

Translating between the representations of an acyclic convex geometry of bounded degree

Oscar Defrain^a Arthur Ohana^{a,b} Simon Vilmin^a

Submitted: Dec 3, 2025; Accepted: July 24, 2026; Published: Oct 23, 2026

© The authors. Released under the CC BY-ND license (International 4.0).

Abstract

We consider the problem of translating between irreducible closed sets and implicational bases in closure systems. To date, the complexity status of this problem is widely open, and it is further known to generalize the notorious hypergraph dualization problem, even in the case of acyclic convex geometries, i.e., closure systems admitting an acyclic implicational base. This paper studies this case with a focus on the degree, which corresponds to the maximal number of implications in which an element occurs. We show that the problem is tractable for bounded values of this parameter, even when relaxed to the notions of premise- and conclusion-degree. Our algorithms rely on structural properties of acyclic convex geometries and involve various techniques from algorithmic enumeration such as solution graph traversal, saturation techniques, and a sequential approach leveraging from acyclicity. They are shown to perform in incremental-polynomial time. Finally, we complete these results by showing that our running times cannot be improved to polynomial delay using the standard framework of flashlight search.

Mathematics Subject Classifications: 68W40, 06A15

1 Introduction

Finite closure systems appear in disguise in several areas of mathematics and computer science, including database theory [42, Chapter 13], Horn logic [18, Chapter 6], or lattice theory [19, Chapter 2], to mention but a few.

A closure system is a family of subsets of some finite ground set being closed by intersection and containing the ground set. The study of the various representations of closure systems and the algorithmic task of translating between them have received lots of attention in these respective fields, and have been the topic of the Dagstuhl Seminar 14201 [1]. The translation task we study in this paper is most importantly known for being

^aAix-Marseille Université, CNRS, LIS, Marseille, France (oscar.defrain@lis-lab.fr, arthur.ohana@lis-lab.fr, simon.vilmin@lis-lab.fr)

^bUniversité Grenoble Alpes, Grenoble, France.

a considerable generalization of the notorious and widely open hypergraph dualization problem, while also not being known to be intractable, contrary to analogous other natural generalizations such as lattice dualization [7, 20, 36]. We redirect the reader to the recent works [6, 8, 21, 23, 47] as well as to the surveys [10, 48, 56] for more information on the topic.

Among the several representations of closure systems, two of them occupy a central role, namely, *irreducible closed sets* and *implicational bases*, whose formal definitions are postponed to Section 2. Informally, irreducible closed sets are particular sets of the closure system, while implicational bases are sets of implications of the form $A \rightarrow B$, where A , the premise, and B , the conclusion, are subsets of the ground set. A closure system has a unique collection of irreducible closed sets, but it can be represented by several equivalent implicational bases. As noted in [38],¹ these two representations are generally incomparable, in size and usefulness. Notably, there exist instances where the minimum cardinality of an implicational base is exponential in the number of irreducible closed sets, and vice versa. Moreover, the computational complexity of some algorithmic tasks such as reasoning, abduction, or recognizing lattice properties can differ significantly depending on the representation at hand [6, 35, 42]. These observations motivate the problem of translating between the two representations, which has two sides: enumerate the irreducible closed sets from an implicational base, and produce a compact implicational base from the collection of irreducible closed sets. These problems are respectively denoted by ICS-ENUM and MIB-GEN throughout the paper and formally defined in the next section.

Because of the possible exponential gap between the sizes of the two representations, input-sensitive complexity is not a meaningful efficiency criterion when analyzing the performance of an algorithm solving ICS-ENUM or MIB-GEN. Instead, the framework of algorithmic enumeration in which an algorithm has to list without repetitions a set of solutions, is usually adopted. The complexity of an algorithm is then analyzed from the perspective of *output-sensitive complexity*, which estimates the running time of the algorithm using both input and output sizes. In that context, an algorithm is considered tractable when it runs in *output-polynomial* time, that is, if its execution time is bounded by a polynomial in the combined sizes of the input and the output. This notion can be further constrained to guarantee some regularity in the enumeration. Namely, we say that an algorithm runs in *incremental-polynomial time* if it outputs the i^{th} solution in a time which is bounded by a polynomial in the input size plus i . It is said to run with *polynomial delay* if the time spent before the first output, after the last output, and between consecutive outputs is bounded by a polynomial in the input size only. We refer the reader to [34, 51] for more details on enumeration complexity.

In this work, we investigate closure systems admitting acyclic implicational bases; see Section 2 for the definition. These closure systems are known as *acyclic convex geometries*, as they satisfy the anti-exchange property which characterizes convex geometries [56]. Acyclic convex geometries have been extensively studied [13, 21, 33, 49, 54]. Even in this

¹Let us point that in the work of Khardon is adopted the language of Horn representations, where an implicational base corresponds to a (definite) Horn formula and the irreducible closed sets correspond to its characteristic models, possibly ignoring the set of all variables.

class though, the two problems ICS·ENUM and MIB·GEN have been shown in [21] to be harder than distributive lattices dualization, a generalization of hypergraph dualization. On the other hand, the authors in [21] show that if, in addition to acyclicity, the implication graph admits a rank function, then the problem can be solved using hypergraph dualization. Using this result, they derive a tractable algorithm for so-called ranked convex geometries admitting an implicational base of bounded premise size. It should be noted that, even under these restrictions, the problem generalizes the one of enumerating all maximal independent sets of a graph (in fact of hypergraphs of bounded edge size) a central and notorious problem in enumeration [34, 52]. These results have then been extended to a more general setting in [47]. Let us finally note that, outside of acyclic convex geometries, the translation task is known to be tractable in modular or k -meet-semidistributive lattices [8, 55], or in closure spaces of bounded Caratheodory number [56].

In this paper, we continue the line of research of characterizing tractable cases of acyclic convex geometries for ICS·ENUM and MIB·GEN, typically, by identifying for which parameters k we can obtain tractable algorithms for these problems for fixed values of k . We note that in this research, finding an algorithm running in quasi-polynomial time should already be considered a success, since no sub-exponential time algorithms are known to date. We focus on the degree, which is defined as the maximum number of implications in which an element participates, and its two obvious relaxations which are the premise- and conclusion-degree; these parameters are defined in Section 2. The principal message of this work is that bounding one of these two relaxed notions of degree suffices to provide not only quasi-polynomial but polynomial-time bounds for the translation. More specifically, as the first of our two main theorems, we obtain the following, whose quantitative formulation is postponed to Sections 5 and 6.

Theorem 1. *There is an incremental-polynomial time algorithm enumerating the irreducible closed sets of an acyclic convex geometry given by an acyclic implicational base of bounded premise- or conclusion-degree.*

Our algorithms for acyclic convex geometries of bounded degree rely on several steps which combine structure and enumeration techniques. We give a rough description which highlights the techniques used, and refer to the appropriate sections for more details.

For generating irreducible closed sets from an acyclic implicational base, we use the fact that in a convex geometry, each irreducible closed set is *attached* to a unique element of the ground set; this is detailed in Section 2. Then, leveraging from acyclicity, our algorithm relies on a sequential procedure which consists of enumerating, incrementally and in a top-down fashion according to a topological ordering $x_1, \dots, x_n$ of the elements of the ground set, all the irreducible closed sets attached to x_i . This step can be seen as enumerating the irreducible closed sets attached to x_i with additional information as input, namely, the irreducible closed sets of ancestor elements. Finally, it remains to solve this latter task, which we can do either by brute forcing the minimal transversals of a carefully designed hypergraph in case of bounded conclusion-degree, or performing a solution graph traversal in case of bounded premise-degree. In the end, the resulting

algorithms run in incremental-polynomial time, where the dependence in the number of solutions arises from the sequential top-down procedure.

We then turn to the opposite translation task of generating a compact implicational base given the irreducible closed sets of an acyclic convex geometry. To do so, we rely on the concept of critical generators [33, 39, 54]. These are known to form an acyclic implicational base of minimum cardinality (in fact, minimum among those with conclusions of size one, or minimum once aggregated, a point that is discussed in Section 3) in the class of acyclic geometries. This implicational base, referred to as the critical base, thus constitutes a candidate of choice to be generated from irreducible closed sets. We obtain the following as our second main theorem.

Theorem 2. *There is a polynomial-time algorithm constructing the critical base of an acyclic convex geometry of bounded premise- or conclusion-degree given by its irreducible closed sets.*

This second algorithm also relies on the top-down approach except that, this time, we generate all the critical generators of x_i at each step. However, since the implicational base is not given, a particular care is devoted to arguing of the choice of the ordering $x_1, \dots, x_n$ and the fact that it can be computed from the irreducible closed sets only. Then, the enumeration of critical generators is conducted using a saturation algorithm relying on the existence of an algorithm for the sequential enumeration of irreducible closed sets we obtained in order to prove Theorem 1. We note that in this direction, our algorithm is much more efficient thanks to the fact that the output has polynomial cardinality.

In order to prove Theorems 1 and 2, we build upon a number of existing properties of acyclic convex geometries coming from different fields. This results in a substantial presentation of acyclic convex geometries which aims to be self-contained and accessible. Besides, we also extend this presentation with a study of the degree in this class.

The remainder of the paper is organized as follows. Section 2 introduces the concepts and notations used throughout the paper. Section 3 continues the state of the art via a focus on acyclic convex geometries, and include additional observations that may be of independent interest. Section 4 presents the top-down approach that we use as the key ingredient in our algorithms. Sections 5 and 6 are devoted to proving Theorem 1 in the respective cases of bounded conclusion- and premise-degree. Theorem 2 is proved in Section 7. Finally, we discuss possible generalizations, limitations, and open directions in Section 8.

2 Preliminaries

All the objects we consider in this paper are finite. If X is a set, $\mathcal{P}(X)$ is its powerset. The size of X is denoted $|X|$. A *set system*, or *hypergraph*, over a ground set X is a pair $(X, \mathcal{F})$ where $\mathcal{F}$ is a collection of subsets of X , that is $\mathcal{F} \subseteq \mathcal{P}(X)$. Its *size* is $|X| + |\mathcal{F}|$; note that the binary length of any standard encoding of a hypergraph is polynomially related to its size. Thus, for simplicity, we shall consider this measure for families as input size in the different problems we will consider. If X is clear from the context, we

may identify $\mathcal{F}$ with $(X, \mathcal{F})$. In examples, if there is no ambiguity, we may write a set as the concatenation of its elements, for instance 123 instead of $\{1, 2, 3\}$. Similarly, to avoid cumbersome notations, we sometimes withdraw brackets when dealing with singletons, e.g., we may write $\phi(x)$ instead of $\phi(\{x\})$.

Independent and hitting sets, hypergraph parameters. Let $\mathcal{H} = (X, \mathcal{E})$ be a hypergraph. We call *vertices* of $\mathcal{H}$ the elements in X , and *edges* the sets in $\mathcal{E}$. We say that $\mathcal{H}$ is *Sperner* if any two of its edges are pairwise incomparable by inclusion. An *independent set* of $\mathcal{H}$ is a set $I \subseteq X$ that contains no edges of $\mathcal{H}$, i.e., $E \not\subseteq I$ for every E in $\mathcal{E}$. The family of inclusion-wise maximal independent sets of $\mathcal{H}$ is denoted $\text{MIS}(\mathcal{H})$. Formally:

$$\text{MIS}(\mathcal{H}) := \max_{\subseteq} \{I : I \subseteq X, E \not\subseteq I \text{ for every } E \text{ in } \mathcal{E}\}.$$

Dually, a *hitting set*, or *transversal*, of $\mathcal{H}$ is a subset T of X that intersects each edge of $\mathcal{H}$. The collection of inclusion-wise minimal hitting sets of $\mathcal{H}$ is called $\text{MHS}(\mathcal{H})$:

$$\text{MHS}(\mathcal{H}) := \min_{\subseteq} \{T : T \subseteq X, T \cap E \neq \emptyset \text{ for every } E \text{ in } \mathcal{E}\}.$$

The hypergraph $\text{MHS}(\mathcal{H})$ is usually called the *dual hypergraph* of $\mathcal{H}$. Observe that $I \subseteq X$ is independent if and only if $T := X \setminus I$ is a hitting set. Consequently, $\text{MHS}(\mathcal{H}) = \{X \setminus I : I \in \text{MIS}(\mathcal{H})\}$. Moreover, if $\mathcal{H}$ is not Sperner, then its minimal hitting sets (resp. maximal independent sets) are precisely those of the hypergraph formed by the inclusion-wise minimal edges of $\mathcal{H}$. The two equivalent generation problems associated to maximal independent sets and minimal hitting sets read as follows:

Maximal Independent Sets Enumeration (MIS-ENUM)

Input: A (Sperner) hypergraph $\mathcal{H} = (X, \mathcal{E})$.

Output: The family $\text{MIS}(\mathcal{H})$.

Minimal Hitting Sets Enumeration (MHS-ENUM)

Input: A (Sperner) hypergraph $\mathcal{H} = (X, \mathcal{E})$.

Output: The family $\text{MHS}(\mathcal{H})$.

To date, the best algorithm for these tasks runs in output-quasi-polynomial time and is due to Fredman and Khachiyan [29]. However, a number of output-polynomial time algorithms have been proposed for particular cases, especially when bounding structural parameters [25]. Among these parameters, we will be interested in the *dimension*, which is the size of the largest edge, or the *dual dimension*, being the dimension of the dual hypergraph.

Closure systems, closure spaces. A *closure system* is a set system $(X, \mathcal{C})$ that satisfies $X \in \mathcal{C}$, and $C \cap C' \in \mathcal{C}$ whenever $C, C' \in \mathcal{C}$. The sets in $\mathcal{C}$ are called *closed (sets)*. A *closure space* is a pair (X, ϕ) where $\phi: \mathcal{P}(X) \rightarrow \mathcal{P}(X)$ is a mapping that satisfies for every $A, B \subseteq X$: (1) $A \subseteq \phi(A)$, (2) $\phi(A) \subseteq \phi(B)$ if $A \subseteq B$, and (3) $\phi(\phi(A)) = \phi(A)$. The map ϕ is a *closure operator*. Closure spaces and closure systems are in one-to-one correspondence. More precisely, if (X, ϕ) is a closure space, the family $\mathcal{C} := \{C : C \subseteq X, \phi(C) = C\}$

defines a closure system. Dually, if $(X, \mathcal{C})$ is a closure system, the map ϕ defined by $\phi(A) := \min_{\subseteq} \{C : C \in \mathcal{C}, A \subseteq C\} = \bigcap \{C : C \in \mathcal{C}, A \subseteq C\}$ is a closure operator whose associated closure system is precisely $(X, \mathcal{C})$. In this paper, if (X, ϕ) is a closure space, we call $(X, \mathcal{C})$ its *corresponding* closure system, and vice-versa. Moreover, we always consider without loss of generality that $\emptyset \in \mathcal{C}$, a standard assumption.

Let $(X, \mathcal{C})$ be a closure system. The poset $(\mathcal{C}, \subseteq)$ of closed sets ordered by inclusion is a (*closure*) *lattice* where $C \wedge C' = C \cap C'$ and $C \vee C' = \phi(C \cup C')$. A *predecessor* of a closed set C is an inclusion-wise maximal closed set F properly contained in C , i.e., $F \subset C$ and for every $C' \in \mathcal{C}$, $F \subseteq C' \subseteq C$ implies $C' = F$ or $C' = C$. *Successors* are defined dually. Let $M \in \mathcal{C}$, $M \neq X$. The closed set M is *irreducible* if it is not the intersection of other distinct closed sets, that is, $M = C \cap C'$ for some $C, C' \in \mathcal{C}$ only if $C = M$ or $C' = M$. Irreducible closed sets are in fact the members of $(\mathcal{C}, \subseteq)$ with a unique successor in the poset. We denote by $\text{irr}(\mathcal{C})$ the family of irreducible closed sets of $(X, \mathcal{C})$. For every $C \in \mathcal{C}$, it is known that $C = \bigcap \{M : M \in \text{irr}(\mathcal{C}), C \subseteq M\}$. Moreover, $\text{irr}(\mathcal{C})$ is the unique minimal subset of $\mathcal{C}$ from which $\mathcal{C}$ can be reconstructed using intersections. Let $x \in X$. We denote by $\text{irr}(x)$ the inclusion-wise maximal closed sets not containing x , i.e.,


$$\text{irr}(x) := \max_{\subseteq} \{M : M \in \mathcal{C}, x \notin M\}.$$

Observe that $\text{irr}(x) \subseteq \text{irr}(\mathcal{C})$ for every $x \in X$. An irreducible closed set M in $\text{irr}(x)$ is said to be *attached* to x . In fact, it is well known that $\text{irr}(\mathcal{C}) = \bigcup_{x \in X} \text{irr}(x)$. A (*non-trivial*) *minimal generator* of x is an inclusion-wise minimal subset A of $X \setminus \{x\}$ satisfying $x \in \phi(A)$. The family of minimal generators of x is denoted $\text{mingen}(x)$. Again, it is well-known [42] that $\text{irr}(x)$ and $\text{mingen}(x)$, seen as hypergraphs over X , are dual to one another, that is:

$$\begin{aligned} \text{irr}(x) &= \text{MIS}(\text{mingen}(x) \cup \{\{x\}\}), \\ \text{mingen}(x) &= \text{MHS}(\{X \setminus (M \cup \{x\}) : M \in \text{irr}(x)\}). \end{aligned}$$

An example of a closure system and its irreducible closed sets is depicted in Figure 1. Let $A \subseteq X$. An *extreme point* of A is an element $x \in A$ such that $x \notin \phi(A \setminus \{x\})$. We denote by $\text{ex}(A)$ the extreme points of A . In particular, if $C \in \mathcal{C}$ and $x \in \text{ex}(C)$, then $C \setminus \{x\} \in \mathcal{C}$. Note that for every $C, C' \in \mathcal{C}$ such that $C \subseteq C'$, the extreme points of C' that belong to C must also be extreme points of C , i.e., $\text{ex}(C') \cap C \subseteq \text{ex}(C)$. Indeed, $x \in \text{ex}(C')$ entails $C' \setminus \{x\} \in \mathcal{C}$ and hence $C \cap C' \setminus \{x\} = C \setminus \{x\} \in \mathcal{C}$, that is, $x \in \text{ex}(C)$. Dually, if $x \notin \text{ex}(C)$ for some $C \in \mathcal{C}$, then $x \notin \text{ex}(C')$ for every $C' \in \mathcal{C}$ such that $C \subseteq C'$. If $C \in \mathcal{C}$, a *spanning set* of C is a subset A of X satisfying $\phi(A) = C$.

Convex geometries. A closure system $(X, \mathcal{C})$ is a *convex geometry* if $\emptyset \in \mathcal{C}$ and for every closed set C distinct from X there exists $x \notin C$ such that $C \cup \{x\}$ is closed. The closure system depicted in Figure 1 is a convex geometry. Equivalently, $(X, \mathcal{C})$ is a convex geometry if the corresponding closure operator ϕ has the *anti-exchange* property: for every closed set C and each distinct $x, y \notin C$, $x \in \phi(C \cup \{y\})$ entails $y \notin \phi(C \cup \{x\})$. Convex geometries have been rediscovered in several fields and, as such, enjoy several characterizations [44]. We summarize some of them in the subsequent statement:



1. $(X, \mathcal{C})$ is a convex geometry
2. the associated closure operator ϕ has the anti-exchange property
3. each irreducible closed set is attached to a unique element, i.e., $\{\text{irr}(x) : x \in X\}$ is a partition of $\text{irr}(\mathcal{C})$
4. each closed set C has a unique minimal spanning set, which is $\text{ex}(C)$

Let $T \subseteq X$. We call *irreducible selection* for T a family of sets obtained by choosing, for each $t \in T$, one $M \in \text{irr}(t)$. Note that since we are dealing with convex geometries, and by Proposition 3, the closed sets M are all distinct and the resulting family has the same cardinal as T . In the following, we note $\mathcal{S}(T)$ the family of all irreducible selections for a given set T . We will use irreducible selections as a way to forbid elements, as stated in the following observation.

Observation 4. For any $T \subseteq X$ and any selection $\mathcal{S} \in \mathcal{S}(T)$, the intersection $\bigcap \mathcal{S}$ of all sets in $\mathcal{S}$ is a closed set disjoint from T .

Before moving to implicational bases, we introduce the notion of critical generators. These have first been defined in terms of circuits of antimatroids (complements of convex geometries) within the context of combinatorial optimization [22, 39, 40] but they also have been studied from the perspective of closure systems [21, 46, 54, 56, 57]. They will play a crucial role in this paper. Recall that $(X, \mathcal{C})$ is a convex geometry. Let $x \in X$ and let $A \in \text{mingen}(x)$. The minimal generator A is a *critical generator* of x if $\phi(A) \setminus \{a, x\} \in \mathcal{C}$ for every $a \in A$. For $x \in X$, we call $\text{critgen}(x)$ the set of critical generators of x . Using the properties of extreme points previously mentioned, we can rephrase the definition of critical generators in a way that will be more convenient to use in our approach:

Proposition 5. *The set A is a critical generator of x if and only if $A = \text{ex}(\phi(A))$ and $\phi(A) \in \min_{\subseteq} \{C : C \in \mathcal{C}, x \in C, x \notin \text{ex}(C)\}$.*

Proof. We start with the if part. By assumption, for every $a \in A$, $x \in \text{ex}(\phi(A) \setminus \{a\})$, i.e., $\phi(A) \setminus \{a, x\} \in \mathcal{C}$ for each $a \in A$. Hence, we only need to show that $A \in \text{mingen}(x)$. Observe first that $x \notin A$ by definition of A . Now, let $a \in A$. We have $\phi(A \setminus \{a\}) \subseteq \phi(A) \setminus \{a\} \in \mathcal{C}$. Moreover, since A is the unique minimal spanning set of $\phi(A)$ by assumption, any subset of A must be the unique minimal spanning set of its closure, so in particular $\text{ex}(\phi(A \setminus \{a\})) = A \setminus \{a\}$. Thus, $x \in \phi(A \setminus \{a\})$ would imply that $x \notin \text{ex}(\phi(A \setminus \{a\}))$, as $x \notin A$, which in turn yields $x \notin \text{ex}(\phi(A) \setminus \{a\})$, a contradiction to the minimality of $\phi(A)$. We deduce $x \notin \phi(A \setminus \{a\})$ and hence that $A \in \text{mingen}(x)$.

As for the only if part, A is a minimal generator of x , hence it must be the unique minimal spanning set of $\phi(A)$. Since $(X, \mathcal{C})$ is a convex geometry, $A = \phi(\text{ex}(A))$ holds by Proposition 3. Thus, the predecessors of $\phi(A)$ are precisely the sets $\phi(A) \setminus \{a\}$ for $a \in A$. Hence, for any $C \in \mathcal{C}$ such that $C \subset \phi(A)$, there is $a \in A$ such that $C \subseteq \phi(A) \setminus \{a\}$. As $x \in \text{ex}(\phi(A) \setminus \{a\})$, $x \in C$ entails $x \in \text{ex}(C)$. We conclude that $\phi(A) \in \min_{\subseteq} \{C : C \in \mathcal{C}, x \in C, x \notin \text{ex}(C)\}$. $\square$

Proposition 5 suggests an algorithm to compute a critical generator of an element x given a set that generates x .

Proposition 6. *Let $(X, \mathcal{C})$ be a convex geometry, $x \in X$ and $Y \subseteq X$ such that $x \notin Y$ and $x \in \phi(Y)$. A critical generator A of x such that $\phi(A) \subseteq \phi(Y)$ can be computed with a number of calls to ϕ being polynomial in $|X|$.*

Proof. Notice that $x \in \phi(Y)$ but $x \notin \text{ex}(\phi(Y))$ by definition. Hence, there exists $C \in \mathcal{C}$ such that $C \subseteq \phi(Y)$ and $C \in \min_{\subseteq} \{C' : C' \in \mathcal{C}, x \in C', x \notin \text{ex}(C')\}$. Moreover, for each such $C \subseteq \phi(Y)$ and each $C' \in \mathcal{C}$ such that $C \subseteq C' \subseteq \phi(Y)$, $x \notin \text{ex}(C')$ must hold.

We describe an algorithm to identify a critical generator A of x included in $\phi(Y)$. As a preprocessing step, we compute $\phi(Y)$. Then, given a closed set C , we compute its extreme points and its predecessors, that are precisely the sets $\{C \setminus \{y\} : y \in \text{ex}(C)\}$, and we identify the element $y \in \text{ex}(C)$ with the smallest index such that $x \notin \text{ex}(C \setminus \{y\})$. If no such y exists, we output $\text{ex}(C)$. Otherwise, we repeat the step with $C \setminus \{y\}$ as input. By Proposition 5, and since we check all the predecessors of the input closed set at each step, this procedure will correctly output a critical generator A of x included in $\phi(Y)$.

Computing $\phi(Y)$ requires one call to ϕ . For a closed set C , finding $\text{ex}(C)$ and checking whether $x \in \text{ex}(C \setminus \{y\})$ for each $y \in \text{ex}(C)$ requires at most $2 \cdot |X|$ calls to ϕ . Since at each step we remove one element from C , we will repeat these operations at most $|X|$ times. Hence, the whole algorithm requires at most $1 + 2 \cdot |X|^2$ calls to ϕ . $\square$

Implicational bases. An *implication* over X is a statement $A \rightarrow B$ where A, B are subsets of X . In $A \rightarrow B$, A is the *premise* and B the *conclusion*. If B is a singleton, that is $B = \{b\}$, the implication $A \rightarrow b$ is (*right*-)unit. An *implicational base* (IB) is a pair (X, Σ) where Σ is a collection of implications over X . The IB (X, Σ) is (*right*-)unit if it consists in unit implications. The *aggregation* of an IB (X, Σ) is the

IB (X, Σ^*) resulting from the unification of the implications with equal premises, i.e., $\Sigma^* := \{A \rightarrow B : B = \bigcup \{B' : A \rightarrow B' \in \Sigma\}\}$. Dually, the *unit-expansion* is obtained by replacing each implication $A \rightarrow B \in \Sigma$ by the set of implications $\{A \rightarrow b : b \in B\}$. Given an IB (X, Σ) , $|\Sigma|$ is the cardinal of Σ , that is the number of implications in Σ . The *size* of the IB is the value $|X| + |\Sigma|$. The *degree* of an element $x \in X$, denoted $\deg(x)$, is the number of implications of Σ in which x appears. We refine this notion by defining the *premise-degree* of x as $\text{pdeg}(x) := |\{A : A \rightarrow B \in \Sigma, x \in A\}|$, and its *conclusion-degree* as $\text{cdeg}(x) := |\{A \rightarrow B : A \rightarrow B \in \Sigma, x \in B\}|$. Then $\deg(x) \leq \text{pdeg}(x) + \text{cdeg}(x)$. Equality may not hold in general because an element can belong to both premise and the conclusion of an implication. In this paper though, we avoid this degenerate case by always considering the premise and conclusion of an implication to be disjoint, which can be ensured in polynomial-time without loss of generality with respect to the closure system. Moreover, we will mainly use the *degree*, *premise-degree*, and *conclusion-degree* of Σ defined by

$$\begin{aligned}\text{pdeg}(\Sigma) &:= \max_{x \in X} \text{pdeg}(x), \\ \text{cdeg}(\Sigma) &:= \max_{x \in X} \text{cdeg}(x), \\ \deg(\Sigma) &:= \max_{x \in X} \deg(x).\end{aligned}$$

Note that $\deg(\Sigma) \leq \text{pdeg}(\Sigma) + \text{cdeg}(\Sigma)$ for any Σ .

Example 7. Let $X = \{1, 2, 3, 4, 5\}$ and consider the IB (X, Σ) given by:

$$\Sigma = \{4 \rightarrow 1, 25 \rightarrow 3, 12 \rightarrow 3, 15 \rightarrow 23\}.$$

We have $\text{pdeg}(2) = 2$, $\text{cdeg}(2) = 1$, $\text{pdeg}(\Sigma) = \text{pdeg}(2) = \text{pdeg}(1) = 2$, and $\text{cdeg}(\Sigma) = \text{cdeg}(3) = 3$.

We now relate IBs to closure systems. Let (X, Σ) be an IB. A set $C \subseteq X$ is *closed* for (X, Σ) if for every implication $A \rightarrow B \in \Sigma$, $A \subseteq C$ implies $B \subseteq C$. The pair $(X, \mathcal{C})$ where $\mathcal{C}$ is the family of closed sets of (X, Σ) is a closure system. An IB (X, Σ) thus induces a closure operator $\phi(\cdot)$ that can be computed in polynomial time using the *forward chaining* algorithm. This procedure, which will be also used in some of our proofs, starts from a set Y and produces a sequence of sets $Y =: Y_0 \subseteq Y_1 \subseteq \dots \subseteq Y_k := \phi(Y)$ such that $Y_i := Y_{i-1} \cup \bigcup \{B : A \rightarrow B \in \Sigma, A \subseteq Y_{i-1}, B \not\subseteq Y_{i-1}\}$ for $1 \leq i \leq k$. On the other hand, any closure system $(X, \mathcal{C})$ can be represented by several *equivalent* IBs. An implication $A \rightarrow B$ is *valid* or *holds* in $(X, \mathcal{C})$ if for all $C \in \mathcal{C}$, $A \subseteq C$ implies $B \subseteq C$. Equivalently, $A \rightarrow B$ holds in $(X, \mathcal{C})$ if and only if $B \subseteq \phi(A)$. For instance, $\Sigma = \{A \rightarrow \phi(A) : A \subseteq X\}$ constitutes an IB for $(X, \mathcal{C})$ as well as $\Sigma' = \{C \cup \{x\} \rightarrow \phi(C \cup \{x\}) : C \in \mathcal{C}, x \notin C\}$ (recall that $\emptyset \in \mathcal{C}$ by assumption). Moreover an IB is always equivalent to its aggregation, and to its unit-expansion. For example, the IB of Example 7 is an IB of the closure system of Figure 1.

The fact that a closure system $(X, \mathcal{C})$ can be represented by several IBs has led to the study of a broad range of different IBs and different measures on IBs, as witnessed by two

surveys on the topic [10, 56]. Among possible IBs, the canonical base [32], the canonical direct base [11, 54] and the D -base [2] have attracted much attention. In this paper, we will be interested instead in some measures that allow to compare two equivalent IBs. Namely, we say that an IB (X, Σ) , or simply Σ , is *minimum* if for any IB (X, Σ') equivalent to (X, Σ) , $|\Sigma| \leq |\Sigma'|$. Similarly, if (X, Σ) is unit, we say that it is *unit-minimum* if it is cardinality-minimum among equivalent unit IBs. If (X, Σ) is an arbitrary IB, then a minimum IB (X, Σ') can be obtained in polynomial time, see e.g., [32, 41, 50, 54]. On the other hand, finding an equivalent unit-minimum IB is computationally hard [3]. Thus, in general, computing a unit-minimum IB is harder than computing a minimum one.

The *degree* of a closure system is the least integer k such that it admits an implicational bases of degree k . The premise- and conclusion-degree are defined analogously.

Translation task. To conclude the preliminaries, we define the two enumeration problems we study in this paper. These are the problems of translating between an IB and irreducible closed sets, two representations of a closure system. Formally, the problems read as follows:

Irreducible Closed Sets Enumeration (ICS·ENUM)

Input: An implicational base (X, Σ) .

Output: The irreducible closed sets of the associated closure system.

Unit-Minimum Implicational Base Generation (MIB·GEN)

Input: A ground set X , the irreducible closed sets $\text{irr}(\mathcal{C})$ of a closure system $(X, \mathcal{C})$.

Output: A unit-minimum implicational base (X, Σ) of $(X, \mathcal{C})$.

We stress the fact that in both problems, only an implicit (and possibly compact) description of the closure system is provided. Namely, for MIB·GEN, only X and $\text{irr}(\mathcal{C})$ are part of the input, not $\mathcal{C}$, and a unit-minimum implicational base is to be computed. Analogously for ICS·ENUM, only (X, Σ) is given, not $\mathcal{C}$. Moreover, we note that in the literature, variants of MIB·GEN have been considered where the goal is to compute a minimum (not unit-minimum) IB. However, and as discussed previously, our formulation is harder; see also [21]. It is also known that these problems are polynomially equivalent in general; see [38] for more details. Finally, we note that in both problems the output may be of exponential size with respect to the input size as suggested by Example 8 below. Hence, these problems are studied through the lens of output-sensitive complexity.

Example 8. Let $X = \{a_1, \dots, a_k\} \cup \{b_1, \dots, b_k\} \cup \{x\}$, and $\Sigma_1 = \{a_i b_i \rightarrow x : 1 \leq i \leq k\}$. Denoting $(X, \mathcal{C}_1)$ the associated closure system, we have $\text{irr}(\mathcal{C}_1) = \{X \setminus \{y\} : y \in X, y \neq x\} \cup \{\{c_1, \dots, c_k\} : c_i \in \{a_i, b_i\}, 1 \leq i \leq k\}$. Therefore, $|\text{irr}(\mathcal{C}_1)| = 2^k + 2k$ while $|\Sigma_1| = k$.

Dually, consider the closure system $(X, \mathcal{C}_2)$ with irreducible closed sets $\text{irr}(\mathcal{C}_2) = \{X \setminus \{y\} : y \in X, x \neq y\} \cup \{X \setminus \{a_i, b_i, x\} : 1 \leq i \leq k\}$. Then, a unit-minimum IB (X, Σ_2) of $(X, \mathcal{C}_2)$ is given by $\Sigma_2 = \{c_1 \cdots c_k \rightarrow x : c_i \in \{a_i, b_i\}, 1 \leq i \leq k\}$. We have $|\text{irr}(\mathcal{C}_2)| = 3k$ while $|\Sigma_2| = 2^k$. The two IBs are pictured in Figure 2.

As mentioned in the introduction, the two problems have been widely studied, but despite being harder than MIS·ENUM (hence MHS·ENUM), their complexity remains




Figure 2: The implicational bases of Example 8 with Σ_1 on the left and Σ_2 on the right. For clarity we only give some of the implications, drawn in grey. In both cases, an irreducible closed set attached to x is highlighted in blue.

unknown to date. In our journey to solve ICS·ENUM, we will use the related problem of generating the family of irreducible closed sets attached to a given element:

Attached Irreducible Closed Sets Enumeration (ACS·ENUM)

Input: An implicational base (X, Σ) and $x \in X$.

Output: The family $\text{irr}(x)$ of irreducible closed sets attached to x .

It is well known that an algorithm solving ACS·ENUM in output-polynomial time can be applied to each element x to produce all irreducible closed sets of a closure system in output-polynomial time. In this process, at most a linear number of repetitions are to be handled, and these repetitions do not occur in convex geometries by Proposition 3. Unfortunately, it was shown by Kavvadias et al. [36] in the language of Horn CNFs that the problem ACS·ENUM admits no output-polynomial time algorithm in general unless $P = NP$. Further results regarding MIB·GEN, ICS·ENUM and ACS·ENUM will be discussed in the next section, which is devoted to acyclic convex geometries, the class of closure systems at the core of our contribution.

3 Acyclic convex geometries

In this section we define and give the main properties of the class of convex geometries we study in this paper: *acyclic convex geometries*. They have been studied under various names such as poset-type convex geometries [49], G -geometries [54], or as acyclic Horn functions from the perspective of Horn CNFs [13, 33].

To define acyclic convex geometries, we will use IBs, thus following the line of propositional logic [33]. To do so, we need the concept of implication graph. Let (X, Σ) be an IB. For the time being, $(X, \mathcal{C})$ needs not be a convex geometry. The *implication graph* [13, 33] of (X, Σ) is the directed graph $G(\Sigma) = (X, E)$ where there is an arc ab in E if there exists $A \rightarrow B$ in Σ such that $a \in A$ and $b \in B$. The IB (X, Σ) is called *acyclic* if $G(\Sigma)$ is. In this case, $G(\Sigma)$ can be seen as a poset, and naturally induces an *ancestor* relation over X defined for all $x \in X$ as follows:

$$\text{anc}(x) := \{y : y \in X, x \neq y, \text{ and there exists a directed path from } y \text{ to } x \text{ in } G(\Sigma)\}.$$

Equivalently, $\text{anc}(x)$ is the set of in-neighbors of x in the transitive closure of $G(\Sigma)$. The descendant relation can be defined dually using out-neighbors. An element x of X is thus maximal if it has no ancestors, and minimal if it has no descendants. As an example, Figure 3 illustrates the implication graph of the IB given in Example 7.

We note that the notions of premise- and conclusion-degree differ from the in- and out-degree of the implication graph. Indeed, on the first hand, a single implication in Σ may lead to several arcs in $G(\Sigma)$. On the other hand, two implications with a same conclusion y and sharing a same element x in their premises yield a single arc (x, y) in $G(\Sigma)$.

A closure system that admits an acyclic IB is called *acyclic*. It is a well-known fact that these closure systems satisfy the properties of convex geometries [49, 54], and thus are also known as *acyclic convex geometries*.




Figure 3: The implication graph $G(\Sigma)$ of the IB (X, Σ) of Example 7, associated to the closure system in Figure 1. This directed graph is acyclic, which makes the closure system an acyclic convex geometry. The ancestors of 2, highlighted in green, are 1, 4, 5, i.e., $\text{anc}(2) = \{1, 4, 5\}$. Maximal elements are 4, 5 and the unique minimal element is 3. A *topological order*, as we will use later, would be 4, 1, 5, 2, 3 in this example.

At first sight, acyclicity seems to be a property of IBs more than a property of the underlying closure system, as when taking two distinct IBs of a closure system $(X, \mathcal{C})$, one may be acyclic and not the other. See Example 9.

Example 9. Let $X = \{1, 2, 3, 4\}$ and consider the two IBs (X, Σ_1) , (X, Σ_2) where:

- $\Sigma_1 = \{1 \rightarrow 4, 12 \rightarrow 3, 23 \rightarrow 4\}$
- $\Sigma_2 = \{1 \rightarrow 4, 124 \rightarrow 3, 23 \rightarrow 4\}$

The two implicational bases are equivalent, i.e., they define the same closure system. However, $G(\Sigma_1)$ is acyclic, while $G(\Sigma_2)$ contains a cycle with arcs 34 and 43.

Consequently, one might expect that acyclicity does not especially reflects in closed sets more than the fact that the closure system is a convex geometry. This however is not true. To see this, consider the implicational base of minimal generators $(X, \Sigma_{\text{mingen}})$ of a closure system $(X, \mathcal{C})$, where $\Sigma_{\text{mingen}} = \{A \rightarrow b : b \in X, A \in \text{mingen}(b)\}$, which has been studied under various names [11]. The implication graph $G(\Sigma_{\text{mingen}})$, where there

is an arc from a to b if some $A \in \text{mingen}(b)$ contains a , turns out to be equal to the (inverse of the) δ -relation of lattice theory [11, 45] where $b\delta a$ holds in $(X, \mathcal{C})$ if there exists $M \in \text{irr}(b)$ such that $a \notin M$. This equality is also a direct consequence of the fact that $\text{irr}(b)$ and $\text{mingen}(b)$ define dual hypergraphs, as explained in Section 2. As was shown by Hammer and Kogan [33], $G(\Sigma_{\text{mingen}})$, or equivalently δ , characterizes the acyclicity of a closure system.

Lemma 10 ([33, Lemma 5.2]). *A closure system $(X, \mathcal{C})$ is acyclic if and only if the implication graph $G(\Sigma_{\text{mingen}})$ of its implicational base of minimal generators $(X, \Sigma_{\text{mingen}})$ is acyclic.*

Lemma 10 has interesting consequences. First, given the correspondence between $G(\Sigma_{\text{mingen}})$ and the δ -relation using irreducible closed sets, it shows that acyclicity indeed reflects in (irreducible) closed sets and can be tested in polynomial time given $\text{irr}(\mathcal{C})$. Furthermore, it follows that $G(\Sigma_{\text{mingen}})$, and hence its associated ancestor relation, can be recovered in polynomial time from $\text{irr}(\mathcal{C})$ too. To do so, one just has to check, for all $a, b \in X$, whether there exists $M \in \text{irr}(b)$ such that $a \notin M$. As another consequence, any subset Σ of Σ_{mingen} produces an acyclic IB. Since such an IB (X, Σ) can always be obtained from an arbitrary IB of $(X, \mathcal{C})$ in polynomial time, checking whether $(X, \mathcal{C})$ is indeed acyclic can be done in polynomial time given an arbitrary IB (see Corollary 5.3 in [33]). Moreover, Boros et al. [12, Theorem 17, Corollary 18] show that any IB (X, Σ) of $(X, \mathcal{C})$ that satisfies $\Sigma \subseteq \Sigma_{\text{mingen}}$ enjoys the same ancestor relation as $(X, \Sigma_{\text{mingen}})$. In other words, $G(\Sigma)$ and $G(\Sigma_{\text{mingen}})$ have the same transitive closure. We thus directly obtain a statement that will be used in a few paragraphs.

Theorem 11. *Let $(X, \mathcal{C})$ be an acyclic convex geometry and let (X, Σ) be an implicational base of $(X, \mathcal{C})$ such that $\Sigma \subseteq \Sigma_{\text{mingen}}$. Then, the ancestor relation associated to Σ can be computed in polynomial time when given X and $\text{irr}(\mathcal{C})$.*

Among the minimal generators of $(X, \mathcal{C})$, we now focus on critical generators, for they play a crucial role in acyclic convex geometries. If $(X, \mathcal{C})$ is an acyclic convex geometry, its *critical base* (X, Σ) of $(X, \mathcal{C})$ is defined as follows:

$$\Sigma_{\text{crit}} := \{A \rightarrow b : b \in X \text{ and } A \text{ is a critical generator of } b\}.$$

As a preliminary remark, $(X, \Sigma_{\text{crit}})$ must be acyclic by Lemma 10, since critical generators are minimal generators. A nice result of Hammer and Kogan [33]² shows that not only $(X, \Sigma_{\text{crit}})$ is indeed a valid IB of $(X, \mathcal{C})$, which is not true for convex geometries in general [2, 46, 54], but also that it is unit-minimum and can be obtained in polynomial time from any other IB of $(X, \mathcal{C})$. More precisely, they show:

Theorem 12 (Corollaries 5.6, 5.7 in [33]). *Let $(X, \mathcal{C})$ be an acyclic convex geometry. Then, its critical base $(X, \Sigma_{\text{crit}})$ is unit-minimum and its aggregation is minimum. Moreover, it can be computed in polynomial time from any implicational base of $(X, \mathcal{C})$.*

²In their terminology, an implication $A \rightarrow b$ where A is a critical generator of b is called an *essential implicate* of the corresponding Horn CNF.

As a consequence, the critical base of an acyclic convex geometry constitutes a valid output to the MIB·GEN problem. Therefore, our approach to MIB·GEN will consist in solving the following problem:

Critical Base Generation (CB·GEN)

Input: A ground set X , and the irreducible closed sets of an acyclic convex geometry $(X, \mathcal{C})$.

Output: The critical base of $(X, \mathcal{C})$.

Given that $(X, \Sigma_{\text{crit}})$ is indeed a valid acyclic IB of $(X, \mathcal{C})$ and $\Sigma_{\text{crit}} \subseteq \Sigma_{\text{mingen}}$, Theorem 11 automatically yields the following corollary that we will use in Section 7.

Corollary 13. *The ancestor relation of the critical base of an acyclic convex geometry $(X, \mathcal{C})$ can be computed in polynomial time from X and $\text{irr}(\mathcal{C})$.*

In the rest of the paper, we will call $\text{critanc}(x)$ the ancestor relation associated to $(X, \Sigma_{\text{crit}})$. As a last property of critical generators, we argue that on top of being unit-minimum, the critical base of $(X, \mathcal{C})$ is also minimum in terms of the different measures we introduced in Section 2: cdeg , pdeg , and deg . To do so, we will rely on a result stating that in a convex geometry $(X, \mathcal{C})$, not necessarily acyclic, any critical generator is a subset of a premise in any IB of $(X, \mathcal{C})$. This is a consequence of [39, Theorem 4.8]. It is also directly stated in [57, Lemma 3.10], as well as in [21, Proposition 4.5] but restricted to acyclic convex geometries. In view of our claim on degrees though, we need to prove a slightly stronger property stating that these premise and critical generator share the same closure. We give the statement below and we include for self-containment a proof which differs from the ones presented in [21, 39, 57].

Lemma 14. *Let $(X, \mathcal{C})$ be a convex geometry and let $b \in X$. For any implicational base (X, Σ) of $(X, \mathcal{C})$ and any critical generator A of b , there exists an implication $D \rightarrow B$ in Σ such that $A \subseteq D$, $b \in B$ and $\phi(A) = \phi(D)$.*

Proof. Let (X, Σ) be an IB of $(X, \mathcal{C})$, let $b \in X$ and A a critical generator of b . Let $A = A_0 \subseteq A_1 \subseteq \dots \subseteq A_k = \phi(A)$ be the sequence of sets obtained by applying the forward chaining on A with (X, Σ) . Since A is a critical generator of b , $b \notin A$ but $b \in \phi(A)$. Therefore, there exists some $1 \leq i \leq k$ such that $b \notin A_{i-1}$ but $b \in A_i$. Hence, there exists an implication $D \rightarrow B \in \Sigma$ such that $b \in B$ and $D \subseteq A_{i-1}$. Since $D \subseteq A_{i-1}$ and $\phi(A_{i-1}) = \phi(A)$, we deduce $\phi(D) \subseteq \phi(A)$. Now, since A is a critical generator of b , we have that $\phi(A) \in \min_{\subseteq} \{C : C \in \mathcal{C}, b \in C, b \notin \text{ex}(C)\}$ by Proposition 5. On the other hand, we have $b \in \phi(D)$ but $b \notin \text{ex}(\phi(D))$. We deduce from the minimality of $\phi(A)$ that $\phi(A) = \phi(D)$. As $(X, \mathcal{C})$ is a convex geometry, the unique minimal spanning set of $\phi(D)$ is $\text{ex}(\phi(D)) = \text{ex}(\phi(A)) = A$ by Propositions 3 and 5. Hence, $A \subseteq D$ holds.

Consequently, $D \rightarrow B$ is an implication of Σ such that $A \subseteq D$, $\phi(D) = \phi(A)$ and $b \in B$ as desired. This concludes the proof. $\square$

We can now proceed to show that the critical base, once aggregated, minimizes all the three measures about degrees.

Lemma 15. *Let $(X, \mathcal{C})$ be an acyclic convex geometry with aggregated critical base $(X, \Sigma_{\text{crit}}^*)$. For any $\mu \in \{\text{cdeg}, \text{pdeg}, \text{deg}\}$, we have:*

$$\mu(\Sigma_{\text{crit}}^*) = \min\{\mu(\Sigma) : (X, \Sigma) \text{ is an implicational base of } (X, \mathcal{C})\}.$$

Proof. We show the statement for each of the three measures. We start with the equality regarding cdeg , the conclusion-degree. Let (X, Σ) be an arbitrary IB of $(X, \mathcal{C})$. We show that $\text{cdeg}(\Sigma_{\text{crit}}^*) \leq \text{cdeg}(\Sigma)$. By definition of cdeg , it is sufficient to argue that for each $b \in X$, $\text{cdeg}_{\Sigma_{\text{crit}}^*}(b) \leq \text{cdeg}_{\Sigma}(b)$. Let $b \in X$ and let $A_1, \dots, A_k$ be the critical generators of b . Then $\text{cdeg}_{\Sigma_{\text{crit}}^*}(b) = k$. By Lemma 14, for each $1 \leq i \leq k$ there exists an implication $D_i \rightarrow B_i$ in Σ such that $b \in B_i$ and $\phi(A_i) = \phi(D_i)$. By Propositions 3 and 5, each A_i is the unique minimal spanning set of $\phi(A_i)$. Hence, $\phi(A_i) \neq \phi(A_j)$ for every $1 \leq i, j \leq k$ such that $i \neq j$. Since for each $1 \leq i \leq k$, $\phi(D_i) = \phi(A_i)$, we deduce $\phi(D_i) \neq \phi(D_j)$ for each $j \neq i$. Consequently, all implications $D_1 \rightarrow B_1, \dots, D_i \rightarrow B_i$ must be distinct, all the while satisfying $b \in B_i$ for each $1 \leq i \leq k$. We deduce $k = \text{cdeg}_{\Sigma_{\text{crit}}^*}(b) \leq \text{cdeg}_{\Sigma}(b)$.

We proceed to analyze pdeg , the premise-degree. Again, we need only to prove that for each $a \in X$, $\text{pdeg}_{\Sigma_{\text{crit}}^*}(a) \leq \text{pdeg}_{\Sigma}(a)$. Let $a \in X$ and let $A_1, \dots, A_k$ be the (distinct) critical generators to which a belongs. We obtain $\text{pdeg}_{\Sigma_{\text{crit}}^*}(a) = k$ because Σ_{crit}^* is aggregated. Applying Lemma 14 once more, there exists $D_1 \rightarrow B_1, \dots, D_k \rightarrow B_k$ in Σ such that, for each $1 \leq i \leq k$, $a \in A_i \subseteq D_i$ and $\phi(A_i) = \phi(D_i)$. Similarly to cdeg , we obtain $\phi(D_i) \neq \phi(D_j)$ for all $1 \leq i, j \leq k$ such that $i \neq j$. Thus, all implications $D_1 \rightarrow B_1, \dots, D_k \rightarrow B_k$ are distinct, with $a \in D_i$ for each $1 \leq i \leq k$. We deduce $k = \text{pdeg}_{\Sigma_{\text{crit}}^*}(a) \leq \text{pdeg}_{\Sigma}(a)$ for all $a \in X$, and $\text{pdeg}(\Sigma_{\text{crit}}^*) \leq \text{pdeg}(\Sigma)$ follows.

As for the degree deg , we have $\text{deg}_{\Sigma_{\text{crit}}^*}(x) = \text{cdeg}_{\Sigma_{\text{crit}}^*}(x) + \text{pdeg}_{\Sigma_{\text{crit}}^*}(x)$ for each $x \in X$ as x cannot belong to one of its critical generators by definition. According to the previous arguments, we thus have $\text{deg}_{\Sigma_{\text{crit}}^*}(x) \leq \text{deg}_{\Sigma}(x)$ for each x . So $\text{deg}(\Sigma_{\text{crit}}^*) \leq \text{deg}(\Sigma)$ also holds, which concludes the proof. $\square$

Moreover, observe that $\text{cdeg}(\Sigma_{\text{crit}}) = \text{cdeg}(\Sigma_{\text{crit}}^*)$ so the critical base, which is not aggregated, is enough for the conclusion degree. We directly deduce the following for acyclic convex geometries.

Corollary 16. *The (premise- or conclusion-) degree of an acyclic convex geometry is exactly the (premise- or conclusion-) degree of its aggregated critical base.*

To conclude this section, we give further insights on the translation tasks within acyclic convex geometries. First, thanks to Theorem 12, one can always assume to be given the critical base as an input to ICS·ENUM in this context. More importantly, the critical base is, by Theorem 12, a correct solution to MIB·GEN, a fact that we will rely upon in Section 7.

It is shown in [21, Theorem 3.1] that in acyclic convex geometries, all three problems ICS·ENUM, ACS·ENUM and MIB·GEN are harder than distributive lattice dualization, a generalization of MIS·ENUM and MHS·ENUM. Here, we briefly discuss a simpler reduction that we will later adapt to obtain further hardness results for acyclic IBs with bounded degrees. Let $\mathcal{H} = (V, \mathcal{E})$ be a Sperner hypergraph as an input to MIS·ENUM.

Let $X := V \cup \{z\}$ where z is a new gadget element and let $\Sigma := \{E \rightarrow z : E \in \mathcal{E}\}$. The IB (X, Σ) is our input to ICS·ENUM and can be computed in polynomial time from $\mathcal{H}$. It is acyclic, and furthermore it has height 2, that is, $G(\Sigma)$ may be seen as a poset of height 2. The reduction is illustrated on an example in Figure 4. We proceed to describe the irreducible closed sets of $(X, \mathcal{C})$, the closure system associated to (X, Σ) . First, each element $v \in V$ is attached to a unique irreducible being $X \setminus \{v\}$, as v is the conclusion of no implications. As for z , by definition of (X, Σ) , $M \subseteq X$ belongs to $\text{irr}(z)$ if and only if M is inclusion-wise maximal for not including any premise of Σ , i.e., if and only if $M \in \text{MIS}(\mathcal{H})$. Therefore, $\text{irr}(\mathcal{C}) = \text{MIS}(\mathcal{H}) \cup \{X \setminus \{v\} : v \in V\}$. Given that $\text{irr}(\mathcal{C})$ has only a linear number of irreducible closed sets not in $\text{MIS}(\mathcal{H})$, which can be identified in polynomial time, an algorithm for ICS·ENUM solves MIS·ENUM with the same time complexity. Similarly, an algorithm solving ACS·ENUM given z would directly solve MIS·ENUM.




Figure 4: An illustration of the reduction from MIS·ENUM to ICS·ENUM for Theorem 17. The initial hypergraph is $\mathcal{H} = (\{1, 2, 3, 4, 5\}, \{123, 234, 45\})$ and the corresponding IB (X, Σ) has implications $123 \rightarrow z$, $234 \rightarrow z$, $45 \rightarrow z$, drawn in grey. The set 125, highlighted in blue, is a maximal independent set of $\mathcal{H}$ and hence an irreducible closed set attached to z .

For the MIB·GEN problem we instead use $\mathcal{H}$ as an input to MHS·ENUM. We consider the closure system with $X := V \cup \{z\}$ and $\text{irr}(\mathcal{C}) := \{V \setminus E : E \in \mathcal{E}\} \cup \{X \setminus \{v\} : v \in V\}$. Note that $\text{irr}(\mathcal{C})$ can be built in polynomial time. The unique unit-minimum IB associated to $(X, \mathcal{C})$ is then (X, Σ) with $\Sigma = \{T \rightarrow z : T \in \text{MHS}(\mathcal{H})\}$. It is also acyclic and has height 2.

We conclude these observations with the following.

Theorem 17. *In acyclic convex geometries, the problems ICS·ENUM, ACS·ENUM and MIB·GEN are at least as hard than MIS·ENUM (hence MHS·ENUM) even for acyclic convex geometries admitting an implicational base of height 2.*

By a small change in the construction of Theorem 17, it can be shown that the result still holds for ACS·ENUM in the case of bounded conclusion-degree. Given $\mathcal{H} = (V, \mathcal{E})$ with $\mathcal{E} = \{E_1, \dots, E_m\}$, we put $X = V \cup \{z_i : 1 \leq i \leq m\}$ where $z_1, \dots, z_m$ are new gadget elements and $\Sigma = \{E_i \rightarrow z_i : 1 \leq i \leq m\} \cup \{z_i \rightarrow z_{i+1} : 1 \leq i < m\}$. Note that Σ has conclusion-degree 2. The element z_m will be the element as input of ACS·ENUM along with (X, Σ) . Again, we illustrate the reduction in Figure 5. In the closure system associated to (X, Σ) , we have $\text{irr}(z_m) = \text{MIS}(\mathcal{H})$. We obtain the subsequent theorem.

Theorem 18. *The problem ACS·ENUM is MIS·ENUM-hard even for acyclic implicational bases of conclusion-degree at most 2.*




Figure 5: The reduction from MIS·ENUM to ACS·ENUM with conclusion-degree at most 2, applied to the hypergraph of Figure 4. The implications are now $123 \rightarrow z_1$, $234 \rightarrow z_2$, $45 \rightarrow z_3$, $z_1 \rightarrow z_2$, $z_2 \rightarrow z_3$. With $z_1 \rightarrow z_2$, $z_2 \rightarrow z_3$, none of z_1, z_2 , highlighted in red, cannot belong to any irreducible closed set attached to z_3 . As compared to the reduction of Theorem 17, irreducible closed sets of z_3 (formerly z) are unchanged.

Notice that this construction does not apply to ICS·ENUM, as the z_i 's have non-trivial attached irreducible closed sets. Indeed, for a given i , $\text{irr}(z_i)$ is in a one-to-one correspondence with the maximal independent sets of the hypergraph consisting of the first i edges $E_1, \dots, E_i$. Hence, the size of $\text{irr}(\mathcal{C})$ might be super-polynomial in the size of $\text{MIS}(\mathcal{H})$. Therefore, the question of the complexity of ICS·ENUM for bounded conclusion-degree arises and will be dealt with in the following section.

4 Top-down procedures

We describe two sequential procedures to solve ICS·ENUM and CB·GEN in acyclic convex geometries. Intuitively, the first procedure consists in enumerating, incrementally and in a top-down fashion according to a topological ordering $x_1, \dots, x_n$ of the elements of the implicational base, all the irreducible closed sets attached to x_i . The second one proceeds analogously and lists all critical generators of x_i at each step.

4.1 Irreducible closed sets

Let us focus on irreducible closed sets first. Recall that the ancestors of an element x in an acyclic implicational base are the elements that participate in a path of implications to x , and which can be identified in polynomial time; see Section 3 for a formal definition. We will be interested in solving the following version of ACS·ENUM in which we are given ancestors' irreducible closed sets as additional information.

Attached Closed Sets Enumeration with Ancestors' solutions (ACS+A·ENUM)

Input: An acyclic implicational base (X, Σ) , an element $x \in X$, and the family $\text{irr}(y)$ for every ancestor y of x in Σ .

Output: The set $\text{irr}(x)$.

Clearly, an algorithm for ACS·ENUM can be used to solve ACS+A·ENUM, while the reverse may not be true. We motivate this latter problem by showing that we have an

incremental-polynomial time algorithm for ICS·ENUM whenever we have one for ACS+A·ENUM. Intuitively, our procedure will iteratively use an algorithm for ACS+A·ENUM on the elements of the ground set, in order, according to an arbitrary topological ordering of the implicational base. Note that the base case corresponds to computing the irreducible closed set of the maximal elements x , i.e., those verifying $\text{irr}(x) = \{X \setminus \{x\}\}$. In terms of difficulty, ACS+A·ENUM may thus be seen as standing as an intermediate problem between ACS·ENUM and ICS·ENUM in acyclic convex geometries.

Theorem 19. *Let $f: \mathbb{N}^2 \rightarrow \mathbb{N}$ be a function. Let us assume that there is an algorithm that, given an instance of ACS+A·ENUM of size N , produces its i^{th} solution in $f(N, i)$ -time. Then there is an algorithm that, given an acyclic instance of ICS·ENUM of size N' , produces its j^{th} solution in $O(\text{poly}(N') + N' \cdot f(N' + j, j))$ time.*

Proof. Let us assume the existence of an algorithm **A** for ACS+A·ENUM as required by the statement. Without loss of generality and for convenience in the rest of the proof, let us assume that f is increasing. We describe a sequential procedure **B** which produces the j^{th} solution of ICS·ENUM within the desired time bound.

Let (X, Σ) be an acyclic instance of ICS·ENUM. Recall that, for convenience, we view the ancestor relation induced by Σ as a partial order on X , and denote sources of $G(\Sigma)$ as maximal elements. First, we identify each maximal element a of X , compute $\text{irr}(a) = \{X \setminus \{a\}\}$ in $\text{poly}(|X|, |\Sigma|)$ time, and mark these a as handled. Every such set is output and stored in memory. Then, iteratively, given a maximal element $x \in X$ among those that are not marked, we identify its ancestors in $\text{poly}(|X|, |\Sigma|)$ time, output and store $\text{irr}(x)$ using **A**, and mark x . The procedure stops when every element of X has been marked. In other words, in this procedure, the elements x are selected with respect to a topological ordering of $G(\Sigma)$.

Note that we can indeed use **A** in the last step of our procedure since the sets $\text{irr}(a)$ for all ancestors a of x have been inductively computed by the choice of x . Furthermore, as by Proposition 3 the sets $\text{irr}(a)$ and $\text{irr}(b)$ are disjoint for any two distinct $a, b \in X$, there is no repetition in the outputs of this procedure: each set output by **A** is a distinct new solution of **B**.

We now turn to the complexity analysis. Observe that, except for the executions of **A**, the different steps of **B** take $\text{poly}(N')$ time for $N' = |X| + |\Sigma|$. Let us now focus on the time spent by **B** in order to output the j^{th} solution in case it is not the irreducible closed set of a maximal element, i.e., when it has been provided by an instance of **A**. By assumption, this instance produces the i^{th} irreducible closed set attached to x , each time being a new solution of **B**, in a time which is bounded by $f(N, i)$, for N being the sum of the sizes of X , Σ and $\bigcup_{a \in \text{anc}(x)} \text{irr}(a)$. Note that the latter family has already been output by **B** at this stage and can be recovered in $\text{poly}(j)$ time from the memory. Thus, and since f is increasing, this is bounded by $f(|X| + |\Sigma| + j, j)$ for j being the total number of solutions output by **B** so far. And finally, **A** will be called only once for each x of the ground set. We conclude to the desired time bound by replacing $|X| + |\Sigma|$ by N' . $\square$

We note that an analogous statement to the one in Theorem 19 holds by relaxing the required and obtained time bounds to output-polynomial times. On the other hand,

improving the complexity of an algorithm for ACS+A·ENUM to polynomial delay (or even input-polynomial as we will provide later for the case of bounded conclusion-degree) does not generally ensure polynomial delay (or incremental-linear) for ICS·ENUM, since the size of an instance of ACS+A·ENUM grows with the number of already obtained solutions for ICS·ENUM, which typically becomes exponential in the size of the implicational base.

4.2 Critical generators

We now turn on enumerating the critical generators of an acyclic convex geometry. Analogously as for irreducible closed sets, we proceed in a top-down fashion according to a topological ordering of the ground set, this time related to the critical base we wish to compute. Only it should be argued that such an ordering can be efficiently computed from the irreducible closed sets, i.e., without the knowledge of the critical base and its implication graph, a fact that follows from Corollary 13.

More formally, we are interested in solving the following enumeration problem, which is an analogue of ACS+A·ENUM to critical generators.

Critical Generators of an Element with Ancestors' solutions (CGE+A·GEN)

Input: A ground set X , the irreducible closed sets $\text{irr}(\mathcal{C})$ of an acyclic convex geometry $(X, \mathcal{C})$, $x \in X$, and the set $\text{critgen}(y)$ for every $y \in \text{critanc}(x)$.

Output: The family $\text{critgen}(x)$.

We stress the fact that $\mathcal{C}$ is not given in this problem; only X and $\text{irr}(\mathcal{C})$ are given as a representation. Note that $\text{critgen}(x)$ is empty for maximal elements of the critanc relation, which corresponds to the base case of our procedure. Moreover, sets lying in $\text{critgen}(y)$ and $\text{critgen}(y')$ for distinct elements $y, y' \in X$ yield distinct implications of the (unit) critical base of $\mathcal{C}$. Then, following a direct analogue of the proof of Theorem 19 we obtain the following result.

Theorem 20. *Let $f: \mathbb{N}^2 \rightarrow \mathbb{N}$ be a function. Let us assume that there is an algorithm that, given any instance of CGE+A·GEN of size N , produces its i^{th} solution in $f(N, i)$ -time. Then there is an algorithm that, given any acyclic instance of CB·GEN of size N' , produces its j^{th} solution in $O(\text{poly}(N') + N' \cdot f(N' + j \cdot |X|, j))$ time.*

4.3 Limitations

Let us end this section by pointing some limitation to the top-down approach. Note that for implicational bases (or critical base) of height at most two, the inputs of ACS+A·ENUM and CGE+A·GEN do not contain any valuable additional information compared to the inputs of ICS·ENUM and CB·GEN. Indeed, as mentioned earlier, for such bases, each maximal element a has a unique, polynomially characterized attached irreducible closed set $X \setminus \{a\}$, due to the fact that a has no minimal generators, and in particular no critical generators. Nevertheless, these instances remain MIS·ENUM-hard by Theorem 17. Thus, in order to expect tractable algorithms using this strategy, one should most probably restrict the closure system to particular cases or bound structural parameters.

5 Bounded conclusion-degree

We show that the top-down approach successfully provides an incremental-polynomial time algorithm for ICS·ENUM whenever the input implicational base has bounded conclusion-degree, proving one of the two cases of Theorem 1.

In the following, let us consider an instance of ACS+A·ENUM, i.e., we fix some acyclic implicational base (X, Σ) , an element $x \in X$, and assume that we are given $\text{irr}(y)$ for all ancestors y of x in Σ . Let us consider the hypergraph of *premises* of x defined by

$$\mathcal{E}_x := \{A : A \rightarrow B \in \Sigma, x \in B\}, \text{ and} \\ \mathcal{H}_x := \left(\bigcup \mathcal{E}_x, \mathcal{E}_x \right).$$

The intuition behind the following lemma is that, in order to get an irreducible closed set attached to x , one should remove at least one element a in the premise of each implication having x in the conclusion, which can be done by considering the intersection with an irreducible closed set attached to a . Recall that $\mathcal{S}(T)$ denotes the family of irreducible selections; see Section 2.

Lemma 21. *For every $M \in \text{irr}(x)$ there exists a minimal transversal T of $\mathcal{H}_x$ and an irreducible selection $\mathcal{S} \in \mathcal{S}(T)$ such that*

$$M = \left(\bigcap \mathcal{S} \right) \setminus \{x\}.$$

Proof. Let $Y := \bigcup \mathcal{E}_x \setminus M$. Note that for each $A \in \mathcal{E}_x$ there exists $a \in A$ such that $a \in Y$, as otherwise $A \subseteq M$ which contradicts the fact that $M \in \text{irr}(x)$ as A implies x . Thus, Y is a transversal of $\mathcal{H}_x$. Let $T \subseteq Y$ be any minimal transversal of $\mathcal{H}_x$ included in Y . Note that for all $a \in T$, as M is a closed set and $a \notin M$, there exists a set $M_a \in \text{irr}(a)$ such that $M \subseteq M_a$. This defines an irreducible selection $\mathcal{S} = (M_a)_{a \in T}$ for T , where each such M_a includes M . Therefore, and as $x \notin M$, we have $M \subseteq \left(\bigcap \mathcal{S} \right) \setminus \{x\}$.

Now, as every member of $\mathcal{S}$ is closed, so is $\bigcap \mathcal{S}$. Moreover, by definition of $\mathcal{S}$, no premise A of $\mathcal{E}_x$ is contained in such an intersection. Thus the set $M' := \left(\bigcap \mathcal{S} \right) \setminus \{x\}$ is closed. By the maximality of $M \in \text{irr}(x)$, and since $M \subseteq M'$, we conclude that $M' = M$, hence to the desired equality. $\square$

Note that, by definition, the number of edges in $\mathcal{H}_x$ is exactly $\text{cdeg}(x)$, which is at most $\text{cdeg}(\Sigma)$. Moreover, since the ground set of $\mathcal{H}_x$ is a subset of the ancestors of x , the selections of Lemma 21 are part of the input we consider for ACS+A·ENUM. We use these observations to derive a polynomial-time algorithm for ACS+A·ENUM whenever the conclusion-degree of x is bounded by a constant.

Theorem 22. *There is an algorithm solving ACS+A·ENUM in $N^{O(k)}$ time on inputs of size N and maximum conclusion-degree k .*

Proof. As $\mathcal{E}_x$ has cardinality at most k , all possible minimal transversals T of $\mathcal{H}_x$ can be computed in $|X|^{O(k)}$ time by brute forcing all subsets of size at most k , and testing

in $\text{poly}(|X| + |\mathcal{H}_x|)$ time, hence $\text{poly}(N)$ time, whether the obtained set is a minimal transversal of $\mathcal{H}_x$. Moreover, all possible irreducible selections $\mathcal{S} \in \mathcal{S}(T)$ for each such T can be computed in $(\max_{a \in T} |\text{irr}(a)|)^{O(k)}$, hence in $N^{O(k)}$ time, since every $a \in T \subseteq \bigcup \mathcal{E}_x$ is an ancestor of x , thus the families $\text{irr}(a)$ are part of the input. As the intersections of Lemma 21 can be computed in $\text{poly}(N)$ time, and as it can be checked in $\text{poly}(N)$ time if they are closed sets attached to x , we conclude that all $M \in \text{irr}(x)$ can be identified in $N^{O(k)}$ time, as desired. $\square$

We conclude to the following, as a corollary of Theorems 19 and 22, and proving the case of bounded-conclusion degree in Theorem 1.

Theorem 23. *There is an incremental-polynomial time algorithm solving ICS·ENUM in acyclic implicational bases of bounded conclusion-degree.*

Let us remark that in our algorithm, we actually do not need the knowledge of $\text{irr}(a)$ for *all* ancestors of x , but only for those in a premise which has x in conclusion. This however makes no asymptotic difference on the final complexity we get for ICS·ENUM.

Also, note that in the proof of Theorem 22 we only need a constant upper bound on the size k of a minimal transversal of $\mathcal{H}_x$, to derive an input-polynomial time algorithm solving ACS+A·ENUM. We immediately derive the following generalization of Theorem 23 dealing with hypergraphs of bounded dual-dimension (i.e., admitting minimal transversals of bounded size) who define complement of conformal hypergraphs [9], which typically arise in geometric settings.

Theorem 24. *There is an incremental-polynomial time algorithm solving ICS·ENUM if, for any element x of the ground set, the hypergraph of premises having x in their conclusion has bounded dual-dimension.*

We end the section by showing the limitations of this approach in order to get output-polynomial time bounds whenever we drop the condition of bounded conclusion-degree.

First, note that the key step of the approach is to guess the correct $T \in \text{MHS}(\mathcal{H}_x)$ and $\mathcal{S} \in \mathcal{S}(T)$ as in the statement of Lemma 21, which we do by testing all possibilities in Theorem 22. However, there are cases where the number of such possibilities is exponential in $|\text{irr}(\Sigma)|$, which typically occurs if minimal transversals of $\mathcal{H}_x$ are of unbounded size, or if $\text{MHS}(\mathcal{H}_x)$ is of exponential size in $\text{irr}(x)$.

To get one such example, consider $k \in \mathbb{N}$ and the implicational base defined by $X := \{a_1, b_1, \dots, a_k, b_k, x\}$ and $\Sigma := \{a_i b_i \rightarrow x : 1 \leq i \leq k\} \cup \{b_i \rightarrow b_{i+1} : 1 \leq i < k\}$; see Figure 6 for an illustration. Then $\text{irr}(x) = \{\{a_1, \dots, a_j\} \cup \{b_{j+1}, \dots, b_k\} : 0 \leq j \leq k\}$ so $|\text{irr}(x)| = k + 1$, and the other elements in $X \setminus \{x\}$ are easily seen to have a single attached irreducible closed set. Therefore, $\text{irr}(\Sigma)$ is of polynomial size in $|X| = 2k + 1$ and $|\Sigma| = 2k - 1$. However, $|\text{MHS}(\mathcal{H}_x)| = 2^k$ and computing this set would fail in providing an output-polynomial time algorithm for enumerating $\text{irr}(x)$.

At last, we note that this example has bounded premise-degree. We will show in the next section that such instances can actually be efficiently solved by solution graph traversal.




Figure 6: An instance where $|\text{MHS}(\mathcal{H}_x)| = 2^k$ but $|\text{irr}(\Sigma)| = 3k + 1$ and $|\Sigma| = 2k - 1$.

6 Bounded premise-degree

We provide an incremental-polynomial time algorithm solving ICS·ENUM whenever the given implicational base has bounded premise-degree. Our algorithm relies on the top-down approach from Section 4 and solves ACS+A·ENUM relying on the solution graph traversal technique.

In the *solution graph traversal* framework, one has to define a digraph called *solution graph*, whose nodes are the solutions we aim to enumerate, and whose arcs (F, F') consists of ways of reconfiguring a given solution F into another solution F' . Then, the goal is not to build the solution graph entirely, but rather to traverse all its nodes efficiently and avoiding repetitions of nodes. This general technique has been fruitfully applied to numerous enumeration problems, see e.g., [14, 37, 58], and generalized to powerful frameworks [4, 16, 17]. In the following we will use the next folklore result which states conditions for such a technique to be efficient; see [26] for a quantitative statement and a proof.

Theorem 25 (Folklore). *Let $\mathcal{F} \subseteq \mathcal{P}(X)$ be a solution set (to be enumerated) over a ground set X , and let us denote by N the size of the input. Then there is a $\text{poly}(N)$ -delay algorithm enumerating $\mathcal{F}$ whenever:*

1. *one member of $\mathcal{F}$ can be found in $\text{poly}(N)$ -time*
2. *there is a function $\mathcal{N} : \mathcal{F} \rightarrow \mathcal{P}(\mathcal{F})$ called neighboring function mapping every solution to a set of solutions that can be computed in $\text{poly}(N)$ time given $F \in \mathcal{F}$*
3. *the directed solution graph $(\mathcal{F}, \{(F, F') : F' \in \mathcal{N}(F)\})$ is strongly connected*

In the following, let us consider an instance of ACS+A·ENUM, i.e., we fix some acyclic implicational base (X, Σ) , an element $x \in X$, and we assume that we are given $\text{irr}(y)$ for all ancestors y of x .

We first argue that Condition 1 holds for ACS+A·ENUM. Note that there is a function GC_x which maps any $Y \subseteq X$ such that $x \notin \phi(Y)$ to a maximal set $M \supseteq Y$ with this property. We call the obtained set the *greedy completion* of Y . Note that $\text{GC}_x(Y)$ can be computed in polynomial time in the size of (X, Σ) , and that the obtained set lies in $\text{irr}(x)$. We derive the following.

Lemma 26. *A first solution $\text{GC}_x(\emptyset) \in \text{irr}(x)$ can be obtained in $\text{poly}(|X| + |\Sigma|)$ time.*

We now turn on showing the existence of a polynomial-time computable neighboring function satisfying Conditions 2 and 3. Given some irreducible closed set $M \in \text{irr}(x)$, and some element $y \in X \setminus M$ we define

$$\begin{aligned}\mathcal{E}_{M,y} &:= \{A : A \rightarrow B \in \Sigma \text{ with } A \setminus M = \{y\}, B \not\subseteq M\}, \\ V_{M,y} &:= \bigcup \mathcal{E}_{M,y}, \\ \mathcal{H}_{M,y} &:= (V_{M,y}, \mathcal{E}_{M,y}).\end{aligned}$$

In other words, $\mathcal{H}_{M,y}$ is the hypergraph of premises of Σ not included in M , triggered by adding y , and with a conclusion not in M . Note that $|\mathcal{E}_{M,y}| \leq \text{pdeg}(y) \leq \text{pdeg}(\Sigma)$. Also, because every premise in $\mathcal{E}_{M,y}$ implies a $b \notin M$, hence an ancestor of x , we derive the following observation by transitivity.

Observation 27. *Every vertex of $\mathcal{H}_{M,y}$ is an ancestor of x .*

Regarding Condition 2, we first need the following statement, which holds the same idea as Lemma 21 for bounded premise-degree.

Lemma 28. *Let M, M^* be two distinct sets in $\text{irr}(x)$. Then there exist $y \in M^* \setminus M$, $T \in \text{MHS}(\mathcal{H}_{M,y})$, and $\mathcal{S} \in \mathcal{S}(T)$ such that*

$$M^* \subseteq \bigcap \mathcal{S}.$$

Proof. Let y be an element of $M^* \setminus M$ which is maximal, with respect to the ancestor order, for this property. Let $A \in \mathcal{E}_{M,y}$ and $A \rightarrow B \in \Sigma$ such that $B \not\subseteq M$. As $y \in A$, every $b \in B$ is a descendant of y so $b \notin M^*$ or $b \in M$ by maximality of y . Moreover, $B \not\subseteq M$ by definition of $\mathcal{E}_{M,y}$ so there is a $b \in B$ such that $b \notin M^*$.

Therefore, for all $A \in \mathcal{E}_{M,y}$, $A \not\subseteq M^*$ (M^* is a closed set). So $V_{M,y} \cap M^*$ is an independent set of $\mathcal{H}_{M,y}$ thus its complement is a transversal. Let T be any minimal transversal of $\mathcal{H}_{M,y}$ included in $V_{M,y} \setminus M^*$. Then, for all $a \in T$, $a \notin M^*$ so there exists $M_a \in \text{irr}(a)$ such that $M^* \subseteq M_a$. For the irreducible selection $\mathcal{S} = (M_a)_{a \in T}$, one gets $M^* \subseteq \bigcap \mathcal{S}$. $\square$

We are now ready to define the aforementioned neighboring function. Given $M \in \text{irr}(x)$ and $y \notin M$, we define

$$\mathcal{N}(M, y) := \left\{ \text{GC}_x \left((M \cap \bigcap \mathcal{S}) \cup \{y\} \right) : \begin{array}{l} \mathcal{S} \in \mathcal{S}(T) \text{ and } T \in \text{MHS}(\mathcal{H}_{M,y}) \\ \text{such that } x \notin \phi((M \cap \bigcap \mathcal{S}) \cup \{y\}) \end{array} \right\}$$

and let $\mathcal{N}(M) := \bigcup_{y \notin M} \mathcal{N}(M, y)$. We prove that Conditions 2 and 3 hold for this function.

Observation 29. *By definition, the greedy completion in the definition of $\mathcal{N}(M, y)$ is well defined and $\mathcal{N}(M) \subseteq \text{irr}(x)$.*

Lemma 30. *Let N be the sum of the sizes of X , Σ , and $\text{irr}(a)$ for all ancestors a of x . Then, for all $M \in \text{irr}(x)$, $\mathcal{N}(M)$ can be computed in $N^{O(k)}$ time where k is the premise-degree of Σ .*

Proof. Let us first note that $\text{GC}_x((M \cap \bigcap \mathcal{S}) \cup \{y\})$ can be computed in $\text{poly}(|X| + |\Sigma| + |\bigcap \mathcal{S}|)$ time, for any y and $\mathcal{S}$. Note that $\mathcal{E}_{M,y}$ has cardinality at most k , hence that all possible minimal transversals T of $\mathcal{H}_{M,y}$ can be computed in $|X|^{O(k)}$ time by brute forcing all subsets of size at most k , and testing whether the obtained set is a minimal transversal of $\mathcal{H}_{M,y}$ in $\text{poly}(|X| + |\mathcal{H}_{M,y}|)$ time, hence in $\text{poly}(|X| + |\Sigma|)$ time. Moreover, all possible irreducible selections $\mathcal{S} \in \mathcal{S}(T)$ for each such T can be computed in $(\max_{a \in T} |\text{irr}(a)|)^{O(k)}$ time, hence in $N^{O(k)}$ time as families $\text{irr}(a)$ are part of the input. Finally, the last union over y multiplies the total time by a factor at most $|X| \leq N$, hence the result. $\square$

Note that Condition 2 is ensured by Observation 29 and Lemma 30. We are thus left to prove Condition 3.

Lemma 31. *Let M, M^* be two distinct sets in $\text{irr}(x)$. Then there exists $M' \in \mathcal{N}(M)$ such that $M \cap M^* \subset M' \cap M^*$.*

Proof. Let $y \in M^* \setminus M$, $T \in \text{MHS}(\mathcal{H}_{M,y})$ and $\mathcal{S} \in \mathcal{S}(T)$ such that $M^* \subseteq \bigcap \mathcal{S}$, as provided in Lemma 28. Note that T may intersect $M \setminus M^*$, but not M^* . Then

$$\begin{aligned} M \cap M^* &\subseteq M \cap \bigcap \mathcal{S} \\ &\subseteq (M \cap \bigcap \mathcal{S}) \cup \{y\} \\ &\subseteq \text{GC}_x((M \cap \bigcap \mathcal{S}) \cup \{y\}) \in \mathcal{N}(M). \end{aligned}$$

So, by setting $M' := \text{GC}_x((M \cap \bigcap \mathcal{S}) \cup \{y\})$ and intersecting with M^* , one gets $M \cap M^* \subseteq M' \cap M^*$.

Finally, let us recall that $y \notin M$ by definition but $y \in M' \cap M^*$ by construction and definition. Hence the inclusion is strict. $\square$

We derive the following, namely Condition 3, which is easily proved by induction since $M' \in \mathcal{N}(M)$ has greater intersection with M^* than M has with M^* .

Corollary 32. *The solution digraph $(\text{irr}(x), \{(M, M') : M' \in \mathcal{N}(M)\})$ is strongly connected.*

Combining Theorem 25, Lemma 26, Lemma 30, and Corollary 32, we conclude to a polynomial-delay algorithm for ACS+A·ENUM by solution graph traversal for acyclic implicational bases of bounded premise-degree. This, together with Theorem 19, in turn implies the following theorem, proving the remaining case of bounded premise-degree in Theorem 1.

Theorem 33. *There is an incremental-polynomial time algorithm solving ICS·ENUM on acyclic implicational bases of bounded premise-degree.*

As in the previous section, notice that we only needed to bound the size of a minimal transversal of $\mathcal{H}_{M,y}$ to derive an input-polynomial time algorithm solving ACS+A·ENUM. This gives the following result, analogous to Theorem 24.

Theorem 34. *There is an incremental-polynomial time algorithm solving ICS·ENUM if, for any irreducible closed set M and $y \notin M$, the hypergraph $\mathcal{H}_{M,y}$ has bounded dual-dimension.*

7 Generating the critical implicational base

We provide an algorithm which constructs the critical base of an acyclic convex geometry given by its irreducible closed sets, namely solving CB·GEN. Our algorithm relies on the top-down approach of Section 4.2, and solves CGE+A·GEN by a saturation procedure which either finds a missing critical generator, or conclude that all critical generators have been found. It performs in polynomial time whenever the aggregated critical base of the closure system has bounded premise- or conclusion-degree, relying on our algorithms from Sections 5 and 6, proving Theorem 2.

In the *saturation technique*, the goal is to decide, starting from the empty set as an initial partial solution set, whether a solution is missing, and to produce one such missing solution if it is the case. This analogue functional problem is usually referred to as the *decision version with counter example* of the enumeration problem. It is folklore and easily seen that there is an incremental-polynomial time algorithm for an enumeration problem whenever its decision version with counterexample can be solved in polynomial time; see e.g., [51] for more details and [15, 29, 43] for typical applications of this technique.

More formally, in the remaining of the section, we are interested in solving the following decision version with counterexample of CGE+A·GEN. Let us recall that critanc, the ancestor relation associated to the critical base, can be computed in polynomial time from the irreducible closed sets by Corollary 13.

Critical Generators of an Element with Ancestors' solutions,

Decision version with counterexample (CGE+A·DEC)

Input: A ground set X , the irreducible closed sets $\text{irr}(\mathcal{C})$ of an acyclic convex geometry $(X, \mathcal{C})$, an element $x \in X$, the set $\text{critgen}(y)$ for every $y \in \text{critanc}(x)$, and a family $\mathcal{F} \subseteq \text{critgen}(x)$.

Output: Yes if $\mathcal{F} = \text{critgen}(x)$, a set $A \in \text{critgen}(x) \setminus \mathcal{F}$ otherwise.

In the following, let us consider an instance of this problem. Let us put

$$\Sigma' := \{A \rightarrow x : A \in \mathcal{F}\} \cup \{A \rightarrow y : A \in \text{critgen}(y), y \in \text{critanc}(x)\}.$$

Note that Σ' defines a subset of the critical base of $(X, \mathcal{C})$. We will note $(X, \mathcal{C}')$ the closure system represented by Σ' , and ϕ' its associated closure operator. Since we are now dealing with two distinct closure systems, we shall write $\text{irr}_{\mathcal{C}}(x)$ and $\text{irr}_{\mathcal{C}'}(x)$ instead of $\text{irr}(x)$ to distinguish between the two closure systems, and we do analogously for critgen .

As $\Sigma' \subseteq \Sigma_{\text{crit}}$, one can notice that $\mathcal{C} \subseteq \mathcal{C}'$. This little fact will be a key element in the following proofs.

To solve CGE+A·DEC, our strategy is to find a set $M \in \text{irr}_{\mathcal{C}'}(x) \setminus \text{irr}_{\mathcal{C}}(x)$ using an algorithm for ACS+A·ENUM. If such a M exists, we argue that it provides a counterexample $A \in \text{critgen}(x) \setminus \mathcal{F}$ (Lemma 35 alongside Proposition 6). If not, we prove that $\mathcal{F} = \text{critgen}(x)$ (Lemma 36).

Lemma 35. *If there exists $M \in \text{irr}_{\mathcal{C}'}(x)$ such that $M \notin \text{irr}_{\mathcal{C}}(x)$, then there exists $A \in \text{critgen}_{\mathcal{C}}(x)$ such that $A \subseteq M$ and $A \notin \mathcal{F}$.*

Proof. First, we argue that $M \notin \mathcal{C}$ and $x \in \phi(M)$. Assume for contradiction that $M \in \mathcal{C}$. Because $x \notin M$ and $M \notin \text{irr}_{\mathcal{C}}(x)$ by assumption, there exists $y \neq x$ such that $M \cup \{y\} \in \mathcal{C}$, as $(X, \mathcal{C})$ is a convex geometry. However, as $\Sigma' \subseteq \Sigma_{\text{crit}}$, we have $\mathcal{C} \subseteq \mathcal{C}'$. Therefore, $M \cup \{y\} \in \mathcal{C}'$ holds, which contradicts $M \in \text{irr}_{\mathcal{C}'}(x)$. We deduce $M \notin \mathcal{C}$. Again as $\mathcal{C} \subseteq \mathcal{C}'$, $\phi(M) \in \mathcal{C}'$ holds. From $M \in \text{irr}_{\mathcal{C}'}(x)$ and $M \subset \phi(M)$, we derive $x \in \phi(M)$ as required.

Since $x \notin M$ but $x \in \phi(M)$ and $(X, \Sigma_{\text{crit}})$ is acyclic, there exists an implication $A \rightarrow b \in \Sigma_{\text{crit}}$ with $A \subseteq M$ but $b \notin M$ such that $A \cup \{b\} \subseteq \text{critanc}_{\mathcal{C}}(x) \cup \{x\}$. If $b \neq x$, $A \rightarrow b \in \Sigma'$ by definition of Σ' . As $M \in \mathcal{C}'$, $A \subseteq M$ thus implies $b \in M$, a contradiction. We deduce that $b = x$ must hold, which concludes the proof. $\square$

Lemma 36. *If $\text{irr}_{\mathcal{C}'}(x) \subseteq \text{irr}_{\mathcal{C}}(x)$, then $\text{irr}_{\mathcal{C}'}(x) = \text{irr}_{\mathcal{C}}(x)$ and so $\mathcal{F} = \text{critgen}(x)$.*

Proof. First, we prove that $\text{irr}_{\mathcal{C}'}(x) = \text{irr}_{\mathcal{C}}(x)$ indeed holds. Let $M \in \text{irr}_{\mathcal{C}}(x)$. Then, $M \in \mathcal{C}'$ as $\mathcal{C} \subseteq \mathcal{C}'$ and there exists $M' \in \text{irr}_{\mathcal{C}'}(x)$ such that $M \subseteq M'$. As $\text{irr}_{\mathcal{C}'}(x) \subseteq \text{irr}_{\mathcal{C}}(x)$ and the members of $\text{irr}_{\mathcal{C}}(x)$ are pairwise incomparable, we deduce $M = M'$ and $M \in \text{irr}_{\mathcal{C}'}(x)$.

Then, we argue that $\mathcal{F} = \text{critgen}_{\mathcal{C}}(x)$ by contradiction. Recall that $\mathcal{F} \subseteq \text{critgen}_{\mathcal{C}}(x)$ by definition of $\mathcal{F}$. Hence, suppose that $\mathcal{F} \subset \text{critgen}_{\mathcal{C}}(x)$ and let $A \in \text{critgen}_{\mathcal{C}}(x) \setminus \mathcal{F}$. We first show that $x \notin \phi'(A)$. Again, suppose for contradiction that $x \in \phi'(A)$. Then, applying forward chaining on A with Σ' will reach x . Since $\Sigma' \subseteq \Sigma_{\text{crit}}$ and $A \rightarrow x \notin \Sigma'$, applying the forward chaining on A with $\Sigma_{\text{crit}} \setminus \{A \rightarrow x\}$ will reach x too. Hence, the two IBs $(X, \Sigma_{\text{crit}})$ and $(X, \Sigma_{\text{crit}} \setminus \{A \rightarrow x\})$ are equivalent. However, this contradicts the fact that Σ_{crit} is unit-minimum by Theorem 12. We conclude that $x \notin \phi'(A)$ must hold as expected. Now, $x \notin \phi'(A)$ implies that there exists $M \in \text{irr}_{\mathcal{C}'}(x)$ such that $\phi'(A) \subseteq M$. By assumption, $M \in \text{irr}_{\mathcal{C}}(x)$, so that $A \subseteq \phi(A) \subseteq \phi(M) = M$. However, $x \in \phi(A)$ holds as $A \in \text{critgen}_{\mathcal{C}}(x)$, which entails $x \in M$, a contradiction with $M \in \text{irr}_{\mathcal{C}}(x)$. We conclude that such A cannot exist, and hence that $\mathcal{F} = \text{critgen}_{\mathcal{C}}(x)$, which was to be proved. $\square$

Before concluding to an algorithm for CGE+A·DEC based on ACS+A·ENUM, there is one last subtlety to handle: we need to guarantee that for each $y \in \text{critanc}(x)$, the irreducible closed sets $\text{irr}(y)$ we provide as an input to ACS+A·ENUM along with Σ' indeed correspond to the irreducible closed sets attached to y in $(X, \mathcal{C}')$, i.e., that $\text{irr}_{\mathcal{C}}(y) = \text{irr}_{\mathcal{C}'}(y)$ for each such y .

Lemma 37. *For every $y \in \text{critanc}(x)$, $\text{irr}_{\mathcal{C}}(y) = \text{irr}_{\mathcal{C}'}(y)$.*

Proof. Because attached irreducible closed sets and minimal generators define dual hypergraphs, it is sufficient to show that $\text{mingen}_{\mathcal{C}}(y) = \text{mingen}_{\mathcal{C}'}(y)$. Since $(X, \Sigma_{\text{crit}})$ is acyclic, so is (X, Σ') and $\text{critanc}_{\mathcal{C}}(y) = \text{anc}_{\Sigma'}(y)$ follows by definition of Σ' . We deduce that for every $A \in \text{mingen}_{\mathcal{C}}(x) \cup \text{mingen}_{\mathcal{C}'}(x)$, $A \subseteq \text{critanc}_{\mathcal{C}}(x)$. Again using acyclicity, for each $Z \subseteq \text{critanc}_{\mathcal{C}}(y)$, $y \in \phi(Z)$ if and only if there exists a subset $\Sigma'' = \{A_1 \rightarrow b_1, \dots, A_k \rightarrow b_k\}$ of implications of Σ_{crit} where $A_i \cup \{b_i\} \subseteq \text{critanc}_{\mathcal{C}}(y) \cup \{y\}$ for all $1 \leq i \leq k$ and $y \in \phi''(Z)$, where ϕ'' is the closure operator associated to (X, Σ'') . Given that any such Σ'' satisfies $\Sigma'' \subseteq \Sigma' \subseteq \Sigma_{\text{crit}}$, we deduce that $y \in \phi(Z)$ if and only if $y \in \phi'(Z)$. Therefore, $A \in \text{mingen}_{\mathcal{C}}(x)$ if and only if $A \in \text{mingen}_{\mathcal{C}'}(x)$, which concludes the proof. $\square$

We conclude to the following as a consequence of Lemmas 35, 37 and 36, and derive its corollary by the saturation technique.

Theorem 38. *There is a polynomial-time algorithm solving CGE+A·DEC in acyclic convex geometries whenever there is an incremental-polynomial time algorithm solving ACS+A·ENUM within the same class.*

Proof. Let us assume the existence of an incremental-polynomial time algorithm **A** solving ACS+A·ENUM on acyclic implicational bases, and let $p : \mathbb{N}^2 \rightarrow \mathbb{N}$ be the polynomial such that **A** outputs the i^{th} solution in time $p(N, i)$ on input (X, Σ) , $x \in X$ and $\bigcup_{y \in \text{anc}(x)} \text{irr}(y)$, whose total size is N . Let us further consider an instance $(X, \text{irr}(\mathcal{C}), x, \{\text{critgen}(y) : y \in \text{critanc}(x)\}, \mathcal{F})$ of CGE+A·DEC, and Σ' as defined above and its aggregation Σ'^* .

In order to solve CGE+A·DEC, we start simulating **A** on input (X, Σ'^*) , x and $\{\text{irr}_{\mathcal{C}'}(y) : y \in \text{critanc}(x)\}$. Note that is indeed a valid input of CGE+A·DEC as the ancestor relation induced by Σ'^* on $\text{critanc}(x)$ is critanc and $\text{irr}_{\mathcal{C}'}(y) = \text{irr}_{\mathcal{C}'}(y)$ for all $y \in \text{critanc}(x)$, according to Lemma 37. Each time **A** outputs a set $M \in \text{irr}_{\mathcal{C}'}(x)$, we check whether $M \in \text{irr}_{\mathcal{C}}(x)$.

If not, we stop **A**: there exists $A \in \text{critgen}(x) \setminus \mathcal{F}$ included in M according to Lemma 35. Thanks to Proposition 6, we can retrieve such a A in polynomial time and return it as a counterexample to CGE+A·DEC.

On the other hand, if **A** terminates only having output sets in $\text{irr}_{\mathcal{C}}(x)$, we derive that $\text{irr}_{\mathcal{C}'}(x) \subseteq \text{irr}_{\mathcal{C}}(x)$. According to Lemma 36, it means in particular that $\text{critgen}(x) = \mathcal{F}$. Hence we can answer **yes** for CGE+A·DEC.

Note that our simulation of **A** may output at most $|\text{irr}_{\mathcal{C}}(x)| + 1$ sets before it either terminates, or produces a counterexample. By assumption this takes a total time of $p(N, |\text{irr}_{\mathcal{C}}(x)| + 1)$ which is polynomial in the input size of CGE+A·DEC. $\square$

Remark 39. Notice that the hypothesis on incremental-polynomial time for ACS+A·ENUM is fundamental, as one could not guarantee that we get the counter example early enough if it performed only in output-polynomial time.

Corollary 40. *There is an incremental polynomial-time algorithm solving CGE+A·GEN in acyclic convex geometries whenever there is an incremental polynomial-time algorithm solving ACS+A·ENUM within the same class.*

Note that, in this paper, we did *not* provide an incremental-polynomial time algorithm for ACS+A·ENUM *in general*, but managed to do so for implicational bases having bounded conclusion- or premise-degree. It thus remains to show that our approach still holds for such class of closure systems in order to prove Theorem 2, which we do now.

First, note that in the proof of Theorem 38, given the irreducible closed sets of a closure system $\mathcal{C}$, we need the existence of an incremental-polynomial time algorithm A solving ACS+A·ENUM on instances Σ'^* obtained by considering an aggregated subset of its critical base. Thus, in order to derive an algorithm solving CGE+A·DEC in acyclic closure systems of bounded (premise- or conclusion-) degree, we need to argue that such an aggregation Σ'^* represents an acyclic closure system that also has bounded (premise- or conclusion-) degree. This follows from the fact that the aggregated critical base of an acyclic convex geometry with bounded degree has bounded degree by Corollary 16, and that any subset Σ' of the critical base Σ_{crit} of an acyclic convex geometry defines an acyclic convex geometry $\mathcal{C}'$.

Now, concerning the complexity, we note that the number of critical generators is bounded by $d \cdot |X|$ if the aggregated critical base has (premise- or conclusion-) degree d . In addition, any incremental-polynomial time algorithm performs in (input) polynomial time on an instance with polynomially many solutions. We derive Theorem 2, that we restate here, as a corollary of Theorem 1 together with these observations, Lemma 15, Corollary 40, and Theorem 20.

Theorem 2. *There is a polynomial-time algorithm constructing the critical base of an acyclic convex geometry of bounded premise- or conclusion-degree given by its irreducible closed sets.*

8 Discussion

In this paper, we provided incremental-polynomial time algorithms for the two translation tasks ICS·ENUM and CB·GEN in acyclic convex geometries of bounded (premise- or conclusion-) degree. We note that the delay of our algorithms is intrinsically not polynomial since they rely on a sequential top-down procedure which reduces the enumeration to an auxiliary enumeration problem whose instance grows with the total number of solutions output so far. This suggests the following question.

Question 41. Can ICS·ENUM and CB·GEN be solved with polynomial delay for acyclic convex geometries of bounded degree?

Toward this direction, we now argue that the classic flashlight search framework may not be used in a straightforward way to obtain a positive answer to this question. In the *flashlight search technique*, the goal is to construct solutions one element at a time, according to a linear order $x_1, \dots, x_n$ on the ground set, and deciding at each step whether x_i should or should not be included in the solution. For this approach to be tractable, the framework requires to solve the so-called “extension problem” in polynomial time: given two disjoint subsets K, F of the ground set, decide whether there exists a solution




Figure 7: An illustration of the construction of Theorem 42. Black edges correspond to the incidence bipartite graphs of ψ and are represented for intuition only: they are not part of the ICS·EXT instance. Rather, to each such edge corresponds a implication having an element of U as conclusion; for readability, only two of these implications are represented.

including K and avoiding F . We refer to e.g., [14, 43] for a more detailed description of this folklore technique. More formally, this technique reduces the enumeration to the following decision problem.

Irreducible Closed Set Extension (ICS·EXT)

Input: An implicational base (X, Σ) and $K, F \subseteq X$.

Question: Does there exist $M \in \text{irr}(\mathcal{C})$ such that $K \subseteq M$ and $M \cap F = \emptyset$?

We show that this problem is NP-complete even for acyclic implicational bases of bounded degree, even for $K = \emptyset$, suggesting that this framework may not be used in a straightforward way to improve our results.

Theorem 42. *The problem ICS·EXT is NP-complete even for acyclic implicational bases with degree, premise-degree, conclusion-degree and dimension (i.e., the maximum size of a premise) simultaneously bounded by 4, 4, 2 and 2, respectively.*

Proof. A polynomial-size certificate for NP is a subset $M \subseteq X$ for which it is easily checked, in polynomial time, whether $M \in \text{irr}(\mathcal{C})$, $K \subseteq M$, and $M \cap F = \emptyset$.

Regarding completeness, it is known that SAT is NP-complete even for instances where each clause contains at most 3 literals, and each variable appears (negatively or positively) in at most 3 clauses [31, Problem L01]. Let us consider an instance ψ of this problem, with variables $v_1, \dots, v_n$ and clauses $C_1, \dots, C_m$. For clause C_j , $1 \leq j \leq m$, let us denote by $\ell_j^1, \dots, \ell_j^k$ its $k \leq 3$ literals, i.e., ℓ_j^i is either a variable or its negation.

We reduce ψ to an instance $((X, \Sigma), \emptyset, F)$ of ICS·EXT as follows. See Figure 7 for an illustration. Let us first describe the ground set X . First, for every variable v_i , $1 \leq i \leq n$ we create three elements x_i, y_i, z_i with x_i representing v_i , y_i representing $\bar{v}_i$, and z_i being a gadget representing a conflicting selection of both v_i and $\bar{v}_i$ in a non-valid

assignment. Then, for every clause C_j , we create one element c_j representing the clause, and $k = |C_j| \leq 3$ additional elements $u_j^1, \dots, u_j^k$ representing the literals of the clause. Finally, we add another gadget element t .

We now describe Σ . For every $1 \leq i \leq n$, we create an implication $\{x_i, y_i\} \rightarrow z_i$. For every clause C_j , and each variable v_i appearing in position k in C_j , $1 \leq k \leq |C_j|$, we add the implication $\{c_j, x_i\} \rightarrow u_j^k$ if it appears positively, i.e., if $\ell_j^k = v_i$, and $\{c_j, y_i\} \rightarrow u_j^k$ if it appears negatively, i.e., if $\ell_j^k = \bar{v}_i$. Let $U := \bigcup_{1 \leq j \leq m} \{u_j^k : 1 \leq k \leq |C_j|\}$ and $u_1, \dots, u_p$ be an arbitrary labeling of U . We complete the construction of Σ by adding the implications $z_i \rightarrow z_{i+1}$ for every $1 \leq i < n$, the implications $u_i \rightarrow u_{i+1}$ for every $1 \leq i < p$, and the two implications $z_n \rightarrow t$ and $u_p \rightarrow t$. Notice that the premise of any implication is a minimal generator of t by construction.

Finally, we put $F = \{c_1, \dots, c_m\} \cup \{z_1, \dots, z_n\} \cup \{u_1, \dots, u_p\}$.

Let us consider a positive instance of ICS-EXT, i.e., some set $M \in \text{irr}(\mathcal{C})$ such that $M \cap F = \emptyset$. Let $a \in X$ such that $M \in \text{irr}(a)$. As $z_n \in F$, it must be an ancestor of a . But note that z_n implies a single element $: t$. Thus, $a = t$ hence $M \in \text{irr}(t)$.

Consider the assignment defined by

$$v_i \mapsto \begin{cases} 1 & \text{if } x_i \in M, \\ 0 & \text{if } y_i \in M. \end{cases}$$

Note that this assignment is valid since $\{x_i, y_i\}$ is a minimal generator of t , hence not both x_i and y_i may belong to M . Let us argue that it satisfies ψ . For every $1 \leq j \leq m$, one variable x_i or y_i representing ℓ_j^k , $1 \leq k \leq |C_j|$ must be in M , as otherwise we can add c_j to M without implying t , a contradiction to $M \in \text{irr}(t)$. Hence each clause is satisfied. This concludes the first direction.

Conversely, suppose that ψ is satisfiable. Consider a satisfying assignment and the set M defined by taking x_i if v_i is set to true, and y_i otherwise, for every $1 \leq i \leq m$. Then $t \notin M$ and M is closed, as it does not contain any premise. Let us argue that it is maximal with that property, hence that $M \in \text{irr}(t) \subseteq \text{irr}(\mathcal{C})$. Clearly, none of $z_1, \dots, z_n$ can be added to M , and similarly for the elements of U without adding t . Suppose toward a contradiction that an element c_j , $1 \leq j \leq m$ can be added to M . Since ψ is satisfiable, M contains either x_i or y_i corresponding to a variable v_i appearing either positively or negatively in C_j . Hence M triggers one of the corresponding implication $\{c_j, \alpha\} \rightarrow u$ for $\alpha = x_i$ or $\alpha = y_i$ and $u \in U$. We derive that M implies t , a contradiction. This concludes the proof of NP-completeness.

Lastly, notice that the implications in Σ all have a premise of size 1 or 2 and that every element $x \in X$ is in at most 2 conclusions (exactly 2 for t , z_i and u_i with $i \geq 2$). Moreover, as a literal participates in at most 3 clauses, each of x_i, y_i , $1 \leq i \leq n$ has total- and premise-degree at most 4, and the other elements have degree at most 3, hence the result. $\square$

Among other directions, we wonder whether our results extend to broader classes of closure systems or to parameters that generalize the degree. As a first step towards this avenue, we show that when we simply drop acyclicity, then the degree is helpless

as far as ICS·ENUM is concerned. More precisely, we show in the next theorem that an algorithm solving ICS·ENUM for closure systems with bounded degree can be used to solve ICS·ENUM in the general case.

Theorem 43. *There is an output-polynomial time algorithm solving ICS·ENUM in general whenever there is an output-polynomial time algorithm solving ICS·ENUM on closure systems with bounded degree.*

Proof. Let (X, Σ) be an instance of ICS·ENUM with $X = \{x_1, \dots, x_n\}$ and $\Sigma = \{A_1 \rightarrow B_1, \dots, A_m \rightarrow B_m\}$. We build an implicational base (X', Σ') of polynomial size and that has bounded degree. We first describe X' . For each $x_i \in X$ and each implication $A_j \rightarrow B_j \in \Sigma$, we create a gadget element x_i^j if $x_i \in A_j \cup B_j$ and we call X_i the set of resulting elements associated with x_i , i.e., $X_i = \{x_i^j : A_j \rightarrow B_j \in \Sigma, x_i \in A_j \cup B_j\}$. Recall that we consider that B_j and A_j are disjoint, so that an element can occur only once in an implication. Therefore, $|X_i| = \deg(x_i)$. We put $X' = \bigcup_{i=1}^n X_i$. We now describe Σ' . For each implication $A_j \rightarrow B_j$ of Σ we create an implication where each element $x_i \in A_j \cup B_j$ is replaced by its corresponding gadget x_i^j , that is, we build the implication $A'_j \rightarrow B'_j$ where $A'_j = \{x_i^j : x_i \in A_j\}$ and $B'_j = \{x_i^j : x_i \in B_j\}$. Finally, for each i , we consider an arbitrary labeling $y_1, \dots, y_k$ of the elements of X_i and make all of them equivalent by putting in Σ' the cycle of implications $\{y_1 \rightarrow y_2, \dots, y_{k-1} \rightarrow y_k, y_k \rightarrow y_1\}$. This completes the description of Σ' . Observe that (X', Σ') can be built in polynomial time in the size of (X, Σ) and that each element in (X', Σ') has degree at most 3. The reduction is illustrated on an example in Figure 8.

Note that the mapping $\psi: \mathcal{C} \rightarrow \mathcal{C}', C \mapsto \bigcup_{x_i \in C} X_i$ is a polynomial-time computable bijection between the two closure systems $(X, \mathcal{C})$ and $(X', \mathcal{C}')$. Hence, with this bijection, an output-polynomial time algorithm solving ICS·ENUM for (X', Σ') can be used to solve ICS·ENUM for (X, Σ) within the same time bound. $\square$




Figure 8: On the left, the IB $\Sigma = \{13 \rightarrow 4, 12 \rightarrow 4, 45 \rightarrow 2\}$. On the right, we apply the reduction of Theorem 43. For readability, we repeat the label of each element in each implication, and we highlight in blue the implications of Σ' encoding the equivalence between each repetition of an element.

Theorem 43 thus suggests to investigate instead more constrained classes of closure systems. The top-down strategy we used for ICS·ENUM requires to devise an ordering on the elements of the ground set that provides a meaningful ancestor relation. Lower-bounded closure systems [2, 30] are a well-known generalization of acyclic convex geometries that appears as a natural candidate to investigate. They are characterized by an acyclic D -relation, a subset of the δ -relation (see Section 3), which suggests that a topological ordering of the elements can be set and used as in our approach. Besides, closure systems arising from graphs, also known as graph convexities [28], may be a source of closure systems where natural orders may be found based on the properties on the underlying graphs. We note, however, that these closure systems are in general not acyclic.

The algorithm we rely on to solve MIB·GEN, through CB·GEN, seems slightly more demanding. On top of a natural ordering of the vertices, it requires the critical base not only to be well defined, but also to be valid. Critical generators are initially defined in convex geometries using extreme points, but in arbitrary closure systems, closed sets need not have extreme points, and critical generators may thus not exist. However, as observed by different authors [46, 56], the critical generators of an element could instead be defined as its minimal generators that have the smallest closure inclusion-wise. This definition makes critical generators of an element exist in any closure system as long as this element admits minimal generators, and it coincides with the original definition in convex geometries. Nevertheless, even under this definition, the corresponding critical base of an arbitrary closure system needs not to be a valid IB (as it already fails in convex geometries), which prevents the strategy of finding critical generators to answer MIB·GEN from applying in general. Identifying classes of closure systems, or convex geometries, where the critical base is valid is thus an interesting question to investigate in order to study the tractability of MIB·GEN.

A possible alternative is to use the lattice theoretic counterpart of critical generators, being E -generators [2, 30]. Still, the corresponding E -base, much as the critical base, does not always constitute a valid IB of its closure system. However, in the class of lower bounded closure systems, the E -base turns out to be valid. Together with the fact that lower bounded closure systems are D -acyclic, this suggests our approach to MIB·GEN based on the problem CB·GEN, adapted to E -generators, might succeed.

As for parameters, an analogue of the notion of degeneracy coming from (hyper)graph theory would be a natural candidate for the generalization of the degree. Degeneracy in implicational bases can be defined as the smallest left-degree among all possible left-to-right ordering of the vertices. Dimension, also known as arity in the language of abstract convexity [53], is another intriguing parameter incomparable to degree. It is the least integer k such that the closure system admits an IB with premises of size at most k , which makes it particularly interesting for the translation task as MIS·ENUM is tractable for hypergraphs of bounded dimension. Other parameters coming from convexity such as the Caratheodory number—for which some tractable algorithms can be obtained [56]—, the Radon number, or the Helly number might also be avenues for future research.

Finally, we note that our algorithms for ICS·ENUM have an XP dependence on the degree k , i.e., it runs in $N^{f(k)}$ time for some computable function f where N is the

size of the input plus the output. This is not due to the top-down approach, that is, Theorem 19, which preserves FPT, i.e., running times of the form $f(k) \cdot N^{O(1)}$ for some computable function. This is neither due to the computation of minimal transversals in Theorem 22, as this can be achieved with FPT delay parameterized by the number of edges [27]. Rather, the bottleneck to FPT time is the guess of irreducible selections from Lemma 21. The same arises in Lemma 28 for the premise-degree. Consequently, and since our algorithm for CB·GEN relies on the ones we devised for ICS·ENUM, they also have an XP dependence on the parameter. Thus, it is natural to ask whether any of these algorithms can be improved to run in FPT times parameterized by the degree.

Note added after revision. During the revision of this paper, Mikhail Babin announced in the preprint [5] the co-NP-completeness of the decision problem associated to ICS·ENUM, even for acyclic implicational bases of dimension 3. This implies that, unless $P = NP$, ICS·ENUM cannot be solved in output-polynomial time for acyclic convex geometries of dimension at most 3, answering a question we asked above. Note that this hardness was not even known for general closure systems. Moreover, it can be seen that this reduction shows the same hardness for the second translation problem MIB·GEN, even in the case of acyclic convex geometries.

Acknowledgments. The authors have been supported by the ANR project PARADUAL (ANR-24-CE48-0610-01).

References

- [1] Kira V. Adaricheva, Giuseppe F. Italiano, Hans Kleine Büning, and György Turán. Horn formulas, directed hypergraphs, lattices and closure systems: related formalisms and applications (Dagstuhl Seminar 14201). *Dagstuhl Reports*, 4(5):1–26, 2014.
- [2] Kira V. Adaricheva, James B. Nation, and Robert Rand. Ordered direct implicational basis of a finite closure system. *Discrete Applied Mathematics*, 161(6):707–723, 2013.
- [3] Giorgio Ausiello, Alessandro D’Atri, and Domenico Saccà. Minimal representation of directed hypergraphs. *SIAM Journal on Computing*, 15(2):418–431, 1986.
- [4] David Avis and Komei Fukuda. Reverse search for enumeration. *Discrete Applied Mathematics*, 65(1-3):21–46, 1996.
- [5] Mikhail Babin. Completeness of Canonical Closure Representations Is coNP-Complete. [arXiv:2607.18534](https://arxiv.org/abs/2607.18534), 2026.
- [6] Mikhail A. Babin and Sergei O. Kuznetsov. Computing premises of a minimal cover of functional dependencies is intractable. *Discrete Applied Mathematics*, 161(6):742–749, 2013.
- [7] Mikhail A. Babin and Sergei O. Kuznetsov. Dualization in lattices given by ordered sets of irreducibles. *Theoretical Computer Science*, 658:316–326, 2017.
- [8] Laurent Beaudou, Arnaud Mary, and Lhouari Nourine. Algorithms for k -meet-semidistributive lattices. *Theoretical Computer Science*, 658:391–398, 2017.

- [9] Claude Berge. *Hypergraphs: combinatorics of finite sets*, volume 45. Elsevier, 1984.
- [10] Karell Bertet, Christophe Demko, Jean-François Viaud, and Clément Guérin. Lattices, closures systems and implication bases: a survey of structural aspects and algorithms. *Theoretical Computer Science*, 743:93–109, 2018.
- [11] Karell Bertet and Bernard Monjardet. The multiple facets of the canonical direct unit implicational basis. *Theoretical Computer Science*, 411(22-24):2155–2166, 2010.
- [12] Endre Boros, Ondřej Čepek, and Alexander Kogan. Horn minimization by iterative decomposition. *Annals of Mathematics and Artificial Intelligence*, 23:321–343, 1998.
- [13] Endre Boros, Ondřej Čepek, Alexander Kogan, and Petr Kučera. A subclass of Horn CNFs optimally compressible in polynomial time. *Annals of Mathematics and Artificial Intelligence*, 57(3-4):249–291, 2009.
- [14] Endre Boros, Khaled Elbassioni, and Vladimir Gurvich. Algorithms for generating minimal blockers of perfect matchings in bipartite graphs and related problems. In *European Symposium on Algorithms*, pages 122–133. Springer, 2004.
- [15] Endre Boros, Khaled Elbassioni, Vladimir Gurvich, and Leonid Khachiyan. Algorithms for enumerating circuits in matroids. In *International Symposium on Algorithms and Computation*, pages 485–494. Springer, 2003.
- [16] Sara Cohen, Benny Kimelfeld, and Yehoshua Sagiv. Generating all maximal induced subgraphs for hereditary and connected-hereditary graph properties. *Journal of Computer and System Sciences*, 74(7):1147–1159, 2008.
- [17] Alessio Conte, Roberto Grossi, Andrea Marino, Takeaki Uno, and Luca Versari. Proximity search for maximal subgraph enumeration. *SIAM Journal on Computing*, 51(5):1580–1625, 2022.
- [18] Yves Crama and Peter L. Hammer. *Boolean functions: Theory, algorithms, and applications*. Cambridge University Press, 2011.
- [19] Brian A. Davey and Hilary A. Priestley. *Introduction to lattices and order*. Cambridge university press, 2002.
- [20] Oscar Defrain and Lhouari Nourine. Dualization in lattices given by implicational bases. *Theoretical Computer Science*, 814:169–176, 2020.
- [21] Oscar Defrain, Lhouari Nourine, and Simon Vilmin. Translating between the representations of a ranked convex geometry. *Discrete Mathematics*, 344(7):112399, 2021.
- [22] Brenda L. Dietrich. A circuit set characterization of antimatroids. *Journal of Combinatorial Theory, Series B*, 43(3):314–321, 1987.
- [23] Felix Distel and Barış Sertkaya. On the complexity of enumerating pseudo-intents. *Discrete Applied Mathematics*, 159(6):450–466, 2011.
- [24] Paul H. Edelman and Robert E. Jamison. The theory of convex geometries. *Geometriae dedicata*, 19(3):247–270, 1985.
- [25] Thomas Eiter, Kazuhisa Makino, and Georg Gottlob. Computational aspects of monotone dualization: A brief survey. *Discrete Applied Mathematics*, 156(11):2035–2049, 2008.

- [26] Khaled Elbassioni. A polynomial delay algorithm for generating connected induced subgraphs of a given cardinality. *Journal of Graph Algorithms and Applications*, 19(1):273–280, 2015.
- [27] Khaled Elbassioni, Matthias Hagen, and Imran Rauf. Some fixed-parameter tractable classes of hypergraph duality and related problems. In *Parameterized and Exact Computation: Third International Workshop, IWPEC 2008, Victoria, Canada, May 14-16, 2008. Proceedings 3*, pages 91–102. Springer, 2008.
- [28] Martin Farber and Robert E. Jamison. Convexity in graphs and hypergraphs. *SIAM Journal on Algebraic Discrete Methods*, 7(3):433–444, 1986.
- [29] Michael L. Fredman and Leonid Khachiyan. On the complexity of dualization of monotone disjunctive normal forms. *Journal of Algorithms*, 21(3):618–628, 1996.
- [30] Ralph Freese, Jaroslav Ježek, and James B. Nation. *Free Lattices*, volume 42 of *Mathematical Surveys and Monographs series*. American Mathematical Society, 1995.
- [31] Michael R. Garey and David S. Johnson. *Computers and intractability*, volume 174. San Francisco: freeman, 1979.
- [32] Jean-Louis Guigues and Vincent Duquenne. Familles minimales d’implications informatives résultant d’un tableau de données binaires. *Mathématiques et sciences humaines*, 95:5–18, 1986.
- [33] Peter L. Hammer and Alexander Kogan. Quasi-acyclic propositional Horn knowledge bases: Optimal compression. *IEEE Transactions on knowledge and data engineering*, 7(5):751–762, 1995.
- [34] David S. Johnson, Mihalis Yannakakis, and Christos H. Papadimitriou. On generating all maximal independent sets. *Information Processing Letters*, 27(3):119–123, 1988.
- [35] Henry A. Kautz, Michael J. Kearns, and Bart Selman. Reasoning with characteristic models. In *Proceedings of the Eleventh National Conference on Artificial Intelligence, AAAI’93*, pages 34—39. AAAI Press, 1993.
- [36] Dimitris J. Kavvadias, Martha Sideri, and Elias C. Stavropoulos. Generating all maximal models of a Boolean expression. *Information Processing Letters*, 74(3-4):157–162, 2000.
- [37] Leonid Khachiyan, Endre Boros, Khaled Elbassioni, and Vladimir Gurvich. On enumerating minimal dicuts and strongly connected subgraphs. *Algorithmica*, 50(1):159–172, 2008.
- [38] Roni Khardon. Translating between Horn representations and their characteristic models. *Journal of Artificial Intelligence Research*, 3:349–372, 1995.
- [39] Bernhard Korte and László Lovász. Shelling Structures, Convexity, and a Happy End. In *Graph Theory and Combinatorics: Proceedings of the Cambridge Combinatorial Conference in Honour of Paul Erdős*, pages 219–232. Academic Press, 1984.
- [40] Bernhard Korte, László Lovász, and Rainer Schrader. *Greedoids*, volume 4 of *Algorithms and Combinatorics series*. Springer Science & Business Media, 2012.

- [41] David Maier. *The Theory of Relational Databases*. Computer software engineering series. Computer Science Press, 1983.
- [42] Heikki Mannila and Kari-Jouko R  ih  . *The design of relational databases*. Addison-Wesley Longman Publishing Co., Inc., 1992.
- [43] Arnaud Mary and Yann Strozecki. Efficient enumeration of solutions produced by closure operations. *Discrete Mathematics & Theoretical Computer Science*, 21(Discrete Algorithms), 2019.
- [44] Bernard Monjardet. A use for frequently rediscovering a concept. *Order*, 1(4):415–417, 1985.
- [45] Bernard Monjardet and Nathalie Caspard. On a dependence relation in finite lattices. *Discrete Mathematics*, 165–166:497–505, 1997.
- [46] Masataka Nakamura and Kenji Kashiwabara. The prime stems of rooted circuits of closure spaces and minimum implicational bases. *The Electronic Journal of Combinatorics*, 20(1):#P22, 2013.
- [47] Lhouari Nourine and Simon Vilmin. Hierarchical decompositions of implicational bases for the enumeration of meet-irreducible elements. *Theoretical Computer Science*, 969:114030, 2023.
- [48] Sebastian Rudolph. Succinctness and tractability of closure operator representations. *Theoretical Computer Science*, 658:327–345, 2017.
- [49] Luigi Santocanale and Friedrich Wehrung. Lattices of regular closed subsets of closure spaces. *International Journal of Algebra and Computation*, 24(07):969–1030, 2014.
- [50] Robert C. Shock. Computing the minimum cover of functional dependencies. *Information Processing Letters*, 22(3):157–159, 1986.
- [51] Yann Strozecki. Enumeration complexity. *Bulletin of EATCS*, 1(129), 2019.
- [52] Shuji Tsukiyama, Mikio Ide, Hiromu Ariyoshi, and Isao Shirakawa. A new algorithm for generating all the maximal independent sets. *SIAM Journal on Computing*, 6(3):505–517, 1977.
- [53] Marcel van de Vel. *Theory of convex structures*. Elsevier, 1993.
- [54] Marcel Wild. A theory of finite closure spaces based on implications. *Advances in Mathematics*, 108(1):118–139, 1994.
- [55] Marcel Wild. Optimal implicational bases for finite modular lattices. *Quaestiones Mathematicae*, 23(2):153–161, 2000.
- [56] Marcel Wild. The joy of implications, aka pure Horn formulas: mainly a survey. *Theoretical Computer Science*, 658:264–292, 2017.
- [57] Hiyori Yoshikawa, Hiroshi Hirai, and Kazuhisa Makino. A representation of antimatroids by horn rules and its application to educational systems. *Journal of Mathematical Psychology*, 77:82–93, 2017.
- [58] Kaiqiang Yu, Cheng Long, Shengxin Liu, and Da Yan. Efficient algorithms for maximal k -biplex enumeration. In *Proceedings of the 2022 International Conference on Management of Data*, pages 860–873, 2022.