Journal of Computer and System Sciences 127 (2022) 1–28

Contents lists available at ScienceDirect

Journal of Computer and System Sciences

www.elsevier.com/locate/jcss

Unlabeled sample compression schemes and corner peelings for ample and maximum classes ☆

Jérémie Chalopin [a,*], Victor Chepoi [a], Shay Moran [b], Manfred K. Warmuth [c,d]

[a] CNRS, Aix-Marseille Université, Université de Toulon, LIS, Marseille, France
[b] Department of Mathematics, Technion and Google Research, Israel
[c] Google Brain, Mountain View, CA, USA
[d] Formerly Computer Science Department, University of California, Santa Cruz, USA

ARTICLE INFO

Article history:
Received 17 December 2020
Received in revised form 17 January 2022
Accepted 23 January 2022
Available online 3 February 2022

Keywords:
VC-dimension
Sample compression
Sauer-Shelah-Perles Lemma
Sandwich Lemma
Maximum class
Ample class
Extremal class
Corner peeling
Unique sink orientation

ABSTRACT

We examine connections between combinatorial notions that arise in machine learning and topological notions in cubical/simplicial geometry. These connections enable to export results from geometry to machine learning. Our first main result is based on a geometric construction by Tracy Hall (2004) [20] of a partial shelling of the cross-polytope which can not be extended. From it, we derive a maximum class of VC dimension 3 without corners. This refutes several previous works in machine learning. In particular, it implies that the previous constructions of optimal unlabeled sample compression schemes for maximum classes are erroneous. On the positive side we present a new construction of an optimal unlabeled sample compression scheme for maximum classes. We leave as open whether our unlabeled sample compression scheme extends to ample classes, which generalize maximum classes. Towards resolving this question, we provide a geometric characterization in terms of unique sink orientations of the associated 1-inclusion graph.

1. Introduction

The Sauer-Shelah-Perles Lemma [38,40,43] is arguably the most basic fact in VC theory; it asserts that any class $C \subseteq \{0,1\}^n$ satisfies $|C| \le \binom{n}{\le d}$, where $d = \text{VC-dim}(C)$. A beautiful generalization of Sauer-Shelah-Perles's inequality asserts that $|C| \le |\overline{X}(C)|$, where $\overline{X}(C)$ is the family of subsets that are shattered by C.[1] The latter inequality is a part of the Sandwich Lemma [3,8,13,33], which also provides a lower bound for $|C|$ (and thus "sandwiches" $|C|$) in terms of the number of its *strongly shattered subsets* (see Section 2). A class C is called *maximum/ample* if the Sauer-Shelah-Perles/Sandwich upper bounds are tight (respectively). Every maximum class is ample, but not vice versa.

Maximum classes were studied mostly in discrete geometry and machine learning, e.g. [15,16,18,23,45]. The history of ample classes is more interesting as they were discovered independently by several works in disparate contexts [3,5,8,13,24,29,46]. Consequently, they received different names such as lopsided classes [24], extremal classes [8,29], and ample classes [5,13]. Lawrence [24] was the first to define them for the investigation of the possible sign patterns realized by

☆ An extended abstract [11] of this paper has appeared in the proceedings of ICALP 2019.
* Corresponding author.
E-mail addresses: jeremie.chalopin@lis-lab.fr (J. Chalopin), victor.chepoi@lis-lab.fr (V. Chepoi), smoran@technion.ac.il (S. Moran), manfred@gmail.com (M.K. Warmuth).

[1] Note that this inequality indeed implies the Sauer-Shelah-Perles Lemma, since $\overline{X}(C) \subseteq \binom{[n]}{\le d}$.

https://doi.org/10.1016/j.jcss.2022.01.003
0022-0000/

points of a convex set of $\mathbb{R}^d$. Interestingly, Lawrence's definition of these classes does not use the notion of shattering nor the Sandwich Lemma. In this context, these classes were discovered by Bollobás and Radcliffe [8] and Bandelt et al. [5], and the equivalence between the two definitions appears in [5]. Ample classes admit a multitude of combinatorial and geometric characterizations [5,6,8,24] and comprise many natural examples arising from discrete geometry, combinatorics, graph theory, and geometry of groups [5,24].

1.1. Main results

1.1.1. Corner peelings

A *corner* in an ample class C is any concept $c \in C$ that belongs to a unique maximal cube of C (equivalently, c is a corner if $C \setminus \{c\}$ is also ample, see Lemma 4.1). A sequence of corner removals leading to a single concept is called a *corner peeling*; corner peeling is a strong version of *collapsibility*. Wiedemann [46] and independently Chepoi (unpublished, 1996) asked whether every ample class has a corner. The machine learning community studied this question independently in the context of *sample compression schemes* for maximum classes: Rubinstein and Rubinstein [35] showed that corner peelings lead to optimal *unlabeled sample compression schemes (USCS)*.

In Theorem 4.5 we refute this conjecture. The crux of the proof is an equivalence between corner peelings and partial shellings of the cross-polytope. This equivalence translates the question whether corners always exist to the question whether partial shellings can always be extended. The latter was an open question in Ziegler's book on polytopes [49], and was resolved in Tracy Hall's PhD thesis [20] where an interesting counterexample is presented. The ample class resulting from Hall's construction yields a maximum class without corners.

1.1.2. Sample compression

Sample compression is a powerful technique to derive generalization bounds in statistical learning. Littlestone and Warmuth [25] introduced it and asked if every class of VC-dimension $d < \infty$ has a sample compression scheme of a finite size. This question was later relaxed by Floyd and Warmuth [16,44] to *whether a sample compression scheme of size $O(d)$ exists*. The first question was recently resolved by [31] who exhibited an $\exp(d)$ sample compression scheme. The second question however remains one of the oldest open problems in machine learning (for more background we refer the reader to [30] and the books [39,47]).

Rubinstein and Rubinstein [35, Theorem 16] showed that the existence of a corner peeling for a maximum class C implies a *representation map* for C (see Section 3 for a definition), which is known to yield an optimal unlabeled sample compression scheme of size VC-dim(C) [23].[2] They claim, using an interesting topological approach, that maximum classes admit corner peelings. Unfortunately, our Theorem 4.5 shows that this does not hold.

While our Theorem 4.5 rules out the program of deriving representation maps from corner peelings, in Theorem 5.1 we provide an alternative derivation of representation maps for maximum classes and therefore also of unlabeled sample compression schemes for them.

1.1.3. Sample compression and unique sink orientations

We next turn to construction of representation maps for ample classes. In Theorem 6.8 we present geometric characterizations of such maps via *unique sink orientations*: an orientation of the edges of a cube B is a *unique sink orientation (USO)* if any subcube $B' \subseteq B$ has a unique sink. Szabó and Welzl [41] showed that any USO of B leads to a representation map for B. We extend this bijection to ample classes C by proving that representation maps are equivalent to orientations r of C such that (i) r is a USO on each subcube $B \subseteq C$, and (ii) for each $c \in C$ the edges outgoing from c belong to a subcube $B \subseteq C$. We further show that any ample class admits orientations satisfying each one of those conditions. However, *the question whether all ample classes admit representation maps remains open.*

1.1.4. Implications on previous works

Our Theorem 4.5 establishes the existence of maximum classes without any corners, thus countering several previous results in machine learning:

- Rubinstein and Rubinstein [35, Theorem 32] showed that any maximum class can be represented by a simple arrangement of piecewise-linear hyperplanes. In [35, Theorem 39], they claim that sweeping such an arrangement leads to a corner peeling of the corresponding maximum class. This is unfortunately false, as witnessed by Theorem 4.5.
- Kuzmin and Warmuth [23] constructed unlabeled sample compression schemes for maximum classes based on the presumed uniqueness of a certain matching (their Theorem 10). This theorem is wrong (as explained in Section 4.2) as it implies the existence of corners and Hall's counterexample does not have corners. However their conclusion is correct: In our Theorem 5.1 we show that such unlabeled compression schemes always exist based on a different construction and proof method.
- Theorem 3 by Samei, Yang, and Zilles [37] is built on a generalization of Theorem 10 from [23] to the multiclass case which is also incorrect.

[2] Pálvölgyi and Tardos [34] recently exhibited a (non-ample) class C with no USCS of size VC-dim(C).

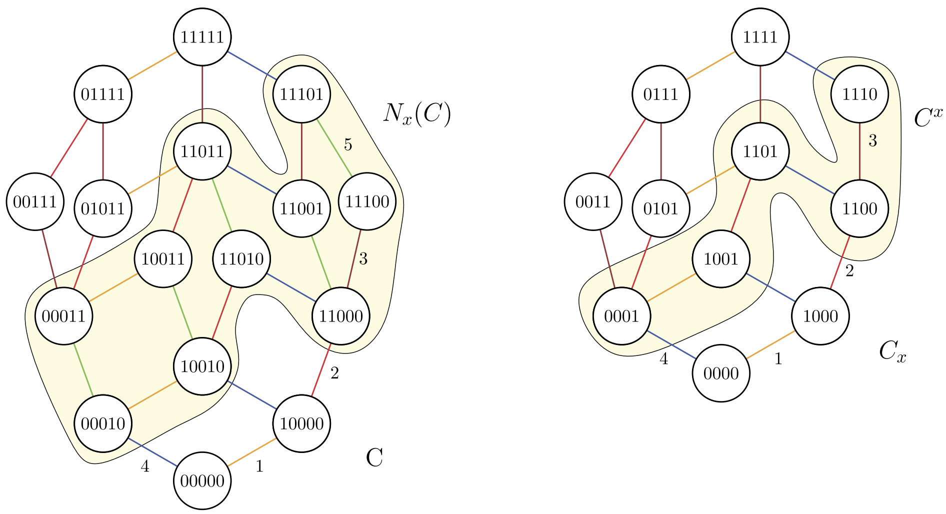

Fig. 1. A 2-dimensional maximum class $C \subseteq 2^{\{1,2,3,4,5\}}$ on the left and the restriction C_x for $x = 5$ on the right. The reduction C^x corresponds to the restriction of the carrier $N_x(C)$.

- Theorem 26 by Doliwa et al. [12] uses the result in [35] to show that the Recursive Teaching Dimension (RTD) of maximum classes equals to their VC dimension. However the VC dimension 3 maximum class from Theorem 4.5 has RTD at least 4. It remains open whether the RTD of every maximum class C is bounded by $O(\text{VC-dim}(C))$.

1.1.5. An optimal proper PAC learner for maximum classes

In a recent work, Bousquet, Hanneke, Moran, and Zhivotovskiy [9] showed that a special type of sample compression schemes, termed *stable compression schemes*, achieve the optimal learning rate in PAC learning. They further noticed that any sample compression scheme which is defined by a representation map is stable. Thus, using the compression scheme constructed in this paper, Bousquet et al. conclude that every maximum class can be properly learned by an algorithm achieving the optimal learning rate.

1.2. Organization

Section 2 presents the main definitions and notations. Section 3 reviews characterizations of ample/maximum classes and characteristic examples. Section 4 demonstrates the existence of the maximum class C_H without corners. Section 5 establishes the existence of representation maps for maximum classes. Section 6 establishes a bijection between representation maps and unique sink orientations for ample classes.

2. Preliminaries

A *concept class* C is a set of subsets (concepts) of a finite ground set U which is called the *domain* of C and denoted $\text{dom}(C)$. We sometimes treat the concepts as characteristic functions rather than subsets. The *support* (or *dimension set*) $\text{supp}(C)$ of C is the set $\{x \in U : x \in c' \setminus c'' \text{ for some } c', c'' \in C\}$.

Let C be a concept class of 2^U. The *complement* of C is $C^* := 2^U \setminus C$. The *twisting* of C with respect to $Y \subseteq U$ is the concept class $C\Delta Y = \{c\Delta Y : c \in C\}$. The *restriction* of a concept $c \in C$ on $Y \subseteq U$ is the concept $c|Y = c \cap Y$. The *restriction* of C on $Y \subseteq U$ is the class $C|Y = \{c|Y : c \in C\}$ whose domain is Y. We use C_Y as shorthand for $C|(U \setminus Y)$; in particular, we write C_x for $C_{\{x\}}$ (see Fig. 1 for an example), and c_x for $c|(U \setminus \{x\})$ for $c \in C$ (note that $c_x \in C_x$). A concept class $B \subseteq 2^U$ is a *cube* if there exists $Y \subseteq U$ such that $B|Y = 2^Y$ and B_Y contains a single concept (denoted by $\text{tag}(B)$). Note that $\text{supp}(B) = Y$ and therefore we say that B is a Y*-cube*; $|Y|$ is called the *dimension* $\dim(B)$ of B. Two cubes B, B' with the same support are called *parallel cubes*. A cube B is *maximal* if there is no cube B' such that $B \subsetneq B'$.

Let Q_n denote the n-dimensional cube where $n = |U|$; $c, c' \in Q_n$ are called adjacent if the symmetric difference $c\Delta c'$ is of size 1. The *1-inclusion graph* of C is the subgraph $G(C)$ of Q_n induced by the vertex-set C when the concepts of C are identified with the corresponding vertices of Q_n. Any cube $B \subseteq C$ is called *a cube of* C. The *cube complex* of C is the set $Q(C) = \{B : B \text{ is a cube of } C\}$. The cubes of C are called the *faces* of $Q(C)$ and the maximal cubes of C are called the facets of $Q(C)$. The *dimension* $\dim(Q(C))$ of $Q(C)$ is the largest dimension $\max_{B \in Q(C)} \dim(B)$ of a cube of $Q(C)$. A concept $c \in C$ is called a *corner* of C if c belongs to a unique maximal cube of C.

The *reduction* C^Y of a concept class C to $Y \subseteq U$ is a concept class on $U \setminus Y$ which has one concept for each Y-cube of C: $C^Y := \{\text{tag}(B) : B \in Q(C) \text{ and } \text{supp}(B) = Y\}$. When $x \in U$ we denote $C^{\{x\}}$ by C^x and call it *the x-hyperplane of* C (see Fig. 1 for an example). Note that a concept c belongs to C^x if and only if c and $c \cup \{x\}$ both belong to C. The union of all cubes of C having x in their support is called the *carrier* of C^x and is denoted by $N_x(C)$. If $c \in N_x(C)$, we also denote $c|(U \setminus \{x\})$ by c^x (note that $c^x \in C^x$).

The *tail* $\mathrm{tail}_x(C)$ of a concept class C on dimension x consists of all concepts that do not have in $G(C)$ an incident edge labeled with x. They correspond to the concepts of $C_x \setminus C^x$, i.e., to the concepts of C_x that have a unique extension in C. The class C can be partitioned as $N_x(C) \uplus \mathrm{tail}_x(C) = 0C^x \uplus 1C^x \uplus \mathrm{tail}_x(C)$, where $\uplus$ denotes the disjoint union and bC^x consists of all concepts in C^x extended with bit b in dimension x.

Given two classes $C \subseteq 2^U$ and $C' \subseteq 2^{U'}$ where U and U' are disjoint, the *Cartesian product* $C \times C' \subseteq 2^{U \uplus U'}$ is the concept class $\{c \uplus c' : c \in C \text{ and } c' \in C'\}$.

A concept class C is *connected* if the graph $G(C)$ is connected. If C is connected, denote by $d_{G(C)}(c, c')$ the graph-distance between c and c' in $G(C)$ and call it the *intrinsic distance* between c and c'. The distance $d(c, c') := d_{Q_n}(c, c')$ between two vertices c, c' of Q_n coincides with the Hamming distance $|c \Delta c'|$ between the 0-1-vectors corresponding to c and c'. Let $B(c, c') = \{t \subseteq U : d(c, t) + d(t, c') = d(c, c')\}$ be the *interval* between c and c' in Q_n; equivalently, $B(c, c')$ is the smallest cube of Q_n containing c and c'. A connected concept class C is called *isometric* if $d(c, c') = d_{G(C)}(c, c')$ for any $c, c' \in C$ and *locally isometric* if $d(c, c') = d_{G(C)}(c, c')$ for any $c, c' \in C$ such that $d(c, c') \leq 2$. Any path of C^Y connecting two concepts $\mathrm{tag}(B)$ and $\mathrm{tag}(B')$ of C^Y can be lifted to a path of parallel Y-cubes connecting B and B' in C; such a path of cubes is called a *gallery*.

A *simplicial complex* X on a set U is a family of subsets of X, called *simplices* or *faces* of X, such that if $\sigma \in X$ and $\sigma' \subseteq \sigma$, then $\sigma' \in X$. The *facets* of X are the maximal (by inclusion) faces of X. The *dimension* d of X is the size of its largest face. A simplicial complex X is a *pure* simplicial complex of dimension d if all its facets have size d.

A set $Y \subseteq U$ is *shattered* by a concept class $C \subseteq 2^U$ if $C|Y = 2^Y$. Furthermore, Y is *strongly shattered* by C if C contains a Y-cube. Denote by $\overline{X}(C)$ and $\underline{X}(C)$ the simplicial complexes consisting respectively of all shattered and of all strongly shattered sets of C. Clearly, $\underline{X}(C) \subseteq \overline{X}(C)$ and both $\overline{X}(C)$ and $\underline{X}(C)$ are closed by taking subsets, i.e., $\overline{X}(C)$ and $\underline{X}(C)$ are simplicial complexes. The classical *VC-dimension* [43] $\mathrm{VC\text{-}dim}(C)$ of a concept class C is the size of the largest set shattered by C, i.e., the dimension of the simplicial complex $\overline{X}(C)$.[3] The fundamental *sandwich lemma* (rediscovered independently by Pajor [33], Bollobás and Radcliffe [8], Dress [13], and Anstee et al. [3]) asserts that $|\underline{X}(C)| \leq |C| \leq |\overline{X}(C)|$. If $d = \mathrm{VC\text{-}dim}(C)$ and $n = |U|$, then $\overline{X}(C)$ cannot contain more than $\Phi_d(n) := \sum_{i=0}^{d} \binom{n}{i}$ simplices, yielding the well-known *Sauer-Shelah-Perles lemma* [38,40,43] that $|C| \leq \Phi_d(n)$.

A *labeled sample* is a set $s = \{(x_1, y_1), \ldots, (x_m, y_m)\}$, where $x_i \in U$ and $y_i \in \{0, 1\}$. An *unlabeled sample* is a set $\{x_1, \ldots, x_m\}$, where $x_i \in U$. A subsample s' of a sample s (labeled or unlabeled) is a subset of s. Given a labeled sample $s = \{(x_1, y_1), \ldots, (x_m, y_m)\}$, the unlabeled sample $\{x_1, \ldots, x_m\}$ is the domain of s and is denoted by $\mathrm{dom}(s)$. A labeled sample s is *realizable* by a concept $c : U \to \{0, 1\}$ (seen as a map) if $c(x_i) = y_i$ for every i, and s is realizable by a concept class C if it is realizable by some $c \in C$. For a concept class C, let $\mathrm{RS}(C)$ be the set of all labeled samples realizable by C.

A *sample compression scheme* for a concept class C is best viewed as a protocol between a *compressor* α and a *reconstructor* β (which both depend of C). The compressor gets a labeled sample s realizable by C from which it picks a small subsample s'. The compressor sends s' to the reconstructor. Based on s', the reconstructor outputs a concept $c \in C$ that needs to be consistent with the entire input sample s. A sample compression scheme has size k if for every realizable input sample s the size of the compressed subsample s' is at most k. An unlabeled sample compression scheme is a sample compression scheme in which the compressed subsample s' is unlabeled. So, the compressor removes the labels before sending the subsample to the reconstructor. An *unlabeled sample compression scheme* of size k for a concept class $C \subseteq 2^U$ is thus defined by a (compressor) function $\alpha : \mathrm{RS}(C) \to \binom{U}{\leq k}$ and a (reconstructor) function $\beta : \mathrm{Im}(\alpha) := \alpha(\mathrm{RS}(C)) \to C$ such that for any realizable sample s of C, the following conditions hold: $\alpha(s) \subseteq \mathrm{dom}(s)$ and $\beta(\alpha(s))|\,\mathrm{dom}(s) = s$.

In the literature, one usually allows the reconstructor β to take values in 2^U, i.e., the reconstructor can return a subset that is not a concept of C. The unlabeled sample compression schemes we consider in this paper, i.e., satisfying the property that $\mathrm{Im}(\beta) \subseteq C$, are usually called *proper* unlabeled sample compression schemes.

3. Ample and maximum classes

In this section, we briefly review the main characterizations and the basic examples of ample classes (maximum classes being one of them).

3.1. Characterizations

A concept class C is called *ample* if $|C| = |\overline{X}(C)|$. Ample classes are closed by taking restrictions, reductions, intersections with cubes, twistings, complements, and Cartesian products.

The following theorem reviews the main combinatorial characterizations of ample classes:

Theorem 3.1 *([5,8,24]). The following conditions are equivalent for a class C:*

[3] Note that the usual definition of the dimension of a simplicial complex X is the size of the largest face of X minus 1. We adopted this convention to have an equality between the VC-dimension of a class C and the dimension of $\overline{X}(C)$.

(1) C is ample;
(2) C^* is ample;
(3) $\underline{X}(C)=\overline{X}(C)$;
(4) $|\underline{X}(C)|=|C|$;
(5) $|\overline{X}(C)|=|C|$;
(6) $C\cap B$ is ample for any cube B;
(7) $(C^Y)_Z=(C_Z)^Y$ for all $Y,Z\subseteq U$ with $Y\cap Z=\varnothing$;
(8) for all partitions $U=Y\cup Z$, either $Y\in\underline{X}(C)$ or $Z\in\underline{X}(C^*)$.[4]

Condition (3) leads to a simple definition of ampleness: C is ample if whenever $Y\subseteq U$ is shattered by C, then there is a Y-cube in C. Thus, if C is ample we will write $X(C)$ instead of $\underline{X}(C)=\overline{X}(C)$. It follows that for ample classes, the VC-dimension of a concept class C, the dimension of the simplicial complex $X(C)$, and the dimension of the cube complex of C are the same. In the following, we talk about a d-dimensional class C when these three dimensions are equal to d.

We continue with metric and recursive characterizations of ample classes:

Theorem 3.2 *([5]). The following are equivalent for a concept class C:*

(1) C is ample;
(2) C^Y is connected for all $Y\subseteq U$;
(3) C^Y is isometric for all $Y\subseteq U$;
(4) C is isometric, and both C_x and C^x are ample for all $x\in U$;
(5) C is connected and all hyperplanes C^x are ample.

Corollary 3.3. *Two maximal cubes of an ample class C have different supports.*

Proof. Indeed, if B and B' are two d-cubes with the same support, by Theorem 3.2(2) B and B' can be connected in C by a gallery, and thus B is contained in a $d+1$-cube. Therefore, B and B' cannot be maximal. □

Lemma 3.4. *Given an ample class C and $x\in U$, for any cube B of C_x, there exists a cube B' of C such that* $\mathrm{supp}(B')=\mathrm{supp}(B)$ *and $B'_x=B$.*

Proof. Consider the cube B^* of 2^U such that $\mathrm{supp}(B^*)=\mathrm{supp}(B)\cup\{x\}$ and $B^*_x=B$. By Theorem 3.1(6), $C\cap B^*$ is ample. Since $\mathrm{supp}(B)$ is shattered by $C\cap B^*$, there exists a cube B' in $C\cap B^*$ such that $\mathrm{supp}(B')=\mathrm{supp}(B)$. Since $B'_x=B^*_x=B$, we are done. □

A concept class $C\subseteq 2^U$ of VC-dimension d is *maximum* if $|C|=\Phi_d(n)=\sum_{i=0}^{d}\binom{n}{i}$, i.e., if $C=\bigcup_{i=0}^{d}\binom{[n]}{i}$ (where $n=|U|$). The Sandwich Lemma and Theorem 3.1(5) imply that maximum classes are ample. Analogously to ample classes, maximum classes are hereditary by taking restrictions, reductions, twistings, and complements. Basic examples of maximum classes are concept classes derived from arrangements of hyperplanes in general position, balls in $\mathbb{R}^n$, and unions of n intervals on the line [15,16,18,21]. The following theorem summarizes some characterizations of maximum classes:

Theorem 3.5 *([15,16,18,45]). The following conditions are equivalent for a concept class C:*

(1) C is maximum;
(2) C_Y is maximum for all $Y\subseteq U$;
(3) C_x and C^x are maximum for all $x\in U$;
(4) C^* is maximum.

Following Kuzmin and Warmuth [23], we define a *representation map* for an ample class C as a bijection $r:C\to X(C)$ satisfying the *non-clashing condition*: $c|(r(c)\cup r(c'))\neq c'|(r(c)\cup r(c'))$, for all $c,c'\in C, c\neq c'$. It was shown in [23] that the existence of a representation map for a maximum class C implies an unlabeled sample compression scheme of size VC-dim(C) for C. In Section 6, we show that this also holds for ample classes. Moreover we show that for ample classes, they are equivalent to Δ-representation maps defined as follows. A Δ-*representation map* for an ample class C is a bijection $r:C\to X(C)$ satisfying the Δ-*non-clashing condition*: $c|(r(c)\Delta r(c'))\neq c'|(r(c)\Delta r(c'))$, for all $c,c'\in C, c\neq c'$.

3.2. Examples

We continue with the main examples of ample classes.

[4] This is the original definition of lopsidedness by Lawrence [24].

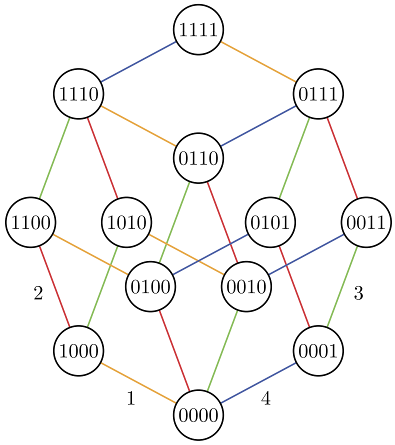

Fig. 2. An ample class which is also a convex geometry.

3.2.1. Simplicial complexes

The set of characteristic functions of simplices of a simplicial complex S can be viewed as a concept class $C(S)$: $C(S)$ is a bouquet of cubes with a common origin $\varnothing$, one cube for each simplex of S. Therefore, $\underline{X}(C(S)) = S$ and since $|S| = |C(S)|$, $C(S)$ is an ample class having S as its simplicial complex [5].

3.2.2. Realizable ample classes

Let $K \subseteq \mathbb{R}^n$ be a convex set. Let $C(K) := \{\mathrm{sign}(v) : v \in K, v_i \neq 0, \forall i \leq n\}$, where $\mathrm{sign}(v) \in \{\pm 1\}^n$ is the sign pattern of v. Lawrence [24] showed that $C(K)$ is ample, and called ample classes representable in this manner *realizable*. Lawrence presented a non-realizable ample class of Q_9 arising from a non-stretchable arrangement of pseudolines. It is shown in [6] that any ample class becomes realizable if instead of a convex set K one considers a Menger ℓ_1-convex set K of $\mathbb{R}^n$.

3.2.3. Median classes

A class C is called *median* if for every three concepts c_1, c_2, c_3 of C their *median* $m(c_1, c_2, c_3) := (c_1 \cap c_2) \cup (c_1 \cap c_3) \cup (c_2 \cap c_3)$ also belongs to C. Median classes are ample by [5, Proposition 2]. Median classes are closed by taking reductions, restrictions, intersections with cubes, and products but not under complementation.

Due to their relationships with other discrete structures, median classes are one of the most important examples of ample classes. Median classes are equivalent to finite median graphs (a well-studied class in metric graph theory, see [4]), to CAT(0) cube complexes, i.e., cube complexes of global nonpositive curvature (central objects in geometric group theory, see [19,36]), and to the domains of event structures (a basic model in concurrency theory [32,48]).

3.2.4. Convex geometries and conditional antimatroids

Let C be a concept class such that (i) $\varnothing \in C$ and (ii) $c, c' \in C$ implies that $c \cap c' \in C$. A point $x \in c \in C$ is called *extremal* if $c \setminus \{x\} \in C$. The set of extremal points of c is denoted by $\mathrm{ex}(c)$. A concept $c \in C$ is *generated* by $s \subseteq c$ if c is the smallest concept of C containing s. A concept class C satisfying (i) and (ii) with the additional property that every concept c of C is generated by its extremal points is called a *conditional antimatroid* [5, Section 3]. If $U \in C$, then we obtain the well-known structure of a *convex geometry* (called also an *antimatroid*) [14] (See Fig. 2 for an example). It was shown in [5, Proposition 1] that if C is a conditional antimatroid, then $\underline{X}(C) = \overline{X}(C)$, since $\underline{X}(C)$ coincides with the sets of extremal points and $\overline{X}(C)$ coincides with the set of all minimal generating sets of sets from C. Hence, any conditional antimatroid is ample. Besides convex geometries, median classes are also conditional antimatroids. Another example of conditional antimatroids is given by the set C of all strict partial orders on a set M. Each partial order is an asymmetric, transitive subset of $U = \{(u, v) : u, v \in M, u \neq v\}$. Then it is shown in [5] that for any $c \in C$, $\mathrm{ex}(c)$ is the set of covering pairs of c (i.e., the pairs (u, v) such that $u < v$ and there is no w with $u < w < v$) and that

$$\overline{X}(C) = \underline{X}(C) = \{H \subseteq U : H \text{ is the Hasse diagram of a partial order on } M\}.$$

Convex geometries comprise many interesting and important examples from geometry, ordered sets, and graphs, see the foundational paper [14]. For example, by the Krein-Milman theorem, any polytope of $\mathbb{R}^n$ is the convex hull of its extremal points. A *realizable convex geometry* is a convex geometry C such that its point set U can be realized as a finite set of $\mathbb{R}^n$ and $c \in C$ if and only if c is the intersection of a convex set of $\mathbb{R}^n$ with U. Acyclic oriented geometries (acyclic oriented matroids with no two point circuits) are examples of convex geometries, generalizing the realizable ones.

We continue with two particular examples of conditional antimatroids.

Example 3.6. Closer to usual examples from machine learning, let U be a finite set of points in $\mathbb{R}^n$, no two points sharing the same coordinate, and let the concept class C_Π consist of all intersections of axis-parallel boxes of $\mathbb{R}^n$ with U. Then C_Π is a convex geometry: for each $c \in C_\Pi$, $\mathrm{ex}(c)$ consists of all points of c minimizing or maximizing one of the n coordinates. Clearly, for any $p \in \mathrm{ex}(c)$, there exists a box Π such that $\Pi \cap U = c \setminus \{p\}$.

Example 3.7. A *partial linear space* is a pair (P, L) consisting of a finite set P whose elements are called *points* and a family L of subsets of P, whose elements are called *lines*, such that any line contains at least two points and any two points belong to at most one line. The projective plane (any pair of points belong to a common line and any two lines intersect in exactly one point) is a standard example, but partial linear spaces comprise many more examples. The concept class $L \subseteq 2^P$ has VC-dimension at most 2 because any two points belong to at most one line. Now, for each line $\ell \in L$ fix an arbitrary total order π_ℓ of its points. Let L^* consist of all subsets of points that belong to a common line ℓ and define an interval of π_ℓ. Then L^* is still a concept class of VC-dimension 2. Moreover, L^* is a conditional antimatroid: if $c \in L^*$ and c is an interval of the line ℓ, then $\mathrm{ex}(c)$ consists of the two end-points of c on ℓ.

3.2.5. Ample classes from graph orientations

Kozma and Moran [22] used the sandwich lemma to derive several properties of graph orientations. They also presented two examples of ample classes related to distances and flows in networks (see also [24, p. 157] for another example of a similar nature). Let $G = (V, E)$ be an undirected simple graph and let o^* be a fixed reference orientation of E. To an arbitrary orientation o of E associate a concept $c_o \subseteq E$ consisting of all edges which are oriented in the same way by o and by o^*. It is proven in [22, Theorem 26] that if each edge of G has a non-negative capacity, a source s and a sink t are fixed, then for any flow-value $w \in \mathbb{R}^+$, the set C_w^{flow} of all orientations of G for which there exists an (s, t)-flow of value at least w is an ample class. An analogous result was obtained if instead of the flow between s and t one considers the distance between those two nodes.

4. Corner peelings and partial shellings

In this section, we prove that corner peelings of ample classes are equivalent to isometric orderings of C as well as to partial shellings of the cross-polytope. This equivalence, combined with a result by Hall [20] yields a maximum class with VC dimension 3 without corners (Theorem 4.5 below). We show that this counterexample also refutes the analysis of the *Tail Matching Algorithm* of Kuzmin and Warmuth [23] for constructing unlabeled sample compression schemes for maximum classes. On the positive side, we prove the existence of corner peelings for conditional antimatroids and 2-dimensional ample classes. Finally we show that the cube complexes of all ample classes are collapsible.

4.1. Corners, isometric orderings, and partial shellings

We first establish some properties satisfied by the corners of an ample class. For $t \notin C$, let $F[t]$ be the smallest cube of Q_n containing t and all neighbors of t in Q_n that are in C. Note that the dimension of $F[t]$ is the number of neighbors of t in $G(C)$.

Lemma 4.1. *Let C be ample. Then:*

(i) If $t \notin C$ then $F[t] \subseteq C \cup \{t\}$.
(ii) If $t \notin C$ and $C' := C \cup \{t\}$ is isometric then C' is ample and t is a corner of C'.
(iii) c is a corner of C if and only if $C \setminus \{c\}$ is ample.

Proof. Item (i): Suppose $F[t] \setminus C \neq \{t\}$. Pick $s \neq t$ that is closest to t in $F[t] \setminus C$ (with respect to the Hamming distance of Q_n). Then t and s are not adjacent (by the definition of $F[t]$). By the choice of s, $B(s, t) \setminus \{s, t\} \subseteq C$, i.e., $B(s, t) \cap C^* = \{t, s\}$. However, by Theorem 3.1, C^* is ample and thus isometric by Theorem 3.2, contradicting that $B(s, t) \cap C^* = \{t, s\}$.

Item (ii): To prove that C' is ample, we use Theorem 3.2(2). First note that by item (i), $F[t] \subseteq C'$. Let $F' \neq F''$ be parallel cubes of C'. If $t \notin F' \cup F''$, then a gallery connecting F' and F'' in C is a gallery in C'. So, assume $t \in F'$. If F' is a proper face of $F[t]$, then F' is parallel to a face F of $F[t]$ not containing t. Since F' and F are connected in $F[t]$ by a gallery and F and F'' are connected in C by a gallery, we obtain a gallery between F' and F'' in C'. Finally, let $F' = F[t]$. In this case, we assert that F'' does not exist. Otherwise, we define a parallelism map π between the concepts of F' and the concepts of F'' as follows: for any $c' \in F'$, $\pi(c')$ is the unique concept $c'' \in F''$ such that $c'|\mathrm{supp}(F') = c''|\mathrm{supp}(F'')$ (recall that $\mathrm{supp}(F') = \mathrm{supp}(F'')$). Note that for any $r \in F'$: $d(t, \pi(t)) = d(r, \pi(r)) = d(F', F'')$. Since C' is isometric, t and $\pi(t)$ can be connected in C' by a path P of length $d(t, \pi(t))$. Let s be the neighbor of t in P. Since $s \in C$ it follows that $s \in F[t] = F'$. So, s is a concept in F' that is closer to $\pi(t)$ than t; this contradicts that $d(t, \pi(t)) = d(F', F'')$.

Item (iii): If $c \in C$ is a corner then there is a unique maximal cube $F \subseteq C$ containing it. Combined with Corollary 3.3, this implies that $\underline{X}(C \setminus \{c\}) = \underline{X}(C) \setminus \{\mathrm{supp}(F)\}$. Next, since $|C| = |\underline{X}(C)|$, we get that $|C \setminus \{c\}| = |\underline{X}(C \setminus \{c\})|$, and by Theorem 3.1

$C \setminus \{c\}$ is ample. Conversely, if $C \setminus \{c\}$ is ample, applying Item (ii) to $C \setminus \{c\}$ and c, since C is ample (and thus isometric), we get that c is a corner of C. $\square$

Let $C_< := (c_1, \ldots, c_m)$ be an ordering of the concepts in C. For any $1 \le i \le m$, let $C_i := \{c_1, \ldots, c_i\}$ denote the *i'th level set*. The ordering $C_<$ is called:

- an *ample* ordering if every level set C_i is ample;
- a *corner peeling* if every c_i is a *corner* of C_i;
- an *isometric* ordering if every level set C_i is isometric;
- a *locally isometric* ordering if every level set C_i is locally isometric.

Proposition 4.2. *The following conditions are equivalent for an ordering $C_< = (c_1, \ldots, c_m)$ of an isometric class C:*

(1) $C_<$ is ample;
(2) $C_<$ is a corner peeling;
(3) $C_<$ is isometric;
(4) $C_<$ is locally isometric.

Proof. Clearly, (3) $\Rightarrow$ (4). Conversely, suppose $C_<$ is locally isometric but one of its levels is not isometric. Hence, there exists $i < j$ such that any shortest (c_i, c_j)-path in C contains some c_k with $k > j$. Additionally, assume that c_i, c_j minimizes the distance $d(c_i, c_j)$ among all such pairs. Since C_j is locally isometric, necessarily $d(c_i, c_j) \ge 3$. Let c_r be the first vertex among $\{c_{j+1}, \ldots, c_m\}$ lying in $B(c_i, c_j) \cap C$. If $d(c_i, c_r) \ge 3$ or $d(c_r, c_j) \ge 3$ (say the first), then one can replace c_i, c_j by c_i, c_r, which contradicts the choice of c_i, c_j. Thus, $d(c_i, c_r), d(c_r, c_j) \le 2$, and at least one of them equals 2 (say $d(c_i, c_r) = 2$). Now, weak isometricity implies that c_i and c_r have a common neighbor c_ℓ with $\ell < \max\{i, r\} = r$. If $\ell < j$ then c_ℓ, c_j contradicts the minimality of c_i, c_j, and if $j < \ell < r$ then c_ℓ contradicts the minimality of c_r. This shows (4) $\Rightarrow$ (3).

If c_i is a corner of C_i, then any two neighbors of c_i in C_i have a second common neighbor in C_i, and therefore $d_{G(C_{i-1})}$ is the restriction of $d_{G(C_i)}$ on C_{i-1}. Since $C_m = C$ is isometric, this proves (2) $\Rightarrow$ (3).

To show (1) $\Rightarrow$ (2), let $C_<$ be an ample ordering of C. We assert that each c_i is a corner of C_i. Indeed, since C_{i-1} is ample and $c_i \notin C_{i-1}$, by Item (i) in Lemma 4.1 the cube $F[c_i]$, defined with respect to C_{i-1}, is included in C_i. Thus, c_i belongs to a unique maximal cube $F[c_i]$ of C_i, i.e., c_i is a corner of C_i. To prove (3) $\Rightarrow$ (1), let $C_<$ be an isometric ordering. The ampleness of each C_i follows by induction from Item (ii) of Lemma 4.1. $\square$

An isometric concept class C is *dismantlable* if it admits an ordering satisfying any of the equivalent conditions (1)-(4) in Proposition 4.2. Isometric orderings of Q_n are closely related to shellings of its dual, the *cross-polytope* O_n (which we define next). Define $\pm U := \{\pm x_1, \ldots, \pm x_n\}$; so, $|\pm U| = 2n$, and we call $-x_i, +x_i$ *antipodal*. The n-dimensional *cross-polytope* is the pure simplicial complex of dimension n whose facets are all $\sigma \subseteq \pm U$ that contain exactly one element in each antipodal pair. Thus, O_n has 2^n facets and each facet σ of O_n corresponds to a vertex c of Q_n ($+x_i \in \sigma$ if and only if $x_i \in c$). Observe that $x_i \in c' \Delta c''$ if and only if $\{+x_i, -x_i\} \subseteq \sigma' \Delta \sigma''$ where σ' correspond to c' and σ'' corresponds to c''.

Let X be a pure simplicial complex of dimension d. Two facets σ, σ' are adjacent if $|\sigma \Delta \sigma'| = 2$. A *shelling* of X is an ordering $\sigma_1, \ldots, \sigma_p$ of all of its facets such that $2^{\sigma_j} \bigcap (\bigcup_{i<j} 2^{\sigma_i})$ is a pure simplicial complex of dimension $d - 1$ for every $2 \le j \le p$ [49, Lecture 8]. A *partial shelling* is an ordering of some facets that satisfies the above condition. Note that $\sigma_1, \ldots, \sigma_m$ is a partial shelling if and only if for every $i < j \le m$ there exists $k < j$ such that $\sigma_i \cap \sigma_j \subseteq \sigma_k \cap \sigma_j$, and $|\sigma_k \cap \sigma_j| = d - 1$, i.e., $\sigma_k \cap \sigma_j$ is a facet of both σ_j and σ_k. A partial shelling $\sigma_1, \ldots, \sigma_m$ of X is *1-step extendable* if there exists a facet τ of X such that $\sigma_1, \ldots, \sigma_m, \tau$ is a partial shelling of X. A partial shelling $\sigma_1, \ldots, \sigma_m$ of X is *extendable* if it can be extended to a shelling of X. A pure simplicial complex X is *extendably shellable* if every partial shelling is extendable. We next establish a relationship between partial shellings and isometric orderings.

Proposition 4.3. *Every partial shelling of the cross-polytope O_n defines an isometric ordering of the corresponding vertices of the cube Q_n. Conversely, if C is an isometric class of Q_n, then any isometric ordering of C defines a partial shelling of O_n.*

Proof. Let $\sigma_1, \ldots, \sigma_m$ be a partial shelling of O_n and $c_1, \ldots, c_m$ be the ordering of the corresponding vertices of Q_n. We need to prove that each level set $C_j = \{c_1, \ldots, c_j\}$ is isometric. It suffices to show that for every $i < j$ there is $k < j$ such that $d(c_k, c_j) = 1$ and $c_k \in B(c_i, c_j)$. Equivalently, for every $i < j$, there is $k < j$ such that $|\sigma_k \Delta \sigma_j| = 2$ and $\sigma_i \cap \sigma_j \subseteq \sigma_k \subseteq \sigma_i \cup \sigma_j$: since $\sigma_1, \ldots, \sigma_m$ is a partial shelling, there is a facet σ_k with $k < j$ such that $|\sigma_k \cap \sigma_j| = n - 1$ and $\sigma_i \cap \sigma_j \subseteq \sigma_k \cap \sigma_j$. We claim that σ_k is the desired facet. It remains to show that (i) $|\sigma_j \Delta \sigma_k| = 2$ and (ii) $\sigma_k \subseteq \sigma_i \cup \sigma_j$. Item (i) follows since $|\sigma_j| = |\sigma_k| = n$, and $|\sigma_k \cap \sigma_j| = n - 1$. For Item (ii), let $\sigma_j \setminus \sigma_k = \{+x\}$ and $\sigma_k \setminus \sigma_j = \{-x\}$. We need to show that $-x \in \sigma_i$, or equivalently that $+x \notin \sigma_i$. The latter follows since $+x \in \sigma_j \setminus \sigma_k$ and $\sigma_j \cap \sigma_i \subseteq \sigma_k$.

Conversely, let $c_1, \ldots, c_m$ be an isometric ordering and $\sigma_1, \ldots, \sigma_m$ be the ordering of the corresponding facets of O_n. We assert that this is a partial shelling. Let $i < j$. It suffices to exhibit $k < j$ such that $|\sigma_k \cap \sigma_j| = n - 1$ and $\sigma_i \cap \sigma_j \subseteq \sigma_k \cap \sigma_j$. Since

C_j is isometric, c_j has a neighbor $c_k \in B(c_i, c_j) \cap C_j$. Since $d(c_j, c_k) = 1$ it follows that $|\sigma_k \cap \sigma_j| = n - 1$. Since $c_k \in B(c_i, c_j)$ it follows that $\sigma_i \cap \sigma_j \subset \sigma_k \subset \sigma_i \cup \sigma_j$ and thus $\sigma_i \cap \sigma_j \subseteq \sigma_k \cap \sigma_j$. $\square$

Corollary 4.4. *For any partial shelling $\sigma_1, \ldots, \sigma_m$, let $(c_1, \ldots, c_m)$ be the ordering of the corresponding vertices of Q_n and let $C = \{c_1, \ldots, c_m\}$. Then the following hold.*

(1) C and $C^* = Q_n \setminus C$ are ample.
(2) the partial shelling $\sigma_1, \ldots, \sigma_m$ is 1-step extendable if and only if C^* has a corner.
(3) the partial shelling $\sigma_1, \ldots, \sigma_m$ is extendable if and only if C^* is dismantlable.

Consequently, if all ample classes are dismantlable, then O_n is extendably shellable.

Proof. By Proposition 4.3, the level sets $\{c_1, \ldots, c_i\}$, $1 \le i \le m$ are isometric, thus C is ample by Proposition 4.2. The complement C^* is also ample, establishing (1).

If C^* has a corner, then by Lemma 4.1(iii), C^* contains a concept t such that $C^* \setminus \{t\}$ is ample. Consequently, its complement $C' := C \cup \{t\}$ is ample. Let τ be the facet of O_n corresponding to t. By Proposition 4.2 $(c_1, \ldots, c_m, t)$ is an isometric ordering of C'. By Proposition 4.3, $\sigma_1, \ldots, \sigma_m, \tau$ is a partial shelling of O_n and thus $\sigma_1, \ldots, \sigma_m$ is 1-step extendable. Conversely, assume that there exists τ such that $\sigma_1, \ldots, \sigma_m, \tau$ is a partial shelling of O_n and let t be the vertex of Q_n corresponding to τ. By Proposition 4.3, $(c_1, \ldots, c_m, t)$ is an isometric ordering of $C' = C \cup \{t\}$. By Proposition 4.2, C' is ample and thus its complement $C^* \setminus \{t\}$ is also ample, establishing that t is a corner of C^* by Lemma 4.1(iii). This establishes (2).

Assertion (3) is a direct consequence of Assertion (2). Consequently, if all ample classes are dismantlable, then any partial shelling of O_n is extendable and thus O_n is extendably shellable. $\square$

It was asked in [49] if any cross-polytope O_n is extendably shellable. In the PhD thesis of H. Tracy Hall from 2004, a nice counterexample to this question is given [20]. Hall's counterexample arises from the 299 regions of an arrangement of 12 pseudo-hyperplanes. These regions are encoded as facets of the 12-dimensional cross-polytope O_{12} and it is shown in [20] that the subcomplex of O_{12} consisting of all other facets admits a shelling which cannot be extended by adding any of the 299 simplices. By Corollary 4.4(ii), the ample concept class C_H defined by those 299 simplices does not have any corner (see Fig. 3 for a picture of C_H).[5] A counting shows that C_H is a maximum class of VC-dimension 3. This completes the proof of our first main result:

Theorem 4.5. *There exists a maximum class C_H of VC-dimension 3 without any corner.*

Remark 4.6. Hall's concept class C_H also provides a counterexample to [30, Conjecture 4.2] asserting that *for any two ample classes $C_1 \subset C_2$ with $|C_2 \setminus C_1| \ge 2$ there exists an ample class C such that $C_1 \subset C \subset C_2$*. As noticed in [30], this conjecture is stronger than the corner peeling conjecture disproved by Theorem 4.5.

Remark 4.7. Notice also that since ample classes are closed by Cartesian products, and any corner in a Cartesian product comes from corners in each factor, one can construct other examples of ample classes without corners by taking the Cartesian product of C_H by any ample class.

4.2. C_H Refutes the analysis by Kuzmin and Warmuth [23]

The algorithm of [23] uses the notion of forbidden labels for maximum classes, introduced in the PhD thesis of Floyd [15] and used in [16,23]; we closely follow [15]. Let C be a maximum class of VC-dimension d on the set U. For any $Y \subseteq U$ with $|Y| = d + 1$, the restriction $C|Y$ is a maximum class of dimension d. Thus $C|Y$ contains $\Phi_d(d+1) = 2^{d+1} - 1$ concepts. There are 2^{d+1} possible concepts on Y. We call the characteristic function of the unique concept that is not a member of $C|Y$ a *forbidden label of size* $d + 1$ for Y. Each forbidden label *forbids* all concepts that contain it from belonging to C. Let c be a concept which contains the forbidden label for Y. Since C is a maximum class, adding c to C would shatter the set Y that is of cardinality $d + 1$.

The algorithm of [23], called the *Tail Matching Algorithm*, recursively constructs a representation map $\tilde{r}$ for C^x, expands $\tilde{r}$ to a map r on the carrier $N_x(C) = 0C^x \cup 1C^x$ of C^x, and extends r to $\mathrm{tail}_x(C)$ using a special subroutine. This subroutine and the correctness proof of the algorithm (Theorem 11) heavily uses that a specially defined bipartite graph (which we will call Γ) has a unique perfect matching. This bipartite graph has the concepts of $\mathrm{tail}_x(C)$ on one side and the forbidden labels of size d for C^x on another side. Both sides have the same size. There is an edge between a concept and a forbidden label

[5] For the interested reader, a file containing the 299 concepts of C_H represented as elements of $\{0,1\}^{12}$ is available at https://arxiv.org/src/1812.02099/anc/CH.txt.

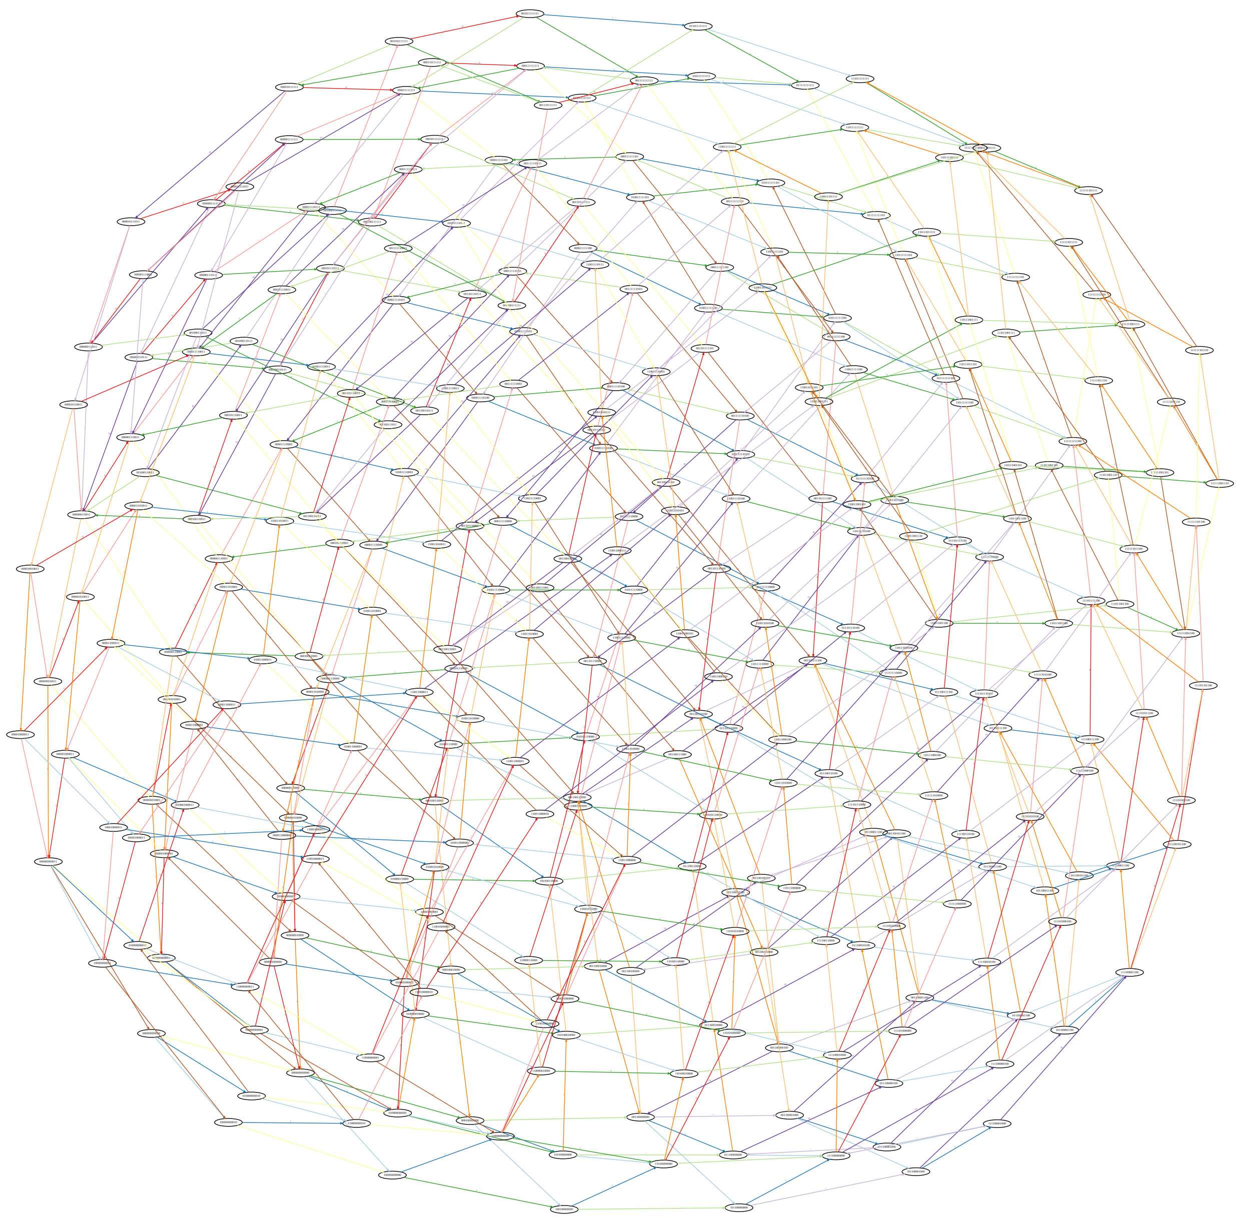

Fig. 3. The maximum class $C_H \subset 2^{12}$ without corners of VC-dimension 3 with $\binom{12}{\leq 3} = 299$ concepts. A different edge color is used for each of the 12 dimensions. The orientation of the edges is a unique sink orientation (defined in Section 6) and defines a representation map r_1 for the class C_H: for each concept c, $r_1(c)$ is the set of the labels of the outgoing edges. This representation map was found using the MiniSat solver. The appendix contains more discussions about this and another representation map of C_H. Best viewed in color.

if and only if the forbidden label is contained in the respective concept. The graph Γ is defined in [23, Theorem 10], which also asserts that Γ *has a unique perfect matching*. In the following, we show that this assertion is false.

We will show that uniqueness of the matching implies that there is a corner in the tail (which is contradicted by Hall's concept class C_H). We use the following lemma about perfect bipartite matchings:

Lemma 4.8 *([10]). Let G be a bipartite graph with bipartition X, Y and unique perfect matching M. Then, there are vertices $x \in X$ and $y \in Y$ with degree one.*

By Lemma 4.8, the uniqueness of the matching claimed in [23, Theorem 10] implies that there is a forbidden labeling that is contained in exactly one concept c of the tail. We claim that c must be a corner: c is the only concept in the tail realizing this forbidden labeling and removing this concept from C reduces the number of shattered sets by at least one. After the removal, the number of concepts is $|C| - 1$. By the Sandwich Lemma, the number of shattered sets is always at least as big as the number of concepts. So removing c reduces the number of shattered sets by exactly one set S and the resulting class is ample. For ample classes, the number of concepts equals the number of supports of cubes of the 1-inclusion graph and for every shattered set there is a cube with this shattered set as its support. Thus c lies in a cube B with support S. There is no other cube with support S because after removing c there is no cube left with support S. Thus B is the unique cube with support S and must be maximal. We conclude that c is a corner of C. Since C_H is maximum and does not contain corners, this leads to a contradiction. This establishes that Theorem 10 from [23] is false.

4.3. Two families of dismantlable ample classes

We continue with two families of dismantlable ample classes: conditional antimatroids and 2-dimensional ample classes.

Proposition 4.9. *Conditional antimatroids are dismantlable.*

Proof. Let $c_1,\dots,c_m$ be an ordering of the concepts of C, where $i<j$ if and only if $|c_i|\le|c_j|$ (breaking ties arbitrarily). Clearly, $c_1=\varnothing$ and the order $c_1,\dots,c_m$ is monotone with respect to distances from c_1. In particular, any order defined by a Breadth First Search (BFS) from c_1 satisfies this condition. We prove that for each i, c_i is a corner of the level set $C_i=\{c_1,\dots,c_{i-1},c_i\}$. The neighbors of c_i in C_i are subsets of c_i containing $|c_i|-1$ elements. From the definition of extremal points of c_i that we denote by $\mathrm{ex}(c_i)$, it immediately follows that $c_i\setminus\{x\}\in C$ if and only if $x\in\mathrm{ex}(c_i)$. For any $s\subseteq\mathrm{ex}(c_i)$, since $c_i\setminus s=\bigcap_{x\in s}(c_i\setminus\{x\})$ and C is closed under intersections, we conclude that $c_i\setminus s\in C$. Therefore, the whole Boolean cube between c_i and $c_i\setminus\mathrm{ex}(c_i)$ is included in C, showing that c_i is a corner of C_i. □

The fact that 2-dimensional maximum classes have corners was proved in [35, Theorem 34] and it was generalized in [28] to 2-dimensional ample classes. We provide here a different proof of this result, originating from 1997-1998 and based on a local characterization of convex sets of general ample classes, which may be of independent interest.

Given two classes $C'\subseteq C$, C' is *convex* in C if $B(c,c')\cap C\subseteq C'$ for any $c,c'\in C'$ and C' is *locally convex* in C if $B(c,c')\cap C\subseteq C'$ for any $c,c'\in C'$ with $d(c,c')=2$.

Lemma 4.10. *A connected subclass C' of an ample class C is convex in C if and only if C' is locally convex.*

Proof. One direction of the statement is trivial. For the other direction, assume that C' is locally convex. For $c,c'\in C'$, recall that $d_{G(C')}(c,c')$ denotes the distance between c and c' in $G(C')$. Recall also that since C is ample, C is isometric and thus $d_{G(C)}(c,c')=d(c,c')$ for any $c,c'\in C$. We prove that for any $c,c'\in C'$, $B(c,c')\cap C\subseteq C'$ by induction on $k=d_{G(C')}(c,c')$, the case $k=2$ being covered by the initial assumption. Pick any $t\in B(c,c')\cap C$ and let L' be a shortest (c,c')-path of C passing via t. By Theorem 3.2, L' is contained in $B(c,c')$. Let L'' be a shortest (c,c')-path in C'. Let c'' be the neighbor of c in L''. Since $d_{G(C')}(c'',c')<k$, by the induction assumption $B(c'',c')\cap C\subseteq C'$. Since $G(C)$ is bipartite, either $c\in C\cap B(c'',c')$ or $c''\in C\cap B(c,c')$. In the first case, $t\in C\cap B(c,c')\subseteq C\cap B(c'',c')\subseteq C'$ and we are done. So suppose that $d_{G(C')}(c,c')=d(c,c')$. This implies that L'' is a shortest (c,c')-path in C. Moreover, we can assume that $t\notin B(c'',c')$ and consequently, $c''\notin L'$ and the edge cc'' is parallel to an edge uv of the path L'. Consider a shortest gallery $e_0:=cc'',e_1,\dots,e_k:=uv$ connecting the edges cc'' and uv in C. It is constituted of two shortest paths $P'=(u_0:=c,u_1,\dots,u_k:=u)$ and $P''=(v_0:=c'',v_1,\dots,v_k=v)$. Then P'' together with the subpath of L' comprised between v and c' constitute a shortest path between c'' and c', thus it belongs to C'. Therefore, if t is comprised in L' between v and c', then we are done. Thus suppose that t belongs to the subpath of L' between c and u. Since u_1 is adjacent to $c,v_1\in C'$, by local convexity of C' we obtain that $u_1\in C'$. Applying this argument several times, we deduce that the whole path P' belongs to C'. In particular, $u\in C'$. Since v is between u and c', $u\neq c'$, thus $d_{G(C')}(c,u)<k$. By induction hypothesis, $B(c,u)\cap C\subseteq C'$. Since $t\in B(c,u)$, we are done. □

Proposition 4.11. *2-Dimensional ample classes are dismantlable.*

Proof. As for conditional antimatroids, the corner peeling for 2-dimensional classes is based on an algorithmic order of concepts. Consider the following total order $c_1,\dots,c_m$ of the concepts of a 2-dimensional ample class C constructed recursively as follows: start with an arbitrary concept and denote it c_1 and at step i, having numbered the concepts $C_{i-1}=\{c_1,\dots,c_{i-1}\}$, select as c_i a concept in $C\setminus C_{i-1}$ which is adjacent to a maximum number of concepts of C_{i-1}. We assert that $C_i:=C_{i-1}\cup\{c_i\}$ is ample and that c_i is a corner of C_i. Since C_{i-1} is ample and C is 2-dimensional, by Lemma 4.1, c_i has at most two neighbors in C_{i-1}. First suppose that c_i has two neighbors u and v in C_{i-1}. By Lemma 4.1, c_i,u and v are included in a 2-cube $F[c_i]$ with a fourth concept t belonging to C_{i-1}. Since $F[c_i]$ is a maximal cube of C, by Corollary 3.3, $F[c_i]$ is not parallel to any other cube of C. On the other hand, the edges c_iu and c_iv are parallel to the edges vt and ut, respectively. This shows that C_i^Y is connected for any $Y\subseteq U$, and therefore by Theorem 3.2(2), C_i is ample and c_i is a corner. Now suppose that c_i has exactly one neighbor c in C_{i-1}. From the choice of c_i at step i, any concept of $C\setminus C_{i-1}$ has at most one neighbor in C_{i-1}, i.e., C_{i-1} is locally convex. By Lemma 4.10, C_{i-1} is convex in C. But then $C_i=C_{i-1}\cup\{c_i\}$ is isometric since c_i has degree 1 in $G(C_i)$. By Lemma 4.1, C_i is thus ample and c_i is a corner of C_i. □

4.4. Collapsibility

A *free face* of a cube complex $Q(C)$ is a face Q of $Q(C)$ strictly contained in only one other face Q' of $Q(C)$. An *elementary collapse* is the deletion of a free face Q (thus also of Q') from $Q(C)$. A cube complex $Q(C)$ is *collapsible* to a vertex v_0 if C can be reduced to v_0 by a sequence of elementary collapses. Namely, there exists an ordered sequence $\Lambda = ((Q_1, Q'_1), \dots, (Q_n, Q'_n))$ where each Q_i is a free face of $Q(C) \setminus \{Q_1, Q'_1, \dots, Q_{i-1}, Q'_{i-1}\}$ contained in Q'_i and $Q(C) \setminus \{Q_1, Q'_1, \dots, Q_n, Q'_n\}$ is $\{v_0\}$. Observe that each face of $Q(C)$ distinct from v_0 appears exactly once in the sequence.

Collapsibility is a stronger version of contractibility. The sequences of elementary collapses of a collapsible cube complex $Q(C)$ can be viewed as discrete Morse functions [17] without critical cells, i.e., acyclic perfect matchings of the face poset of $Q(C)$. From the definition it follows that if C has a corner peeling, then the cube complex $Q(C)$ is collapsible: the sequence of elementary collapses follows the corner peeling order (in general, detecting if a finite complex is collapsible is NP-complete [42]). Theorem 3.2(5) implies that the cube complexes of ample classes are contractible (see also [7] for a more general result). In fact, the cube complexes of ample classes are collapsible (this extends the collapsibility of finite CAT(0) cube complexes established in [1]):

Proposition 4.12. *If $C \subseteq 2^U$ is an ample class, then the cube complex $Q(C)$ is collapsible.*

Proof. We proceed by induction on the size of C. If C contains only one concept, then the statement is trivially true. Otherwise, let $x \in U$ and suppose by induction hypothesis that $Q(C_x)$ is collapsible to v_0. Let $\Lambda := ((Q_1, Q'_1), \dots, (Q_n, Q'_n))$ be the corresponding collapsing sequence of $Q(C_x)$, i.e., a partition of the faces of $Q(C_x) \setminus \{v_0\}$ into pairs (Q_i, Q'_i) such that Q_i is a free face in the current subcomplex $Q(C_x) \setminus \{Q_1, Q'_1, \dots, Q_{i-1}, Q'_{i-1}\}$ of C_x and Q'_i is the unique face properly containing Q_i.

Each cube Q^* of C is mapped to a cube Q of C_x with $\mathrm{supp}(Q) = \mathrm{supp}(Q^*) \setminus \{x\}$. Observe that $\dim(Q) = \dim(Q^*) - 1$ if $x \in \mathrm{supp}(Q^*)$ and $\dim(Q) = \dim(Q^*)$ otherwise. In the first case, Q is contained in C^x and Q^* is entirely contained in the carrier $N_x(C)$. Conversely, by Lemma 3.4, any cube Q of C_x is the image of at least one cube Q^* of C (with $\mathrm{supp}(Q) = \mathrm{supp}(Q^*) \setminus \{x\}$). If there exists Q^* such that $\mathrm{supp}(Q^*) = \mathrm{supp}(Q) \cup \{x\}$, then there exist exactly three cubes of C that are mapped to Q: Q^* and the two opposite facets P^* and R^* of Q^* such that $\mathrm{supp}(P^*) = \mathrm{supp}(R^*) = \mathrm{supp}(Q) = \mathrm{supp}(Q^*) \setminus \{x\}$. Otherwise, there exists a unique cube Q^* of C that is mapped to Q.

We derive a collapsing sequence Λ^* for $Q(C)$ by replacing each elementary collapse of Λ by one, two, or three elementary collapses in $Q(C)$ and when $v_0 \in C^x$, we add a last elementary collapse. Consider a pair $(Q_i, Q'_i) \in \Lambda$. If neither Q_i nor Q'_i are contained in C^x, let Q^*_i and $Q^{*\prime}_i$ be the unique preimages of Q_i and Q'_i in $Q(C)$ and insert the pair $(Q^*_i, Q^{*\prime}_i)$ in Λ^*.

If both Q_i and $Q_i{}'$ are contained in C^x, let P^*_i, R^*_i, Q^*_i be the cubes of C mapped to Q and $P^{*\prime}_i, R^{*\prime}_i, Q^{*\prime}_i$ the ones mapped to Q' such that $\dim(P^*_i) = \dim(R^*_i) = \dim(Q^*_i) - 1 = \dim(Q_i)$, $\dim(P^{*\prime}_i) = \dim(R^{*\prime}_i) = \dim(Q^{*\prime}_i) - 1 = \dim(Q'_i)$, P^*_i is contained in $P^{*\prime}_i$, R^*_i is contained in $R^{*\prime}_i$, and Q^*_i is contained in $Q^{*\prime}_i$. We insert the three pairs $(Q^*_i, Q^{*\prime}_i), (P^*_i, P^{*\prime}_i), (R^*_i, R^{*\prime}_i)$ in Λ^* in this order.

Suppose now that Q_i is included in C^x and Q'_i is not included in C^x. Let $Q^{*\prime}_i$ be the unique cube of C mapped to Q'_i and let P^*_i, R^*_i, Q^*_i be the cubes of C mapped to Q such that $\dim(P^*_i) = \dim(R^*_i) = \dim(Q^*_i) - 1 = \dim(Q_i)$. Assume without loss of generality that P^*_i is a facet of both Q^*_i and $Q^{*\prime}_i$. We insert the two pairs $(R^*_i, Q^*_i), (P^*_i, Q^{*\prime}_i)$ in Λ^* in this order.

Finally, we consider the vertex v_0. If $v_0 \notin C^x$, then the preimage of v_0 contains only one vertex that we denote by v^*_0. Suppose now that $v_0 \in C^x$. Let $v^*_0, u^*_0 \in C$ such that $v^*_0 = v_0$ and $u^*_0 = v_0 \cup \{x\}$. Then the preimage of v_0 in Q is of size 3: it contains the vertices v^*_0, u^*_0 and the edge $\{u^*_0, v^*_0\}$. In this case, we add to Λ^* the elementary collapse $(u^*_0, \{u^*_0, v^*_0\})$.

Each cube Q^* of $Q(C)$ is in the preimage of a single cube of $Q(C_x)$. Moreover, each cube of $Q(C_x) \setminus \{v_0\}$ appears in exactly one collapsing pair of the sequence Λ. Consequently, each cube of $Q(C)$ whose image is not v_0 appears exactly once in the sequence Λ^*. Notice also that when the preimage of v_0 is of size 3, both u^*_0 and $\{u^*_0, v^*_0\}$ appear in the last pair of Λ^*. Consequently, each cube of $Q(C) \setminus \{v^*_0\}$ appears in exactly one pair of the sequence Λ^*. Since each pair of Λ^* (except potentially $(u^*_0, \{u^*_0, v^*_0\})$) is derived from a collapsing pair of Λ and since Λ is a sequence of elementary collapses of $Q(C_x)$, we can deduce that Λ^* is a sequence of elementary collapses of the cube complex $Q(C)$ to v^*_0. □

Example 4.13. We illustrate the construction in the proof of Proposition 4.12 with the ample class C from Fig. 4 and its restriction C_x. The names of the concepts are depicted in the figure. The faces will be denoted by the list of their vertices.

Consider the following two collapsing sequences of C_x:

- $\Lambda_1 = (\{a\}, \{a, b\}), (\{b\}, \{b, c\})$, and
- $\Lambda_2 = (\{a\}, \{a, b\}), (\{c\}, \{b, c\})$.

Note that Λ_1 collapses $Q(C_x)$ to c and Λ_1 collapses $Q(C_x)$ to b.

The corresponding collapsing sequences for C obtained by the construction described in the proof of Proposition 4.12 are the following:

- $\Lambda^*_1 = (\{a, a'\}, \{a, b, b', a\}), (\{a'\}, \{a', b'\}), (\{a\}, \{a, b\}), (\{b'\}, \{b, b'\}), (\{b\}, \{b, c\})$ collapses $Q(C)$ to c, and

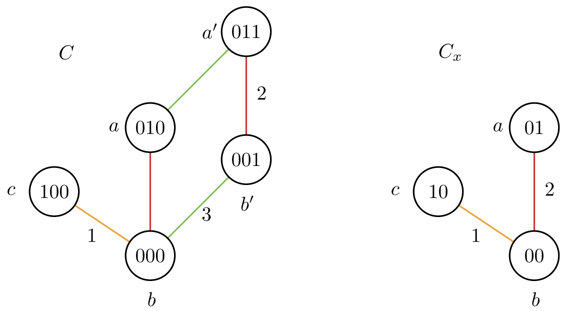

Fig. 4. The ample class C used in Example 4.13 and its restriction C_x for $x = 3$.

- $\Lambda_2^* = (\{a, a'\}, \{a, b, b', a\}), (\{a'\}, \{a', b'\}), (\{a\}, \{a, b\}), (\{c\}, \{b, c\}), (\{b'\}, \{b, b'\})$ collapses $Q(C)$ to b.

In both cases, the subsequence $(\{a, a'\}, \{a, b, b', a\}), (\{a'\}, \{a', b'\}), (\{a\}, \{a, b\})$ corresponds to the elementary collapse $(\{a\}, \{a, b\})$. In Λ_1^*, the subsequence $(\{b'\}, \{b, b'\}), (\{b\}, \{b, c\})$ corresponds to the elementary collapse $(\{b\}, \{b, c\})$ in Λ_1. In Λ_2^*, the elementary collapse $(\{c\}, \{b, c\})$ corresponds to the elementary collapse $(\{c\}, \{b, c\})$ in Λ_2 and the elementary collapse $(\{b'\}, \{b, b'\})$ is the elementary collapse associated to the last vertex b of Λ_2.

5. Representation maps for maximum classes

In this section, we prove that maximum classes admit representation maps, and therefore optimal unlabeled sample compression schemes.

Theorem 5.1. *Any maximum class $C \subseteq 2^U$ of VC-dimension d admits a representation map, and consequently, an unlabeled sample compression scheme of size d.*

The crux of the proof of Theorem 5.1 is the following proposition. Let C be a d-dimensional maximum class and let $D \subseteq C$ be a $(d-1)$-dimensional maximum subclass. A *missed simplex* for the pair (C, D) is a simplex $\sigma \in X(C) \setminus X(D)$. Note that since C and D are maximum, any missed simplex has size d. An *incomplete cube* Q for (C, D) is a cube of C such that $\operatorname{supp}(Q)$ is a missed simplex. For any incomplete cube Q with $\sigma = \operatorname{supp}(Q)$, $C|\sigma$ is a d-cube and $D|\sigma$ is a maximum class of dimension $d-1$. Observe that any incomplete cube for (C, D) is a maximal cube in C and by Corollary 3.3, there is a bijection between the missed simplices for (C, D) and the incomplete cubes for (C, D).

Consider a missed simplex σ for (C, D) and the incomplete cube Q for (C, D) such that $\operatorname{supp}(Q) = \sigma$. Since $|\sigma| = d$, we have $|C|\sigma| = \binom{d}{\leq d} = \binom{d}{\leq d-1} + 1 = |D|\sigma| + 1$. Since $Q|\sigma = C|\sigma$, there exists a unique concept $c \in Q$ such that $c|\sigma \notin D|\sigma$. We call c the *source* of Q and we consider the *source-map* s from the set of incomplete cubes for (C, D) to $C \setminus D$ where $s(Q)$ is the source of Q. In fact, we show in the following proposition that the source-map is a bijection between the incomplete cubes for (C, D) and the concepts of $C \setminus D$.

Proposition 5.2. *Each $c \in C \setminus D$ is the source of a unique incomplete cube for (C, D). Moreover, if $r' : D \to X(D)$ is a representation map for D and $r : C \to X(C)$ extends r' by setting $r(c) = \operatorname{supp}(s^{-1}(c))$ for each $c \in C \setminus D$, then r is a representation map for C.*

Proof of Theorem 5.1. Following the general recursive construction idea of [23], we derive a representation map for C by induction on $|U|$. If $|U| = 0$, then C is a 0-dimensional maximum class containing a unique concept c. In this case, we set $r(c) = \varnothing$ and r is clearly a representation map for C. For the induction step, (see Fig. 5), pick $x \in U$ and consider the maximum classes C_x and $C^x \subset C_x$ with domain $U \setminus \{x\}$. By induction, C^x has a representation map r^x. Use Proposition 5.2 to extend r^x to a representation map r_x of C_x. Define a map r on C as follows:

$$r(c) = \begin{cases} r_x(c_x) & \text{if } c_x \notin C^x \text{ or } x \notin c, \\ r_x(c_x) \cup \{x\} & \text{if } c_x \in C^x \text{ and } x \in c. \end{cases}$$

It is easy to verify that r is non-clashing: indeed, if $c' \neq c'' \in C$ satisfy $c'_x \neq c''_x$ then $c'_x|r_x(c'_x) \cup r_x(c''_x) \neq c''_x|r_x(c'_x) \cup r_x(c''_x)$. Since $r_x(c'_x) \subseteq r(c')$, $r_x(c''_x) \subseteq r(c'')$, it follows that also c', c'' disagree on $r(c') \cup r(c'')$. Else, $c'_x = c''_x \in C^x$ and $c'(x) \neq c''(x)$. In this case, $x \in r(c') \cup r(c'')$ and therefore c', c'' disagree on $r(c') \cup r(c'')$.

It remains to show that r is a bijection between C and $X(C) = \binom{U}{\leq d}$. It is easy to verify that r is injective. So, it remains to show that $|r(c)| \leq d$, for every $c \in C$. This is clear when $c_x \notin C^x$ or $x \notin c$. If $c_x \in C^x$ and $x \in c$, then $r(c) = r^x(c_x) \cup \{x\}$ and $|r^x(c_x)| \leq d-1$ (since C^x is $(d-1)$-dimensional). Hence, $|r(c)| \leq d$ as required, concluding the proof. $\square$

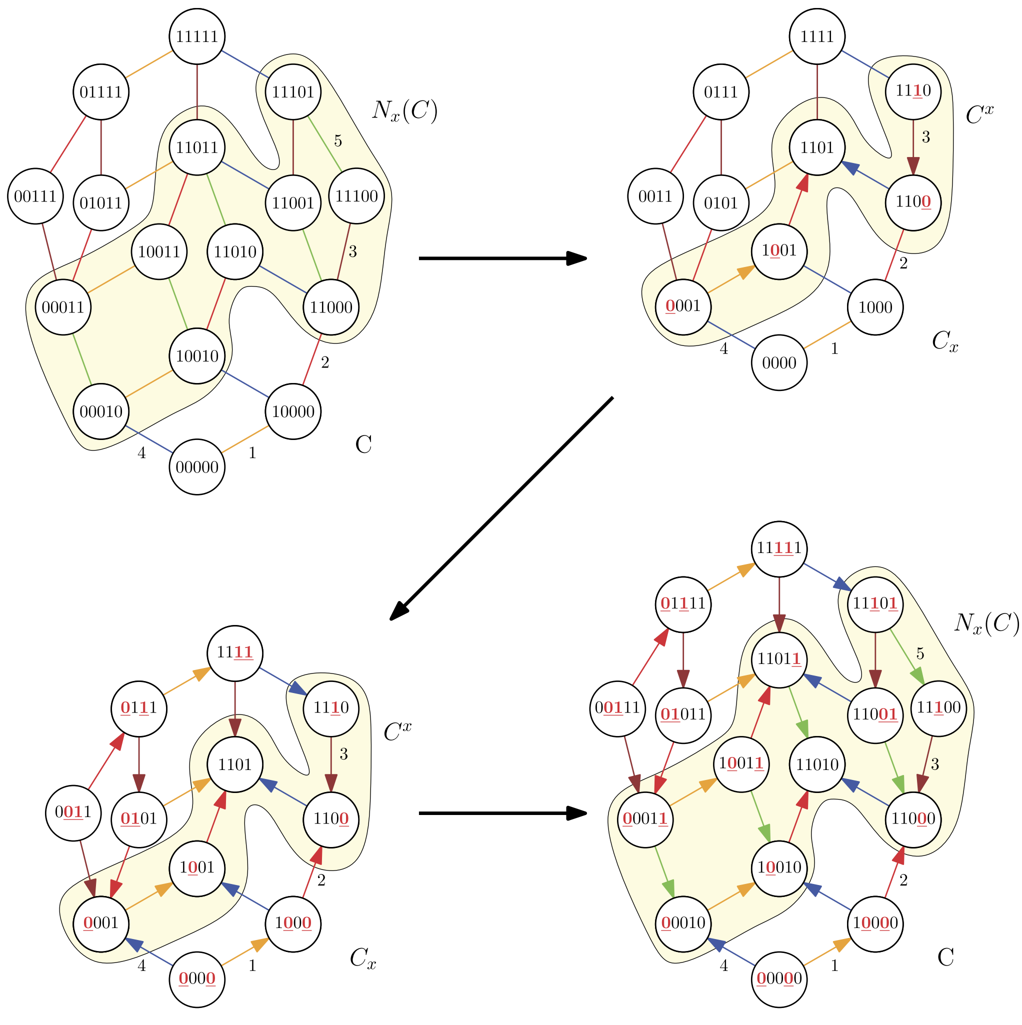

Fig. 5. Illustrating the proof of Theorem 5.1 (when $x = 5$): to construct a representation map for C, we inductively construct a representation map r^x for C^x, extend it to a representation map r_x for C_x using Proposition 5.2 with $D = C^x$, and finally extend it to a representation map r for C. The representation maps r^x, r_x, and r are defined by the coordinates of the underlined bits or by the labels of the outgoing edges (see Theorem 6.8).

Proof of Proposition 5.2. To prove the proposition, we first prove that incomplete cubes and their sources are preserved by restrictions and reductions (Lemma 5.4). To show that each concept of $C \setminus D$ is the source of one incomplete cube, we consider a minimal counterexample and we establish that in this counterexample each concept is the source of at most 2 incomplete cubes. Moreover, if a concept is the source of 0 (respectively, 1, 2) incomplete cubes, then any of its neighbors in $G(C)$ is the source of 2 (respectively, 1, 0) incomplete cubes (Lemma 5.5). Using this and the notions of galleries and the associated trees defined below, we establish the first assertion of the proposition. Then using Lemma 5.4 and the first assertion, we establish the second part of the proposition. We also give a geometric characterization of sources (Lemma 5.3) that is used later to estimate the complexity of computing the representation map (see Remark 5.7).

Call a maximal cube of C a *chamber* and a facet of a chamber a *panel* (a *σ'-panel* if its support is σ'). Any σ'-panel in C satisfies $|\sigma'| = d - 1$ and $\sigma' \in X(D)$. Recall that a gallery between two parallel cubes Q', Q'' (say, two σ'-cubes) is any simple path of σ'-cubes $(Q_0 := Q', Q_1, \ldots, Q_k := Q'')$, where $Q_i \cup Q_{i+1}$ is a d-cube. By Theorem 3.2(3), any two parallel cubes of C are connected by a gallery in C. Since D is a maximum class, any panel of C is parallel to a panel that is a maximal cube of D. Also for any maximal simplex $\sigma' \in X(D)$, the class $C^{\sigma'}$ is a maximum class of dimension 1 and $D^{\sigma'}$ is a maximum class of dimension 0 (single concept). Thus $C^{\sigma'}$ is a tree (e.g. [18, Lemma 7]) which contains the unique concept $c \in D^{\sigma'}$. We call c *the root* of $C^{\sigma'}$ and we denote by $P(\sigma')$ the unique σ'-panel P of D such that $P^{\sigma'} = c$.

The next result provides a geometric characterization of sources:

Lemma 5.3. *Let σ be a missed simplex of the pair (C, D). A concept $c \in Q$ is the source of the unique σ-cube Q if and only if for any $x \in \sigma$, if $\sigma' := \sigma \setminus \{x\}$ and P', P'' are the two σ'-panels of Q with $c \in P''$, then $(P')^{\sigma'}$ is on the path between $(P'')^{\sigma'}$ and the root $(P(\sigma'))^{\sigma'}$ of the tree $C^{\sigma'}$.*

Proof. Observe that there exists a unique concept $c \in Q$ such that for any $x \in \sigma$, if $\sigma' := \sigma \setminus \{x\}$ and P', P'' are the two σ'-panels of Q with $c \in P''$, then $(P')^{\sigma'}$ is on the path between $(P'')^{\sigma'}$ and the root $(P(\sigma'))^{\sigma'}$ of the tree $C^{\sigma'}$. Since Q contains a unique source, it is enough to show that this unique concept c is the source.

Assume by way of contradiction that this is not the case, i.e., that c is the source of Q and that there exists $x \in \sigma$ and two σ'-panels P', P'' with $\sigma' = \sigma \setminus \{x\}$ and $c \in P''$ such that the unique gallery L between P' and the root $P(\sigma')$ passes via P'', i.e., $L = (P_0 = P(\sigma'), P_1, \dots, P_{m-1} = P'', P_m = P')$. Since Q is a maximal cube, by Corollary 3.3, x is not in the domain of the chamber $P_i \cup P_{i+1}$ for $i < m-1$. This implies that there exists $c_0 \in P(\sigma') \subseteq D$ such that $c_0|\sigma = c|\sigma$, and consequently, $c|\sigma$ is not the missed sample for σ. $\square$

In the next lemma, we show that incomplete cubes and their sources are preserved by restrictions and reductions.

Lemma 5.4. *Let Q be an incomplete cube for (C, D) with source s and support σ, and let $x, y \in U$ such that $x \notin \sigma$ and $y \in \sigma$. Then, the following holds:*

(i) Q_x is an incomplete cube for (C_x, D_x) whose source is s_x.
(ii) Q^y is an incomplete cube for (C^y, D^y) whose source is s^y.

Proof. Item (i): C_x and D_x are maximum classes on $U \setminus \{x\}$ of VC-dimensions d and $d-1$, and $\operatorname{supp}(Q_x) = \sigma$. Therefore, Q_x is an incomplete cube for (C_x, D_x). By definition, s is the unique concept $c \in Q$ such that $c|\sigma \notin D|\sigma$. Since $x \notin \sigma$, $D|\sigma = D_x|\sigma$ and s_x is the unique concept c of Q_x so that $c|\sigma \notin D_x|\sigma$, i.e., s_x is the source of Q_x.

Item (ii): C^y and D^y are maximum classes on $U \setminus \{y\}$ of VC-dimensions $d-1$ and $d-2$. Since $y \in \operatorname{supp}(Q)$, $\dim(Q^y) = d-1$ and Q^y is an incomplete cube for (C^y, D^y). Let $\sigma' = \sigma \setminus \{y\}$. It remains to show that $s^y|\sigma' \notin D^y|\sigma'$. Indeed, otherwise both extensions of s^y in σ, namely $s, s\Delta\{y\}$, are in $D|\sigma$ which contradicts that $s = s(Q)$. $\square$

Next we prove that each concept of $C \setminus D$ is the source of a unique incomplete cube. Since there is a bijection between incomplete cubes and missed simplices, the number of incomplete cubes is $|X(C) \setminus X(D)| = |C \setminus D|$. Therefore, it is sufficient to show that each concept of $C \setminus D$ is the source of at most one incomplete cube. Assume the contrary and let (C, D) be a counterexample minimizing the size of U. First, if a concept $c \in C \setminus D$ is the source of two incomplete cubes Q_1, Q_2, then $\operatorname{dom}(C) = \operatorname{supp}(Q_1) \uplus \operatorname{supp}(Q_2)$. Indeed, let $\sigma_1 = \operatorname{supp}(Q_1)$ and $\sigma_2 = \operatorname{supp}(Q_2)$. By Lemma 5.4(i) and minimality of (C, D), $\operatorname{dom}(C) = \sigma_1 \cup \sigma_2$. Indeed, if there exists $x \notin \sigma_1 \cup \sigma_2$, c_x is the source of the incomplete cubes $(Q_1)_x$ and $(Q_2)_x$ for (C_x, D_x), contrary to minimality of (C, D). By Lemma 5.4(ii) and minimality of (C, D), $\sigma_1 \cap \sigma_2 = \varnothing$. Indeed, if there exists $x \in \sigma_1 \cap \sigma_2$, c^x is the source of the incomplete cubes Q_1^x and Q_2^x for (C^x, D^x), contrary to minimality of (C, D).

Next we assert that any $c \in C \setminus D$ is the source of at most 2 incomplete cubes. Indeed, let c be the source of incomplete cubes Q_1, Q_2, Q_3. Then $\operatorname{dom}(C) = \operatorname{supp}(Q_1) \uplus \operatorname{supp}(Q_2)$, i.e., $\operatorname{supp}(Q_2) = \operatorname{dom}(C) \setminus \operatorname{supp}(Q_1)$. For similar reasons, $\operatorname{supp}(Q_3) = \operatorname{dom}(C) \setminus \operatorname{supp}(Q_1) = \operatorname{supp}(Q_2)$. Thus, by Corollary 3.3, $Q_2 = Q_3$.

Lemma 5.5. *Let $c', c'' \in C \setminus D$ be neighbors and let $c' \Delta c'' = \{x\}$. Then, c' is the source of 2 incomplete cubes if and only if c'' is the source of 0 incomplete cubes. Consequently, every connected component in $G(C \setminus D)$ either contains only concepts c with $|s^{-1}(c)| \in \{0, 2\}$, or only concepts c with $|s^{-1}(c)| = 1$.*

Proof. By minimality of (C, D), $(c')^x = (c'')^x$ is the source of a unique incomplete cube for (C^x, D^x) and $c'_x = c''_x$ is the source of a unique incomplete cube for (C_x, D_x). Let Q_1 be the incomplete cube for (C, D) such that $(c')_x$ is the source of $(Q_1)_x$. Let Q_2 be the incomplete cube for (C, D) such that $(c')^x$ is the source of $(Q_2)^x$. Since $(Q_1)_x$ has a unique source $c'_x = c''_x$, by Lemma 5.4(i), $(s(Q_1))_x = c'_x$ and consequently, $s(Q_1) \in \{c', c''\}$. Similarly, since Q_2^x has a unique source $(c')^x = (c'')^x$, by Lemma 5.4(ii), $(s(Q_2))^x = (c')^x$ and thus $s(Q_2) \in \{c', c''\}$. Consequently, c' is the source of 2 incomplete cubes (Q_1 and Q_2) if and only if c'' is the source of 0 incomplete cubes. $\square$

Pick $c \in C \setminus D$ that is the source of two incomplete cubes for (C, D) and an incomplete cube Q such that $c = s(Q)$. Let $\sigma = \operatorname{supp}(Q)$, $x \in \sigma$, and $\sigma' = \sigma \setminus \{x\}$. The concept c belongs to a unique σ'-panel P. Let $L = (P_0 = P(\sigma'), P_1, \dots, P_{m-1}, P_m = P)$ be the unique gallery between the root $P(\sigma')$ of the tree $C^{\sigma'}$ and P. For $i = 1, \dots, m$, denote the chamber $P_{i-1} \cup P_i$ by Q_i. Since $Q_i \cap D$ is ample and Q_i is not contained in D, it follows that the complement $Q_i \setminus D$ is a nonempty ample class. Hence $Q_i \setminus D$ induces a nonempty connected subgraph of $G(C \setminus D)$. Therefore, it follows that c and each concept $c' \in Q_i \setminus D$ are connected by a path in $G(C \setminus D)$, and by Lemma 5.5 it follows that

For each i, each $c' \in Q_i \setminus D$ is the source of either 0 or 2 incomplete cubes. (5.1)

Consider the chamber $Q_1 = P_0 \cup P_1$ and its source $s = s(Q_1)$. By the definition of the source, necessarily $s \in P_1$ and $s \notin D$. Therefore, Property (5.1) implies that there must exist another cube Q' such that $s = s(Q')$. Let s' be the neighbor of s in $P_0 = P(\sigma')$; note that $s' \in D$. Since $\operatorname{supp}(Q_1) \cap \operatorname{supp}(Q') = \varnothing$, it follows that $s|\operatorname{supp}(Q') = s'|\operatorname{supp}(Q') \in D|\operatorname{supp}(Q')$, contradicting that $s = s(Q')$. This establishes the first assertion of Proposition 5.2.

We prove now that the map r defined in Proposition 5.2 is a representation map for C. It is easy to verify that it is a bijection between C and $X(C)$, so it remain to establish the non-clashing property: $c|(r(c) \cup r(c')) \neq c'|(r(c) \cup r(c'))$ for all distinct pairs $c, c' \in C$. This holds when $c, c' \in D$ because r' is a representation map. Next, if $c \in C \setminus D$ and $c' \in D$, this holds because $c|r(c) \notin D|r(c)$ by the properties of s.

Thus, it remains to show that every distinct $c, c' \in C \setminus D$ satisfy the non-clashing condition. Assume towards contradiction that this does not hold and consider a counterexample with minimal domain size $|U|$. Consequently, there exist distinct $c, c' \in C \setminus D$ such that $c(z) = c'(z)$ for any $z \in \operatorname{supp}(Q) \cup \operatorname{supp}(Q')$, where $Q = s^{-1}(c)$ and $Q' = s^{-1}(c')$. Since $c \neq c'$, there exists $x \in U \setminus (\operatorname{supp}(Q) \cup \operatorname{supp}(Q'))$ such that $c(x) \neq c(x')$. If there exists $y \in U \setminus (\operatorname{supp}(Q) \cup \operatorname{supp}(Q'))$ distinct from x, then from Lemma 5.4(1), Q_y and Q'_y are incomplete cubes for (C_y, D_y) whose respective sources are c_y and c'_y. Since $r(c_y) = r(c)$, $r(c'_y) = r(c')$, and $c_y \neq c'_y$ (since they differ on x), c_y and c'_y clash in C_y, contradicting the minimality of the counterexample (C, D). Consequently, we can assume that $U = \operatorname{supp}(Q) \cup \operatorname{supp}(Q') \cup \{x\}$ and that c and c' differ only on x. In this case, by Lemma 5.4(1), $c_x = c'_x$ is the source of both incomplete cubes Q_x and Q'_x. Since Q and Q' are distinct incomplete cubes, we have $\operatorname{supp}(Q_x) = \operatorname{supp}(Q) \neq \operatorname{supp}(Q') = \operatorname{supp}(Q'_x)$ and thus, $c_x = c'_x$ is the source of two different incomplete cubes of C_x, contradicting the first statement of the proposition. This ends the proof of Proposition 5.2. $\square$

Remark 5.6. Proposition 5.2 relies on a canonical bijection between missed simplices and incomplete cubes (that allows to define the source-map). This is due to the fact that any missed simplex for (C, D) is a maximal simplex of C and is thus the support of a unique cube in C. This property does not longer hold for general ample classes since there are missed simplices that are not maximal and therefore there are the supports of several cubes in C. A first step to generalize Proposition 5.2 to ample classes could be to find a way to define a source-map between the missed simplices in $X(C) \setminus X(D)$ and the concepts in $C \setminus D$.

Remark 5.7. To compute the representation map for a maximum class C of dimension d, we make d recursive calls. For each call, the costliest operation is to compute the source-map that can be done in $O(|C|^3)$ time as follows. First we naively compute in $O(|C|^3)$ the cubes of dimension d in C and the cubes of dimension $d-1$ in C^x. Then we compute in $O(d|C|)$ time the trees $C^{\sigma'}$ for any maximal simplex $\sigma' \in X(C^x)$ (the roots of those trees are the cubes of C^x of dimension $d-1$) and we can then compute the source of each cube by Lemma 5.3. Consequently, one can compute a representation map for a maximum class C of dimension d in $O(d|C|^3)$.

Remark 5.8. In the appendix, we give two representation maps for Hall's concept class C_H presented in Fig. 3. One of this representation map is obtained by the method described in the proof of Theorem 5.1. The other one is obtained by transforming the problem into a SAT formula and using a SAT solver.

6. Representation maps for ample classes

In this section, we provide combinatorial and geometric characterizations of representation maps of ample classes. We first show that representation maps lead to optimal unlabeled sample compression schemes and that they are equivalent to unique sink orientations (USO) of $G(C)$ (see below for the two conditions defining USOs). We also show that corner peelings of ample classes are equivalent to the existence of acyclic USOs. In Section 6.3, we show how from representation maps for an ample class C to derive representation maps for substructures of C (intersections with cubes, restrictions and reductions). In Section 6.4, we show that there exist maps satisfying each one of the two conditions defining USOs (but not both). Finally, using the geometric characterization of representation maps as USOs, we show that constructing a representation map for an ample concept class can be reduced to solving an instance of the Independent System of Representatives problem [2].

6.1. Unlabeled sample compression schemes and representation maps

In the next theorem, we prove that, analogously to maximum classes, representation maps for ample classes lead to unlabeled sample compression schemes of size VC-dim(C). This also shows that the representation maps for ample classes are equivalent to Δ-representation maps.

Theorem 6.1. *Let $C \subseteq 2^U$ be an ample class and let $r : C \to X(C)$ be a bijection. The following conditions are equivalent:*

(R1) $\cup$-non-clashing: For all distinct concepts $c', c'' \in C$, $c'|r(c') \cup r(c'') \neq c''|r(c') \cup r(c'')$.
(R2) Reconstruction: For every realizable sample s of C, there is a unique $c \in C$ that is consistent with s and $r(c) \subseteq \operatorname{dom}(s)$.
(R3) Cube injective: For every cube B of 2^U, the map $c \mapsto r(c) \cap \operatorname{supp}(B)$ is an injection from $C \cap B$ to $X(C \cap B)$.
(R4) Δ-non-clashing: For all distinct concepts $c', c'' \in C$, $c'|r(c') \Delta r(c'') \neq c''|r(c') \Delta r(c'')$.

Moreover, any Δ-non-clashing map $r : C \to X(C)$ is bijective and is therefore a representation map. Furthermore, if r is a representation map for C, then there exists an unlabeled sample compression scheme for C of size VC-dim(C).

Proof. Fix $Y \subseteq U$ and partition C into equivalence classes where two concepts c, c' are equivalent if $c|Y = c'|Y$. Thus, each equivalence class corresponds to a sample of C with domain Y, i.e., a concept in $C|Y$. We first show that the number of such equivalence classes equals the number of concepts whose representation set is contained in Y:

$$\begin{aligned} |C|Y| &= |\overline{X}(C|Y)| && \text{(Since } C|Y \text{ is ample)} \\ &= |\overline{X}(C) \cap 2^Y| \\ &= |\{c : r(c) \subseteq Y\}| && \text{(Since } r : C \to \overline{X}(C) = \underline{X}(C) \text{ is a bijection)} \end{aligned}$$

Condition (R2) asserts that in each equivalence class there is exactly one concept c such that $r(c) \subseteq Y$.

(R1) $\Rightarrow$ (R2): Assume $\neg(R2)$ and consider a sample s for which the property does not hold. This implies that there exists an equivalence class with either zero or (at least) two equivalent concepts c for which $r(c) \subseteq Y$ with $Y = \mathrm{dom}(s)$. Note that since the number of equivalence classes equals the number of concepts whose representation set is contained in Y, if some equivalence class has no concept c for which $r(c) \subseteq Y$, then there must be another equivalence class with two distinct concepts $c', c'' \in C$ for which $r(c'), r(c'') \subseteq Y$. Therefore, in both cases, there exist two equivalent concepts $c, c' \in C$ such that $r(c), r(c') \subseteq Y$. Since $c|Y = c'|Y$, we have $c|r(c) \cup r(c') = c'|r(c) \cup r(c')$, contradicting $(R1)$.

(R2) $\Rightarrow$ (R1): Assume $\neg(R1)$, i.e., for two distinct concepts $c', c'' \in C$, we have $c'|r(c') \cup r(c'') = c''|r(c') \cup r(c'')$. Now for the sample $s = c'|r(c') \cup r(c'')$, we have $\mathrm{dom}(s) = r(c') \cup r(c'')$. Furthermore, $c'|\mathrm{dom}(s) = c''|\mathrm{dom}(s)$ and $r(c'), r(c'') \subseteq \mathrm{dom}(s)$. This implies $\neg(R2)$.

(R1) & (R2) $\Rightarrow$ (R3): Since $C \cap B$ is ample, it suffices to show that for every $Y \in X(C \cap B)$ there is some $c \in C \cap B$ with $r(c) \cap \mathrm{supp}(B) = Y$. This is established by the following claim.

Claim 6.2. *Conditions (R1) and (R2) together imply that for any $Y \in X(C \cap B)$, there exists a unique concept $c_Y \in C \cap B$ such that $r(c_Y) \cap \mathrm{supp}(B) = Y$.*

Proof. For any concept $c \in C \cap B$, let $r_B(c) := r(c) \cap \mathrm{supp}(B)$. Let $Z = U \setminus \mathrm{supp}(B)$. Note that all concepts in B agree on domain Z: indeed, $B|Z$ is the single sample $\mathrm{tag}(B)$. We prove the claim by induction on $|Y|$.

Suppose first that $Y = \varnothing$. Since $X(C \cap B) \neq \varnothing$, $C \cap B \neq \varnothing$ and $\mathrm{tag}(B)$ is a sample of C. By condition (R2), there is a unique concept $c \in C$ such that (i) $c|Z = \mathrm{tag}(B)$ (i.e., $c \in B$), and (ii) $r(c) \subseteq Z$ (i.e., $r_B(c) = \varnothing$). Thus choosing $c_\varnothing = c$ settles this case.

Assume now that $Y \neq \varnothing$ and that for every $V \subsetneq Y$, there is a unique $c_V \in C \cap B$ with $r_B(c_V) = r(c_V) \cap \mathrm{supp}(B) = V$.

We assert that for any two distinct concepts $c, c' \in C \cap B$ such that $r_B(c) \cup r_B(c') \subseteq Y$, we have $c|Y \neq c'|Y$. Indeed, suppose by way of contradiction that $c|Y = c'|Y$ and note that $r_B(c) \cup r_B(c') = (r(c) \cup r(c')) \cap \mathrm{supp}(B) \subseteq Y$. If $c|Y = c'|Y$, then

$$c|(r(c) \cup r(c')) \cap \mathrm{supp}(B) = c'|(r(c) \cup r(c')) \cap \mathrm{supp}(B).$$

Since $c|Z = c'|Z = \mathrm{tag}(B)$, this implies that $c|r(c) \cup r(c') = c'|r(c) \cup r(c')$, contradicting condition (R1).

Consequently, all the samples $c_V|Y$ are pairwise distinct. There are $2^{|Y|} - 1$ such samples and each sample $c_V|Y$ corresponds to a proper subset V of Y. By the previous assertion, there exist at most $2^{|Y|}$ concepts c such that $r_B(c) \subseteq Y$, and thus there exists at most one concept $c_Y \in C \cap B$ such that $r_B(c_Y) = Y$.

Thus, it remains to establish the existence of $c_Y \in C \cap B$ such that $r_B(c_Y) = r(c_Y) \cap \mathrm{supp}(B) = Y$. Since there are $2^{|Y|} - 1$ c_V's for $V \subsetneq Y$, it follows that there is a unique sample s' with $\mathrm{dom}(s') = Y$ that is not realized by any of the c_V's. Consider the sample s with domain $Y \cup Z$ defined by

$$s(x) = \begin{cases} s'(x) & \text{if } x \in Y, \\ \mathrm{tag}(B)(x) & \text{if } x \in Z. \end{cases}$$

Since Y is shattered by $C \cap B$, it follows that s is realized by C. By condition (R2) there is a unique concept $c \in C$ that agrees with s such that $r(c) \subseteq Y \cup Z$ (i.e. $r(c) \cap \mathrm{supp}(B) \subseteq Y$). We claim that c is the desired concept c_Y. First notice that $c \in C \cap B$, because c agrees with $\mathrm{tag}(B)$ on Z. Since s is not realized by any c_V, $V \subsetneq Y$, and since c_V is the unique concept of B such that $r_B(c_V) = V$ (by induction hypothesis), necessarily we have that $c \neq c_V$ for any $V \subsetneq Y$. Consequently, $r_B(c) = Y$, concluding the proof of the claim. □

(R3) $\Rightarrow$ (R4): For any distinct concepts $c', c'' \in C$, consider the minimal cube $B := B(c', c'')$ which contains both c', c''. This means that $c'(x) \neq c''(x)$ for every $x \in \mathrm{supp}(B)$, and that $c'(x) = c''(x)$ for every $x \notin \mathrm{supp}(B)$. Condition (R3) guarantees that the map $r(c) \mapsto r(c) \cap \mathrm{supp}(B)$ is an injection from $C \cap B$ to $X(C \cap B)$. Therefore $r(c') \cap \mathrm{supp}(B) \neq r(c'') \cap \mathrm{supp}(B)$. It follows that there must be some $x \in \mathrm{supp}(B)$ such that $x \in (r(c') \cap \mathrm{supp}(B)) \Delta (r(c'') \cap \mathrm{supp}(B))$. Since $x \in \mathrm{supp}(B)$, $c'(x) \neq c''(x)$ and therefore $c'|r(c') \Delta r(c'') \neq c''|r(c') \Delta r(c'')$ and condition (R4) holds for c' and c''.

(R4) $\Rightarrow$ (R1): This is immediate because if two concepts clash on their symmetric difference, then they also clash on their union.

Moreover, observe that for any map $r: C \to X(C)$, if $r(c) = r(c')$ for $c \neq c'$, then $r(c)\Delta r(c') = \varnothing$ and r is not Δ-non-clashing. Consequently, any Δ-non-clashing map $r: C \to X(C)$ is injective and thus bijective since $|C| = |X(C)|$.

We now show that if $r: C \to X(C)$ is a representation map for C then there exists an unlabeled sample compression scheme for C. Indeed, by (R2), for each realizable sample $s \in \mathrm{RS}(C)$, let $\gamma(s)$ be the unique concept $c \in C$ such that $r(c) \subseteq \mathrm{dom}(s)$ and $c|\mathrm{dom}(s) = s$. Then consider the compressor $\alpha : \mathrm{RS}(C) \to X(C)$ such that for any $s \in \mathrm{RS}(C)$, $\alpha(s) = r(\gamma(s))$ and the reconstructor $\beta : X(C) \to C$ such that for any $Z \in X(C)$, $\beta(Z) = r^{-1}(Z)$. Observe that by the definition of $\gamma(s)$, $\alpha(s) \subseteq \mathrm{dom}(s)$ and $\beta(\alpha(s)) = \gamma(s)$ coincides with s on $\mathrm{dom}(s)$. Consequently, α and β defines an unlabeled sample compression scheme for C of size $\dim(X(C)) = \text{VC-dim}(C)$. This concludes the proof of the theorem. $\square$

By applying (R3) to the 1-dimensional cubes of B, we get the following corollary.

Corollary 6.3. *For any representation map r of an ample class C, for any $c \in C$ and $x \in r(c)$, we have $c\Delta\{x\} \in C$.*

6.2. Representation maps as unique sink orientations

We call a map $r: C \to 2^U$ *edge-non-clashing* if for any x-edge cc', $x \in r(c)\Delta r(c')$. Note that r defines an orientation o_r of the edges of $G(C)$: an x-edge cc' is oriented from c to c' if and only if $x \in r(c) \setminus r(c')$. Conversely, given an orientation o of the edges of $G(C)$, the *out-map* r_o of o associates to each $c \in C$ the coordinate set of the edges outgoing from c. Note that the out-map of any orientation is edge-non-clashing. Note also that any Δ-representation map $r: C \to X(C)$ is edge-non-clashing and thus defines an orientation o_r of $G(C)$. Moreover, by Corollary 6.3, the out-map of o_r coincides with r.

An orientation o of the edges of $G(C)$ (or the corresponding out-map r_o) is a *unique sink orientation (USO)* if it satisfies the following two conditions.

(C1) For any $c \in C$, the cube of 2^U which has support $r_o(c)$ and contains c is a cube of C, i.e., all outgoing neighbors of c belong to a cube of C;
(C2) For any cube B of C, there exists a unique $c \in C \cap B$ such that $r_o(c) \cap \mathrm{supp}(B) = \varnothing$, i.e., c is a sink in $G(C \cap B)$.

Observe also that a map $r: C \to 2^U$ satisfying Condition (C2) is necessarily an edge-non-clashing map. If C is a cube, then Condition (C1) trivially holds and Condition (C2) corresponds to the usual definition of USOs on cubes [41]. In the following, we use the characterization of USOs for cubes given in [41].

Lemma 6.4 *([41]). For a cube B and a map $r: B \to 2^{\mathrm{supp}(B)}$, the following are equivalent:*

(1) r is the out-map of a unique sink orientation of B;
(2) r is Δ-non-clashing;
(3) for any subcube B' of B, the map $c \mapsto r(c) \cap \mathrm{supp}(B')$ is a bijection between B' and $2^{\mathrm{supp}(B')}$;
(4) r is the out-map of a unique source orientation *of B, i.e., for any subcube B' of B, there exists a unique $c \in B'$ such that $r(c) \cap \mathrm{supp}(B') = \mathrm{supp}(B)$.*

The equivalences (1) ⇔ (2) and (1) ⇔ (4) are respectively [41, Lemma 2.3] and [41, Lemma 2.1]. The implications (2) ⇒ (3) and (3) ⇒ (1) are trivial.

Remark 6.5. In view of Lemma 6.4, Condition (C2) looks similar to Condition (R3) of Theorem 6.1. Note however that (C2) is about the cubes of 2^U contained in C while (R3) is about all cubes of 2^U. Similarly, Condition (C2) implies that r_o is Δ-non-clashing on each cube of 2^U contained in C while condition (R4) of Theorem 6.1 requires that r_o is Δ-non-clashing on C.

Remark 6.6. If C is an ample class and o is a USO of $G(C)$, then for any cube B of C, the restriction of o to the edges of $G(B)$ trivially satisfies (C1) and (C2) and is thus a USO of $G(B)$. Consequently, its out-map $r_B : c \mapsto r_o(c) \cap \mathrm{supp}(B)$ from B to $2^{\mathrm{supp}(B)}$ satisfies the conditions of Lemma 6.4.

We now show that representation maps for ample classes give rise to USOs.

Corollary 6.7. *If $r: C \to X(C)$ is a representation map for an ample class $C \subseteq 2^U$, then o_r is a unique sink orientation.*

Proof. Pick some concept $c \in C$ and consider the unique cube B of 2^U that contains c and has support $\mathrm{supp}(B) = r(c)$. Since r is a representation map, by Condition (R3) of Theorem 6.1, the map $c' \mapsto r(c') \cap \mathrm{supp}(B)$ is a bijection between the ample set $C \cap B$ and $X(C \cap B)$. Consequently, $r(c) = \mathrm{supp}(B) \in X(C \cap B)$, which is possible only if $C \cap B = B$. This proves (C1). Condition (C2) follows from Condition (R4) of Theorem 6.1 applied to the cubes of C and Lemma 6.4. $\square$

We continue with a characterization of representation maps of ample classes as out-maps of USOs, extending a similar result of Szabó and Welzl [41] for cubes. This characterization is "local-to-global", since (C1) and (C2) are conditions on the cubes around each concept $c \in C$.

Theorem 6.8. *For an ample class C and a map $r : C \to 2^U$, the following are equivalent:*

(1) r is a representation map;
(2) r is the out-map of a USO;
(3) $r(c) \in X(C)$ for any $c \in C$ and r satisfies (C2).

Before proving the theorem, starting from an edge-non-clashing map for an ample class C, we show how to derive maps for restrictions C_x, reductions C^x and intersections $C \cap B$ with cubes of 2^U. Consider an edge-non-clashing map $r : C \to X(C)$ for an ample class C. Given a cube B of 2^U, define $r_B : C \cap B \to X(C \cap B)$ by setting $r_B(c) := r(c) \cap \mathrm{supp}(B)$ for any $c \in C \cap B$. Note that r_B is the out-map of the orientation o_r restricted to the edges of $G(C \cap B)$.

Given $x \in U = \mathrm{dom}(C)$, define $r^x : C^x \mapsto 2^{U \setminus \{x\}}$ by

$$r^x(c) = \begin{cases} r(c) \setminus \{x\} & \text{if } x \in r(c), \\ r(c \cup \{x\}) \setminus \{x\} & \text{otherwise.} \end{cases}$$

Hence, $r^x(c \setminus \{x\}) = r(c) \setminus \{x\}$ for each $c \in C$ with $x \in r(c)$. Consequently, for an x-edge of C between c and $c \cup \{x\}$, $r^x(c)$ is the label of the origin of this edge minus x; we call r^x the *x-out-map* of r.

Given $x \in U = \mathrm{dom}(C)$, define $r_x : C_x \mapsto 2^{U \setminus \{x\}}$ by

$$r_x(c) = \begin{cases} r(c) & \text{if } x \notin r(c), \\ r(c \cup \{x\}) & \text{otherwise.} \end{cases}$$

Hence, $r_x(c \setminus \{x\}) = r(c)$ for each $c \in C$ with $x \notin r(c)$. Consequently, for an x-edge of C between c and $c \cup \{x\}$, $r_x(c)$ is the label of the destination of this edge; we call r_x the *x-in-map* of r.

In the next lemma, we show that if the map r is the out-map of a USO, then r_B and r^x are also outmaps of USOs. A similar result holds for r_x but it will be proved in Section 6.3. The first assertion of the lemma generalizes Remark 6.6.

Lemma 6.9. *Consider a map $r : C \to 2^U$ that is the out-map of a USO of an ample class C. For any cube $B \subseteq 2^U$ and any $x \in U = \mathrm{dom}(C)$, the following hold:*

(1) r_B is the out-map of a USO of $C \cap B$;
(2) r^x is the out-map of a USO of C^x.

Proof. Item (1): Consider a cube $B \subseteq 2^U$. Since (C1) holds for C and r, for any concept $c \in C \cap B$, the cube $B(c)$ that has support $r(c)$ and contains c is a cube of C. Since the intersection of two cubes is a cube, the cube $B'(c) = B(c) \cap B$ has support $r(c) \cap \mathrm{supp}(B) = r'(c)$ and contains c. Consequently, (C1) holds for C' and r'. Since any cube of $C' = C \cap B$ is also a cube of C and since (C2) holds for C and r, (C2) also holds for C' and r'.

Item (2): Consider $c \in C^x$ and let $c' \in \{c, c \cup \{x\}\}$ such that $r^x(c) = r(c') \setminus \{x\}$. By (C1), there exists an $r(c')$-cube B' containing c' in C. Thus, there exists an $(r(c') \setminus \{x\})$-cube containing c in C^x. Consequently, r^x satisfies (C1).

Suppose that there exists a cube B in C^x violating (C2), i.e., there exist $c_1, c_2 \in C^x \cap B$ such that $r^x(c_1) \cap \mathrm{supp}(B) = r^x(c_2) \cap \mathrm{supp}(B)$. By definition of r^x, there exist $c'_1 \in \{c_1, c_1 \cup \{x\}\}, c'_2 \in \{c_2, c_2 \cup \{x\}\}$ such that $r(c'_1) = r^x(c_1) \cup \{x\}$ and $r(c'_2) = r^x(c_2) \cup \{x\}$. Since B is a cube of C^x containing c_1 and c_2, there exists a cube B' of C containing c'_1 and c'_2 such that $\mathrm{supp}(B') = \mathrm{supp}(B) \cup \{x\}$. Consequently, $r(c'_1) \cap \mathrm{supp}(B') = (r^x(c_1) \cup \{x\}) \cap (\mathrm{supp}(B) \cup \{x\}) = (r^x(c_1) \cap \mathrm{supp}(B)) \cup \{x\}$ and similarly, $r(c'_2) \cap \mathrm{supp}(B') = (r^x(c_2) \cap \mathrm{supp}(B)) \cup \{x\}$. Therefore, $r(c'_1) \cap \mathrm{supp}(B') = r(c'_2) \cap \mathrm{supp}(B')$ and r is not injective on the cube B' of C, contradicting Lemma 6.4 (see Remark 6.6). □

Proof of Theorem 6.8. The implication (1) ⇒ (2) is established in Corollary 6.7. Now, we prove (2) ⇒ (1). Clearly, property (C1) implies that $r(c) \in X(C)$ for any $c \in C$, whence r is a map from C to $X(C)$.

Let C be an ample class of smallest size admitting a non-representation map $r : C \to X(C)$ satisfying (C1) and (C2). Hence there exist $u_0, v_0 \in C$ such that $u_0|(r(u_0)\Delta r(v_0)) = v_0|(r(u_0)\Delta r(v_0))$, i.e., $(u_0 \Delta v_0) \cap (r(u_0)\Delta r(v_0)) = \varnothing$; (u_0, v_0) is called a *clashing pair.*

Claim 6.10. *If (u_0, v_0) is a clashing pair, then $C = C \cap B(u_0, v_0)$ and $r(u_0) = r(v_0) = \varnothing$.*

Proof. Since $C \cap B(u_0, v_0)$ is ample and $(u_0 \Delta v_0) \cap (r(u_0)\Delta r(v_0)) = \varnothing$, (u_0, v_0) is a clashing pair for $C \cap B(u_0, v_0)$ and the restriction r_B of r to $\mathrm{supp}(B(u_0, v_0))$. By Lemma 6.9, r_B is the out-map of a USO of $C \cap B(u_0, v_0)$. Consequently, by minimality of C, $C = C \cap B(u_0, v_0)$ and thus $\mathrm{dom}(C) = \mathrm{supp}(B(u_0, v_0)) = u_0 \Delta v_0$. Moreover, if $r(u_0) \neq r(v_0)$, then there is $x \in r(u_0)\Delta r(v_0)$ and $x \in \mathrm{supp}(B(u_0, v_0)) = u_0 \Delta v_0$, contradicting that (u_0, v_0) is a clashing pair.

Suppose $r(u_0) \neq \varnothing$ and pick $x \in r(u_0) = r(v_0)$. Consider the x-out-map $r^x : C^x \to 2^{U\setminus\{x\}}$. By Lemma 6.9, r^x is the out-map of a USO for C^x. Suppose without loss of generality that $x \in v_0 \setminus u_0$. Let $u'_0 = u_0$ and $v'_0 = v_0 \setminus \{x\}$. Then $r^x(u'_0) = r(u_0) \setminus \{x\} = r(v_0) \setminus \{x\} = r^x(v'_0)$, and consequently (u'_0, v'_0) is a clashing pair for r^x on C^x. Since C^x is ample, smaller than C, and since r^x is the out-map of a USO of C^x, this contradicts the minimality of C. □

Claim 6.11. *C is a cube minus a vertex.*

Proof. By (C2) and Claim 6.10, C is not a cube. If C is not a cube minus a vertex, since the complement $C^* = 2^U \setminus C$ is also ample (thus $G(C^*)$ is connected), $G(C^*)$ contains an x-edge ww' with $x \notin w$ and $x \in w'$. Consider the x-in-map $r_x : C_x \mapsto 2^{U\setminus\{x\}}$.

For any $u' \in C_x$, let $u \in \{u', u' \cup \{x\}\}$ such that $r_x(u') = r(u)$. Since C satisfies (C1), u belongs to an $r(u)$-cube in C and consequently, u' belongs to an $r_x(u')$-cube in C_x. Thus C_x and r_x satisfy (C1).

Suppose that C_x and r_x violate (C2). Then there exists a cube B' of C_x and $u', v' \in B'$ such that $(u'\Delta v') \cap (r_x(u')\Delta r_x(v')) = \varnothing$. Without loss of generality, we can assume that $B' = B(u', v')$. Let $u \in \{u', u' \cup \{x\}\}$ such that $r(u) = r_x(u')$ and let $v \in \{v', v' \cup \{x\}\}$ such that $r(v) = r_x(v')$. The restriction C'_x of the ample class $C' := C \cap B(u, v)$ is the cube B'. By Lemma 6.9, $r_{B(u,v)}$ is the out-map of an USO of C'. Since $w, w' \notin C$ and ww' is an x-edge, $w \notin C'_x$. Thus there exists $y \in \mathrm{supp}(C)$ such that C' and the edge ww' of C^* belong to different y-half-spaces $C^- = \{c \in C : y \notin c\}$ and $C^+ = \{c \in C : y \in c\}$ of the cube 2^U. Since $y \in \mathrm{supp}(C)$, the half-space containing ww' also contains a concept of C. Hence, C' is a proper ample subset of C. Since $u \subseteq u' \cup \{x\}$, $v \subseteq v' \cup \{x\}$, $x \notin r(u) = r_x(u')$, $x \notin r(v) = r_x(v')$, we deduce that $u \cap (r(u)\Delta r(v)) = u' \cap (r_x(u')\Delta r_x(v'))$ and $v' \cap (r_x(u')\Delta r_x(v')) = v \cap (r(u)\Delta r(v))$. Since $(u'\Delta v') \cap (r_x(u')\Delta r_x(v')) = \varnothing$, (u, v) is a clashing pair for the restriction of r on C', contrary to the minimality of C. Consequently, C_x and r_x satisfy (C2) and r_x is the outmap of an USO of C_x. By minimality of C, r_x is a representation map for C_x.

Consider a clashing pair (u_0, v_0) for C and r, and let $u'_0 = u_0 \setminus \{x\}$ and $v'_0 = v_0 \setminus \{x\}$. By Claim 6.10, $r_x(u'_0) = r(u_0) = r(v_0) = r_x(v'_0) = \varnothing$. Since r_x is a representation map for C_x, necessarily $u'_0 = v'_0$. Consequently, $u_0 \Delta v_0 = \{x\}$, i.e., $u_0 v_0$ is an x-edge of $G(C)$. This is impossible since C satisfies (C2). Therefore, C is necessarily a cube minus a vertex. □

Now, we complete the proof of the implication (2) ⇒ (1). By Claim 6.10, $r(u_0) = r(v_0) = \varnothing$. By condition (C1) and Claim 6.11, $r(c) \neq U$ for any $c \in C$. Thus there exists a set $s \in X(C) = 2^U \setminus \{U, \varnothing\}$ such that $s \neq r(c)$ for any $c \in C$. Every s-cube B of C contains a source $p(B)$ for o_{r_B} (i.e., $s \subseteq r(p(B))$). For each s-cube B of C, let $t(B) = r(p(B)) \setminus s$. Notice that $\varnothing \subsetneq t(B) \subsetneq U \setminus s$ since $s \subsetneq r(p(B)) \subsetneq U$. Consequently, the number of distinct sets $t(B)$ (when B runs over the s-cubes of C) is at most $2^{|U|-|s|} - 2$. On the other hand, since C is a cube minus one vertex by Claim 6.11, there are $2^{|U|-|s|} - 1$ s-cubes in C. Consequently, there exist two s-cubes B, B' such that $t(B) = t(B')$. Thus $\varnothing \subsetneq s \subsetneq r(p(B)) = r(p(B'))$ and $(p(B), p(B'))$ is a clashing pair for C and r, contradicting Claim 6.10.

The implication (2) ⇒ (3) is trivial. To prove (3) ⇒ (2), we show by induction on $|U|$ that a map $r : C \to X(C)$ satisfying (C2) also satisfies (C1). For any $x \in U$, consider the x-out-map r^x. Recall that if cc' is an x-edge directed from c to c', then $x \in r(c)$ and r^x maps $c^x = (c')^x \in C^x$ to $r(c) \setminus \{x\} \in X(C^x)$. Thus r^x maps C^x to $X(C^x)$. Moreover, each cube B^x of C^x is contained in a unique cube B of C such that $\mathrm{supp}(B) = \mathrm{supp}(B^x) \cup \{x\}$. If there exist $c^x_1, c^x_2 \in B^x$ such that $r^x(c^x_1) = r(c^x_2)$, then there exist $c_1, c_2 \in B$ such that $r(c_1) = r^x(c^x_1) \cup \{x\} = r^x(c^x_2) \cup \{x\} = r(c_2)$, contradicting (C2). Consequently, o_{r^x} satisfies (C2). By induction hypothesis, o_{r^x} satisfies (C1) for any $x \in U$.

For any concept $c \in C$, pick $x \in r(c)$. Since r^x satisfies (C1), c^x belongs to a σ'-cube in C^x with $\sigma' = r^x(c^x) = r(c) \setminus \{x\}$. This implies that c belongs to a σ-cube in C with $\sigma = \sigma' \cup \{x\} = r(c)$. Thus o_r satisfies (C1), concluding the proof of Theorem 6.8. □

A consequence of Theorems 6.1 and 6.8 is that corner peelings correspond *exactly* to acyclic unique sink orientations.

Proposition 6.12. *An ample class C admits a corner peeling if and only if there exists an acyclic orientation o of the edges of $G(C)$ that is a unique sink orientation.*

Proof. Suppose that $C_< = (c_1, \ldots, c_m)$ is a corner peeling and consider the orientation o of $G(C)$ where an edge $c_i c_j$ is oriented from c_i to c_j if and only if $i > j$. Clearly, this orientation is acyclic. For any i, since c_i is a corner in $C_i = \{c_1, \ldots, c_i\}$, the outgoing neighbors of c_i belong to a cube of C_i, i.e., o satisfies (C1). For any cube B of C, assume that c_{i_B} is the first concept of B in the ordering $C_<$. Observe that c_{i_B} is a sink of B for the orientation o. Note that for each $i > i_B$, $\varnothing \subsetneq C_{i-1} \cap B \subseteq C_i \cap B$. By Theorem 3.1(6), $C_i \cap B$ is ample and thus connected. Consequently, for each $i > i_B$, there exists $c_i c_j$ in $G(C)$ with $i_B \leq j < i$ such that $c_j \in B$. Consequently, since $c_i c_j$ is oriented from c_i to c_j, c_i is not a sink of B. Therefore every cube B has a unique sink for the orientation o and o is an acyclic unique sink orientation of $G(C)$.

Suppose now that $G(C)$ admits an acyclic unique sink orientation o. Consider a concept c that is a source for o. By (C1) all the neighbors of c belong to a cube of C and consequently, c is a corner of C. Since C is ample, by Lemma 4.1,

$C' = C \setminus \{c\}$ is ample. Clearly the restriction o' of the orientation o to C' is acyclic. We claim that o' is also a USO for C'. Observe that any cube B of C' is also a cube of C and consequently, o' satisfies (C2) since o satisfies (C2). Suppose now that there exists $c' \in C'$ such that the outgoing neighbors of c' in C' do not belong to a cube of C'. Then since o is a USO for C, necessarily the outgoing neighbors of c' in C belong to a cube B of C that contains c. But then c and c' are both sources of B for the orientation o. By Lemma 6.4, this implies that o does not satisfy (C2) on B and thus on C (see Remark 6.6), a contradiction. Therefore, o' satisfies (C1) and (C2) and is a unique sink orientation of C'. Applying the previous argument inductively we obtain a corner peeling of C. □

Remark 6.13. A consequence of Proposition 6.12 is that for any representation map for Hall's concept class C_H, the corresponding USO of $G(C_H)$ contains a directed cycle and thus at least one non-trivial strongly connected component.

The USOs corresponding to the two representation maps of C_H presented in the appendix contain one or two non-trivial strongly connected components.

6.3. Representation maps for substructures

In this subsection, from a representation map r for an ample class C, we show how to derive representation maps for restrictions C_Y, reductions C^Y and intersections $C \cap B$ with cubes of 2^U.

Given a subset $Y \subseteq U = \mathrm{dom}(C)$, define $r^Y : C^Y \to X(C^Y)$ as follow. For any $c \in C^Y$, there exists a unique Y-cube B in C such that $\mathrm{tag}(B) = c$. By Lemmas 6.4 and 6.9(1), there exists a unique $c^B \in B$ such that $r_B(c^B) = r(c^B) \cap Y = Y$ (c^B is the source of B for o_r). We set $r^Y(c) := r(c^B) \setminus Y$. If $Y = \{x\}$, r^Y coincides with the x-out-map r^x defined in Section 6.2.

Given a subset $Y \subseteq U = \mathrm{dom}(C)$, define $r_Y : C_Y \to X(C_Y)$ as follow. For any $c \in C_Y$, there exists a unique Y-cube B in 2^U such that $\mathrm{tag}(B) = c$. Since $c \in C_Y$, $C \cap B \neq \varnothing$. By Claim 6.2, there exists a unique $c_B \in C \cap B$ such that $r_B(c_B) = r(c_B) \cap Y = \varnothing$ (c_B is the unique sink of $C \cap B$ for o_r). We set $r_Y(c) := r(c_B)$. If $Y = \{x\}$, r_Y coincides with the x-in-map r_x defined in Section 6.2.

Proposition 6.14. *For a representation map r for an ample class C, any cube B of 2^U, and any $Y \subseteq U$, the following hold:*

(1) r_B is a representation map for $C \cap B$;
(2) r^Y is a representation map for C^Y;
(3) r_Y is a representation map for C_Y.

Proof. By Theorem 6.8, Assertion (1) follows from Assertion (1) of Lemma 6.9 and Assertion (2) follows by iteratively applying Assertion (2) of Lemma 6.9 to the elements of Y.

By Theorem 6.8, to prove Assertion (3), it is enough to show that r_Y is the out-map of a USO of C_Y. Recall that for any $c \in C_Y$, we have $r_Y(c) = r(c_B)$, where B is the unique Y-cube of 2^U such that $\mathrm{tag}(B) = c$ and c_B is the unique sink of $C \cap B$ for o_r. By (C1), there exists an $r(c_B)$-cube B' containing c_B in C. Since $r(c_B) \cap Y = \varnothing$, there exists an $r(c_B)$-cube of C_Y containing c. Consequently, C_Y and r_Y satisfy (C1).

Suppose there exists a cube B in C_Y violating (C2), i.e., there exist $c_1, c_2 \in C_Y \cap B$ such that $r_Y(c_1) \cap \mathrm{supp}(B) = r_Y(c_2) \cap \mathrm{supp}(B)$. Any cube B in C_Y extends to a unique cube B' of 2^U such that $\mathrm{supp}(B') = \mathrm{supp}(B) \cup Y$. By definition of r_Y, there exist $c'_1, c'_2 \in B'$ such that $r(c'_1) = r_Y(c_1)$ and $r(c'_2) = r_Y(c_2)$. Consequently, $r(c'_1) \cap \mathrm{supp}(B') = r_Y(c_1) \cap \mathrm{supp}(B)$ since $r(c'_1) \cap Y = \varnothing$ and similarly $r(c'_2) \cap \mathrm{supp}(B') = r_Y(c_2) \cap \mathrm{supp}(B)$. Consequently, the map $c \mapsto r(c) \cap \mathrm{supp}(B')$ is not injective on $C \cap B'$, contradicting Condition (R3) of Theorem 6.1. □

6.4. Pre-representation maps

We now show that we can find maps satisfying each of the conditions (C1) and (C2). Nevertheless, we were not able to find a map satisfying (C1) and (C2). It is surprising that, while each d-cube has at least $d^{\Omega(2^d)}$ USOs [27], it is so difficult to find a single USO for ample classes.

Proposition 6.15. *For any ample class C there exists a bijection $r' : C \to X(C)$ and an injection $r'' : C \to 2^U$ such that r' satisfies the condition (C1) and r'' satisfies the condition (C2).*

Proof. First we prove the existence of the bijection r'. For $s \in X(C)$, denote by $N_s(C)$ the union of all s-cubes included in C and call $N_s(C)$ the *carrier* of s. For $S \subseteq X(C)$, denote by $N_S(C)$ the union of all carriers $N_s(C), s \in S$. Define a bipartite graph $\Gamma(C) = (C \cup X(C), E)$, where there is an edge between a concept $c \in C$ and a strongly shattered set $s \in X(C)$ if and only if c belongs to the carrier $N_s(C)$. We assert that $\Gamma(C)$ admits a perfect matching M. By the definition of the edges of $\Gamma(C)$, if the edge cs is in M, then c belongs to $N_s(C)$ and thus the unique s-cube containing c is included in C. Thus $r' : C \to X(C)$ defined by setting $r'(c) = s$ if and only if $cs \in M$ is a bijection satisfying (C1).

Since C is ample, we have $|X(C)| = |C|$, and thus to prove the existence of M, we show that the graph Γ satisfies the conditions of Philip Hall's theorem [26]: if S is an arbitrary subset of simplices of $X(C)$, then $|N_S(C)| \geq |S|$. Indeed, since C

is ample, for any $s \in S$ the carrier $N_s(C)$ contains at least one s-cube, thus $S \subseteq \underline{X}(N_S(C))$. Consequently, $|S| \leq |\underline{X}(N_S(C))| \leq |N_S(C)|$ by the Sandwich Lemma applied to the class $N_S(C)$.

We now prove that there exists an injection $r'' : C \to 2^U$ satisfying (C2). We prove the existence of r'' by induction on the size of U. If $|U| = 0$, r'' trivially exists. Consider now $x \in X$ and note that there exists an injection $r''_x : C_x \to 2^{U\setminus\{x\}}$ satisfying (C2) by induction hypothesis. We define r'' by:

$$r''(c) = \begin{cases} r''_x(c) & \text{if } x \notin c, \\ r''_x(c \setminus \{x\}) & \text{if } x \in c \text{ and } c \setminus \{x\} \notin C, \\ r''_x(c \setminus \{x\}) \cup \{x\} & \text{otherwise.} \end{cases}$$

It means that the orientation of the edges of $G(C)$ is obtained by keeping the orientation of the edges of $G(C_x)$ and orienting all x-edges of $G(C)$ from $c \cup \{x\}$ to c. It is easy to verify that r'' is injective.

Consider two distinct concepts $c, c' \in C$ belonging to a common cube of C and let $B = B(c, c')$ that is the minimal cube of C containing c and c'. If $c\Delta c' \neq \{x\}$, then since $r''(c) \setminus \{x\} = r''_x(c \setminus \{x\})$, since $r''(c') \setminus \{x\} = r''_x(c' \setminus \{x\})$, and since r''_x satisfies (C2) on the cube B_x of C_x, there exists $y \in (c \setminus \{x\} \cup c' \setminus \{x\})\Delta(r''_x(c \setminus \{x\}) \cup r''_x(c' \setminus \{x\})) \subseteq (c \cup c')\Delta(r''(c) \cup r''(c'))$. Suppose now that $c' = c \cup \{x\}$. In this case, $x \in r''(c') \setminus r''_x(c)$ and $x \in (c\Delta c') \cap (r''(c)\Delta r''(c'))$. □

One can try to find representation maps for ample classes by extending the approach for maximum classes: given ample classes C and D with $D \subset C$, a representation map r for C is called *D-entering* if all edges cd with $c \in C \setminus D$ and $d \in D$ are directed by o_r from c to d. The representation map defined in Proposition 5.2 is D-entering. Given $x \in \text{dom}(C)$, suppose that r_x is a C^x-entering representation map for C_x.

We can extend the orientation o_{r_x} to an orientation o of $G(C)$ as follows. Each x-edge cc' of $G(C)$ is directed arbitrarily, while each other edge cc' is directed as the edge $c_x c'_x$ is directed by o_{r_x}. Since o_{r_x} satisfies (C1), (C2) and r_x is C^x-entering, o also satisfies (C1), (C2), thus the map r_o is a representation map for C. So, ample classes would admit representation maps, if *for any ample classes $D \subseteq C$, any representation map r' of D extends to a D-entering representation map r of C.*

6.5. Representation maps as ISRs

Our next result formulates the construction of representation maps for ample classes as an instance of the Independent System of Representatives problem. A system $(G, (V_i)_{1\leq i\leq n})$ consisting of a graph G and a partition $V_1, \ldots, V_n$ of its vertex set $V(G)$ is called an *ISR-system*. An independent set in G of the form $\{v_1, \ldots, v_n\}$, where $v_i \in V_i$ for each $1 \leq i \leq n$, is called an *Independent System of Representatives*, or *ISR* for short [2].

Consider an ample class $C \subseteq 2^U$. We build an ISR-system $(G, (V_c)_{c\in C})$ as follows. For each concept c and each set $Y \subseteq X(C)$ such that c belongs to a Y-cube of C, there is a vertex (c, Y) in $V(G)$. For each concept $c \in C$, set $V_c := \{(c, Y) \in V(G)\}$. Finally, $E(G)$ is defined as follows: there is an edge between two vertices (c_1, Y_1), (c_2, Y_2) in G if c_1, c_2 are distinct vertices belonging to a common cube B such that $Y_1 \cap \text{supp}(B) = Y_2 \cap \text{supp}(B)$.

Proposition 6.16. *There is an ISR for $(G, (V_c)_{c\in C})$ if and only if C admits a representation map.*

Proof. Assume first that r is a representation map for C. We show that $\{(c, r(c))\}_{c\in C}$ is an ISR for $(G, (V_c)_{c\in C})$. By (C1), for every $c \in C$, c belongs to an $r(c)$-cube of C and thus $(c, r(c)) \in V(G)$. Moreover, if $\{(c, r(c))\}_{c\in C}$ is not an independent set of G, there exist two concepts c_1, c_2 in a cube B such that $r(c_1) \cap \text{supp}(B) = r(c_2) \cap \text{supp}(B)$, contradicting (C2).

Conversely, if $\{(c, Y_c)\}_{c\in C}$ is an ISR for $(G, (V_c)_{c\in C})$, then the map $r : c \to X(C)$ defined by $r(c) = Y_c$ is a representation map. Indeed, (C1) is satisfied by the definition of the vertices of $V(G)$, (C2) is satisfied by the definition of the edges of $E(G)$ and since $\{(c, Y_c)\}_{c\in C}$ is an independent set of G. □

CRediT authorship contribution statement

The four authors have equally participated in the thinking, the formalizing, and the writing of the results presented in this paper.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We are very grateful to the referees of this submission and to some of the referees of previous versions of this paper for useful remarks that helped improving the presentation of this work. We are also grateful to Olivier Bousquet for several insightful discussions. The research of J. C. and V. C. was supported by the ANR project DISTANCIA (ANR-17-CE40-0015). S. M. is a Robert J. Shillman Fellow and was supported by the ISF, grant no. 1225/20, by an Azrieli Faculty Fellowship, and by BSF grant 2018385; part of this project was carried while S.M. was at the Institute for Advanced Study and was supported by the NSF grant CCF-1412958. The research of M. W. was partially supported by the NSF grant IIS-1619271.

Appendix A. Two representation maps for Hall's concept class C_H

We give two representation maps for Hall's concept class C_H given in Fig. 3. The representation map r_1 presented in Table 1 was found by constructing Boolean clauses for a matching between the concepts and the dimension sets of size ≤ 3 that satisfies the non-clashing condition. A satisfying assignment was then found with the open source MiniSat solver. The representation map r_2 presented in Table 2 was found by following the steps of the proof of Theorem 5.1 by recursing on dimensions 12, 11, ... In both cases, the representation map r_i, $i = 1, 2$ of C_H is described as follows: we list all the 299 concepts as bit vectors and for each concept $c \in C_H$, we underline the bits of $r_i(c)$ (a subset of size ≤ 3 from $\{1, 2, \ldots, 12\}$).

The USOs corresponding to these two representation maps are given in Figs. 6 and 7; they were obtained by using Sage and Graphviz. In both cases, the non-trivial strongly connected components are represented by the blue boxes. Except from the arcs within these strongly connected components, the other arcs are downward arcs.

Table 1
The representation r_1 map for C_H obtained using MiniSat.

000000000000	000000000001	000000000010	000000000011	000000010000	000000010001	000000010011
000000110000	000000110001	000000110011	000000110111	000000111011	000000111111	000001000000
000001000001	000001000011	000001010000	000001010001	000001010011	000001110000	000001110001
000001110011	000001110111	000001111011	000001111111	000010000011	000010010011	000010110011
000010111011	000010111111	000011000000	000011000001	000011000011	000011010000	000011010001
000011010011	000011110000	000011110001	000011110011	000011110100	000011110101	000011110111
000011111011	000011111100	000011111101	000011111111	000100000000	000100010000	000100110000
000100110001	000100110011	000100110111	000100111111	000101110000	000101110001	000101110011
000101110111	000101111111	000110111111	000111110000	000111110001	000111110011	000111110100
000111110101	000111110111	000111111100	000111111101	000111111111	001000110011	001000110111
001000111011	001000111111	001010111111	001011111111	001100000000	001100000100	001100001000
001100001100	001100010000	001100010100	001100011100	001100100000	001100100100	001100101000
001100101100	001100110000	001100110001	001100110011	001100110100	001100110101	001100110111
001100111000	001100111001	001100111011	001100111100	001100111101	001100111111	001101110000
001101110001	001101110011	001101110100	001101110101	001101110111	001101111100	001101111101
001101111111	001110111111	001111110000	001111110001	001111110011	001111110100	001111110101
001111110111	001111111100	001111111101	001111111111	010000000000	010000000001	010000000010
010000000011	010000010000	010000110000	010001000000	010001000001	010001000011	010001010000
010001110000	010011000000	010011000001	010011000011	010011010000	010011010001	010011010011
010011110000	010011110001	010011110011	010011110100	010011110101	010011110111	010011111100
010011111101	010011111111	010100000000	010100010000	010100110000	010101110000	010111110000
010111110100	010111111100	010111111101	010111111111	011011111111	011100000000	011100000100
011100001000	011100001100	011100010000	011100010100	011100011100	011100110000	011100110100
011100111100	011101110000	011101110100	011101111100	011111110000	011111110100	011111111100
011111111101	011111111111	100000000000	100000000010	100000000011	100001000011	100010000011
100011000011	100100000000	101100000000	101100001000	101100001100	101100011100	101100101100
101100111100	110000000000	110000000001	110000000010	110000000011	110000010000	110000110000
110001000000	110001000001	110001000011	110001010000	110001110000	110010000000	110010000001
110010000010	110010000011	110011000000	110011000001	110011000010	110011000011	110011000100
110011000101	110011000110	110011000111	110011001100	110011001101	110011001110	110011001111
110011010000	110011010001	110011010011	110011010100	110011010101	110011010111	110011011100
110011011101	110011011111	110011101111	110011110000	110011110001	110011110011	110011110100
110011110101	110011110111	110011111100	110011111101	110011111111	110100000000	110100010000
110100110000	110101000000	110101010000	110101110000	110111000000	110111000100	110111001100
110111001101	110111001111	110111010000	110111010100	110111011100	110111011101	110111011111
110111101111	110111110000	110111110100	110111111100	110111111101	110111111111	111011001100
111011001101	111011001110	111011001111	111011101111	111011111111	111100000000	111100000100
111100001000	111100001100	111100010000	111100010100	111100011100	111100101100	111100110000
111100110100	111100111100	111101000000	111101000100	111101001100	111101010000	111101010100
111101011100	111101101100	111101110000	111101110100	111101111100	111111000000	111111000100
111111001100	111111001101	111111001110	111111001111	111111010000	111111010100	111111011100
111111011101	111111011110	111111011111	111111101100	111111101110	111111101111	111111110000
111111110100	111111111100	111111111101	111111111110	111111111111		

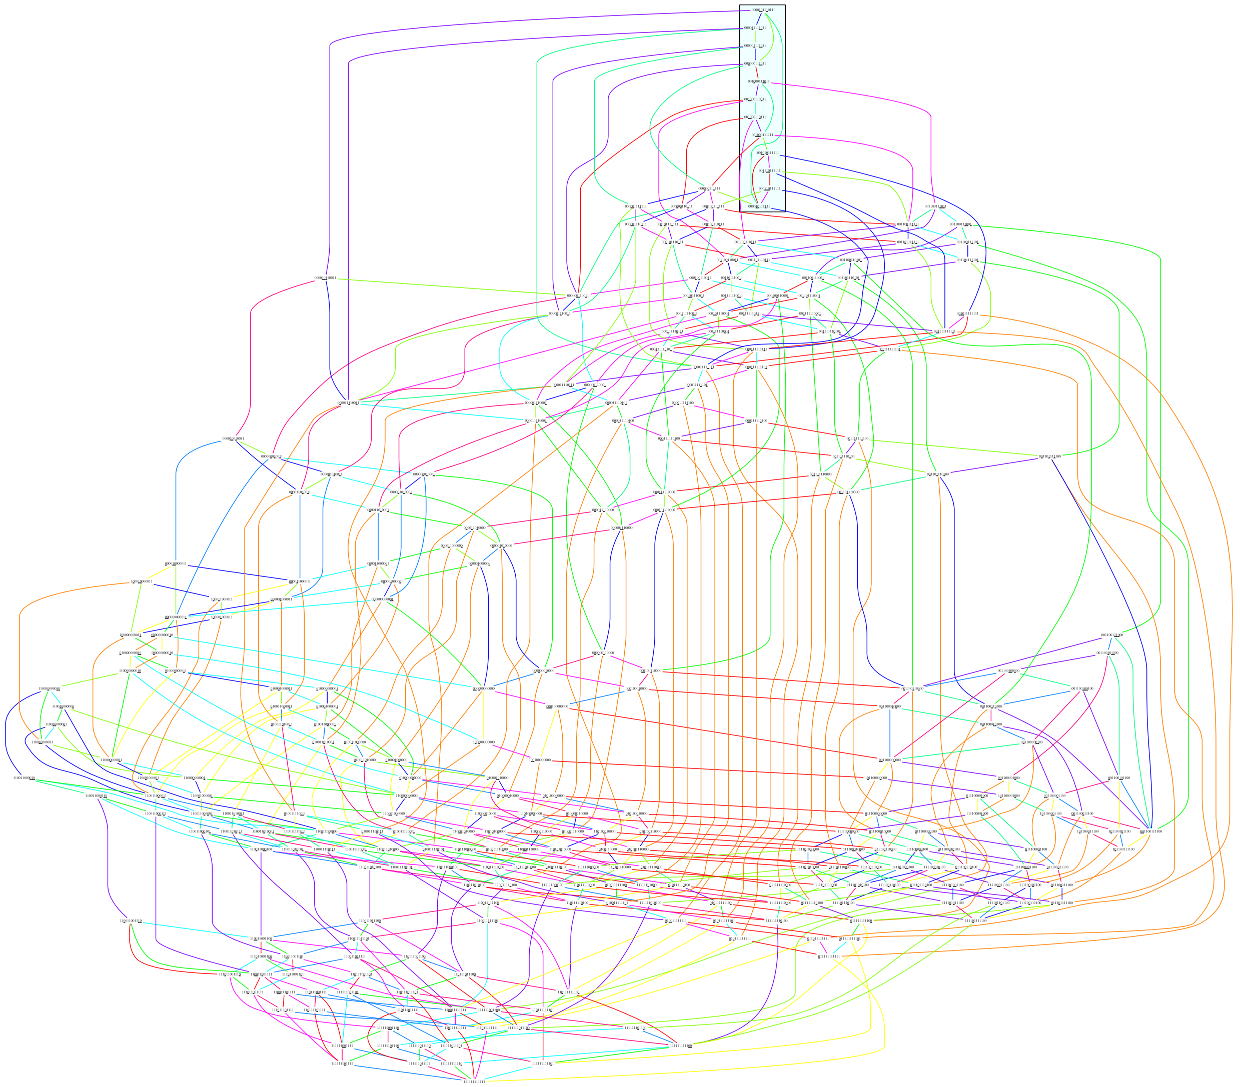

Fig. 6. The USO corresponding to the representation map r_1 of C_H described in Table 1. The 12-vertex NSCC is in the blue box. Best viewed by zooming in.

The USO from Fig. 6 has a unique non-trivial strongly connected component (NSCC) K while the one from Fig. 7 has two such components K' and K''. Since those components are small (each of them contains 12 concepts), these USOs are not "far" from being acyclic and C_H is (not "far" from being) a minimal corner-free example.

Given a concept class C and an orientation o of the edges of $G(C)$, we denote by $\overrightarrow{G_o}(C)$ the resulting directed graph. From the following proposition, we deduce that any representation map for C_H obtained by the algorithm of Theorem 5.1 contains at least two NSCCs. In particular, this shows that the representation map r_1 cannot be obtained in such a way.

Proposition A.1. *For any maximum class C without corners, any representation map r of C computed by Theorem 5.1, and the USO o corresponding to r, the directed graph $\overrightarrow{G}_o(C)$ contains at least two non-trivial strongly connected components.*

Proof. We prove the result by induction on the dimension of C. If C is of dimension at most 2, then by Proposition 4.11 C contains a corner and we are done.

Now suppose that the dimension of C is at least 3. Recall that the first step is to contract C along one of the coordinates $x \in U$ (to construct r_2, we contracted C_H along coordinate 12, see Fig. 8) and to compute recursively a USO o^x (and the corresponding representation map r^x) of C^x. Then we extend o^x to a USO o_x (and find the corresponding representation map r_x) of C_x. Finally, we obtain o from o_x by keeping the orientation of all edges that do not correspond to coordinate x and by orienting all x-edges from $1C^x$ to $0C^x$.

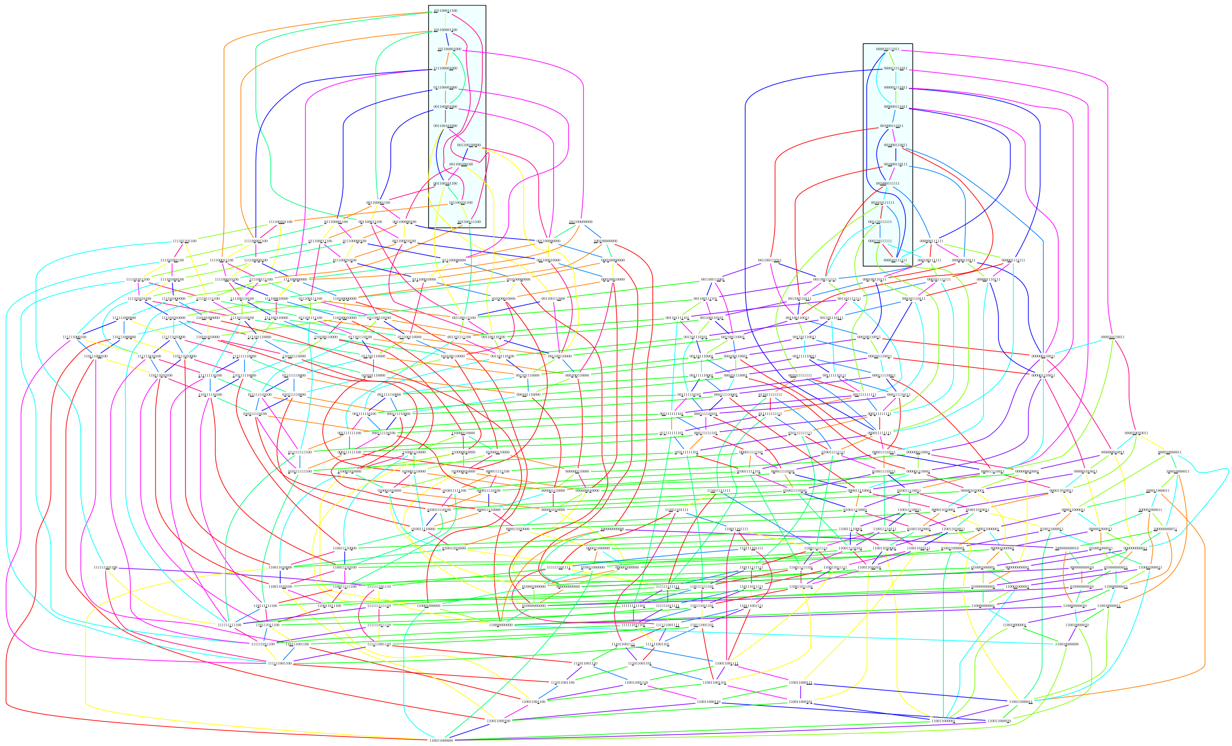

Fig. 7. The USO corresponding to the representation map r_2 of C_H described in Table 2. This map has two NSCCs of size 12 (in blue).

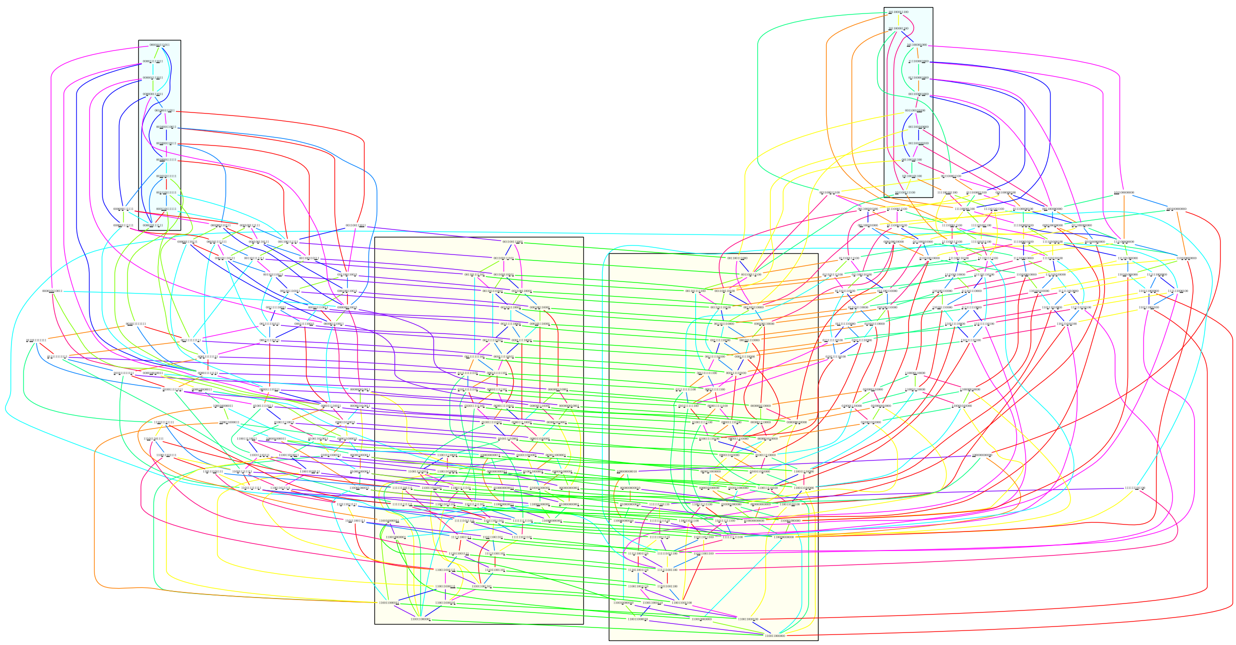

Fig. 8. Another representation of the USO corresponding to r_2 illustrating the top recursion (on dimension 12) of the algorithm from Theorem 5.1: $0C^x$ and $1C^x$ are represented by yellow boxes.

Table 2
The representation map r_2 for C_H obtained using Theorem 5.1 by recursing on dimensions 12, 11, ….

000000000000	000000000001	000000000010	000000000011	000000010000	000000010001	000000010011
000000110000	000000110001	000000110011	000000110111	000000111011	000000111111	000001000000
000001000001	000001000011	000001010000	000001010001	000001010011	000001110000	000001110001
000001110011	000001110111	000001111011	000001111111	000010000011	000010010011	000010110011
000010111011	000010111111	000011000000	000011000001	000011000011	000011010000	000011010001
000011010011	000011110000	000011110001	000011110011	000011110100	000011110101	000011110111
000011111011	000011111100	000011111101	000011111111	000100000000	000100010000	000100110000
000100110001	000100110011	000100110111	000100111111	000101110000	000101110001	000101110011
000101110111	000101111111	000110111111	000111110000	000111110001	000111110011	000111110100
000111110101	000111110111	000111111100	000111111101	000111111111	001000110011	001000110111
001000111011	001000111111	001010111111	001011111111	001100000000	001100000100	001100001000
001100001100	001100010000	001100010100	001100011100	001100100000	001100100100	001100101000
001100101100	001100110000	001100110001	001100110011	001100110100	001100110101	001100110111
001100111000	001100111001	001100111011	001100111100	001100111101	001100111111	001101110000
001101110001	001101110011	001101110100	001101110101	001101110111	001101111100	001101111101
001101111111	001110111111	001111110000	001111110001	001111110011	001111110100	001111110101
001111110111	001111111100	001111111101	001111111111	010000000000	010000000001	010000000010
010000000011	010000010000	010000110000	010001000000	010001000001	010001000011	010001010000
010001110000	010011000000	010011000001	010011000011	010011010000	010011010001	010011010011
010011110000	010011110001	010011110011	010011110100	010011110101	010011110111	010011111100
010011111101	010011111111	010100000000	010100010000	010100110000	010101110000	010111110000
010111110100	010111111100	010111111101	010111111111	011011111111	011100000000	011100000100
011100001000	011100001100	011100010000	011100010100	011100011100	011100110000	011100110100
011100111100	011101110000	011101110100	011101111100	011111110000	011111110100	011111111100
011111111101	011111111111	100000000000	100000000010	100000000011	100001000011	100010000011
100011000011	100100000000	101100000000	101100001000	101100001100	101100011100	101100101100
101100111100	110000000000	110000000001	110000000010	110000000011	110000010000	110000110000
110001000000	110001000001	110001000011	110001010000	110001110000	110010000000	110010000001
110010000010	110010000011	110011000000	110011000001	110011000010	110011000011	110011000100
110011000101	110011000110	110011000111	110011001100	110011001101	110011001110	110011001111
110011010000	110011010001	110011010011	110011010100	110011010101	110011010111	110011011100
110011011101	110011011111	110011101111	110011110000	110011110001	110011110011	110011110100
110011110101	110011110111	110011111100	110011111101	110011111111	110100000000	110100010000
110100110000	110101000000	110101010000	110101110000	110111000000	110111000100	110111001100
110111001101	110111001111	110111010000	110111010100	110111011100	110111011101	110111011111
110111101111	110111110000	110111110100	110111111100	110111111101	110111111111	111011001100
111011001101	111011001110	111011001111	111011101111	111011111111	111100000000	111100000100
111100001000	111100001100	111100010000	111100010100	111100011100	111100101100	111100110000
111100110100	111100111100	111101000000	111101000100	111101001100	111101010000	

Observe that any directed cycle of $\overrightarrow{G_o}(C)$ corresponds either to a cycle of C^x or to a directed cycle of a connected component of $G(C_x \setminus C^x)$. If C^x does not contain a corner, then by induction hypothesis, C^x contains at least two NSCCs and thus there are two NSCCs in $0C^x$ and thus in C and we are done.

Suppose now that C^x contains a corner. Consider a connected component A of $G(C_x \setminus C^x)$. If A does not contain any directed cycle, then the concept class A is acyclic and it contains a source s. By (C1) applied to s and r, we conclude that s is a corner of C. Consequently, each connected component of $G(C_x \setminus C^x)$ contains at least one NSCC.

If C^x is not a separator of $G(C_x)$, i.e., if $N_x(C)$ is not a separator of $G(C)$, then without loss of generality, we can assume that there is no edge from $0C^x$ to $C \setminus N_x(C)$. Consider any corner c of C^x that is contained in a unique maximal cube B^x of C^x. In C there exists a unique cube B containing c such that $\mathrm{supp}(B) = \mathrm{supp}(B^x) \cup \{x\}$. Since c has no neighbor outside the carrier $N_x(C)$, B is the unique maximal cube containing c and thus c is a corner of C, a contradiction.

We can thus assume that C^x is a separator of $G(C_x)$ and thus $G(C_x \setminus C^x)$ contains at least two connected components, and by the previous assertion, each of them contains at least one NSCC and we are done. □

A strongly connected component (SCC) S of a directed graph $\overrightarrow{G}$ is called a *source-component* if there is no arc from u to v with $v \in S$ and $u \notin S$.

In the directed graph $\overrightarrow{G_o}$ corresponding to a representation map r of an ample class C, if a source-component S is reduced to a single concept c (i.e., S is trivial), then c is a corner of C. Thus, one can view the source-components of $\overrightarrow{G_o}$ as a generalization of corners. However, the definition of source-components depends on a given representation map r and the corresponding orientation o_r of $G(C)$, while the corners are defined in the undirected graph $G(C)$. In the following proposition, we prove that, similarly to corners, removing a source component S from an ample class C results into an ample class $C \setminus S$. Moreover, the restriction of the representation map r to $C \setminus S$ is still a representation map.

Proposition A.2. *Let C be an ample class C, r be a representation map of C, $o = o_r$ the USO defined by r, and $\overrightarrow{G_o}$ be the corresponding directed graph. Then for any source-component S of $\overrightarrow{G_o}$, $C \setminus S$ is an ample class and the restriction of r to $C \setminus S$ is a representation map of $C \setminus S$.*

Proof. Set $C' := C \setminus S$ and denote by r' the restriction of r to C'. Since r satisfies (C2) on C, r' satisfies (C2) on C'. Now, we show that r' satisfies (C1) on C'. Pick any concept $c \in C'$ and let B be the cube of 2^U containing c and defined by $r(c)$. By (C1), B is included in C. If B is also included in C', then we are done. So, suppose that $B \cap S \neq \varnothing$ and pick a concept $c' \in B \cap S$ closest to c. Let $B' = B(c, c')$ be the cube spanned by c and c'. Since $B' \subset B$, B' is a cube of C. By the definition of B, c is a source of B and thus of B'. On the other hand, from the choice of c' all neighbors of c' in B' belong to $C \setminus S$. Since $c' \in S$ and S is a source-component of $\overrightarrow{G_o}$, all these edges have c' as a source. Consequently, c' is a second source of B', contrary to condition (C2) applied to C. This shows that B is included in C', establishing that r' satisfies (C1) on C'.

Since r is a bijective map from C to $X(C)$, from (C1) applied to r' we conclude that r' is an injective map from C' to $\underline{X}(C')$, yielding $|C'| \leq |\underline{X}(C')|$. Since $|\underline{X}(C')| \leq |C'|$ by the sandwich lemma, we deduce that $|C'| = |\underline{X}(C')|$ and thus C' is ample by Theorem 3.1(4). Since the out-map of r' satisfies (C1) and (C2), r' is a representation map of C'. □

Corollary A.3. *$C_H \setminus K$ and $C_H \setminus (K' \cup K'')$ are ample classes admitting corner peelings.*

Proof. By Proposition A.2, the concept classes $C_H \setminus K$ and $C_H \setminus K'$ are ample classes and the restrictions of r to them are representation maps. Applying Proposition A.2 we conclude that $C_H \setminus (K' \cup K'')$ is an ample class and the restriction of r to $C_H \setminus (K' \cup K'')$ is a representation map. Moreover, since all other NSCCs of $\overrightarrow{G}(C)$ are trivial, the orientations of the edges of graphs $G(C \setminus K)$ and $C_H \setminus (K' \cup K'')$ defined by these representation maps are acyclic USOs. By Proposition 6.12, $C_H \setminus K$ and $C_H \setminus (K' \cup K'')$ admit corner peelings. □

References

[1] Karim Adiprasito, Bruno Benedetti, Collapsibility of CAT(0) spaces, Geom. Dedic. 206 (2019) 181–199, https://doi.org/10.1007/s10711-019-00481-x.
[2] Ron Aharoni, Eli Berger, Ran Ziv, Independent systems of representatives in weighted graphs, Combinatorica 27 (3) (2007) 253–267, https://doi.org/10.1007/s00493-007-2086-y.
[3] Richard P. Anstee, Lajos Rónyai, Attila Sali, Shattering news, Graphs Comb. 18 (1) (2002) 59–73, https://doi.org/10.1007/s003730200003.
[4] Hans-Jürgen Bandelt, Victor Chepoi, Metric graph theory and geometry: a survey, in: J.E. Goodman, J. Pach, R. Pollack (Eds.), Surveys on Discrete and Computational Geometry: Twenty Years Later, in: Contemp. Math., vol. 453, Amer. Math. Soc., Providence, RI, 2008, pp. 49–86.
[5] Hans-Jürgen Bandelt, Victor Chepoi, Andreas W.M. Dress, Jack H. Koolen, Combinatorics of lopsided sets, Eur. J. Comb. 27 (5) (2006) 669–689, https://doi.org/10.1016/j.ejc.2005.03.001.
[6] Hans-Jürgen Bandelt, Victor Chepoi, Andreas W.M. Dress, Jack H. Koolen, Geometry of lopsided sets, 2012, unpublished.
[7] Hans-Jürgen Bandelt, Victor Chepoi, Kolja Knauer, COMs: complexes of oriented matroids, J. Comb. Theory, Ser. A 156 (2018) 195–237, https://doi.org/10.1016/j.jcta.2018.01.002.
[8] Béla Bollobás, Andrew J. Radcliffe, Defect Sauer results, J. Comb. Theory, Ser. A 72 (2) (1995) 189–208, https://doi.org/10.1016/0097-3165(95)90060-8.
[9] Olivier Bousquet, Steve Hanneke, Shay Moran, Nikita Zhivotovskiy, Proper learning, Helly number, and an optimal SVM bound, in: COLT 2020, in: Proceedings of Machine Learning Research, vol. 125, 2020, pp. 582–609, http://proceedings.mlr.press/v125/bousquet20a.html.
[10] Katarína Cechlárová, The uniquely solvable bipartite matching problem, Oper. Res. Lett. 10 (4) (1991) 221–224, https://doi.org/10.1016/0167-6377(91)90062-T.
[11] Jérémie Chalopin, Victor Chepoi, Shay Moran, Manfred K. Warmuth, Unlabeled sample compression schemes and corner peelings for ample and maximum classes, in: ICALP 2019, in: LIPIcs, vol. 132, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019, pp. 29:1–29:15.
[12] Thorsten Doliwa, Gaojian Fan, Hans Ulrich Simon, Sandra Zilles, Recursive teaching dimension, VC-dimension and sample compression, J. Mach. Learn. Res. 15 (1) (2014) 3107–3131, http://jmlr.org/papers/v15/doliwa14a.html.
[13] Andreas W.M. Dress, Towards a theory of holistic clustering, in: Mathematical Hierarchies and Biology, in: DIMACS Ser. Discrete Math. Theoret. Comput. Sci., vol. 37, Amer. Math. Soc., 1996, pp. 271–290.
[14] Paul H. Edelman, Robert E. Jamison, The theory of convex geometries, Geom. Dedic. 19 (3) (1985) 247–270, https://doi.org/10.1007/BF00149365.
[15] Sally Floyd, On space bounded learning and the Vapnik-Chervonenkis dimension, PhD thesis International Computer Science Institut, Berkeley, CA, 1989, https://www.icsi.berkeley.edu/icsi/node/2289.
[16] Sally Floyd, Manfred K. Warmuth, Sample compression, learnability, and the Vapnik-Chervonenkis dimension, Mach. Learn. 21 (3) (1995) 269–304, https://doi.org/10.1007/BF00993593.
[17] Robin Forman, Morse theory for cell complexes, Adv. Math. 134 (1) (1998) 90–145, https://doi.org/10.1006/aima.1997.1650.
[18] Bernd Gärtner, Emo Welzl, Vapnik-Chervonenkis dimension and (pseudo-)hyperplane arrangements, Discrete Comput. Geom. 12 (4) (1994) 399–432, https://doi.org/10.1007/BF02574389.
[19] Mikhaïl Gromov, Hyperbolic groups, in: S.M. Gersten (Ed.), Essays in Group Theory, in: Math. Sci. Res. Inst. Publ., vol. 8, Springer, New York, 1987, pp. 75–263.
[20] H. Tracy Hall, Counterexamples in discrete geometry, PhD thesis, University of California, 2004.
[21] Hunter R. Johnson, Some new maximum VC classes, Inf. Process. Lett. 114 (6) (2014) 294–298, https://doi.org/10.1016/j.ipl.2014.01.006.
[22] László Kozma, Shay Moran, Shattering, graph orientations, and connectivity, Electron. J. Comb. 20 (3) (2013) P44, http://www.combinatorics.org/ojs/index.php/eljc/article/view/v20i3p44.
[23] Dima Kuzmin, Manfred K. Warmuth, Unlabeled compression schemes for maximum classes, J. Mach. Learn. Res. 8 (2007) 2047–2081, http://www.jmlr.org/papers/v8/kuzmin07a.html.
[24] James F. Lawrence, Lopsided sets and orthant-intersection of convex sets, Pac. J. Math. 104 (1) (1983) 155–173, https://doi.org/10.2140/pjm.1983.104.155.
[25] Nick Littlestone, Manfred K. Warmuth, Relating data compression and learnability, Technical report, Department of Computer and Information Sciences, Santa Cruz, CA, 1986.
[26] László Lovász, Michael D. Plummer, Matching Theory, AMS Chelsea Publishing, Providence, RI, 2009, Corrected reprint of the 1986 original.

[27] Jiří Matoušek, The number of unique-sink orientations of the hypercube, Combinatorica 26 (1) (2006) 91–99, https://doi.org/10.1007/s00493-006-0007-0.

[28] Tamás Mészáros, Lajos Rónyai, Shattering-extremal set systems of VC dimension at most 2, Electron. J. Comb. 21 (4) (2014) P4.30, http://www.combinatorics.org/ojs/index.php/eljc/article/view/v21i4p30.

[29] Shay Moran, Shattering-extremal systems, preprint, arXiv:1211.2980, 2012.

[30] Shay Moran, Manfred K. Warmuth, Labeled compression schemes for extremal classes, in: ALT 2016, in: Lecture Notes in Comput. Sci., vol. 9925, Springer, 2016, pp. 34–49.

[31] Shay Moran, Amir Yehudayoff, Sample compression schemes for VC classes, J. ACM 63 (3) (2016) 21:1–21:10, https://doi.org/10.1145/2890490.

[32] Mogens Nielsen, Gordon D. Plotkin, Glynn Winskel, Petri nets, event structures and domains, part I, Theor. Comput. Sci. 13 (1) (1981) 85–108, https://doi.org/10.1016/0304-3975(81)90112-2.

[33] Alain Pajor, Sous-Espaces ℓ_1^n des Espaces de Banach, Travaux en Cours, Hermann, Paris, 1985.

[34] Dömötör Pálvölgyi, Gábor Tardos, Unlabeled compression schemes exceeding the VC-dimension, Discrete Appl. Math. 276 (2019) 102–107, https://doi.org/10.1016/j.dam.2019.09.022.

[35] Benjamin I.P. Rubinstein, J. Hyam Rubinstein, A geometric approach to sample compression, J. Mach. Learn. Res. 13 (2012) 1221–1261, http://www.jmlr.org/papers/v13/rubinstein12a.html.

[36] Michah Sageev, CAT(0) cube complexes and groups, in: M. Bestvina, M. Sageev, K. Vogtmann (Eds.), Geometric Group Theory, in: IAS/Park City Math. Ser., vol. 21, Amer. Math. Soc., Inst. Adv. Study, 2012, pp. 6–53.

[37] Rahim Samei, Boting Yang, Sandra Zilles, Generalizing labeled and unlabeled sample compression to multi-label concept classes, in: ALT 2014, in: Lecture Notes in Comput. Sci., vol. 8776, Springer, 2014, pp. 275–290.

[38] Norbert Sauer, On the density of families of sets, J. Comb. Theory, Ser. A 13 (1) (1972) 145–147, https://doi.org/10.1016/0097-3165(72)90019-2.

[39] Shai Shalev-Shwartz, Shai Ben-David, Understanding Machine Learning: From Theory to Algorithms, Cambridge Univ. Press, 2014.

[40] Saharon Shelah, A combinatorial problem, stability and order for models and theories in infinitary languages, Pac. J. Math. 41 (1) (1972) 247–261, https://doi.org/10.2140/pjm.1972.41.247.

[41] Tibor Szabó, Emo Welzl, Unique sink orientations of cubes, in: FOCS 2001, IEEE Computer Society, 2001, pp. 547–555.

[42] Martin Tancer, Recognition of collapsible complexes is NP-complete, Discrete Comput. Geom. 55 (1) (2016) 21–38, https://doi.org/10.1007/s00454-015-9747-1.

[43] Vladimir N. Vapnik, Alexey Y. Chervonenkis, On the uniform convergence of relative frequencies of events to their probabilities, Theory Probab. Appl. 16 (2) (1971) 264–280, https://doi.org/10.1137/1116025.

[44] Manfred K. Warmuth, Compressing to VC dimension many points, in: COLT/Kernel 2003, in: Lecture Notes in Comput. Sci., vol. 2777, Springer, 2003, pp. 743–744.

[45] Emo Welzl, Complete range spaces, 1987, unpublished notes.

[46] Douglas H. Wiedemann, Hamming geometry, PhD thesis, University of Waterloo, 1986, re-typeset July, 2006.

[47] Avi Wigderson, Mathematics and Computation, Princeton Univ. Press, 2019.

[48] Glynn Winskel, Events in computation, PhD thesis, Edinburgh University, 1980.

[49] Günter M. Ziegler, Lectures on Polytopes, Grad. Texts in Math., vol. 152, Springer–Verlag, New York, 1995.