programming (GULP-PRODE'94), Valencia, Spain, Vol. 2, Servicio de Publicaciones, Univ. Politéc. Valencia, 1994, pp. 32–46, Preliminary version (invited paper) in Proc. DAISD 1993, pp. 267–286.
[14] J. Dix, Semantics of logic programs: their intuitions and formal properties, An overview, in: A. Fuhrmann, H. Rott (Eds.), Logic, Action and Information. Proceedings of the Konstanz Colloquium in Logic and Information (LogIn'92), DeGruyter, 1995, pp. 241–329.
[15] P. Doherty, W. Lukaszewicz, A. Szalas, Computing circumscription revisited: a reduction algorithm, J. Automated Reasoning 18 (3) (1997) 297–336.
[16] T. Eiter, G. Gottlob, Propositional circumscription and extended closed world reasoning are Π_2^P-complete, Theoret. Comput. Sci. 114 (2) (1993) 231–245 (Addendum 118:315).
[17] T. Eiter, G. Gottlob, On the computational cost of disjunctive logic programming: propositional case, Ann. Math. Artif. Intell. 15 (3/4) (1995) 289–323.
[18] T. Eiter, G. Gottlob, On the complexity of theory curbing, in: M. Parigot, A. Voronkov (Eds.), Proc. Seventh Internat. Conf. on Logic for Programming and Automated Reasoning (LPAR 2000), Reunion Island, France, Lecture Notes in Computer Science, Vol. 1955, Springer, Berlin, 2000, pp. 1–19.
[19] T. Eiter, G. Gottlob, Y. Gurevich, Curb your theory! A circumscriptive approach for inclusive interpretation of disjunctive information, in: R. Bajcsy (Ed.), Proc. 13th Internat. Joint Conf. on Artificial Intelligence (IJCAI-93), Morgan Kaufman, Los Altos, CA, 1993, pp. 634–639.
[20] CIRC2DLP software for translating parallel circumscription to disjunctive logic programming, 2005. ⟨www.tcs.hut.fi/Software/circ2dlp/⟩.
[21] D.W. Etherington, Reasoning with Incomplete Information, Morgan Kaufmann, Los Altos, CA, 1988.
[22] M.L. Ginsberg, A circumscriptive theorem prover, Artif. Intell. 39 (2) (1989) 209–230.
[23] K. Inoue, N. Helft, On theorem provers for circumscription, in: P. Patel-Schneider (Ed.), Proc. Eighth Biennial Conf. of the Canadian Society for Computational Studies of Intelligence (CSCSI 1990), Morgan Kaufman, CA, 1990, pp. 212–219.
[24] T. Janhunen, E. Oikarinen, Capturing parallel circumscription with disjunctive logic programs, in: J.J. Alferes, J. Leite (Eds.), Proc. Ninth European Conf. on Artificial Intelligence (JELIA 2004), Lecture Notes in Computer Science/Lecture Notes in Artificial Intelligence, Vol. 3229, Springer, Berlin, 2004, pp. 134–146.
[25] D. Le Berre, M. Narizzano, L. Simon, A. Tacchella, The second QBF solvers comparative evaluation, Available at ⟨http://www.qbflib.org/⟩, 2004.
[26] J. Lee, F. Lin, Loop formulas for circumscription, in: Proc. 19th National Conf. on Artificial Intelligence, Sixteenth Conf. on Innovative Applications of Artificial Intelligence, July 25–29, 2004, San Jose, CA, USA, AAAI Press/The MIT Press, Cambridge, MA, 2004, pp. 281–286.
[27] P. Liberatore, The complexity of iterated belief revision, in: F. Afrati, P. Kolaitis (Eds.), Proc. Sixth Internat. Conf. on Database Theory (ICDT-97), Lecture Notes in Computer Science, Vol. 1186, January 1997, pp. 276–290.
[28] P. Liberatore, The complexity of belief update, Artif. Intell. 119 (1–2) (2000) 141–190.
[29] V. Lifschitz, Computing circumscription, in: A.K. Joshi (Ed.), Proc. Ninth Internat. Joint Conf. on Artificial Intelligence (IJCAI-85), Morgan Kaufmann, Los Altos CA, 1985, pp. 121–127.
[30] V. Lifschitz, Circumscription, in: D. Gabbay, C. Hogger, J. Robinson (Eds.), Handbook of Logic in Artificial Intelligence and Logic Programming, Vol. III, Clarendon Press, Oxford, 1994, pp. 297–352.
[31] V. Lifschitz, Nested abnormality theories, Artif. Intell. 74 (2) (1995) 351–365.
[32] J. Lobo, J. Minker, A. Rajasekar, Foundations of Disjunctive Logic Programming, MIT Press, Cambridge, MA, 1992.
[33] R. Manthey, F. Bry, Satchmo: a theorem prover implemented in prolog, in: E.L. Lusk, R.A. Overbeek (Eds.), Proc. Ninth Internat. Conf. on Automated Deduction, (CADE 1988), Lecture Notes in Computer Science, Vol. 310, Springer, Berlin, 1988, pp. 415–434.
[34] J. McCarthy, Circumscription—A form of non-monotonic reasoning, Artif. Intell. 13 (1980) 27–39.
[35] J. McCarthy, Applications of circumscription to formalizing common-sense knowledge, Artif. Intell. 28 (1986) 89–116.
[36] J. Minker, D. Seipel, Disjunctive logic programming: a survey and assessment, in: A. Kakas, F. Sadri (Eds.), Computational Logic: from Logic Programming into the Future, Lecture Notes in Computer Science/Lecture Notes in Artificial Intelligence, Vol. 2407, Springer, Berlin, 2002, pp. 472–511, Festschrift in honour of Bob Kowalski.
[37] A. Nerode, R. Ng, V. Subrahmanian, Computing circumscriptive databases. I: theory and algorithms, Inform. Comput. 116 (1995) 58–80.
[38] I. Niemelä, A tableau calculus for minimal model reasoning, in: P. Miglioli, U. Moscato, D. Mundici, M. Ornaghi (Eds.), Proc. Fifth Internat. Workshop on Theorem Proving with Analytic Tableaux and Related Methods (TABLEAUX '96), Lecture Notes in Computer Science, Vol. 1071, Springer, Berlin, 1996, pp. 278–294.
[39] C.H. Papadimitriou, Computational Complexity, Addison-Wesley, Reading, MA, 1994.
[40] K. Ross, The well-founded semantics for disjunctive logic programs, in: W. Kim, J.-M. Nicholas, S. Nishio (Eds.), Proc. First Internat. Conf. on Deductive and Object-Oriented Databases (DOOD-89), Elsevier, North-Holland, Amsterdam, 1990, pp. 352–369.
[41] K. Ross, R. Topor, Inferring negative information from disjunctive databases, J. Automated Reasoning 4 (2) (1988) 397–424.
[42] C. Sakama, Possible model semantics for disjunctive databases, in: W. Kim, J.-M. Nicholas, S. Nishio (Eds.), Proc. First Internat. Conf. on Deductive and Object-Oriented Databases (DOOD-89), Elsevier, North-Holland, Amsterdam, 1990, pp. 337–351.

[43] C. Sakama, K. Inoue, Negation in disjunctive logic programs, in: D.S. Warren (Ed.), Proc. Tenth Internat. Conf. on Logic Programming (ICLP-93), Budapest, Hungary, MIT-Press, Cambridge, MA, June 1993, pp. 703–719.
[44] C. Sakama, K. Inoue, Embedding circumscriptive theories in general disjunctive programs, in: W. Marek, A. Nerode, M. Truszczyński (Eds.), Proc. Third Internat. Conf. on Logic Programming and Nonmonotonic Reasoning (LPNMR-95), Lecture Notes in Computer Science, Vol. 982, LNAI, Springer, Berlin, 1995, pp. 344–357.
[45] F. Scarcello, N. Leone, L. Palopoli, Curbing theories: fixpoint semantics and complexity issues, in: M. Alpuente, M.I. Sessa (Eds.), Proc. 1995 Joint Conf. on Declarative Programming (GULP-PRODE'95), Marina di Vietri, Italy, 11–14 September 1995, Palladio Press, 1995, pp. 545–554.
[46] A. vanGelder, K. Ross, J. Schlipf, The well-founded semantics for general logic programs, J. ACM 38 (3) (1991) 620–650.
[47] T. Wakaki, K. Inoue, Compiling prioritized circumscription into answer set programming, in: B. Demoen, V. Lifschitz (Eds.), Proc. 20th Internat. Conf. on Logic Programming (ICLP'04), Lecture Notes in Computer Science, Vol. 3132, Springer, Berlin, 2004, pp. 356–370.
[48] M. Winslett, Reasoning about action using a possible models approach, in: Proc. Seventh National Conf. on Artificial Intelligence (AAAI-88), AAAI Press/The MIT Press, Cambridge, MA, 1988, pp. 89–93.
[49] Y. Zhang, N.Y. Foo, Updating knowledge bases with disjunctive information, in: Proc. 13th National Conf. on Artificial Intelligence and Eighth Innovative Applications of Artificial Intelligence Conference (AAAI/IAAI), Vol. 1, AAAI Press/The MIT Press, Cambridge, MA, 1996, pp. 562–568.
[50] Y. Zhang, N.Y. Foo, Updates with disjunctive information: from syntactical and semantical perspectives, Comput. Intell. 16 (1) (2000) 29–52.