1id1 Preprint | 20 November 2026
1id1:core:20261120-0001@v1 Version 1
https://1id1.net/id/core/20261120-0001/v1/

Learning Orthant Comparability: Proper–Improper Separation and Compression without Side Information

Wuming Pan PANWUMING@SCU.EDU.CN
College of Computer Science, Sichuan University
Chengdu 610065, China

Abstract

We study learning coordinate-reversed product orders from comparability feedback: an example specifies whether two ordinal feature vectors are comparable, without revealing which precedes the other. Directions are therefore identifiable only up to simultaneous reversal. For the resulting class of VC dimension $d-1$, we establish a computational separation between proper and improper realizable learning. Proper PAC learning is impossible in polynomial time unless $\mathsf{RP} = \mathsf{NP}$, even when negative examples have support three. In contrast, the deductive closure of positive examples is an efficient improper learner with at most $d-1$ online mistakes and sample complexity $O((d + \log(1/\delta))/\varepsilon)$. Every realizable mixed sample compresses to at most $d-1$ positive examples, without auxiliary bits, with closure as the reconstruction rule. The algorithm exploits an affine positive version space represented by a signed graph. This representation also yields computable PAC-Bayes bounds despite hard exact inference from mixed feedback. As a complementary identification analysis, we prove Poisson and coupon-collector thresholds for dependent pairs drawn from continuous product distributions, using multiplicative estimates for joint antichain probabilities.

Keywords: comparability learning, improper learning, sample compression, signed graphs, PAC-Bayes bounds

1 Introduction

A product order treats each feature as a benefit or a cost. When these polarities are unknown, predicting comparability differs from reconstructing one explanatory order: we show that efficient prediction is possible outside the order-induced class, although finding an accurate member of that class is hard.

On a finite set of records, choosing independently which coordinate orders to reverse gives a family of preorders that we call a *finite orthant-order (FOO) system*. An unknown direction $\sigma \in \{-,+\}^d$ selects one member. An example is an unordered pair $z = \{x, y\}$, labelled 1 if x and y are comparable in the resulting order, and 0 otherwise. The class $\mathcal{C}_d = \{h_\sigma : \sigma \in \Sigma_d\}$ has 2^{d-1} distinct members: $h_\sigma = h_{-\sigma}$. Thus comparability feedback identifies an order only up to global reversal. The learning target is this Boolean pairwise predicate, not a real-valued monotone regression function.

Main result. A positive pair imposes parity relations among the coordinates on which its endpoints differ. Their joint solution set $\mathcal{V}(P)$ is an affine space, represented by a signed graph on d coordinates. Instead of selecting one direction from the mixed-feedback version

space, define the *closure predictor*

$$g_P(z) = \mathbf{1}[\mathcal{V}(P) \subseteq C(z)], \qquad C(z) = \{\sigma : h_\sigma(z) = 1\}.$$

This fits every realizable mixed sample $S = (P, N)$: positives are implied, while a true target in $\mathcal{V}(P)$ prevents any negative from being implied. The affine representation yields four main results.

(i) *Proper–improper separation.* The closure rule is an efficient distribution-free realizable PAC learner with sample complexity $O((d + \log(1/\delta))/\varepsilon)$, and its online version makes at most $d - 1$ mistakes, all false negatives (Theorems 4.2 and 4.4). In contrast, efficient proper PAC learning with negative supports of size three would imply $\mathsf{RP} = \mathsf{NP}$ (Corollary 5.3). Negative supports at most two admit efficient proper learning.

(ii) *Compression without auxiliary bits.* Every realizable mixed sample admits a compression to at most $d - 1$ positive examples. Reconstruction is the same closure rule, uses no negative examples or auxiliary bits, and is generally improper (Theorem 4.3). The online mistake bound is optimal; the compression theorem is an upper bound, not a claim of minimal compression size among all possible schemes.

(iii) *Efficient representation and risk certificates.* Positive consistency, counting, and small inconsistency certificates follow from signed-graph balance (Theorems 3.6 and 4.5). A single graph preprocessing pass permits linear-in-support closure queries and evaluation of the positive-version-space Gibbs risk. Its KL term is exactly the number of independent positive parity constraints, times $\ln 2$ (Theorem 8.2).

(iv) *Complementary identification theory.* For complete feedback on continuous random instances, one positive pair identifies the projective direction; negative pairs require a coupon-collector scale. We handle the dependence between shared-endpoint pairs using joint multiplicative probability estimates (Theorems 9.3 and 9.4). This sampling protocol is distinct from i.i.d. labelled-pair PAC learning.

Scope and motivation. Monotone learning (Barlow et al., 1972; Ben-David, 1995; Sill, 1998; Daniels and Velikova, 2010), multi-objective comparison (Deb, 2001), preference learning (Fürnkranz and Hüllermeier, 2010), and skyline queries (Börzsönyi et al., 2001) motivate order-induced hypotheses, but do not automatically supply our feedback: *one dominates the other* versus *trade-off*, without a dominance arrow. Scalar response labels do not generally determine comparability: monotone responses may be ordered on incomparable inputs. Conversely, comparability does not specify a scalar monotonicity inequality. Once the projective direction is identified, one direction-revealing comparison of distinct comparable instances resolves global reversal.

Distinct nominal labels, including machine-learning strings, motivate preserving incomparability instead of imposing arbitrary rankings. They fall outside our product-of-chains assumption: general partially ordered coordinates require separate analysis. Our results concern ordinal features, not empirical training of general-purpose AI models.

1id1:core:20261120-0001@v1 1id1.net

Prediction from compact relational evidence. Positive closure supports prediction and compression without selecting one explanatory order, suggesting a potential role in structured evaluation of language-model outputs and agent plans. Preference modelling beyond scalar rewards (Munos et al., 2024) and multi-objective language-model alignment (Mukherjee et al., 2024) motivate retaining trade-offs rather than forcing a ranking. Our model instead concerns undirected comparability under unknown coordinate polarities, not preference strength or policy optimisation. These are potential settings, not demonstrated applications.

1.1 Related work and contribution boundaries

Signed-graph balance and its linear-time test are classical (Harary, 1953; Harary and Kabell, 1980; Zaslavsky, 1982). The parity/NAE tractability boundary follows from established constraint-satisfaction theory (Schaefer, 1978; Creignou and Hermann, 1996), and the reduction from hard consistency to hard proper learning is the Pitt–Valiant argument (Pitt and Valiant, 1988). We use these results explicitly, rather than presenting them as new graph or CSP theorems.

The closure rule is a most-specific consistent predictor in the intersection closure of $\mathcal{C}_d$. Conservative prediction, monotonic error regions and compression have established guarantees (Hanneke, 2016); compression may reconstruct outside the target class (Littlestone and Warmuth, 1986; Floyd and Warmuth, 1995). Our contribution combines an efficient orthant-comparability closure representation, $d-1$-example compression without auxiliary bits, an exact online mistake bound, and proper-learning hardness within the same finite class. Our direct high-probability PAC proof does not require a separate compression theorem.

The VC, query and teaching parameters use standard models (Valiant, 1984; Angluin, 1988; Goldman and Kearns, 1995). Agnostic empirical optimisation connects to signed agreement and maximum cut (Kearns et al., 1994; Garey et al., 1976; Goemans and Williamson, 1995; Håstad, 2001). Our PAC-Bayes analysis applies established inequalities (McAllester, 1999; Seeger, 2002; Maurer, 2004; Langford, 2005) to an exactly computable affine distribution; we do not propose a new general PAC-Bayes inequality. Random product orders have been studied by Winkler (1985); Bollobás and Brightwell (1995). Our identification analysis uses Chen–Stein (Arratia et al., 1989) at bounded intensity and the two-sided estimates of Suen (1990); Janson (1998) at the coupon-collector scale.

Coordinate-reversed orders appeared in earlier work on local Hasse graphs for instance-based learning (Pan, 2008). Here the direction is hidden: the learner observes whether two vectors are comparable, but not which one precedes the other. This distinction is what makes prediction possible without reconstructing a single order.

2 Preliminaries

2.1 Oriented product orders and sign vectors

Throughout, $d \geq 2$ and $[d] = \{1, \ldots, d\}$. The instance space is a product of chains

$$\mathcal{X} = X_1 \times \cdots \times X_d,$$

1id1:core:20261120-0001@v1 1id1.net

for instance $\mathbb{R}^d$ or a product of finite ordinal scales. We assume that every chain has at least two elements. Let $\Sigma_d = \{-,+\}^d$ be the set of *directions*; signs are multiplied in the usual way, with + as identity.

Definition 2.1 (Oriented product order and finite orthant-order system) *For $\sigma \in \Sigma_d$ and $x, y \in \mathcal{X}$,*

$$x \preceq_\sigma y \iff (\sigma_i = + \Rightarrow x_i \le y_i) \text{ and } (\sigma_i = - \Rightarrow x_i \ge y_i) \quad \text{for every } i \in [d].$$

Given a finite set of records D and a feature map $\varphi : D \to \mathcal{X}$, put $a \preceq_\sigma^D b$ if and only if $\varphi(a) \preceq_\sigma \varphi(b)$. The finite orthant-order (FOO) system *is the indexed family $\mathcal{O}(D, \varphi) = (D, (\preceq_\sigma^D)_{\sigma \in \Sigma_d})$.*

Each $\preceq_\sigma$ is a partial order on $\mathcal{X}$, with $\preceq_{-\sigma}$ its opposite. On D the pullbacks are preorders: records with identical vectors precede one another in both directions for every σ. Identifying such records gives the ordered image $\varphi(D)$ for each direction. We learn on pairs of distinct feature vectors in the ambient product; on a particular finite image, even nonopposite directions may give identical comparability labels. Write $x \parallel_\sigma y$ for incomparability under $\preceq_\sigma$.

Definition 2.2 (Pairwise sign vector) *For $x, y \in \mathcal{X}$, let $\chi(x, y) \in \{-, 0, +\}^d$ be given by $\chi(x, y)_i = \operatorname{sgn}(y_i - x_i)$, that is, + if $x_i < y_i$, 0 if $x_i = y_i$, and − if $x_i > y_i$. Its* support *is $S(x, y) = \{i : \chi(x, y)_i \neq 0\}$, and we write $e(x, y) = d - |S(x, y)|$ for the number of tied coordinates.*

For $v \in \{-, 0, +\}^d$ the *face* of the cube Σ_d determined by v is

$$F(v) = \{\sigma \in \Sigma_d : \sigma_i = v_i \text{ for all } i \in \operatorname{supp}(v)\},$$

a set of $2^{d - |\operatorname{supp}(v)|}$ directions.

Every nonzero support–sign pattern is realised on any product of chains with two available levels per coordinate: choose the two endpoints in the indicated order on its support, and equal endpoints elsewhere. Consequently, constructions written with numeric coordinates below use only their sign patterns and transfer to these products.

Proposition 2.3 (Sign representation) *For all $x, y \in \mathcal{X}$,*

$$\{\sigma : x \preceq_\sigma y\} = F(\chi(x, y)),$$
$$C(x, y) := \{\sigma : x \preceq_\sigma y \text{ or } y \preceq_\sigma x\} = F(\chi(x, y)) \cup F(-\chi(x, y)).$$

If $x \neq y$, the two faces are disjoint, and $|C(x, y)| = 2^{\,e(x,y)+1}$.

Proof If $x_i < y_i$, then $x \preceq_\sigma y$ forces $\sigma_i = +$; if $x_i > y_i$ it forces $\sigma_i = -$; a tie imposes nothing. These are the defining conditions of $F(\chi(x, y))$. Apply this to (y, x), noting $\chi(y, x) = -\chi(x, y)$. When $\chi(x, y) \neq 0$, no direction agrees with both χ and $-\chi$ on their common nonempty support, so the faces are disjoint, each of size $2^{e(x,y)}$. ∎

1id1:core:20261120-0001@v1 1id1.net

2.2 The hypothesis class

Definition 2.4 (Comparability concepts) *For $\sigma \in \Sigma_d$, define $h_\sigma \colon \binom{\mathcal{X}}{2} \to \{0,1\}$ by $h_\sigma(\{x,y\}) = 1$ iff $\sigma \in C(x,y)$. Let $\mathcal{C}_d = \{h_\sigma : \sigma \in \Sigma_d\}$.*

Since $C(x,y)$ is invariant under $\sigma \mapsto -\sigma$, we have $h_\sigma = h_{-\sigma}$. Let $\mathbb{P}\Sigma_d = \Sigma_d/\{\sigma \sim -\sigma\}$ denote the set of 2^{d-1} *projective directions*, written $[\sigma]$. By taking x and y with $\chi(x,y) = \sigma$ (possible because every chain has two elements), h_σ and $h_{\sigma'}$ differ unless $\sigma' = \pm\sigma$; hence $|\mathcal{C}_d| = 2^{d-1}$ exactly.

A *labelled example* is a pair $(\{x,y\}, \ell)$ with $\ell \in \{0,1\}$; it is *positive* if $\ell = 1$ (the pair is comparable) and *negative* if $\ell = 0$ (the pair is a trade-off). A sample $S = (P, N)$ is a finite sequence, partitioned into its positive and negative subsequences. Version spaces ignore repetitions, but empirical risks and sample sizes count them with multiplicity. The *version space* of S is

$$\mathcal{V}(S) = \{\sigma \in \Sigma_d : h_\sigma(p) = 1 \ \forall p \in P,\ h_\sigma(n) = 0 \ \forall n \in N\} = \bigcap_{p \in P} C(p) \ \cap \ \bigcap_{n \in N} (\Sigma_d \setminus C(n)),$$

and S is *realizable* (or *consistent*) if $\mathcal{V}(S) \neq \varnothing$. The *support* of an example $\{x,y\}$ is $S(x,y)$; we say that a sample has *support at most s* if every example differs in at most s coordinates. Tied coordinates are the norm for ordinal and discrete features, and they are precisely what makes the problem nontrivial.

Remark 2.5 (What is and is not observed) *If $x \preceq_{\sigma^*} y$ is observed, the signs of σ^* on $S(x,y)$ are fixed. Such directed comparisons add unary constraints. Our feedback does not contain this information. All identification claims below therefore concern $[\sigma^*] = \{\sigma^*, -\sigma^*\}$. The distribution-free learning results use independent draws of labelled pairs; the complete-pair sample generated by N random instances in Section 9 has dependent observations. In particular, one cannot substitute $m = \binom{N}{2}$ into an i.i.d. PAC bound for that second sampling protocol.*

2.3 Invariance under feature transformations

Assume the chains are real intervals and let $f(x) = (f_1(x_1), \dots, f_d(x_d))$ with each f_i strictly monotone. Let $\varepsilon(f) \in \Sigma_d$ record which f_i are increasing $(+)$ or decreasing $(-)$.

Proposition 2.6 (Covariance) *Coordinatewise sign transformation gives*

$$\chi(f(x), f(y)) = \varepsilon(f)\, \chi(x,y), \qquad h_\sigma(f(x), f(y)) = h_{\varepsilon(f)\sigma}(x,y).$$

Consequently $\mathcal{V}(f(S)) = \varepsilon(f)\, \mathcal{V}(S)$. In particular, realizability, version-space size, and all quantities studied below are invariant under coordinatewise strictly monotone reparametrisation of the features.

Proof A strictly increasing map preserves $<, =, >$, a strictly decreasing one interchanges $<$ and $>$ and preserves $=$. Apply Proposition 2.3. ∎

1id1:core:20261120-0001@v1 1id1.net

3 Version spaces from positive examples

3.1 Parity form

Lemma 3.1 (One positive example) *Let $x \neq y$ and $\alpha = \chi(x,y)|_{S(x,y)}$. Then $\sigma \in C(x,y)$ if and only if there is $\tau \in \{-,+\}$ with $\sigma_i = \tau\alpha_i$ for all $i \in S(x,y)$; equivalently,*

$$\sigma_i\sigma_j = \alpha_i\alpha_j \qquad \text{for all } i,j \in S(x,y).$$

Proof By Proposition 2.3, $\sigma \in C(x,y)$ iff σ agrees on $S(x,y)$ with α or with $-\alpha$. Eliminating the common sign τ gives the pairwise equations; conversely, the pairwise equations say that $\sigma_i\alpha_i$ is constant on $S(x,y)$. ■

A positive example therefore does not select a direction. It imposes parity constraints between the coordinates on which its two instances differ. Identify Σ_d with $\mathbb{F}_2^d$ via $+ \mapsto 0$, $- \mapsto 1$, so that sign multiplication becomes addition. For $x \neq y$ let

$$W_{xy} = \{b \in \mathbb{F}_2^d : \ b_i = b_j \text{ for all } i,j \in S(x,y)\} = \langle e_i : i \notin S(x,y)\rangle + \langle \mathbf{1}\rangle.$$

Proposition 3.2 (Coset structure) *$C(x,y)$ is a coset of W_{xy}. Hence for every set P of positive examples, $\mathcal{V}(P) = \bigcap_{p\in P} C(p)$ is either empty or a coset of $W_P = \bigcap_{p\in P} W_p$, and $\dim_{\mathbb{F}_2} W_P$ is the number of connected components of the support hypergraph $H(P) = ([d], \{S(p) : p \in P\})$, isolated vertices included. Since $\mathbf{1} \in W_P$, a nonempty $\mathcal{V}(P)$ is closed under $\sigma \mapsto -\sigma$.*

Proof Let $\tilde{\beta} \in \mathbb{F}_2^d$ agree with the binary form of α on $S(x,y)$ and vanish elsewhere. By Lemma 3.1, $b \in C(x,y)$ iff $b + \tilde{\beta}$ is constant on $S(x,y)$, i.e. $b \in \tilde{\beta} + W_{xy}$. A nonempty intersection of cosets is a coset of the intersection of the subspaces. A vector lies in W_P iff it is constant on every support, i.e. on every component of $H(P)$; one free bit per component gives the dimension. ■

3.2 The signed constraint graph

A signed multigraph has edges labelled $\pm$; the sign of a cycle is the product of its edge signs, and the graph is *balanced* if every cycle is positive (Harary, 1953; Zaslavsky, 1982). A vertex signing $\sigma \in \Sigma_d$ *satisfies* the graph if $\sigma_i\sigma_j = \lambda(ij)$ for every edge ij of sign $\lambda(ij)$.

Definition 3.3 (Signed constraint graph) *For a set P of positive examples, $G(P)$ has vertex set $[d]$. For each $p = \{x,y\} \in P$ with $|S(p)| \geq 2$, fix a root $r_p \in S(p)$ and add, for every $i \in S(p) \setminus \{r_p\}$, an edge $r_p i$ of sign $\alpha_{r_p}\alpha_i$, where $\alpha = \chi(x,y)$.*

The graph depends on the roots, but by Lemma 3.1 its set of satisfying signings equals $\mathcal{V}(P)$ for any choice, and its connected components are those of $H(P)$. Let $c(P)$ denote their number.

Theorem 3.4 (Version-space theorem) *For a set P of positive examples the following are equivalent:*

1id1:core:20261120-0001@v1 1id1.net

(a) P is realizable;

(b) $G(P)$ admits a satisfying signing;

(c) $G(P)$ is balanced.

If they hold, $\mathcal{V}(P)$ is an affine space over $\mathbb{F}_2$ of dimension $c(P)$, so $|\mathcal{V}(P)| = 2^{c(P)}$, and the number of consistent hypotheses in $\mathcal{C}_d$ is $2^{c(P)-1}$.

Proof (a)⇔(b) is Lemma 3.1. If σ satisfies all edges, the sign of any cycle $v_0 v_1 \cdots v_k = v_0$ is $\prod_j \sigma_{v_j}\sigma_{v_{j+1}} = +$, as each vertex sign occurs twice; so (b)⇒(c). Conversely, if $G(P)$ is balanced, choose a root in each component, assign it an arbitrary sign, and propagate along a spanning tree by $\sigma_j = \lambda(ij)\sigma_i$. A violated non-tree edge would close a negative cycle with the tree path, so the signing satisfies every edge. The two root choices per component are independent and everything else is forced, giving $2^{c(P)}$ satisfying signings; this matches Proposition 3.2. Antipodal pairs collapse to $2^{c(P)-1}$ hypotheses. ■

Thus $d - c(P)$ is the number of independent bits a realizable positive sample has revealed about σ^*, and $c(P) - 1$ is the number of bits still unknown about the hypothesis h_{σ^*}.

Corollary 3.5 (General position) *If $P \neq \varnothing$ and no example in P has a tied coordinate, then P is realizable iff all sign vectors $\chi(p)$, $p \in P$, coincide up to global sign, and in that case exactly one hypothesis is consistent.*

Proof Full supports connect all vertices, so $c(P) = 1$, and Lemma 3.1 forces $\sigma = \pm\chi(p)$ for every p. ■

3.3 Algorithmic consequences

Write $L(P) = \sum_{p \in P} \max\{|S(p)| - 1, 0\}$ for the number of edges of $G(P)$.

Theorem 3.6 (Linear-time consistency, counting, and certificates) *From the support–sign representation, one can decide positive consistency in $O(d + |P| + L(P))$ time. The algorithm outputs the exponent $c(P)$ of $|\mathcal{V}(P)| = 2^{c(P)}$ and a consistent direction when it is, and output a negative cycle of $G(P)$—a set of at most d examples that is already unrealizable—when it is not. For instances given as coordinate vectors the total time is $O(d(1 + |P|))$, linear in the input. The test is maintainable online under insertion of examples by a disjoint-set structure with parity labels, in amortised inverse-Ackermann time per edge.*

Proof Build $G(P)$ and run the propagation in the proof of Theorem 3.4 by depth-first search; a disagreeing edge together with the tree path between its endpoints is a negative cycle. A simple cycle in a multigraph on d vertices has at most d edges, and each edge belongs to one example. Union–find with parity merges components or detects a parity conflict per edge. ■

1id1:core:20261120-0001@v1 1id1.net

For fixed observed pairs and a noiseless likelihood, a uniform prior on Σ_d gives a uniform posterior on this coset after positive labels. The count is represented by its exponent; running times for graph operations assume unit-cost access to support indices and signs. Extracting the support–sign representation from dense inputs incurs the coordinate-comparison cost stated above.

4 Closure learning and compression without side information

Theorem 4.1 (Compression of positive samples) *For every realizable set P of positive examples there is $P' \subseteq P$ with $|P'| \leq d-1$ and $\mathcal{V}(P') = \mathcal{V}(P)$. The bound $d-1$ is attained for every $d \geq 2$.*

Proof Process the examples in any order and keep p iff its support $S(p)$ merges at least two current components of the support hypergraph. Each kept example reduces the number of components by at least one, from d to $c(P) \geq 1$, so at most $d-1$ examples are kept, and the kept set P' has the same components as P. By Proposition 3.2, $W_{P'} = W_P$; both $\mathcal{V}(P)$ and $\mathcal{V}(P') \supseteq \mathcal{V}(P)$ are nonempty cosets of this subspace, hence equal.

For tightness take $q = 0 \in \mathbb{R}^d$ and $p_i = \{q, e_i + e_{i+1}\}$, $1 \leq i \leq d-1$. The constraints $\sigma_i = \sigma_{i+1}$ are consistent with $c(P) = 1$; deleting any p_i disconnects the path and doubles the version space. ■

4.1 The closure predictor

For a consistent positive sample P, define

$$g_P(z) = \mathbf{1}[\mathcal{V}(P) \subseteq C(z)].$$

This is the pointwise intersection of all concepts consistent with P. It need not belong to $\mathcal{C}_d$. Already for $P = \varnothing$, g_P accepts exactly the pairs of support one, whereas each h_σ accepts some full-support pair.

Theorem 4.2 (Efficient closure learning and optimal mistake bound) *Let the labels be given by some $h_{\sigma^*} \in \mathcal{C}_d$.*

(a) For every positive sample P, $g_P \leq h_{\sigma^*}$ pointwise. It is consistent with every realizable mixed sample $S = (P, N)$.

(b) After $O(d + |P| + L(P))$ preprocessing, $g_P(z)$ is evaluated in $O(|S(z)|)$ time from the support–sign representation.

(c) Starting with $P = \varnothing$, predict by g_P and retain each mistaken positive example. This online algorithm makes at most $d-1$ mistakes on every realizable sequence, all false negatives. No deterministic online learner has a smaller worst-case mistake bound.

1id1:core:20261120-0001@v1 1id1.net

Proof (a) The target direction belongs to $\mathcal{V}(P)$, so every pair implied by $\mathcal{V}(P)$ is positive under the target. Each observed positive is implied by definition. This also proves consistency on the negative examples, without using them in the construction.

(b) Choose a satisfying signing a of $G(P)$, and store the component index $t(i)$ of each coordinate. Then $\mathcal{V}(P) = \{\sigma : \sigma_i = a_i u_{t(i)},\ u_t \in \{-,+\}\}$. For a query with signs α_i on support T, it is implied iff all coordinates of T lie in the same component and $a_i\alpha_i$ is constant on T. Indeed, these conditions make $\sigma_i\alpha_i$ constant for every choice of the free component signs. If two components occur, their free signs can violate this equality; if the within-component signs disagree, no direction satisfies the query. A single scan suffices, including the support-one case.

(c) There are no false positives by (a). A mistaken positive z satisfies $\varnothing \neq \mathcal{V}(P \cup \{z\}) \subsetneq \mathcal{V}(P)$. These are affine spaces, so their dimensions drop by at least one. The dimension starts at d and stays at least one because of global sign reversal. Thus at most $d-1$ mistakes occur. A correctly predicted positive adds no constraint, and negative observations need no update; the retained examples therefore generate the same positive version space as the full history. Parity union–find implements insertions in amortised inverse-Ackermann time per support edge; query scans using its representatives have the same inverse-Ackermann overhead online.

For the lower bound, use $d-1$ pairs supported on $\{1, i\}$, $i = 2, \ldots, d$, with both nonzero signs positive. Every labelling of these pairs is realised by setting $\sigma_1 = +$ and choosing each σ_i independently. An adversary presents these pairs once each and chooses the opposite of a deterministic learner's prediction. The resulting sequence is realizable and causes $d-1$ mistakes. ■

Theorem 4.3 (Mixed-sample compression without auxiliary bits) *The class $\mathcal{C}_d$ admits an efficient sample compression scheme of size at most $d-1$, with no auxiliary bits, for arbitrary realizable mixed samples. The compressed sample consists solely of positive examples, and reconstruction returns their closure predictor.*

Proof Given $S = (P, N)$, select $P' \subseteq P$ as in Theorem 4.1, so $|P'| \leq d-1$ and $\mathcal{V}(P') = \mathcal{V}(P)$. Fixing the processing order, for example the original sample order, makes the selection deterministic. Reconstruction receives only the selected examples and returns $g_{P'}$. Their order is immaterial to the reconstruction, no negative examples are stored, and their labels are all known to be positive. The equality of version spaces gives $g_{P'} = g_P$, which fits the entire mixed sample by Theorem 4.2(a). For inputs to the reconstruction map that are not consistent positive samples, define its output arbitrarily; such inputs never arise in the claimed scheme. The selected subset can be found with parity union–find, and reconstruction uses Theorem 3.6. ■

This is an *improper* compression scheme in the usual sense (Floyd and Warmuth, 1995): reconstruction is a Boolean predictor, not necessarily an order-induced concept. The tightness statement in Theorem 4.1 concerns preservation of the entire positive version space, and does not prove a lower bound for every mixed-sample compression scheme.

1id1:core:20261120-0001@v1 1id1.net

4.2 A distribution-free PAC guarantee

Theorem 4.4 (PAC guarantee for the final closure predictor) *Let S_m consist of m i.i.d. pairs drawn from any distribution $\mathcal{D}$, labelled by h_{σ^*}, and let P_m be its positive subsequence. Put $k = d - 1$ and $R_m = \mathbb{P}_{z\sim\mathcal{D}}(g_{P_m}(z) \neq h_{\sigma^*}(z))$. For $0 < \varepsilon < 1$,*

$$\mathbb{P}(R_m > \varepsilon) \leq \mathbb{P}(\mathrm{Bin}(m, \varepsilon) \leq k) \leq \exp\bigl(k - (1 - e^{-1})m\varepsilon\bigr).$$

In particular, with probability at least $1 - \delta$,

$$R_m \leq \min\left\{1, \frac{2(d - 1 + \ln(1/\delta))}{m}\right\}.$$

Thus $\mathcal{C}_d$ is efficiently improperly PAC-learnable with $O((d + \ln(1/\delta))/\varepsilon)$ examples.

Proof Run the online rule of Theorem 4.2 on the sample. Its predictor after t examples equals g_{P_t}. The positive version spaces decrease, so the accepted regions of g_{P_t} increase. They are all contained in the target's positive region; hence R_t is nonincreasing. Let I_t be the mistake indicator at step t. Conditional on the past, $\mathbb{P}(I_t = 1) = R_{t-1}$, and $\sum_{t=1}^{m} I_t \leq k$.

For completeness we give the binomial comparison, including the stopping step. On an enlarged probability space construct a Bernoulli variable B_t as follows. If $R_{t-1} > \varepsilon$, retain a mistake independently with probability ε/R_{t-1}, and put $B_t = 1$ exactly when that mistake is retained. Otherwise draw B_t independently with probability ε. Given the enlarged past, $\mathbb{P}(B_t = 1) = \varepsilon$; therefore the B_t's are independent Bernoulli(ε) variables. On $\{R_m > \varepsilon\}$, monotonicity gives $R_{t-1} > \varepsilon$ for all t, so $B_t \leq I_t$ and $\sum_t B_t \leq k$. This proves the first inequality. If $B \sim \mathrm{Bin}(m, \varepsilon)$, Markov's inequality gives

$$\mathbb{P}(B \leq k) \leq e^k \mathbb{E}e^{-B} = e^k(1 - \varepsilon + \varepsilon e^{-1})^m \leq e^{k-(1-e^{-1})m\varepsilon}.$$

Substitute $\varepsilon = 2(k + \ln(1/\delta))/m$ when this is less than one. Since $2(1 - e^{-1}) > 1$, the failure probability is at most δ; otherwise the risk bound is trivial. ∎

The rate matches the finite-class upper bound and the usual VC-dimension dependence (Theorem 6.1 and Corollary 6.2), but has an efficient improper implementation even in the hard proper regime. It also illustrates the role of monotonic error regions (Hanneke, 2016); an online mistake bound alone does not, for an arbitrary learning algorithm, imply this guarantee for its final iterate.

4.3 Certificates of positive inconsistency

Theorem 4.5 (Helly number of positive consistency) *Every unrealizable set P of positive examples contains an unrealizable subset of at most d examples. For every $d \geq 2$ there is an unrealizable set of d examples all of whose proper subsets are realizable. If no example has a tied coordinate, the bound drops to 2.*

Proof The sign of a cycle is additive over symmetric differences, and the cycle space of a multigraph is spanned by simple cycles, so an unbalanced signed multigraph contains a simple negative cycle; on d vertices it has at most d edges (two parallel edges of opposite

1id1:core:20261120-0001@v1 1id1.net

sign included). By Theorem 3.4, $G(P)$ has such a cycle Z. Select one example per edge of Z; with the same roots, the graph of the selected examples contains Z and is unbalanced.

For attainment with $d \geq 3$, take $q = 0$ and $x_i = e_i + e_{i+1}$ for $i < d$, $x_d = e_d - e_1$; the examples $\{q, x_i\}$ impose $\sigma_i\sigma_{i+1} = +$ for $i < d$ and $\sigma_d\sigma_1 = -$, a d-cycle with one negative edge. Deleting any example leaves a forest. For $d = 2$ take $\{0, (1, 1)\}$ and $\{0, (1, -1)\}$.

In general position, pairwise realizability means pairwise agreement of sign vectors up to global sign (Corollary 3.5), which is an equivalence relation, so pairwise realizability implies realizability. ■

Theorem 4.5 says that a negative-cycle certificate of inconsistency is small—at most d examples—but not of constant size: no fixed number of local checks suffices in growing dimension. Combined with Theorem 3.6, a learner can always exhibit at most d examples that no orthant comparability concept can satisfy simultaneously.

5 Negative examples: a complexity dichotomy

Lemma 5.1 (Signed not-all-equal form) *Let $x \neq y$, $\alpha = \chi(x, y)|_{S(x,y)}$. Then $\sigma \notin C(x, y)$ iff $\{\sigma_i\alpha_i : i \in S(x, y)\} = \{-, +\}$, that is, iff the binary variables $b_i + \beta_i$, $i \in S(x, y)$, are not all equal. If $|S(x, y)| \leq 1$, then $C(x, y) = \Sigma_d$ and a negative label on $\{x, y\}$ is unrealizable.*

Proof Negate Lemma 3.1. ■

Positive feedback is thus a parity (affine) constraint language, negative feedback a not-all-equal language. Schaefer's dichotomy (Schaefer, 1978) places these on opposite sides of the tractability frontier, and the frontier is crossed exactly at support three.

Theorem 5.2 (Consistency dichotomy) *Consider the consistency problem for $\mathcal{C}_d$: given a sample $S = (P, N)$ by its support–sign pairs, decide whether $\mathcal{V}(S) \neq \varnothing$. Let s bound the supports of the negative examples; the positive supports are unrestricted.*

(a) If $s \leq 2$, then $\mathcal{V}(S)$ is empty or a coset of a subspace of $\mathbb{F}_2^d$ containing $\mathbf{1}$; consistency and $|\mathcal{V}(S)|$ are computed in $O\big(d + |P| + L(P) + |N|\big)$ time, and $|\mathcal{V}(S)| = 2^c$ with c the number of components of a signed multigraph on $[d]$. Every inconsistent sample contains an inconsistent subsample of at most d examples.

(b) If $s \geq 3$, the problem is NP*-complete, already for $P = \varnothing$, all negative examples of the form $\{0, x\}$ with $x \in \{-1, 0, 1\}^d$, and $|S(x)| = 3$.*

Proof (a) A negative example of support at most one is unrealizable by Lemma 5.1. A negative example with $S = \{i, j\}$ states $\sigma_i\sigma_j = -\alpha_i\alpha_j$: a single signed edge. Add these edges to $G(P)$. All constraints are now parity equations on pairs of variables, and the proof of Theorem 3.4 applies verbatim to the combined multigraph, including the count 2^c, the linear-time algorithm, and the Helly bound (a negative cycle uses at most d edges, each from one example; a support-≤ 1 negative example is a one-element certificate).

1id1:core:20261120-0001@v1 1id1.net

(b) Membership in NP is clear. Reduce from NAE-3-SAT with signed literals, which is NP-complete (Schaefer, 1978). Represent Boolean variables by $\sigma_1, \ldots, \sigma_d$, and a clause requiring $\gamma_i\sigma_i, \gamma_j\sigma_j, \gamma_k\sigma_k$ to be not all equal, $\gamma \in \{-, +\}$, by the negative example $\{0, x\}$ with $x_i = \gamma_i$, $x_j = \gamma_j$, $x_k = \gamma_k$, and $x_r = 0$ otherwise. By Lemma 5.1, $\sigma \notin C(0, x)$ iff the clause holds. ■

Corollary 5.3 (Proper–improper separation) *For distributions whose realizable negative examples have support at most two, $\mathcal{C}_d$ is efficiently properly PAC-learnable. This includes distributions supported entirely on positive examples. For arbitrary realizable distributions with negative supports of size three, no polynomial-time proper PAC learner exists unless* RP = NP. *Nevertheless Theorem 4.4 gives an efficient improper learner without any support restriction.*

Proof In the first regime, output a satisfying direction from the combined signed graph of Theorem 5.2(a) and apply the realizable finite-class bound of Corollary 6.2. The positive-only statement assumes that the sampling distribution itself has no negative labels; it is not a guarantee for learning an arbitrary distribution from a label-conditioned positive-only oracle.

For hardness, take an instance with m examples of the consistency problem in Theorem 5.2(b), sample uniformly from these examples, and run a putative proper learner with $\varepsilon = 1/(2m)$ and $\delta = 1/4$. Verify its returned direction on every input example and accept only if all labels agree. On a realizable input it succeeds with probability at least $3/4$; error below $1/m$ forces exact consistency. On an unrealizable input no direction passes verification. The polynomial running-time requirement is enforced on all inputs, by truncation at the learner's guaranteed polynomial bound if necessary. This yields a one-sided randomized polynomial-time decision procedure for an NP-complete problem (Pitt and Valiant, 1988), implying RP = NP. Improper closure predictions cannot be used as such certificates: fitting a finite negative sample does not certify that an explaining direction exists. ■

Remark 5.4 *Hardness requires ties. In general position every pair has full support, a negative example $\{x, y\}$ excludes exactly the projective class $[\chi(x, y)]$, and consistency is the trivial problem of avoiding a list of at most $|N|$ classes among 2^{d-1}. Ordinal and discrete features, the setting in which monotonicity constraints are most often imposed, are exactly where the dichotomy bites.*

Counting is governed by the same boundary.

Proposition 5.5 (Inclusion–exclusion) *For every finite set N of negative examples,*

$$|\mathcal{V}(\varnothing, N)| = \sum_{B \subseteq N} (-1)^{|B|} \, |\mathcal{V}(B, \varnothing)| = \sum_{\substack{B \subseteq N \\ B \text{ realizable as positives}}} (-1)^{|B|} \, 2^{c(B)},$$

with $c(\varnothing) = d$.

1id1:core:20261120-0001@v1 1id1.net

Proof $\mathcal{V}(\varnothing, N) = \Sigma_d \setminus \bigcup_{n \in N} C(n)$; apply inclusion–exclusion and Theorem 3.4. ■

Corollary 5.6 (Counting dichotomy) *Computing $|\mathcal{V}(S)|$ takes $O(d + |P| + L(P) + |N|)$ time when negative supports are at most two, and is #P-complete under polynomial-time Turing reductions when supports of size three are allowed, already for $P = \varnothing$ and examples $\{0, x\}$, $x \in \{-1, 0, 1\}^d$.*

Proof The first part is Theorem 5.2(a). For the second, the reduction in Theorem 5.2(b) is parsimonious—directions in $\mathcal{V}$ correspond bijectively to satisfying assignments, unused variables being free on both sides—and counting signed NAE-3-SAT solutions is #P-complete because the not-all-equal relation is not affine (Creignou and Hermann, 1996). ■

For fixed designs and noiseless labels, the positive-only posterior has a closed form, whereas exact normalisation of the full mixed-feedback posterior is #P-hard. This counting statement alone does not establish hardness of every approximate sampling task. Proposition 5.5 shows that enumerating subsets of the negative examples gives an exact $2^{|N|}$-times-polynomial algorithm, with polynomial work in d and the total support input length for each subset.

6 Statistical parameters: VC dimension, queries, teaching

Theorem 6.1 (VC dimension) $\mathrm{VC}(\mathcal{C}_d) = d - 1$.

Proof Upper bound: $|\mathcal{C}_d| = 2^{d-1}$, and a class shattering m points has at least 2^m members. For the lower bound, choose two ordered levels per coordinate and let q use the lower levels. For $i = 2, \dots, d$, an instance $x^{(i)}$ that differs from q exactly in coordinates 1 and i, with $\chi(q, x^{(i)})_1 = \chi(q, x^{(i)})_i = +$. By Lemma 3.1, $h_\sigma(\{q, x^{(i)}\}) = 1$ iff $\sigma_1 \sigma_i = +$. Given any labels $\ell_2, \dots, \ell_d \in \{0, 1\}$, take $\sigma_1 = +$ and $\sigma_i = +$ iff $\ell_i = 1$. The $d - 1$ pairs are shattered. ■

Corollary 6.2 (PAC sample complexity) *Let $\mathcal{D}$ be any distribution on $\binom{\mathcal{X}}{2}$. In the realizable case, any hypothesis in $\mathcal{C}_d$ consistent with $m = O\big((d + \log(1/\delta))/\varepsilon\big)$ i.i.d. examples has error at most ε with probability $1 - \delta$; in the agnostic case, $m = O\big((d + \log(1/\delta))/\varepsilon^2\big)$ examples suffice for empirical risk minimisation to be ε-optimal. Both bounds are tight up to constants in d.*

Proof Standard bounds for finite classes with $\log_2 |\mathcal{C}_d| = d - 1$, and VC-dimension lower bounds (Vapnik and Chervonenkis, 1971; Blumer et al., 1989; Shalev-Shwartz and Ben-David, 2014). ■

These statistical bounds do not ensure efficient proper learning. The closure algorithm in Section 4 achieves the realizable rate improperly; Section 5 gives the contrasting proper hardness. The agnostic computational problem is considered next.

1id1:core:20261120-0001@v1 1id1.net

Theorem 6.3 (Query complexity) *In Angluin's model with membership queries "is $\{x, y\}$ comparable under $\preceq_{\sigma^*}$?", exact identification of h_{σ^*} requires exactly $d-1$ queries in the worst case.*

Proof Upper bound: query the pairs $\{q, x^{(i)}\}$ of the proof of Theorem 6.1; answer i reveals $\sigma_1^*\sigma_i^*$, and the $d-1$ products determine σ^* up to global sign. Lower bound: there are 2^{d-1} hypotheses and each answer is one bit, so an adversary answering so as to keep the larger half of the current version space forces at least $\lceil \log_2 2^{d-1} \rceil = d-1$ queries. ■

Proposition 6.4 (Teaching dimension) *Every $h_\sigma \in \mathcal{C}_d$ has a teaching set of size one, hence* $\mathrm{TD}(\mathcal{C}_d) = 1$.

Proof Choose x, y with $\chi(x, y) = \sigma$, a full-support pair. By Proposition 2.3, $h_{\sigma'}(\{x, y\}) = 1$ iff $\sigma' \in \{\sigma, -\sigma\}$, so the single positive example $(\{x, y\}, 1)$ is consistent with h_σ only. ■

The gap between teaching dimension 1 and VC dimension $d-1$ is the largest possible for a class of size 2^{d-1}. It reflects the two sources of information in a positive example: a full-support pair is maximally informative (it pins down the hypothesis), whereas a pair with many ties contributes only parity relations among the few coordinates where it differs. Ordinal samples can contain the second kind; the signed-graph representation handles them without a general-position assumption.

7 Agnostic learning

In the agnostic setting (Kearns et al., 1994) the feedback need not be realizable, and the goal is to minimise the empirical risk $\mathrm{err}_S(h) = |\{e \in S : h(e) \neq \ell(e)\}|$ over $h \in \mathcal{C}_d$. By Corollary 6.2 a sample of size $O((d + \log(1/\delta))/\varepsilon^2)$ makes the empirical risk minimiser ε-optimal; the obstacle is computational.

Theorem 7.1 (Hardness of empirical risk minimisation) *Deciding whether some $h \in \mathcal{C}_d$ has $\mathrm{err}_S(h) \leq k$ is* NP*-complete, already for samples consisting of positive examples only, all of the form $\{0, x\}$ with $x \in \{-1, 0, 1\}^d$ and $|S(x)| = 2$. Maximising the number of satisfied examples is* APX*-hard, and no polynomial-time algorithm approximates the maximum number of satisfied positive examples within $16/17 + \varepsilon$ unless* P = NP*. On the positive side, for positive-only samples of support two the maximum number of satisfied examples is approximable within* 0.878 *in polynomial time.*

Proof A positive example $\{0, x\}$ with $S(x) = \{u, v\}$ is one signed edge uv with sign $\chi_u\chi_v$, and a subset of such examples is realizable iff the corresponding signed subgraph is balanced (Theorem 3.4). Maximising satisfied examples is therefore the maximum balanced subgraph problem, and when all edges are negative it is exactly Max Cut: a set of negative edges is balanced iff it is bipartite, and a bipartite edge set extends to a cut containing it. Given a simple graph $H = ([d], E)$, the sample $\{(\{0, e_u - e_v\}, 1) : uv \in E\}$ realises this: Simple Max Cut is NP-complete (Garey et al., 1976); its agreement objective is inapproximable

1id1:core:20261120-0001@v1 1id1.net

beyond 16/17 (Håstad, 2001). Membership in NP is clear. The approximation guarantee is the Goemans–Williamson algorithm for Max Res Cut, which handles equality and inequality edge constraints simultaneously (Goemans and Williamson, 1995). ∎

Examples with larger support bundle several edges into one unit that must be satisfied or violated as a whole; the approximability of this bundled version, and of mixed samples, is left open. We note that the hardness in Theorem 7.1 is for the empirical problem; whether $\mathcal{C}_d$ is *improperly* agnostically learnable in polynomial time, for instance by a convex surrogate over $\{-1,1\}^d$ relaxations, is likewise open.

8 PAC-Bayes bounds and the Gibbs classifier

The counting results of Sections 3 and 5 have a direct statistical use. For fixed observed pair inputs (or an input distribution independent of the target), a uniform prior π on Σ_d and a noiseless label likelihood give a posterior for a realizable sample S that is the uniform distribution ρ_S on $\mathcal{V}(S)$. The Gibbs classifier G_ρ draws $\sigma \sim \rho$ and predicts with h_σ; its risk under a distribution $\mathcal{D}$ on labelled pairs is $\operatorname{err}_{\mathcal{D}}(G_\rho) = \mathbb{E}_{\sigma\sim\rho} \operatorname{err}_{\mathcal{D}}(h_\sigma)$, and its empirical risk on a sample S of size m is $\operatorname{err}_S(G_\rho) = \mathbb{E}_{\sigma\sim\rho} \operatorname{err}_S(h_\sigma)/m$. PAC-Bayes theory (McAllester, 1999; Seeger, 2002; Maurer, 2004; Langford, 2005) bounds the former by the latter plus a complexity term $\mathrm{KL}(\rho \,\|\, \pi)/m$. We use the form of Maurer (Maurer, 2004): for $m \geq 8$, with probability at least $1-\delta$ over $S \sim \mathcal{D}^m$, simultaneously for all distributions ρ on Σ_d,

$$\mathrm{kl}\big(\operatorname{err}_S(G_\rho) \,\|\, \operatorname{err}_{\mathcal{D}}(G_\rho)\big) \leq \frac{\mathrm{KL}(\rho \,\|\, \pi) + \ln\big(2\sqrt{m}/\delta\big)}{m}, \tag{1}$$

where $\mathrm{kl}(a\|b) = a \ln \frac{a}{b} + (1-a) \ln \frac{1-a}{1-b}$ is the binary relative entropy. Write $\mathrm{kl}^{-1}(a,\varepsilon) = \sup\{b \in [a,1] : \mathrm{kl}(a\|b) \leq \varepsilon\}$; by Pinsker's inequality $\mathrm{kl}^{-1}(a,\varepsilon) \leq a + \sqrt{\varepsilon/2}$, and $\mathrm{kl}^{-1}(0,\varepsilon) = 1 - e^{-\varepsilon} \leq \varepsilon$.

The two ingredients of (1) are exactly the quantities computed in Section 3. For a realizable positive set P and an example $n = \{x, y\}$ define the *rank increment*

$$\mathrm{rk}_P(n) = c(P) - c(P \cup \{n\}),$$

the dimension reduction if the added positive constraint is consistent. The component counts are defined even for an inconsistent augmentation: the increment is one less than the number of components meeting $S(n)$, and $0 \leq \mathrm{rk}_P(n) \leq |S(n)| - 1$.

Lemma 8.1 (Posterior mass of a comparability event) *Let P be realizable and ρ_P uniform on $\mathcal{V}(P)$. For every example n,*

$$\rho_P\big(C(n)\big) = \mathbb{P}_{\sigma\sim\rho_P}\big(h_\sigma(n) = 1\big) = \begin{cases} 2^{-\mathrm{rk}_P(n)}, & P \cup \{n\} \text{ realizable}, \\ 0, & \text{otherwise}. \end{cases}$$

In particular $\rho_P(C(n)) \in \{0\} \cup \{2^{-r} : 0 \leq r \leq |S(n)| - 1\}$, and $\rho_P(C(n)) = 1$ iff the comparability of n is implied by P.

1id1:core:20261120-0001@v1 1id1.net

Proof $\mathcal{V}(P) \cap C(n) = \mathcal{V}(P \cup \{n\})$, which by Theorem 3.4 is empty or has size $2^{c(P\cup\{n\})}$; divide by $|\mathcal{V}(P)| = 2^{c(P)}$. ∎

Theorem 8.2 (Computable PAC-Bayes bounds for direction learning) *Let $\mathcal{D}$ be any distribution on labelled pairs, let π be uniform on Σ_d, and let $S = (P, N) \sim \mathcal{D}^m$ with $m \geq 8$. With probability at least $1 - \delta$ the following hold simultaneously.*

(a) (Realizable affine case.) *If S is realizable and either $N = \varnothing$ or all negative supports have size at most two, let ρ_S be uniform on $\mathcal{V}(S)$ and $c(S)$ the component count of the combined signed graph of Theorem 5.2(a). Then* $\mathrm{KL}(\rho_S \| \pi) = (d - c(S)) \ln 2$, $\mathrm{err}_S(G_{\rho_S}) = 0$, *and*

$$\mathrm{err}_{\mathcal{D}}(G_{\rho_S}) \leq 1 - \exp\Big(-\frac{(d - c(S)) \ln 2 + \ln(2\sqrt{m}/\delta)}{m}\Big) \leq \frac{(d - c(S)) \ln 2 + \ln(2\sqrt{m}/\delta)}{m}.$$

(b) (Mixed case, arbitrary negative supports.) *If the positive part P is realizable, let ρ_P be uniform on $\mathcal{V}(P)$. Then* $\mathrm{KL}(\rho_P \| \pi) = (d - c(P)) \ln 2$,

$$\mathrm{err}_S(G_{\rho_P}) = \frac{1}{m} \sum_{\substack{n \in N \\ P \cup \{n\} \text{ realizable}}} 2^{-\mathrm{rk}_P(n)},$$

and

$$\begin{aligned} \mathrm{err}_{\mathcal{D}}(G_{\rho_P}) &\leq \mathrm{kl}^{-1}\Big(\mathrm{err}_S(G_{\rho_P}),\ \frac{(d - c(P)) \ln 2 + \ln(2\sqrt{m}/\delta)}{m}\Big) \\ &\leq \mathrm{err}_S(G_{\rho_P}) + \sqrt{\frac{(d - c(P)) \ln 2 + \ln(2\sqrt{m}/\delta)}{2m}}. \end{aligned}$$

The graph statistics and powers of two on the right-hand sides require

$$O\left(d + |P| + L(P) + |N| + \sum_{n \in N} |S(n)|\right)$$

operations after support extraction. Exact rational arithmetic and numerical inversion of kl *have their usual additional bit/precision costs. The positive rank is maintainable online by parity union–find.*

Proof Both posteriors are uniform on a coset of size 2^c inside $|\Sigma_d| = 2^d$, so $\mathrm{KL} = \ln(2^d/2^c) = (d - c) \ln 2$. In (a) every $\sigma \in \mathcal{V}(S)$ is consistent with S, so the empirical Gibbs risk is zero and (1) gives $-\ln(1 - \mathrm{err}_{\mathcal{D}}) \leq \varepsilon$, i.e. $\mathrm{err}_{\mathcal{D}} \leq 1 - e^{-\varepsilon} \leq \varepsilon$. In (b) every $\sigma \in \mathcal{V}(P)$ is correct on P, and on a negative example n it errs iff $\sigma \in C(n)$, an event of ρ_P-probability $2^{-\mathrm{rk}_P(n)}$ or 0 by Lemma 8.1. Summing and applying (1) with Pinsker's inequality gives the bound. For the running time, preprocess a satisfying signing a_i and component indices $t(i)$ once. Scan the support of each n, recording one value $a_i \alpha_i$ for every component encountered. A disagreement within a component means $P \cup \{n\}$ is inconsistent; otherwise, if q components are encountered, $\mathrm{rk}_P(n) = q - 1$. An array

1id1:core:20261120-0001@v1 1id1.net

indexed by components with a query timestamp avoids clearing all d entries for each query. This takes $O(|S(n)|)$ operations per negative example without rebuilding the graph. The bound (1) holds for all ρ at once, so (a) and (b) share the same probability-$1-\delta$ event. ■

Three features of Theorem 8.2 deserve comment.

The complexity term is the information actually extracted. The union bound behind Corollary 6.2 charges $\log_2 |\mathcal{C}_d| = d - 1$ bits for every sample. The PAC-Bayes bound charges $d - c(S)$ bits, the rank of the parity system the sample has imposed: a sample of pairs with many ties, which pins down few bits, is charged little. For a full-support positive pair the rank is $d-1$, so the complexity term reaches the logarithm of the number of hypotheses. This observation concerns the formula, not the applicability of an i.i.d. bound to the dependent complete-pair sampling protocol.

Exact inference is hard, but the bound is not. By Corollary 5.6, when negative examples have support three the true posterior ρ_S has a #P-hard normalising constant, and even deciding whether it exists is NP-complete (Theorem 5.2). Part (b) sidesteps this by conditioning on the positive examples only and paying for the negative examples through an explicit, linear-time empirical term. Each negative example costs $2^{-\mathrm{rk}_P(n)}$: nothing if its comparability is already refuted by P, one full unit if its comparability is already implied by P (the sample is then not realizable), and a fractional amount otherwise.

Unrefuted negative examples contribute fractional Gibbs risk. Since $\mathrm{rk}_P(n) \le |S(n)| - 1$, an incomparable pair of support s contributes at least 2^{1-s} unless it is refuted outright. A large support does not by itself guarantee a small contribution: the relevant quantity is the number of newly linked components. When those components are mostly distinct, the contribution is small; when the comparability is already implied, it is one.

The Gibbs classifier can be derandomised. The *Bayes classifier* B_ρ predicts 1 on a pair n iff $\rho(C(n)) > \frac{1}{2}$, with prediction (0) at equality. This tie rule is essential below.

Corollary 8.3 (The Bayes classifier is the deductive closure) *Let P be realizable and ρ_P uniform on $\mathcal{V}(P)$. Then $B_{\rho_P}(n) = 1$ if and only if $P \cup \{n\}$ is realizable and $\mathrm{rk}_P(n) = 0$, that is, if and only if every direction consistent with P makes n comparable. Consequently, $B_{\rho_P} = g_P$ is computable in $O(|S(n)|)$ time per query after the preprocessing in Theorem 4.2, and $\mathrm{err}_{\mathcal{D}}(B_{\rho_P}) \le 2\,\mathrm{err}_{\mathcal{D}}(G_{\rho_P})$, so the bounds of Theorem 8.2(b) apply to B_{ρ_P} with an extra factor of two.*

Proof By Lemma 8.1, $\rho_P(C(n))$ is 0, 1, or $2^{-r} \le \frac{1}{2}$ with $r \ge 1$, so it exceeds $\frac{1}{2}$ iff it equals 1, iff $\mathcal{V}(P) \subseteq C(n)$. Whenever B_{ρ_P} errs on a labelled pair, the ρ_P-mass of hypotheses that err on it is at least $\frac{1}{2}$, so the Gibbs risk on that pair is at least $\frac{1}{2}$ the Bayes risk; integrate over $\mathcal{D}$ (Langford, 2005). ■

Thus the closure learner is also the majority vote of ρ_P when ties are resolved as incomparable. In a realizable model it has no false positives, and its false-negative region decreases as true positives are added, as used in Theorem 4.4. For arbitrary labelled distributions, a positive query can also be rejected because its comparability is inconsistent with P; no monotonic improvement of both error types is asserted in that setting.

1id1:core:20261120-0001@v1 1id1.net

Remark 8.4 (Chosen distribution versus a full posterior) *The distribution ρ_P is a valid data-dependent choice for PAC-Bayes, even when it is not the Bayesian posterior for the complete observation protocol. In particular, selecting positives from a mixed sample and discarding negatives need not correspond to conditioning on the full sampling and selection mechanism. Formula* (1) *needs no such Bayesian interpretation.*

If the full sample is realizable, the uniform distribution ρ_S on $\mathcal{V}(S)$ has zero empirical risk and

$$\mathrm{KL}(\rho_S \| \pi) = d \ln 2 - \ln |\mathcal{V}(S)| \;\geq\; (d - c(P)) \ln 2.$$

It pays a larger*, not a smaller, complexity penalty than ρ_P. The resulting risk bound may improve because its empirical risk is zero, but is not uniformly tighter. Exact evaluation is* #P*-hard for negative supports of size three (Corollary 5.6).*

9 The high-dimensional regime

The hardness results above exploit ties. We now turn to the opposite extreme—continuous data in general position—where the computational questions trivialise and the limiting factor is information.

Let $X^{(1)}, \dots, X^{(N)}$ be independent random instances in $\mathbb{R}^d$ whose coordinates are independent and, within each coordinate, identically distributed with a continuous law. Fix $\sigma \in \Sigma_d$ and let K_σ be the number of unordered pairs comparable under $\preceq_\sigma$. Ties have probability zero, so every pair has full support and is comparable under σ iff $\chi = \pm\sigma$, an event of probability 2^{1-d}. Hence $\mathbb{E}K_\sigma = \binom{N}{2}2^{1-d}$ and, by Markov's inequality, the sample is a σ-antichain with probability $\geq 1 - \binom{N}{2}2^{1-d}$.

Theorem 9.1 (Poisson limit) *Suppose $N \to \infty$, $d = d(N) \to \infty$, and $\lambda_N = \binom{N}{2}2^{1-d} \to \lambda \in [0, \infty)$. Then K_σ converges in distribution to* Poisson(λ)*, and the total-variation distance between the law of K_σ and* Poisson(λ_N) *is $O(N^3 3^{-d})$, uniformly in σ. In particular, if $d - 2\log_2 N \to t \in \mathbb{R}$, then $\mathbb{P}(K_\sigma = 0) \to \exp(-2^{-t})$.*

Proof For an unordered pair π let I_π indicate comparability under σ, $p_\pi = 2^{1-d}$, and let the dependence neighbourhood B_π consist of the pairs meeting π. Since I_π depends only on the two instances of π and the indicators outside B_π depend only on the other instances, the Chen–Stein term b_3 vanishes (Arratia et al., 1989). Next, $b_1 = \binom{N}{2}(2N-3)4^{1-d} = O(N^3 4^{-d})$. For b_2, fix pairs $\{u, v\}, \{u, w\}$ and let $s_i = \mathrm{sgn}(X_i^{(v)} - X_i^{(u)})$, $t_i = \mathrm{sgn}(X_i^{(w)} - X_i^{(u)})$. The six equiprobable rankings of three continuous values give $\mathbb{P}(s_i = a, t_i = b) = 1/3$ if $a = b$ and $1/6$ otherwise. Both pairs are comparable iff $s = \tau\sigma$, $t = \tau'\sigma$ for some signs τ, τ', so the joint probability is $2 \cdot 3^{-d} + 2 \cdot 6^{-d}$ and $b_2 = \binom{N}{2}\, 2(N-2)(2 \cdot 3^{-d} + 2 \cdot 6^{-d}) = O(N^3 3^{-d})$. With $N^2 2^{-d}$ bounded, $N^3 4^{-d} \to 0$ and $N^3 3^{-d} = (N^2 2^{-d})^{3/2}(2^{3/2}/3)^d \to 0$ since $2^{3/2} < 3$. The Chen–Stein theorem bounds the total-variation distance by $2(b_1 + b_2)$, and none of the quantities depends on σ. Finally $\lambda_N = (1 - N^{-1})2^{-(d - 2\log_2 N)} \to 2^{-t}$. ∎

The threshold is $d \approx 2\log_2 N$ for a *fixed* direction. The question of whether *some* direction makes the sample an antichain is a covering problem: in general position every pair hits exactly one of the 2^{d-1} projective classes, and $Z_N := \frac{1}{2}|\,\mathrm{Ant}(A_N)|$, where $\mathrm{Ant}(A) =$

1id1:core:20261120-0001@v1 1id1.net

$\{\sigma : A \text{ is a } \preceq_\sigma \text{-antichain}\}$ and $A_N = \{X^{(1)}, \dots, X^{(N)}\}$, counts the classes missed by all $\binom{N}{2}$ pairs. Heuristically the hits are nearly uniform, so one expects a coupon-collector threshold at $\binom{N}{2} \approx 2^{d-1} \ln 2^{d-1}$. The additive bound used above does not establish this: at that scale $\mathbb{P}(K_\sigma = 0) \approx 2^{1-d}$, far below the additive error $O(N^3 3^{-d})$. Suen's correlation inequality, in Janson's two-sided form (Suen, 1990; Janson, 1998), provides the required *multiplicative* control. Write $\lambda_N = \binom{N}{2} 2^{1-d}$, $c_N = \lambda_N - (d-1)\ln 2$, $\nu_N = 2^{d-1} e^{-\lambda_N} = e^{-c_N}$, and $\eta_N = N^3 3^{-d} + N 2^{-d} + \lambda_N 2^{-d}$.

Lemma 9.2 (Multiplicative Poisson estimate) *There are absolute constants $C, c_0 > 0$ such that for any k distinct projective classes $[\sigma_1], \dots, [\sigma_k]$ with $k 2^{1-d} \le \frac{1}{2}$ and $k^2 \eta_N \le c_0$,*

$$\mathbb{P}(K_{\sigma_1} = \cdots = K_{\sigma_k} = 0) = e^{-k\lambda_N}(1+\theta), \qquad |\theta| \le C k^2 \eta_N.$$

Proof *Outline; full details appear in the supplement.* Let I_π indicate that the pair π is comparable under some σ_j; these events are disjoint over j, so $p_\pi = k 2^{1-d}$, and pairs are dependent only when they share a point. For the sum $S = \sum_\pi I_\pi$, Janson's Theorems 2 and 8 (Janson, 1998) give

$$(1 - \Delta' e^{4\delta} \exp(\Delta e^{4\delta})) \prod_\pi (1 - p_\pi) \le \mathbb{P}(S = 0) \le e^{-\mu + \Delta e^{2\delta}},$$

where $\mu = k\lambda_N$, $\delta = \max_\pi \sum_{\pi' \sim \pi} p_{\pi'} \le 4kN2^{-d}$, and Δ, Δ' sum $\mathbb{E}(I_\pi I_{\pi'})$ and $p_\pi p_{\pi'}$ over unordered adjacent pairs. The ranking computation in the proof of Theorem 9.1 gives $\mathbb{E}(I_\pi I_{\pi'}) \le 4k^2 3^{-d}$, so $\Delta, \Delta' \le 2k^2 N^3 3^{-d}$; and $\prod_\pi (1-p_\pi) = e^{-k\lambda_N}(1 + O(k^2 \lambda_N 2^{-d}))$. All error terms are $O(k^2 \eta_N)$. Full details are in the supplement. ∎

Theorem 9.3 (Antichain-direction threshold) *Let $d \to \infty$ and $c_N = \lambda_N - (d-1)\ln 2$.*

(a) *If $c_N \to +\infty$, then $\mathrm{Ant}(A_N) = \emptyset$ with probability $\to 1$.*

(b) *If $c_N \to -\infty$, then $Z_N / \nu_N \to 1$ in probability with $\nu_N \to \infty$; in particular $\mathrm{Ant}(A_N) \neq \emptyset$ with probability $\to 1$.*

(c) *If $c_N \to c \in \mathbb{R}$, then $Z_N \to \mathrm{Poisson}(e^{-c})$ in distribution, so $\mathbb{P}(\mathrm{Ant}(A_N) = \emptyset) \to \exp(-e^{-c})$.*

The critical window is specified by $\binom{N}{2} 2^{1-d} = (d-1)\ln 2 + c + o(1)$. It implies $d = 2\log_2 N - \log_2 \ln N - 1 + o(1)$, but this coarser dimension relation alone does not determine a critical-window limit.

Proof In (b) and (c), $\lambda_N = O(d)$, hence $\eta_N \to 0$, and Lemma 9.2 applied to all ordered k-tuples of distinct classes gives $\mathbb{E}[(Z_N)_k] = (2^{d-1})_k e^{-k\lambda_N}(1 + O(k^2\eta_N)) = \nu_N^k (1 + o(1))$ for each fixed k, where $(m)_k$ is the falling factorial. In (c) $\nu_N \to e^{-c}$ and the method of moments yields the Poisson limit. In (b), $k = 1, 2$ give $\mathrm{Var}\, Z_N = \nu_N(1 + o(1)) + o(\nu_N^2)$ and Chebyshev's inequality gives concentration. For (a), $\mathrm{Ant}(A_N)$ is decreasing in N. Choose $\omega_N \to \infty$ with $\omega_N \le \min\{c_N/2, \sqrt{d}\}$, and take the least $N' \le N$ with $\lambda_{N'} \ge (d-1)\ln 2 + \omega_N$. Then $N' = O(\sqrt{d 2^d})$, consecutive intensities differ by $O(\sqrt{d}\, 2^{-d/2}) = o(1)$, and $c_{N'} = \omega_N + o(1)$. Thus $\eta_{N'} \to 0$ and $\mathbb{P}(Z_N \ge 1) \le \mathbb{E} Z_{N'} = e^{-\omega_N + o(1)} \to 0$. ∎

1id1:core:20261120-0001@v1 1id1.net

The existence of some antichain direction therefore has a sharp threshold $\log_2 \ln N + 1$ dimensions *below* the fixed-direction threshold, with a Gumbel-type profile $\exp(-e^{-c})$ in place of $\exp(-2^{-t})$.

Theorem 9.4 (Identification; negative feedback alone) *In the setting of Theorem 9.1 with hidden direction σ^*, suppose the learner receives the labels of all $\binom{N}{2}$ pairs. Let $\mathcal{I}_N$ be the event that the version space is $\{\pm\sigma^*\}$, and let $\mathcal{V}_N^-$ be the version space determined by the negative examples alone.*

(a) *If $\lambda_N \to \lambda \in [0,\infty]$, then $\mathbb{P}(\mathcal{I}_N) \to 1 - e^{-\lambda}$. In particular, if $d - 2\log_2 N \to t$, then $\mathbb{P}(\mathcal{I}_N) \to 1 - \exp(-2^{-t})$.*

(b) *Almost surely $\frac{1}{2}|\mathcal{V}_N^-| = 1 + Z_N - \mathbf{1}[K_{\sigma^*} = 0]$. Hence negative examples alone identify σ^* with probability $\to 1$ if $c_N \to +\infty$, with probability $\to \exp(-e^{-c})$ if $c_N \to c$, and with probability $\to 0$ if $c_N \to -\infty$; in the critical window $\frac{1}{2}|\mathcal{V}_N^-| - 1 \to \mathrm{Poisson}(e^{-c})$.*

Proof Almost surely every pair has full support, so a pair is negative iff its class differs from $[\sigma^*]$, and it then eliminates exactly that class; this gives (b)'s identity, and the limits follow from Theorem 9.3 because $\mathbb{P}(K_{\sigma^*} = 0) \to 0$ when c_N is bounded below (use the thinning argument below if λ_N grows too quickly for Lemma 9.2), while $\mathbf{1}[K_{\sigma^*} = 0] \le 1 \ll \nu_N$ when $c_N \to -\infty$. For (a), a positive pair identifies $\pm\sigma^*$ by Corollary 3.5, so $\mathcal{I}_N = \{K_{\sigma^*} \ge 1\} \sqcup \{K_{\sigma^*} = 0, Z_N = 1\}$. If $\lambda < \infty$, then $c_N \to -\infty$ and $\mathbb{P}(Z_N = 1) \le \mathbb{P}(Z_N \le \nu_N/2) \to 0$ by Theorem 9.3(b), so $\mathbb{P}(\mathcal{I}_N) = \mathbb{P}(K_{\sigma^*} \ge 1) + o(1) \to 1 - e^{-\lambda}$ by Theorem 9.1; if $\lambda = \infty$, fix $T > 0$ and take the least $N_T \le N$ with $\binom{N_T}{2} 2^{1-d} \ge T$. Its intensity tends to T, so the bounded-intensity theorem and monotonicity in N give $\limsup \mathbb{P}(K_{\sigma^*} = 0) \le e^{-T}$. Letting $T \to \infty$ proves the claim. ∎

Theorem 9.4 does not contradict the distribution-free bound of Corollary 6.2: under the product distribution the constant predictor "incomparable" already has error 2^{1-d}, so PAC learning to accuracy $\varepsilon > 2^{1-d}$ is trivial, and one positive pair is sufficient for identification. Without a positive pair, identification can still occur when exactly one class remains. The two kinds of feedback reveal the $d-1$ bits at very different rates: one positive example reveals all of them, whereas negative examples reveal them at the coupon-collector rate, so that identification from negatives alone needs $\binom{N}{2} \ge 2^{d-1}((d-1)\ln 2 + \omega(1))$ pairs, a factor $(d-1)\ln 2$ above the counting lower bound $2^{d-1} - 1$.

Finally, Z_N also measures how much monotonicity itself constrains a sample: under $\sigma \in \mathrm{Ant}(A)$, *every* labelling of A is σ-monotone.

Corollary 9.5 (Vacuous monotone bias) *Let $F_N = |\mathrm{Ant}(A_N)|/2^d$ be the fraction of directions under which the monotonicity assumption imposes no constraint on the sample. If $\lambda_N \to \lambda < \infty$ (in particular if $d - 2\log_2 N \to t$, with $\lambda = 2^{-t}$), then $F_N \to e^{-\lambda}$ in probability; if $c_N \to c$, then $2^{d-1}F_N \to \mathrm{Poisson}(e^{-c})$; and if $c_N \to +\infty$, then $F_N = 0$ with probability $\to 1$.*

Proof $2^{d-1}F_N = Z_N$; apply Theorem 9.3, noting that $\lambda_N \to \lambda < \infty$ forces $c_N \to -\infty$ and $\nu_N = 2^{d-1}e^{-\lambda_N}$. ∎

1id1:core:20261120-0001@v1 1id1.net

This is a statement about inequalities among the observed sample points. It does not imply that a global monotone function class has no restrictions outside that sample, or that monotonicity is useless for prediction under other data distributions.

Small-dimensional witnesses and a general upper bound. For A in general position, $|\operatorname{Ant}(A)| \geq 2^d - 2\binom{N}{2}$, so $\binom{N}{2} < 2^{d-1}$ guarantees an antichain direction, i.e. a direction realizing the all-negative labelling. Let $N_0(d) = \min\{N : \binom{N}{2} \geq 2^{d-1}\}$ and let $N^*(d)$ be the least N for which some N-point set in general position has $\operatorname{Ant}(A) = \emptyset$.

Proposition 9.6 (Bounds on the minimum covering sample) *$N_0(d) \leq N^*(d)$ for all d, with equality for $2 \leq d \leq 6$ ($N^* = 3, 4, 5, 7, 9$); $N^*(7) \in \{12, 13\}$, $N^*(8) \in \{17, 18\}$; and $N^*(d) \leq (1 + o(1))\sqrt{2^d(d-1)\ln 2}$, so $N^*(d)/N_0(d) \leq (1 + o(1))\sqrt{(d-1)\ln 2}$.*

Proof The first inequality is the union bound. $\operatorname{Ant}(A)$ depends only on the d coordinate rankings, so witnesses are d-tuples of permutations of $[N]$; explicit witnesses for $2 \leq d \leq 8$ are listed in the supplement, with a standalone verification program. Their verification is exhaustive over the pairs of each displayed witness; no search-optimality claim is needed. The asymptotic upper bound is Theorem 9.3(a) applied with $\omega_d \to \infty$ and $\omega_d = o(d)$: a random sample of $N = \lceil\sqrt{2^d((d-1)\ln 2 + \omega_d)}\rceil$ points has $\operatorname{Ant}(A_N) = \emptyset$ with positive probability. ■

The counting lower bound is attained for $d \leq 6$. In general the displayed upper and lower bounds still differ by a factor of order $\sqrt{d}$ in the number of instances. Thus we do not claim worst-case sharpness in all dimensions. Whether $N^*(d) = N_0(d)$ for every d remains open.

10 Conclusion and open problems

Orthant comparability separates prediction from parameter reconstruction: affine positive closure yields efficient improper learning, an exact $d - 1$ online mistake bound, compression without auxiliary bits, and an optimal-order distribution-free PAC bound, despite proper-learning hardness when negative supports reach three. It also supports explicit Gibbs-risk certificates. The dependent-pair analysis distinguishes fixed-accuracy prediction from exact identification and positive from negative feedback. These results do not cover noise or ordinary response labels that fail to reveal comparability.

Potential applications and their limits. A potential AI application is a checkable relational layer over ordinal attributes, not a replacement for a language model or training objective. Under realizability, every direction consistent with the positive evidence shares a positive closure prediction; a negative prediction only means comparability is not implied, not that true incomparability is certified. Compression preserves the version space and closure, not arbitrary dialogue; the bound counts examples, not storage bits. The noiseless analysis gives no robustness guarantee for erroneous AI-generated annotations.

Applications must justify unknown polarities and comparability feedback: known polarities make comparability directly computable, while one full-support positive identifies the projective direction. Structured comparisons with ties or controlled attribute changes

1id1:core:20261120-0001@v1 1id1.net

may therefore be more informative than dense continuous representations. Neither a suitable feedback protocol nor empirical advantages for language-model or agent systems are established here.

The main open problems are efficient improper *agnostic* learning; approximation of the all-or-nothing bundled constraints induced by larger positive supports; useful tractable distributions incorporating negative constraints for PAC-Bayes inference; random satisfiability thresholds with ties; and the minimum covering sample size in dimensions $d \geq 7$. Extending the model to genuinely unordered categorical coordinates or general partially ordered features requires a new analysis rather than an arbitrary numeric encoding.

Formal verification scope. The accompanying Lean development checks the mathematical core of Theorem 4.2(a),(c) and Theorem 4.3, but not their running-time claims or an efficient compression implementation. Its additional verification of Theorem 4.4 is restricted to finitely supported distributions. Section 1 of the supplement records the precise correspondence and exclusions.

Acknowledgments

AI-assisted tools were used for manuscript organisation, language editing, LaTeX preparation, and assistance with proof and code checks. The author is responsible for the final mathematical claims and submission.

References

Dana Angluin. Queries and concept learning. *Machine Learning*, 2(4):319–342, 1988. doi: 10.1007/BF00116828.

Richard Arratia, Larry Goldstein, and Louis Gordon. Two moments suffice for Poisson approximations: The Chen–Stein method. *The Annals of Probability*, 17(1):9–25, 1989. doi: 10.1214/aop/1176991491.

Richard E. Barlow, David J. Bartholomew, J. M. Bremner, and H. Daniel Brunk. *Statistical Inference under Order Restrictions*. Wiley, London, 1972.

Arie Ben-David. Monotonicity maintenance in information-theoretic machine learning algorithms. *Machine Learning*, 19(1):29–43, 1995. doi: 10.1007/BF00994659.

Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K. Warmuth. Learnability and the Vapnik–Chervonenkis dimension. *Journal of the ACM*, 36(4):929–965, 1989. doi: 10.1145/76359.76371.

Béla Bollobás and Graham Brightwell. Random high-dimensional orders. *Advances in Applied Probability*, 27(1):161–184, 1995. doi: 10.2307/1428102.

Stephan Börzsönyi, Donald Kossmann, and Konrad Stocker. The skyline operator. In *Proceedings of the 17th International Conference on Data Engineering*, pages 421–430. IEEE Computer Society, 2001. doi: 10.1109/ICDE.2001.914855.

Nadia Creignou and Miki Hermann. Complexity of generalized satisfiability counting problems. *Information and Computation*, 125(1):1–12, 1996. doi: 10.1006/inco.1996.0016.

1id1:core:20261120-0001@v1 1id1.net

Hennie Daniels and Marina Velikova. Monotone and partially monotone neural networks. *IEEE Transactions on Neural Networks*, 21(6):906–917, 2010. doi: 10.1109/TNN.2010.2044803.

Kalyanmoy Deb. *Multi-Objective Optimization Using Evolutionary Algorithms*. Wiley, Chichester, 2001.

Sally Floyd and Manfred K. Warmuth. Sample compression, learnability, and the Vapnik–Chervonenkis dimension. *Machine Learning*, 21(3):269–304, 1995. doi: 10.1007/BF00993593.

Johannes Fürnkranz and Eyke Hüllermeier, editors. *Preference Learning*. Springer, Berlin, Heidelberg, 2010. doi: 10.1007/978-3-642-14125-6.

Michael R. Garey, David S. Johnson, and Larry Stockmeyer. Some simplified NP-complete graph problems. *Theoretical Computer Science*, 1(3):237–267, 1976. doi: 10.1016/0304-3975(76)90059-1.

Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42(6):1115–1145, 1995. doi: 10.1145/227683.227684.

Sally A. Goldman and Michael J. Kearns. On the complexity of teaching. *Journal of Computer and System Sciences*, 50(1):20–31, 1995. doi: 10.1006/jcss.1995.1003.

Steve Hanneke. Refined error bounds for several learning algorithms. *Journal of Machine Learning Research*, 17(135):1–55, 2016. URL `https://www.jmlr.org/papers/v17/15-655.html`.

Frank Harary. On the notion of balance of a signed graph. *Michigan Mathematical Journal*, 2(2):143–146, 1953. doi: 10.1307/mmj/1028989917.

Frank Harary and Jerald A. Kabell. A simple algorithm to detect balance in signed graphs. *Mathematical Social Sciences*, 1(1):131–136, 1980. doi: 10.1016/0165-4896(80)90010-4.

Johan Håstad. Some optimal inapproximability results. *Journal of the ACM*, 48(4):798–859, 2001. doi: 10.1145/502090.502098.

Svante Janson. New versions of Suen's correlation inequality. *Random Structures & Algorithms*, 13(3–4):467–483, 1998.

Michael J. Kearns, Robert E. Schapire, and Linda M. Sellie. Toward efficient agnostic learning. *Machine Learning*, 17(2–3):115–141, 1994. doi: 10.1007/BF00993468.

John Langford. Tutorial on practical prediction theory for classification. *Journal of Machine Learning Research*, 6:273–306, 2005.

Nick Littlestone and Manfred K. Warmuth. Relating data compression and learnability. Unpublished manuscript, University of California, Santa Cruz, 1986.

Andreas Maurer. A note on the PAC Bayesian theorem. arXiv:cs/0411099, 2004.

David A. McAllester. Some PAC-Bayesian theorems. *Machine Learning*, 37(3):355–363, 1999. doi: 10.1023/A:1007618624809.

Subhojyoti Mukherjee, Anusha Lalitha, Sailik Sengupta, Aniket Deshmukh, and Branislav Kveton. Multi-objective alignment of large language models through hypervolume maximization, 2024. URL `https://arxiv.org/abs/2412.05469`. arXiv:2412.05469.

Rémi Munos, Michal Valko, Daniele Calandriello, Mohammad Gheshlaghi Azar, Mark Rowland, Zhaohan Daniel Guo, Yunhao Tang, Matthieu Geist, Thomas Mesnard, Andrea

1id1:core:20261120-0001@v1 1id1.net

Michi, Marco Selvi, Sertan Girgin, Nikola Momchev, Olivier Bachem, Daniel J. Mankowitz, Doina Precup, and Bilal Piot. Nash learning from human feedback, 2024. URL `https://arxiv.org/abs/2312.00886`. arXiv:2312.00886, version 4.

Wuming Pan. Fast generation of local Hasse graphs for learning from structurally connected instances. In *Rough Sets and Knowledge Technology*, volume 5009 of *Lecture Notes in Computer Science*, pages 548–555, Berlin, Heidelberg, 2008. Springer. doi: 10.1007/978-3-540-79721-0_74.

Leonard Pitt and Leslie G. Valiant. Computational limitations on learning from examples. *Journal of the ACM*, 35(4):965–984, 1988. doi: 10.1145/48014.63140.

Thomas J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the Tenth Annual ACM Symposium on Theory of Computing*, pages 216–226, New York, 1978. Association for Computing Machinery. doi: 10.1145/800133.804350.

Matthias Seeger. PAC-Bayesian generalisation error bounds for Gaussian process classification. *Journal of Machine Learning Research*, 3:233–269, 2002.

Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, New York, 2014.

Joseph Sill. Monotonic networks. In *Advances in Neural Information Processing Systems 10*, pages 661–667. MIT Press, 1998.

W.-C. Stephen Suen. A correlation inequality and a Poisson limit theorem for nonoverlapping balanced subgraphs of a random graph. *Random Structures & Algorithms*, 1(2):231–242, 1990.

Leslie G. Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984. doi: 10.1145/1968.1972.

Vladimir N. Vapnik and Alexey Ya. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and Its Applications*, 16 (2):264–280, 1971. doi: 10.1137/1116025.

Peter Winkler. Random orders. *Order*, 1(4):317–331, 1985. doi: 10.1007/BF00582738.

Thomas Zaslavsky. Signed graphs. *Discrete Applied Mathematics*, 4(1):47–74, 1982. doi: 10.1016/0166-218X(82)90033-6.

1id1:core:20261120-0001@v1 1id1.net