

# Efficient Algorithms for Learning Monophonic Halfspaces in Graphs

**Marco Bressan**

*Università degli Studi di Milano, Milan, Italy*

MARCO.BRESSAN@UNIMI.IT

**Emmanuel Esposito**

*Università degli Studi di Milano, Milan, Italy*

*Istituto Italiano di Tecnologia, Genoa, Italy*

EMMANUEL@EMMANUELESPOSITO.IT

**Maximilian Thiessen**

*TU Wien, Vienna, Austria*

MAXIMILIAN.THIESSEN@TUWIEN.AC.AT

**Editors:** Shipra Agrawal and Aaron Roth

## Abstract

We study the problem of learning a binary classifier on the vertices of a graph. In particular, we consider classifiers given by *monophonic halfspaces*, partitions of the vertices that are convex in a certain abstract sense. Monophonic halfspaces, and related notions such as geodesic halfspaces, have recently attracted interest, and several connections have been drawn between their properties (e.g., their VC dimension) and the structure of the underlying graph  $G$ . We prove several novel results for learning monophonic halfspaces in the supervised, online, and active settings. Our main result is that a monophonic halfspace can be learned with near-optimal passive sample complexity in time polynomial in  $n = |V(G)|$ . This requires us to devise a polynomial-time algorithm for consistent hypothesis checking, based on several structural insights on monophonic halfspaces and on a reduction to 2-satisfiability. We prove similar results for the online and active settings. We also show that the concept class can be enumerated with delay  $\text{poly}(n)$ , and that empirical risk minimization can be performed in time  $2^{\omega(G)} \text{poly}(n)$  where  $\omega(G)$  is the clique number of  $G$ . These results answer open questions from the literature (González et al., 2020), and show a contrast with geodesic halfspaces, for which some of the said problems are NP-hard (Seiffarth et al., 2023).

**Keywords:** graph convexity, online learning, active learning, polynomial time

## 1. Introduction

We study the problem of binary classification of the vertices of a graph. With the advent of social networks and the use of graph-based techniques in machine learning, this problem has received considerable attention in all the most common learning settings, including supervised learning (Hanneke, 2006; Pelckmans et al., 2007), active learning (Afshani et al., 2007; Guillory and Bilmes, 2009; Cesa-Bianchi et al., 2010; Dasarathy et al., 2015), and online learning (Herbster et al., 2005; Cesa-Bianchi et al., 2013; Herbster et al., 2015). Most results obtained so far rely on the homophily principle, that is, they assume that adjacent vertices tend to belong to the same class. We take a different perspective on this problem and assume that the ground truth classes are *monophonic halfspaces*, a notion related to linear separability and convexity in Euclidean spaces. In this way, we hope to exploit the intuition and technical machinery behind classes of convex concepts, which are often at the heart of machine learning models (e.g., think of intervals, halfspaces, or polytopes).

Let us spend a few words to introduce the notion of graph convexity. In  $\mathbb{R}^d$ , a set  $C$  is convex if for all  $x, y \in C$  the connecting segment  $I(x, y)$  lies in  $C$ . One can port this intuition to graphs in the following way. Given two vertices  $x$  and  $y$  in a graph  $G$ , one defines the *interval*  $I(x, y)$ . For

instance,  $I(x, y)$  can be the set of all nodes  $z \in V(G)$  that lie on some shortest path between  $x$  and  $y$ , that is, such that  $d(x, y) = d(x, z) + d(z, y)$  where  $d: V \times V \rightarrow \mathbb{N}$  is the shortest-path distance in  $G$ . In this way  $I(x, y)$  becomes the graph analogue of a segment in  $\mathbb{R}^d$ . This yields the well-known notion of *geodesic* convexity: a set  $C \subseteq V$  is geodesically convex if the induced subgraph  $G[C]$  is connected and  $I(x, y) \subseteq C$  for all  $x, y \in C$ , and is a geodesic halfspace if  $V \setminus C$  is geodesically convex too. Thus,  $C$  is the graph analogue of a halfspace in  $\mathbb{R}^d$ . Interestingly, several bounds that hold for halfspaces in  $\mathbb{R}^d$  continue to hold: for example, the VC dimension of geodesic halfspaces is bounded by their Radon number. One can therefore port several learning-theoretical results to the realm of graphs.

Starting from decade-old classic results, in recent years there has been a revived interest in the study of graph halfspaces and their learning properties (Duchet and Meyniel, 1983; Chepoi, 1986, 1994; Farber and Jamison, 1986; Pelayo, 2013; Thiessen and Gärtner, 2021; Bressan et al., 2021; Chalopin et al., 2022c,a; Seiffarth et al., 2023). Most works, however, give purely existential bounds on the invariants (the Radon number, the Helly number, the VC dimension, etc.) and study the connection with properties of the graph (maximum degree, clique number, etc.). Little is known, instead, about the existence of efficient algorithms for learning a halfspace in the most common supervised learning settings. Note that the existence of efficient algorithms is not obvious: already for the class of geodesic halfspaces described above, even just deciding if there exists a hypothesis consistent with a given labeled sample of vertices is NP-hard (Seiffarth et al., 2023). In this work we study efficient learning algorithms for another variant of abstract graph halfspaces, called *monophonic halfspaces*, which are defined through what is called minimal convexity or induced-path convexity (Farber and Jamison, 1986; Duchet, 1988).

Let us introduce some notation. Given two distinct vertices  $x, y \in V$ , the *monophonic interval*  $I(x, y)$  between  $x$  and  $y$  is the set of all vertices that lie on some induced path between  $x$  and  $y$ . A set  $C \subseteq V$  is *monophonically convex* (m-convex) if  $G[C]$  is connected and  $I(x, y) \subseteq C$  for all  $x, y \in C$ . A set  $H \subseteq V$  is a *monophonic halfspace* (m-halfspace) if both  $H$  and  $\overline{H} = V \setminus H$  are m-convex. For instance, if  $G$  is a tree, then the connected components left by deleting an edge are m-halfspaces; if  $G$  is a clique, then any subset is a m-halfspace, see Figure 1 for another example. In real-world networks, communities and clusters often tend to be geodesically convex, e.g., in gene similarity networks (Zhou et al., 2002), protein-protein interaction networks (Li et al., 2013), community detection benchmark datasets (Thiessen and Gärtner, 2021), and collaboration networks (Šubelj et al., 2019). In the latter and many other cases, the set of monophonic and geodesic convexity actually coincides (Malvestuto et al., 2012).

Monophonic halfspaces are among the most studied graph halfspaces (Bandelt, 1989; Changat et al., 2005; Dourado et al., 2010), second only to geodesic halfspaces. Nonetheless, only a couple of facts are known; for instance, a connection between the Radon number of m-convex sets and the clique number  $\omega(G)$ , or that computing a  $k$ -partition of  $V$  into m-convex subsets is hard for  $k \geq 3$  (González et al., 2020). The present work fills this gap by providing several concrete results.

## Our Contributions

For a graph  $G = (V, E)$ , let  $n = |V|$  and  $m = |E|$ , and let  $\mathcal{H}_{\text{MP}} = \mathcal{H}_{\text{MP}}(G)$  denote the concept class consisting of all m-halfspaces of  $G$ .

1. In the realizable PAC setting, we give an algorithm that with probability  $1 - \delta$  learns a monophonic halfspace with accuracy  $\varepsilon$  by using  $\mathcal{O}(\varepsilon^{-1} (\omega(G) \log(1/\varepsilon) + \log(1/\delta)))$  labeled samples

and  $\text{poly}(n)$  time. The heart of this result is a polynomial-time consistency checker, i.e., an algorithm that in time  $\text{poly}(n)$  finds a  $m$ -halfspace  $H \in \mathcal{H}_{\text{MP}}$  consistent with a given labeled subset of vertices, if one exists. This polytime consistency checker is nontrivial; it requires us to exploit heavily the structure of monophonic halfspaces, and to devise a polynomial-time reduction to 2-satisfiability by carefully constructing boolean formulas that express constraints on  $G$ .

2. In the agnostic PAC setting, we give an algorithm that with probability  $1 - \delta$  learns a monophonic halfspace with accuracy  $\varepsilon$  by using  $\mathcal{O}(\varepsilon^{-2}(\omega(G) + \log(1/\delta)))$  samples and  $|\mathcal{H}_{\text{MP}}| \text{poly}(n) \leq 2^{\omega(G)} \text{poly}(n)$  time. To this end we list  $\mathcal{H}_{\text{MP}}$  in time  $|\mathcal{H}_{\text{MP}}| \text{poly}(n)$  using the polynomial-time checker above; then, we prove that  $|\mathcal{H}_{\text{MP}}| \leq \frac{4m 2^{\omega(G)}}{\omega(G)} + 2$  by a careful analysis of a second listing algorithm tailored to the purpose.
3. In the realizable active setting, we give an algorithm that learns  $H \in \mathcal{H}_{\text{MP}}(G)$  in time  $\text{poly}(n)$  using  $\mathcal{O}(\omega(G) + \log \text{diam}(G) + h(G))$  queries, where  $\text{diam}(G)$  and  $h(G)$  are respectively the diameter and the “monophonic hull number” of  $G$ . We leave open the problem of obtaining a similar result for agnostic active learning.
4. In the realizable online setting, we give an algorithm that learns  $H \in \mathcal{H}_{\text{MP}}(G)$  in time  $\text{poly}(n)$  by making  $\mathcal{O}(\omega(G) \log n)$  mistakes. We obtain this algorithm by showing a decomposition of  $m$ -halfspaces into a small disjunction of simpler concepts, and applying the Winnow algorithm. We give a similar result for the agnostic online setting. Additionally we achieve an improved mistake bound  $\mathcal{O}(\omega(G) + \log(n/\omega(G)))$  using Halving and a runtime of  $2^{\omega(G)} \text{poly}(n)$ .
5. As a byproduct of the results above, we resolve the open case  $k = 2$  of the partitioning problem of [González et al. \(2020\)](#), which showed that for  $k \geq 3$  it is NP-hard to decide whether the vertex set  $V$  of a graph  $G$  admits a  $k$ -partition  $V_1, \dots, V_k$  where each  $V_i$  is monophonically convex. We also prove that  $|\mathcal{H}_{\text{MP}}| \leq 2^{\omega(G)} \text{poly}(n)$ , significantly improving on the only previously known and trivial bound  $|\mathcal{H}_{\text{MP}}| = \mathcal{O}(n^d)$ , where  $d = \mathcal{O}(\omega(G))$  is the VC dimension of  $\mathcal{H}_{\text{MP}}$ .

From a technical point of view, we rely on several technical results on the structure of  $m$ -halfspaces, including characterizations of their cutsets and efficient computation of so-called “shadows”. It should be noted that ours are among the few *constructive* results on the efficiency of learning abstract halfspaces (or convex concepts) in graphs. Our work belongs to a line of research on graph hypothesis spaces, their learning properties, and fixed-parameter learnability ([Chepoi et al., 2007, 2021](#); [Chalopin et al., 2022b](#); [Le and Wulff-Nilsen, 2024](#); [Brand et al., 2023](#)). Note that our runtime bounds for agnostic PAC learning and Halving are indeed polynomial for any family of graphs with bounded  $\omega(G)$ . This is for example the case for triangle-free graphs or planar graphs and a much weaker condition than more commonly used parameters such as treewidth ([Thiessen and Gärtner, 2021](#)) and maximum *clique-minor* size ([Duchet and Meyniel, 1983](#); [Chalopin et al., 2022b](#); [Le and Wulff-Nilsen, 2024](#)) to achieve polynomial runtime or bounds on the VC dimension.

**Organization of the manuscript.** Section 2 reviews related work. Sections 4-6 discuss our main results, their significance, and their technical ingredients. Section 7 presents structural properties of monophonic halfspaces that are needed by our algorithms. Section 8 presents our polynomial-time consistency checker, and Section 9 presents our empirical risk minimization (ERM) algorithm. All missing details and proofs can be found in the appendix.

## 2. Further Related Work

The complexity of consistency checking for monophonic halfspaces was unknown. [González et al. \(2020\)](#) show that deciding whether  $V$  can be partitioned into  $k \geq 3$  nonempty  $m$ -convex sets is NP-hard. They leave open the case  $k = 2$ , i.e., deciding if there is a nontrivial  $m$ -halfspace of  $G$ , or in other words whether  $\mathcal{H}_{\text{MP}}(G) \neq \emptyset$ . Our polynomial-time consistency checker proves that the case  $k = 2$  is in polynomial time: for every pair  $\{u, v\} \in \binom{V}{2}$  use the checker to verify whether there is a halfspace that separates  $u$  from  $v$ . This result was recently independently achieved by [Elaroussi et al. \(2024\)](#) and [Chepoi \(2024\)](#). This result should be contrasted with consistency checking in geodesic halfspaces, which is known to be NP-hard ([Seiffarth et al., 2023](#)). As monophonic halfspaces are a sub-family of geodesic halfspaces, our result can be seen as pushing the boundary of tractability of halfspace separation problems on graphs, see also [Chepoi \(2024\)](#) and [Elaroussi et al. \(2024\)](#).

Empirical risk minimization can be reduced to listing the concept class  $\mathcal{H}_{\text{MP}}(G)$ . [Duchet \(1988\)](#) observes that, if  $H$  is  $m$ -convex, then the vertices of  $\bar{H}$  that are adjacent to  $H$  form a clique. Therefore one can list  $\mathcal{H}_{\text{MP}}(G)$  by listing all pairs of cliques of  $G$  and checking if the edges between them form a cut of  $G$ , for a running time of  $n^{2\omega(G)} \text{poly}(n)$ . A better bound can be achieved if one is given a polynomial-time consistency checker: in that case, by a folklore algorithm one can list  $\mathcal{H}_{\text{MP}}(G)$  in time  $|\mathcal{H}_{\text{MP}}(G)| \text{poly}(n)$ . Our work gives both a polynomial-time consistency checker and a tight bound on  $|\mathcal{H}_{\text{MP}}(G)|$ ; neither one was known before. In particular, bounds on  $|\mathcal{H}_{\text{MP}}(G)|$  given by standard VC-dimension arguments suffer an exponential dependence on the cutsize  $c$  (i.e., the number of edges) of the halfspace ([Kleinberg, 2004](#)). In our case  $c$  can be as large as  $\Theta(\omega(G)^2)$ , which yields  $|\mathcal{H}_{\text{MP}}| \leq n^{\mathcal{O}(\omega(G)^2)}$ . This is significantly beaten by our novel bound  $|\mathcal{H}_{\text{MP}}| \leq \frac{4m 2^{\omega(G)}}{\omega(G)} + 2$ . [Glantz and Meyerhenke \(2017\)](#) give polynomial time algorithms for enumerating geodesic halfspaces of bipartite and planar graphs, but do not have results for general graphs. By contrast, we can enumerate all  $m$ -halfspaces in optimal time  $|\mathcal{H}_{\text{MP}}| \text{poly}(n)$  up to polynomial factors.

For active learning, [Thiessen and Gärtner \(2021\)](#) give lower and upper bounds, but for *geodesic* halfspaces. Their algorithm requires computing the geodesic hull number, which is APX-hard and without constant-approximation algorithms; our algorithm runs in polynomial time. [Bressan et al. \(2021\)](#) also studied active learning on graphs under a geodesic convexity assumption. They achieved polynomial time however with additional assumptions on the convex sets, such as margin. For online learning, [Thiessen and Gärtner \(2022\)](#) give again results for geodesic halfspaces. Their algorithms, however, are computationally inefficient and/or loose in terms of mistakes. We instead rewrite monophonic halfspaces as a union of a small number of  $m$ -shadows, which enables an algorithm (Winnow) that is both computationally efficient and near-optimal in terms of mistakes.

## 3. Preliminaries

Let  $G = (V, E)$  be a simple undirected graph and let  $n = |V|$  and  $m = |E|$ . Without loss of generality we assume  $G$  is connected.<sup>1</sup> For any  $v \in V$  let  $N(v) = \{u \in V \mid \{u, v\} \in E\}$ . The *cutset* induced by  $X \subseteq V$  is  $\delta(X) = \{\{u, v\} \in E \mid u \in X, v \notin X\}$ . For an edge  $\{u, v\} \in \delta(X)$  we may write  $uv \in \delta(X)$  to specify that  $u \in X$ . The *border* of  $X$  is  $\Gamma(X) = \{u \mid uv \in \delta(X)\}$ , the set of vertices of  $X$  with a neighbour in  $V \setminus X$ . We denote by  $\bar{G}$  or  $\neg G$  the complement  $(V, \binom{V}{2} \setminus E)$  of  $G$ . We let  $\omega(G)$  and  $\alpha(G)$  be respectively the clique number and the independence number of

---

1. If  $G$  is not connected then it has no  $m$ -halfspace, unless it consists of precisely two connected components, which are then the only two halfspaces.

$G$ , and  $\tilde{\omega}(G) = \max\{\omega(G), 3\}$ . For any  $X \subseteq V$  we denote by  $G[X]$  the subgraph of  $G$  induced by  $X$ . If  $P$  is a path in  $G$ , any edge in  $E(G) \setminus E(P)$  is called a *chord*; a path is induced if and only if it has no chords. Any shortest path is an induced path. We denote by  $\text{diam}_g(G)$  the diameter of  $G$ , and by  $\text{diam}_m(G)$  the maximum number of edges in any induced path. If  $a, b \in V$  and  $X \subseteq V$ , then  $X$  is an  $(a, b)$ -separator if in  $G$  every path between  $a$  and  $b$  intersects  $X$ . An algorithm is FPT (fixed-parameter tractable) with respect to some parameter  $k$  (in our case  $k = \omega(G)$ ) if its runtime is bounded by  $f(k) \text{poly}(n)$  for some computable function  $f$ .




Figure 1: A toy graph  $G$  whose vertex set  $V$  is partitioned into a monophonic halfspace  $H$  (in blue) and its complement  $\bar{H}$  (in orange). The diamond-shaped vertices form the borders  $\Gamma(H)$  and  $\Gamma(\bar{H})$ , see below. Both  $G[\Gamma(H)]$  and  $G[\Gamma(\bar{H})]$  are cliques, see Lemma 16.

Let  $V$  be a set and  $\mathcal{C} \subseteq 2^V$ . The set system  $(V, \mathcal{C})$  is a *convexity space* if (i)  $\emptyset, V \in \mathcal{C}$  and (ii)  $(\bigcap_{C \in \mathcal{C}'} C) \in \mathcal{C}$  for every  $\mathcal{C}' \subseteq \mathcal{C}$ .<sup>2</sup> A set  $C \subseteq V$  is said to be *convex* if  $C \in \mathcal{C}$ . Convexity spaces abstract standard Euclidean convexity (see, e.g., van de Vel (1993)). The *convex hull* of  $A \subseteq V$  is  $\text{conv}(A) = \bigcap_{C \in \mathcal{C}: A \subseteq C} C \in \mathcal{C}$ . The *hull number* of  $(V, \mathcal{C})$  is the size of the smallest  $A$  such that  $\text{conv}(A) = V$ . A set  $R \subseteq V$  is *Radon-independent* if  $\text{conv}(R') \cap \text{conv}(R \setminus R') = \emptyset$  for all  $R' \subseteq R$ . The *Radon number* of  $(V, \mathcal{C})$  is the size of its largest Radon-independent set. Many convexity spaces are defined by *intervals*. A map  $I: V \times V \rightarrow 2^V$  is an *interval map* if (i)  $u, v \in I(u, v)$  and (ii)  $I(u, v) = I(v, u)$  for all  $u, v \in V$ . An example is the geodesic interval of a metric space  $(V, d)$  defined by  $I_d(u, v) = \{z \in V \mid d(u, z) + d(z, v) = d(u, v)\}$ . Any interval map  $I$  defines a convexity space  $(V, \mathcal{C})$  where  $C \in \mathcal{C}$  if and only if  $I(u, v) \subseteq C$  for all  $u, v \in C$ . A *graph convexity space* for  $G = (V, E)$  is any convexity space  $(V, \mathcal{C})$  where  $G[C]$  is connected for all  $C \in \mathcal{C}$  (Pelayo, 2013). We denote by  $r(G)$  the Radon number and by  $h(G)$  the hull number of  $(V, \mathcal{C})$ .

For any distinct  $u, v \in V$  the *monophonic interval* (m-interval)  $I(u, v)$  between  $u$  and  $v$  is the set of all vertices that lie on some induced path between  $u$  and  $v$ ; while  $I(u, u) = \emptyset$  for all  $u \in V$ . A set  $X \subseteq V$  is *monophonically convex* (m-convex) if  $I(u, v) \subseteq X$  for all  $u, v \in X$ . A set  $X \subseteq V$  is a *monophonic halfspace* (m-halfspace), if both  $X$  and its complement  $\bar{X} = V \setminus X$  are m-convex. We denote by  $\mathcal{H}_{\text{MP}}(G)$  the set of all m-halfspaces of  $G$ . The *monophonic shadow* (m-shadow) of  $u$  with respect to  $v$  is the set  $u/v = \{z \in V : u \in I(z, v)\}$ . It is known that computing  $I$  is NP-hard in general (Dourado et al., 2010); as  $I(u, v) = \{z \in V : u \in z/v\}$ , computing m-shadows is NP-hard, too. Note that if  $X \subseteq V$  is m-convex and  $vu \in \delta(X)$  then  $u/v \subseteq \bar{X}$ . The *monophonic convex hull* (m-hull) of  $X \subseteq V$ , denoted  $\text{conv}(X)$ , is the smallest m-convex set  $X'$  such that  $X \subseteq X' \subseteq V$ . It is known that one can compute  $\text{conv}(X)$  in time  $\text{poly}(n)$  for any given  $X$  (Dourado et al., 2010). To discuss some tightness results of our achieved bounds we rely on so-called *separation axioms*, which determine the separability of vertices by halfspaces.

**Definition 1 (van de Vel (1993))** A convexity space  $(X, \mathcal{C})$  is:

2. When  $V$  is not finite one needs additional constraints, see van de Vel (1993); clearly this is not our case.

- $S_1$  if and only if each singleton  $x \in X$  is convex.
- $S_2$  if and only if each pair of distinct elements  $x, y \in X$  is halfspace separable.
- $S_3$  if and only if each convex set  $C \in \mathcal{C}$  and elements  $x \in X \setminus C$  are halfspace separable.
- $S_4$  if and only if any two disjoint convex sets are halfspace separable.

By slightly abusing notation, we write  $S_i$  for the family of graphs whose m-convexity space is  $S_i$  (for  $i = 1, \dots, 4$ ). All graphs are in  $S_1$ , and  $S_2 \supsetneq S_3 \supsetneq S_4$ . The graph families  $S_2$ ,  $S_3$ , and  $S_4$  are exactly characterized (Jamison-Waldner, 1981; Bandelt, 1989), see also Chepoi (1994, 2024).

#### 4. Supervised Learning

We study the standard PAC (probably approximately correct) setting with a given finite instance space  $V$  and a known hypothesis space  $\mathcal{H} \subseteq 2^V$ . In our case  $\mathcal{H} = \mathcal{H}_{\text{MP}}(G)$  is defined implicitly by the input graph  $G$ . A *labeled sample* is a pair  $(S, y)$  where  $S \subseteq V$  is a finite (multi)set of vertices and  $y: S \rightarrow \{0, 1\}$  is a binary labeling of the sample. We say that it is possible to  $(\varepsilon, \delta)$ -PAC learn  $(V, \mathcal{H}_{\text{MP}})$  if we can provide an m-halfspace  $H = H(S, y)$  that has error bounded by  $\varepsilon$  with probability at least  $1 - \delta$ , for any  $\varepsilon, \delta \in (0, 1)$ , under some unknown joint distribution over vertices and their labels. The aim of PAC learning is to guarantee this property with the smallest sample complexity possible. Note that, while standard PAC learning formulations typically measure the running time as a function of the sample size (see, e.g., Valiant (1984); Shalev-Shwartz and Ben-David (2014)), we measure running time as a function of  $n$ . Note that a dependence on  $n$  is unavoidable: even just to check whether two vertices  $u, v \in V$  are separated by some  $H \in \mathcal{H}_{\text{MP}}(G)$ , one needs to distinguish between  $G$  being a path or a cycle, which takes time  $\Omega(n)$ . We are in the *realizable setting* if for any given sample  $(S, y)$  there exists a  $H \in \mathcal{H}_{\text{MP}}$  such that for all  $s \in S$  it holds that  $y(s) = 1$  if and only if  $s \in H$ . Our main result for realizable PAC learning is:

**Theorem 2 (Realizable PAC is poly-time)** *There exists an algorithm that, for any  $\varepsilon, \delta \in (0, 1)$ , can  $(\varepsilon, \delta)$ -PAC learn  $(V, \mathcal{H}_{\text{MP}})$  in the realizable setting using  $\mathcal{O}\left(\frac{\omega(G)\log(1/\varepsilon)+\log(1/\delta)}{\varepsilon}\right)$  labeled samples and time  $\text{poly}(n)$ .*

The bound on the sample complexity of our algorithm follows by a bound on the VC dimension of  $\mathcal{H}_{\text{MP}}$  in terms of  $\omega(G)$  (Proposition 4), coupled with standard PAC learning bounds (Blumer et al., 1989). The bound is near-optimal save possibly for the  $\log(1/\varepsilon)$  factor (see below). The bound on the running time is instead given by Theorem 5 (see below) which provides a polynomial-time algorithm for computing an m-halfspace consistent with a given realizable sample. Similarly to the realizable case, we obtain the following result for agnostic PAC learning:

**Theorem 3 (Agnostic PAC is FPT)** *There exists an algorithm that, for any  $\varepsilon, \delta \in (0, 1)$ , can  $(\varepsilon, \delta)$ -PAC learn  $(V, \mathcal{H}_{\text{MP}})$  in the agnostic setting using  $\mathcal{O}\left(\frac{\omega(G)+\log(1/\delta)}{\varepsilon^2}\right)$  labeled samples and time  $|\mathcal{H}_{\text{MP}}| \text{poly}(n) \leq 2^{\omega(G)} \text{poly}(n)$ .*

Similarly to the realizable case, the sample complexity bound follows by coupling VC dimension bounds with standard PAC bounds, and is near-optimal. The result follows by using the ERM from Corollary 10, which uses the version space enumeration of Theorem 8 (see below).

To prove that the sample complexity bounds of Theorem 2 and Theorem 3 are near-optimal, it suffices to note that the VC dimension of  $\mathcal{H}_{\text{MP}}$  is  $\omega(G)$  for clique graphs. We prove that this is true for all  $S_4$  graphs (see Section 3), and essentially optimal. Let  $\tilde{\omega}(G) = \max\{\omega(G), 3\}$ . Then:

**Proposition 4** *If  $G \in S_4$  then  $\text{VC}(\mathcal{H}_{\text{MP}}(G)) \geq \omega(G)$ . Moreover,  $\text{VC}(\mathcal{H}_{\text{MP}}(G)) \leq \tilde{\omega}(G)$  for all  $G$ .*

#### 4.1. Polynomial-Time Consistency Checking and Version Space Enumeration

The algorithms of Theorem 2 and Theorem 3 rely on the standard empirical risk minimization (ERM) approach: given a labeled sample  $(S, y)$ , we find a hypothesis  $H \in \mathcal{H}_{\text{MP}}$  that has minimal disagreement with  $(S, y)$ . This, we need efficient algorithms for ERM. In the realizable case, where  $(S, y)$  is realizable, we need an algorithm to find a hypothesis consistent with the sample. Formally, an  $m$ -halfspace  $H \in \mathcal{H}_{\text{MP}}$  is *consistent* with a given labeled sample  $(S, y)$  if  $y^{-1}(0) \subseteq \bar{H}$  and  $y^{-1}(1) \subseteq H$ . Our main technical contribution is MH-CHECK, an algorithm that gives:

**Theorem 5 (Consistency check is in polynomial time)** *Let  $G = (V, E)$  and  $(S, y)$  a labeled sample. Then,  $\text{MH-CHECK}(G, (S, y))$  (Algorithm 1) runs in time  $\text{poly}(n)$  and returns an  $H \in \mathcal{H}_{\text{MP}}(G)$  consistent with  $(S, y)$ , or reports that no such  $H$  exists.*

We note that Elaroussi et al. (2024) and Chepoi (2024) independently obtained an analogous result recently using different techniques.

The intuition behind MH-CHECK is to guess an edge in the cut of  $H$ , and then construct a set of constraints that reduces the search space to precisely those subsets  $H \subseteq V$  that are  $m$ -halfspaces and that are consistent with  $(S, y)$ . It turns out that, by carefully choosing those constraints, one actually obtains in time  $\text{poly}(n)$  an instance of 2-satisfiability, which in turn is solvable in time  $\text{poly}(n)$ . Choosing the constraints requires defining several specific subsets of vertices, as well as proving several structural results on  $m$ -halfspaces and their cutsets. Another computational ingredient is showing that computing  $m$ -shadows over edges is easy, although hard in general.

**Lemma 6** *Let  $\{z, v\} \in E$ . Then,  $z/v = \{x \in V \mid N(v) \setminus \{z\} \text{ is not a } (z, x)\text{-separator in } G\}$ . As a consequence one can compute  $z/v$  in time  $\text{poly}(n)$ .*

As an immediate byproduct of Theorem 5 we answer an open problem of González et al. (2020), who proved that deciding if  $V$  admits a proper  $k$ -partition into  $m$ -convex sets is NP-hard for  $k \geq 3$ .

**Corollary 7** *In time  $\text{poly}(n)$  one can decide if  $V$  admits a proper 2-partition into  $m$ -convex sets in  $G$  (that is, if there exists  $H \in \mathcal{H}_{\text{MP}}$  such that  $\emptyset \subsetneq H \subsetneq V$ ).*

A detailed description of MH-CHECK and a sketch of the proof of Theorem 5 are given in Section 8. Using MH-CHECK, we obtain an algorithm for ERM. For a given labeled sample  $(S, y)$ , the *version space*  $\text{VS}((S, y), \mathcal{H}_{\text{MP}})$  consists of all  $H \in \mathcal{H}_{\text{MP}}$  consistent with  $(S, y)$ . Thus, one can find an empirical risk minimizer for  $(S, y)$  by just listing  $\text{VS} = \text{VS}((S, y), \mathcal{H}_{\text{MP}})$  and picking a  $m$ -halfspace with minimum disagreement. We say that an algorithm lists  $\text{VS}$  with *delay*  $t$  if it spends time at most  $t$  to output the next element of  $\text{VS}$ . A listing algorithm is considered to be efficient if it has polynomial delay, that is,  $t = \text{poly}(n)$ ; this also implies that  $\text{VS}$  is listed in total time  $|\text{VS}| \text{poly}(n)$ . Using MH-CHECK and a folklore technique, we prove that such an algorithm exists.

**Theorem 8 (Listing the version space with polynomial delay)** *There exists an algorithm MH-LIST that, given a graph  $G$  and a labeled sample  $(S, y)$ , lists  $\text{VS} = \text{VS}(S, \mathcal{H}_{\text{MP}})$  with delay  $\text{poly}(n)$  and total time  $|\text{VS}| \text{poly}(n)$ .*

As a last ingredient for our ERM bounds, we bound the size of  $\mathcal{H}_{\text{MP}}(G)$ . To this end we develop another listing algorithm, Algorithm 3, which by a careful analysis yields the following upper bound. Further details on how we accomplish this are deferred to Section 9.

**Theorem 9**  $|\mathcal{H}_{\text{MP}}(G)| \leq \frac{4m2^{\omega(G)}}{\omega(G)} + 2$  for all graphs  $G$ .

Together this yields the fact that ERM is FPT for  $m$ -halfspaces.

**Corollary 10 (ERM is FPT)** *Given  $G = (V, E)$  and a labeled sample  $(S, y)$ , one can compute  $H_{\text{ERM}} \in \arg \min_{H \in \mathcal{H}_{\text{MP}}} \frac{1}{|S|} \sum_{s \in S} \mathbb{1}\{y(s) \neq \mathbb{1}_{s \in H}\}$  in time  $|\mathcal{H}_{\text{MP}}| \text{poly}(n) \leq 2^{\omega(G)} \text{poly}(n)$ .*

## 5. Active Learning

In the active learning setting, the algorithm is given a graph  $G = (V, E)$ , and nature selects a concept  $H \in \mathcal{H}_{\text{MP}}(G)$ . The algorithm can query any vertex  $x \in V$  for its label  $\mathbb{1}_{x \in H}$ . The goal of the algorithm is to output  $H$  by making as few queries as possible. This problem is a special case of realizable transductive active learning on a graph (Afshani et al., 2007; Guillory and Bilmes, 2009; Cesa-Bianchi et al., 2010; Dasarathy et al., 2015), and can be seen as a variant of query learning (Angluin, 1988; Hegedűs, 1995) with an additional fixed graph given where vertices correspond to the instance space. Now let  $\mathcal{H} \subseteq 2^V$  be a concept class. The *query complexity* of  $(V, \mathcal{H})$  is the maximum number of queries an optimal algorithm makes over  $H \in \mathcal{H}$ . More precisely, for any algorithm  $A$  and any  $H \in \mathcal{H}$ , let  $\text{qc}(A, H)$  denote the number of queries  $A$  make on  $G$  when nature chooses  $H$ . The query complexity of  $A$  on  $\mathcal{H}$  is  $\text{qc}(A, \mathcal{H}) = \max_{H \in \mathcal{H}} \text{qc}(A, H)$ . The query complexity of  $\mathcal{H}$  is  $\text{qc}(\mathcal{H}) = \min_A \text{qc}(A, \mathcal{H})$ . Our main result is stated as follows.

**Theorem 11 (Poly-time active learning)** *It is possible to actively learn  $(V, \mathcal{H}_{\text{MP}})$  in time  $\text{poly}(n)$  with query complexity  $\mathcal{O}(h(G) + \log \text{diam}_g(G) + \omega(G))$ .*

The full proof of Theorem 11 is given in Appendix E; here we provide some intuition. First, we compute a minimum monophonic hull set in polynomial time relying on Dourado et al. (2010). If any cut-edge  $\{u, v\}$  exists we can then find it with  $\mathcal{O}(\log \text{diam}_g(V))$  queries. Using some structural results (Lemma 17 and Lemma 19, see Section 7) we can infer the remaining halfspace by querying at most  $\omega(G)$  from the set  $\Delta_{uv} = N(u) \cup N(v) \cup \{u, v\}$ . The required  $m$ -shadows can be computed efficiently by Lemma 6. Theorem 11 should be contrasted with the active learning algorithm provided by Thiessen and Gärtner (2021) for *geodesic* halfspaces; that algorithm does not guarantee polynomial running time, as it requires solving the minimum geodesic hull set problem, which is APX-hard.

Along the previously mentioned separation axioms from Definition 1 we achieve increasingly tighter lower bounds on the query complexity, eventually matching our algorithmic upper bound from Theorem 11 for all  $S_4$  graphs, the strongest separability assumption.

**Theorem 12** *Let  $G$  be a graph. The following holds for the query complexity  $\text{qc}(\mathcal{H}_{\text{MP}}(G))$ :*

- if  $G \in S_2$ , then  $\text{qc}(\mathcal{H}_{\text{MP}}(G)) \geq \Omega(\log \text{diam}_m(G))$ ,
- if  $G \in S_3$ , then  $\text{qc}(\mathcal{H}_{\text{MP}}(G)) \geq \Omega(\log \text{diam}_m(G) + h(G))$ , and
- if  $G \in S_4$ , then  $\text{qc}(\mathcal{H}_{\text{MP}}(G)) \geq \Omega(\log \text{diam}_m(G) + h(G) + \omega(G))$ .

## 6. Online Learning

The classical (realizable) online learning problem of Littlestone (1988) can be modelled as an iterative game between a learner and the environment over a finite number  $T$  of rounds. The instance



space  $V$  and a hypothesis space  $\mathcal{H} \subseteq 2^V$  is known and fixed. First, the environment chooses a hypothesis  $H$  from a hypothesis space  $\mathcal{H}$ . Then, in each round  $t = 1, \dots, T$ :

1. the environment chooses a point  $v_t \in V$ ,
2. the learner predicts the label  $\hat{y}_t \in \{0, 1\}$  of  $v_t$ ,
3. the environment reveals the true label  $y_t = \mathbb{1}_{v_t \in H}$ ,
4. the learner made a mistake if  $\hat{y}_t \neq y_t$ .

The goal of the learner is to minimize the total number of mistakes. More precisely let  $A$  be an algorithm for this online learning problem. Then, let  $M(A, H)$  for  $H \in \mathcal{H}$  denote the worst-case number of mistakes  $A$  would make on any sequence labeled by  $H$  over  $T$  rounds. The mistake bound of  $A$  on  $\mathcal{H}$  is thus defined as  $M(A, \mathcal{H}) = \max_{H \in \mathcal{H}} M(A, H)$ . We are interested in the optimal mistake bound  $M(\mathcal{H}) = \min_A M(A, \mathcal{H})$ , also known as the Littlestone dimension of  $\mathcal{H}$ . This setting can be extended to the agnostic/non-realizable case, see [Ben-David et al. \(2009\)](#).

The node classification (or node labeling) variant of this problem is well studied ([Herbster et al., 2005](#); [Cesa-Bianchi et al., 2013](#); [Herbster et al., 2015](#)). As in the active learning variant, the main parameter is the (potentially effective resistance weighted) cutsizes, linearly determining the mistake bounds. In this section, we mainly study the variant of the above realizable online learning problem over our hypothesis class  $\mathcal{H}_{\text{MP}}$ .

**Theorem 13 (Poly-time online learning)** *Realizable online learning of monophonic halfspaces is possible in time  $\text{poly}(n)$  with a mistake bound of  $\mathcal{O}(\omega(G) \log n)$ . It is also possible in time  $|\mathcal{H}_{\text{MP}}| \text{poly}(n) \leq 2^{\omega(G)} \text{poly}(n)$  with a mistake bound of  $\mathcal{O}\left(\omega(G) + \log \frac{n}{w(G)}\right)$ .*

The first mistake bound is achieved by the Winnow algorithm ([Littlestone, 1988](#)) and relies on a novel representation of  $m$ -halfspaces as a sparse disjunction of  $m$ -shadows (Lemma 18). The second mistake bound is achieved by the Halving algorithm ([Barzdin, 1972](#); [Littlestone, 1988](#)), together with our version space listing algorithm (Theorem 8) and our upper bound on  $|\mathcal{H}_{\text{MP}}|$  (Theorem 9). In Appendix F.3 we additionally discuss agnostic online learning of  $m$ -halfspaces relying again on our decomposition by Lemma 18 and known results for Winnow ([Blum, 1996](#)). We achieve the following mistake bound efficiently almost matching standard expert-based techniques ([Cesa-Bianchi et al., 1997](#); [Ben-David et al., 2009](#)), which typically require  $\mathcal{O}(|\mathcal{H}_{\text{MP}}|)$  time per round.

**Theorem 14** *A mistake bound  $\mathcal{O}(\omega(G)(M^* + \log n))$  for agnostic online learning over  $\mathcal{H}_{\text{MP}}(G)$  is possible in time  $\text{poly}(n)$ , where  $M^*$  is the minimum number of mistakes of any  $m$ -halfspace.*

**Lower bounds.** The achieved mistake bounds are near-optimal by the following result, whose proof is in Appendix F.2.

**Proposition 15** *For any  $S_4$  graph  $G$ , it holds that  $M(\mathcal{H}_{\text{MP}}(G)) \geq \omega(G)$ .*

## 7. Structural Lemmas regarding Monophonic Halfspaces

This section provides some structural results on monophonic halfspaces used by our algorithms. Loosely speaking, the main message is that monophonic halfspaces can be expressed as unions of a small number of “simpler” sets, whose labels can be inferred from a small number of labeled vertices. First, however, we give a basic result about the border of the cut.

**Lemma 16**  $H \in \mathcal{H}_{\text{MP}}(G)$  if and only if  $G[\Gamma(H)]$  and  $G[\Gamma(\overline{H})]$  are cliques.

The statement follows by Duchet (1988, Proposition 4.2) applied to both  $H$  and  $\overline{H}$ ; a self-contained proof can be found in Appendix A. Figure 1 shows an example. For an edge  $uv \in E$  let  $\Delta_{uv} = (N(u) \cap N(v)) \cup \{u, v\}$ . Our first result shows that every halfspace  $H \in \mathcal{H}_{\text{MP}}$  can be expressed as the union of m-shadows over  $H \cap \Delta_{uv}$ , for every  $uv \in \delta(H)$ .

**Lemma 17** Let  $H \in \mathcal{H}_{\text{MP}}(G)$ , let  $uv \in \delta(H)$ , and let  $\Delta_{uv} = (N(u) \cap N(v)) \cup \{u, v\}$ . Then,  $H$  consists precisely of those vertices whose label can be inferred from  $u, v$ , and a vertex in  $H \cap \Delta_{uv}$ . Formally:

$$H = \bigcup_{z \in H \cap \Delta_{uv}} z/v .$$

Note that Lemma 17 does not yet give a *sparse* representation, in the sense that it does not bound the number of m-shadows  $|H \cap \Delta_{uv}|$  used to obtain  $H$ . Combining Lemma 17 with Lemma 16, however, we can prove that  $H$  can be expressed as the union of at most  $\omega(G)$  such m-shadows.

**Lemma 18** Let  $H \in \mathcal{H}_{\text{MP}}(G)$ . Then, there exists a subset  $C \subseteq \delta(H)$  with  $|C| \leq \omega(G)$  such that

$$H = \bigcup_{\{z,v\} \in C} z/v .$$

Lemma 18 will be crucial to enable efficient online learning using Winnow. We shall also note the following fact, used by our bounds on  $|\mathcal{H}_{\text{MP}}|$  and by our polynomial-time active learning algorithm.

**Lemma 19** Let  $H \in \mathcal{H}_{\text{MP}}$  and  $uv$  be a cut-edge of  $H$ . Then, the number  $N_{uv}$  of connected components in  $\neg G[\Delta_{uv}]$  satisfies  $N_{uv} \leq \omega(G)$ .

## 8. A Polynomial-Time Consistency Checker

This section describes our main technical contribution: MH-CHECK (Algorithm 1), an algorithm to find an m-halfspace consistent with a given a labeled sample if one exists. We show that such an algorithm runs in polynomial time; this proves Theorem 5, which we restate for convenience.

**Theorem 5 (Consistency check is in polynomial time)** Let  $G = (V, E)$  and  $(S, y)$  a labeled sample. Then, MH-CHECK( $G, (S, y)$ ) (Algorithm 1) runs in time  $\text{poly}(n)$  and returns an  $H \in \mathcal{H}_{\text{MP}}(G)$  consistent with  $(S, y)$ , or reports that no such  $H$  exists.

Let us illustrate the idea behind MH-CHECK. Let  $(S, y)$  be the labeled sample, and suppose a halfspace  $H$  consistent with  $(S, y)$  exists. For every  $uv \in E$ , we guess that  $uv$  is the cut  $\delta(H)$ , and we compute the set  $\Delta_{uv} = (N(u) \cap N(v)) \cup \{u, v\}$ . By exploiting the results of Section 7, we know that from  $\Delta_{uv} \cap H$  one can infer the rest of  $H$ . Unfortunately, there seems to be no easy way of learning  $\Delta_{uv} \cap H$ , apart from exhaustively guessing all subsets of  $\Delta_{uv}$ . Thus, we need a more sophisticated strategy. We use the results of Section 7, as well as the closure properties of monophonic halfspaces, to construct a set of constraints on the space of all possible subsets  $H \subseteq V$ . The hope is that those constraints identify precisely the subsets  $H \subseteq V$  that satisfy  $H \in \mathcal{H}_{\text{MP}}(G)$  and are consistent with  $(S, y)$ .

Formally, we will perform a reduction to 2-SAT. For every orientation  $uv$  of every edge of  $G$ , we construct in time  $\text{poly}(|V|)$  a 2-SAT formula  $\varphi_{uv}$  whose literals are in the form “ $x \in H$ ” or “ $x \notin H$ ”, where  $x \in V$ . Any  $H \subseteq V$  yields an evaluation of  $\varphi_{uv}$  in the obvious way. We can prove that the solutions to  $\varphi_{uv}$  are precisely the  $m$ -halfspaces  $H \in \mathcal{H}_{\text{MP}}(G)$  such that  $uv \in \delta(H)$ . Constructing the 2-SAT formula  $\varphi_{uv}$  in polynomial time is the challenging part of the algorithm. Once we have computed  $\varphi_{uv}$ , it is straightforward to extend it to a 2-SAT formula  $\varphi_{uv}(S, y)$  that is satisfied precisely by the solutions to  $\varphi_{uv}$  that are consistent with  $(S, y)$ :

$$\varphi_{uv}(S, y) = \varphi_{uv} \wedge \bigwedge_{x \in y^{-1}(0)} (x \notin H) \wedge \bigwedge_{x \in y^{-1}(1)} (x \in H) . \quad (1)$$

Algorithm 1 gives the pseudocode of MH-CHECK. The rest of this section describes the construction of  $\varphi_{uv}$  and sketches the proof of Theorem 5. All missing details can be found in Appendix B.

---

**Algorithm 1:** MH-CHECK
 

---

**Input:**  $G = (V, E)$ , and a labeled sample  $(S, y)$

**Output:**  $H \in \mathcal{H}_{\text{MP}}(G)$  consistent with  $(S, y)$ , or failure if no such  $H$  exists

**for**  $uv \in E$  **do**

compute the constraints given in Appendix B, and their conjunction  $\varphi_{uv}$   
 compute  $\varphi_{uv}(S, y)$  as given by Equation (1)  
 compute a solution  $H$  to  $\varphi_{uv}(S, y)$  and return it, if one exists

**end**

**return** “ $H$  does not exist”

---

**Constructing  $\varphi_{uv}$ .** To construct  $\varphi_{uv}$ , we start by computing the following subsets:

- $\Delta_{uv}^- = N(u) \cap N(v)$ . Note that  $u, v \notin \Delta_{uv}^-$ .
- $\square_{uv} = \{x \in V : x \text{ appears in a 4-cycle having } uv \text{ as an edge}\}$ . Note that  $u, v \in \square_{uv}$ .
- $A = \square_{uv} \setminus \Delta_{uv}^- = \{x \in \square_{uv} : d(x, u) \neq d(x, v)\}$ .
- $A^u = \{x \in A : d(x, u) < d(x, v)\}$ .
- $A^v = \{x \in A : d(x, u) > d(x, v)\} = A \setminus A^u$ .
- $T = G \setminus E(G[\square_{uv} \cup \Delta_{uv}^-])$ .

Note that all those sets are computable in time  $\text{poly}(|V|)$ .

**A set of constraints.** As said,  $\varphi_{uv}$  is a conjunction of several constraints (i.e., 2-SAT formulas). Each constraint is based on some of the sets defined above. We describe some of those constraints to convey the idea. The complete list is given in Appendix B, where the full proof of Theorem 5 checks that all of them are satisfied if and only if  $H \in \mathcal{H}_{\text{MP}}(G)$ .

The first constraint we present here, given by the formula  $\varphi_{\text{conv}}$ , ensures that  $A^u$  and  $A^v$  are closed under monophonic convex hulls:

$$\varphi_{\text{conv}} = \bigwedge_{x \in \text{conv}(A^u)} (x \in H) \wedge \bigwedge_{x \in \text{conv}(A^v)} (x \notin H) . \quad (2)$$

Note that, as computing monophonic convex hulls in an  $n$ -vertex graph takes time  $\text{poly}(n)$  (Dourado et al., 2010), we can then compute  $\varphi_{\text{conv}}$  in time  $\text{poly}(n)$ , too.

The next constraint,  $\varphi_{\Delta_{uv}^-}$ , ensures that  $G[\Delta_{uv}^- \cap H] \subseteq G[\Gamma(H)]$  and  $G[\Delta_{uv}^- \cap \overline{H}] \subseteq G[\Gamma(\overline{H})]$  are both cliques, as required by Lemma 16. In fact it ensures a slightly stronger condition: that the edge complement of  $G[\Delta_{uv}^-]$  is a bipartite graph where adjacent vertices have opposite labels under  $H$ , as imposed by Lemma 22:

$$\varphi_{\Delta_{uv}^-} = \bigwedge_{\{x,y\} \in E(\neg G[\Delta_{uv}^-])} (x \in H \vee y \in H) \wedge (x \notin H \vee y \notin H) . \quad (3)$$

Denote by  $\text{cc}(T)$  the set of connected components of  $T$ . It is not hard to prove that every  $T_i \in \text{cc}(T)$  is either the  $m$ -shadow  $z/v$  or in the  $m$ -shadow  $z/u$  for some  $z \in \Delta_{uv}$ . Therefore, either  $T_i \subseteq H$  or  $T_i \subseteq \overline{H}$ . This is captured by the next constraint:

$$\varphi_T = \bigwedge_{T_i \in \text{cc}(T)} \bigwedge_{x,y \in V(T_i)} (x \in H \vee y \notin H) \wedge (x \notin H \vee y \in H) . \quad (4)$$

Finally, we compute constraints ensuring that there is no induced path starting from  $A^u$  or  $\text{conv}(A_u)$  that violates monophonic convexity. To compute them, we first need to define what follows: for  $X, Y \subseteq V$  and  $k \in \{3, 4\}$  denote by  $\Pi_k(X, Y)$  the set of all induced paths on  $k$  vertices that connect  $X$  to  $Y$ . We report only one of such constraints them as an example:

$$\varphi_{u,3} = \bigwedge_{\substack{\pi=(x,y,z) \\ \pi \in \Pi_3(A^u, \Delta_{uv}^-)}} (y \in H \vee z \notin H) . \quad (5)$$

It is easy to see that  $\varphi_{\Delta_{uv}^-}$ ,  $\varphi_T$ , and  $\varphi_{u,3}$  can be computed in time  $\text{poly}(n)$ , too.

## 9. Bounding the Number of Monophonic Halfspaces

This section describes an algorithm to enumerate all  $m$ -halfspaces in FPT time and, through an additional charging argument, a near-tight bound on  $|\mathcal{H}_{\text{MP}}(G)|$ . The first intuition behind our algorithm comes from Lemma 16, as follows. First, we guess each edge  $uv$  of  $G$  as a cut edge of some  $H \in \mathcal{H}_{\text{MP}}$ . Starting from the endpoints  $u, v$ , we then try to find the vertices in  $U = \Gamma(H) \cup \Gamma(\overline{H})$ . Since by Lemma 16 both  $G[\Gamma(H)]$  and  $G[\Gamma(\overline{H})]$  are cliques, then  $|U| \leq 2\omega(G)$ . At that point we exhaustively guess the partition of  $U$  into  $\Gamma(H)$  and  $\Gamma(\overline{H})$ , and we check if that partition is indeed an  $m$ -halfspace. This would result in a total running time of  $2^{2\omega(G)} \text{poly}(n)$  per edge and, therefore, in total.

Unfortunately, this idea does not work straight away. The reason is that there is no efficient way to find all of  $U$ . However, by Lemma 17, we can show that this approach still works if instead of  $U$  we look at the set  $\Delta_{uv} = (N(u) \cap N(v)) \cup \{u, v\}$ . Thus, the idea is to compute  $\Delta_{uv}$ , iterate over all possible labelings of it, and for each such labeling use Lemma 17 to infer  $H$  and check whether  $H$  is a halfspace. This still does not yield the desired bound, as  $\Delta_{uv}$  could be much larger than  $\omega(G)$ ; in fact we could have  $|\Delta_{uv}| = \Omega(n)$  even though  $\omega(G) = \mathcal{O}(1)$ . Thus, we first check if  $\Delta_{uv}$  is a subset of  $U$ , which can be done by just testing if  $\neg G[\Delta_{uv}]$  is bipartite. If this is the case, then we already know  $|\Delta_{uv}| \leq 2\omega(G)$ . By a careful analysis of  $\neg G[\Delta_{uv}]$  we can then show that we need only to test certain labelings of  $\Delta_{uv}$ , whose number can be bounded by roughly  $2^{\omega(G)}$ . The pseudocode of MH-FPT-LIST is given in Appendix C, together with the proof of:

**Theorem 20**  $\text{MH-FPT-LIST}(G)$  lists  $\mathcal{H}_{\text{MP}}(G)$  in time  $2^{\omega(G)} \text{poly}(n)$ .

The proof of Theorem 20 immediately implies  $|\mathcal{H}_{\text{MP}}(G)| \leq 2^{\omega(G)} 2m$ . However, starting from the analysis of MH-FPT-LIST, and using a charging argument that takes into account that an edge  $uv$  can appear in the cut  $\delta(H)$  of many  $m$ -halfspaces, we can prove the improved bound stated in Theorem 9 (see Appendix C.1 for the complete proof).

## Acknowledgments

MB and EE acknowledge the financial support from the FAIR (Future Artificial Intelligence Research) project, funded by the NextGenerationEU program within the PNRR-PE-AI scheme and the the EU Horizon CL4-2022-HUMAN-02 research and innovation action under grant agreement 101120237, project ELIAS (European Lighthouse of AI for Sustainability). MT acknowledges support from a DOC fellowship of the Austrian academy of sciences (ÖAW) and partial support from FFG through the AI4SAR project (885324) in the ASAP programme.

## References

- Peyman Afshani, Ehsan Chiniforooshan, Reza Dorigiv, Arash Farzan, Mehdi Mirzazadeh, Narges Simjour, and Hamid Zarrabi-Zadeh. On the complexity of finding an unknown cut via vertex queries. In *Computing and Combinatorics: 13th Annual International Conference*, 2007.
- Dana Angluin. Queries and concept learning. *Machine learning*, 2:319–342, 1988.
- Hans-Jürgen Bandelt. Graphs with intrinsic  $s_3$  convexities. *Journal of graph theory*, 13(2):215–228, 1989.
- Ja M. Barzdin. On the prediction of general recursive functions. In *Soviet Mathematics Doklady*, volume 13, pages 1224–1228, 1972.
- Shai Ben-David, Dávid Pál, and Shai Shalev-Shwartz. Agnostic online learning. In *COLT*, volume 3, page 1, 2009.
- Avrim Blum. On-line algorithms in machine learning. In *Online Algorithms, The State of the Art*, volume 1442 of *Lecture Notes in Computer Science*, pages 306–325. Springer, 1996.
- Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K. Warmuth. Learnability and the Vapnik-Chervonenkis dimension. *Journal of the ACM (JACM)*, 36(4):929–965, 1989.
- Cornelius Brand, Robert Ganian, and Kirill Simonov. A parameterized theory of PAC learning. In *AAAI*, 2023.
- Marco Bressan, Nicolò Cesa-Bianchi, Silvio Lattanzi, and Andrea Paudice. Exact recovery of clusters in finite metric spaces using oracle queries. In *COLT*, 2021.
- Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006.
- Nicolò Cesa-Bianchi, Yoav Freund, David Haussler, David P. Helmbold, Robert E. Schapire, and Manfred K. Warmuth. How to use expert advice. *Journal of the ACM*, 44(3):427–485, 1997.

- Nicolò Cesa-Bianchi, Claudio Gentile, Fabio Vitale, and Giovanni Zappella. Active learning on trees and graphs. In *COLT*, 2010.
- Nicolò Cesa-Bianchi, Claudio Gentile, Fabio Vitale, and Giovanni Zappella. Random spanning trees and the prediction of weighted graphs. *Journal of Machine Learning Research*, 14(1):1251–1284, 2013.
- Jérémie Chalopin, Manoj Changat, Victor Chepoi, and Jeny Jacob. First-order logic axiomatization of metric graph theory. *arXiv preprint arXiv:2203.01070*, 2022a.
- Jérémie Chalopin, Victor Chepoi, Fionn Mc Inerney, Sébastien Ratel, and Yann Vaxès. Sample compression schemes for balls in graphs. In *47th International Symposium on Mathematical Foundations of Computer Science (MFCS 2022)*, 2022b.
- Jérémie Chalopin, Victor Chepoi, Shay Moran, and Manfred K Warmuth. Unlabeled sample compression schemes and corner peelings for ample and maximum classes. *Journal of Computer and System Sciences*, 127:1–28, 2022c.
- Manoj Changat, Henry Martyn Mulder, and Gerard Sierksma. Convexities related to path properties on graphs. *Discrete Mathematics*, 290(2-3):117–131, 2005.
- Victor Chepoi. Some properties of domain finite convexity structures (in russian), res. *Algebra, Geometry and Appl.(Moldova State University)*, pages 142–148, 1986.
- Victor Chepoi. Separation of two convex sets in convexity structures. *Journal of Geometry*, 50: 30–51, 1994.
- Victor Chepoi. Separation axiom  $S_3$  for geodesic convexity in graphs. *arXiv preprint arXiv:2405.07512*, 2024.
- Victor Chepoi, Bertrand Estellon, and Yann Vaxes. Covering planar graphs with a fixed number of balls. *Discrete & Computational Geometry*, 37:237–244, 2007.
- Victor Chepoi, Kolja Knauer, and Manon Philibert. Labeled sample compression schemes for complexes of oriented matroids. *arXiv preprint arXiv:2110.15168*, 2021.
- Gautam Dasarathy, Robert Nowak, and Xiaojin Zhu. S2: An efficient graph based active learning algorithm with application to nonparametric classification. In *COLT*, 2015.
- Mitre C. Dourado, Fábio Protti, and Jayme L. Szwarcfiter. Complexity results related to monophonic convexity. *Discrete Applied Mathematics*, 158(12):1268–1274, 2010.
- Pierre Duchet. Convex sets in graphs, II. Minimal path convexity. *Journal of Combinatorial Theory, Series B*, 44(3):307–316, 1988.
- Pierre Duchet and Henry Meyniel. Ensemble convexes dans les graphes i: Théorèmes de Helly et de Radon pour graphes et surfaces. *European Journal of Combinatorics*, 4(2):127–132, 1983.
- Mohammed Elaroussi, Lhouari Nourine, and Simon Vilmin. Half-space separation in monophonic convexity. *arXiv preprint arXiv:2404.17564*, 2024.

- Martin Farber and Robert E. Jamison. Convexity in graphs and hypergraphs. *SIAM Journal on Algebraic Discrete Methods*, 7(3):433–444, 1986.
- Claudio Gentile, Mark Herbster, and Stephen Pasteris. Online similarity prediction of networked data from known and unknown graphs. In *COLT*, 2013.
- Roland Glantz and Henning Meyerhenke. On finding convex cuts in general, bipartite and plane graphs. *Theoretical Computer Science*, 695:54–73, 2017.
- Lucía M. González, Luciano N. Grippo, Martín D. Safe, and Vinicius F. dos Santos. Covering graphs with convex sets and partitioning graphs into convex sets. *Information Processing Letters*, 158, 2020.
- Andrew Guillory and Jeff A. Bilmes. Label selection on graphs. In *NIPS*, 2009.
- Steve Hanneke. An analysis of graph cut size for transductive learning. In *Proceedings of the 23rd International Conference on Machine Learning*, pages 393–399, 2006.
- Tibor Hegedűs. Generalized teaching dimensions and the query complexity of learning. In *COLT*, 1995.
- Mark Herbster, Massimiliano Pontil, and Lisa Wainer. Online learning over graphs. In *ICML*, 2005.
- Mark Herbster, Stephen Pasteris, and Shaona Ghosh. Online prediction at the limit of zero temperature. *NIPS*, 2015.
- Robert E. Jamison-Waldner. *Convexity and block graphs*. Clemson University. Department of Mathematical Sciences, 1981.
- Jon Kleinberg. Detecting a network failure. *Internet Mathematics*, 1(1):37–55, 2004.
- Hung Le and Christian Wulff-Nilsen. VC set systems in minor-free (di) graphs and applications. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 5332–5360, 2024.
- Bi-Qing Li, Jin You, Lei Chen, Jian Zhang, Ning Zhang, Hai-Peng Li, Tao Huang, Xiang-Yin Kong, and Yu-Dong Cai. Identification of lung-cancer-related genes with the shortest path approach in a protein-protein interaction network. *BioMed research international*, 2013(1):267375, 2013.
- Nick Littlestone. Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm. *Machine learning*, 2:285–318, 1988.
- Francesco M Malvestuto, Mauro Mezzini, and Marina Moscarini. Characteristic properties and recognition of graphs in which geodesic and monophonic convexities are equivalent. *Discrete Mathematics, Algorithms and Applications*, 4(04):1250063, 2012.
- Shay Moran and Amir Yehudayoff. On weak  $\varepsilon$ -nets and the Radon number. *Discrete & Computational Geometry*, 64(4):1125–1140, 2020.
- Ignacio M. Pelayo. *Geodesic convexity in graphs*. Springer Briefs in Mathematics. Springer, 2013. doi: 10.1007/978-1-4614-8699-2.

- Kristiaan Pelckmans, John Shawe-Taylor, Johan AK Suykens, and Bart De Moor. Margin based transductive graph cuts using linear programming. In *Artificial Intelligence and Statistics*, pages 363–370, 2007.
- Florian Seiffarth, Tamás Horváth, and Stefan Wrobel. Maximal closed set and half-space separations in finite closure systems. *Theoretical Computer Science*, 973:114105, 2023.
- Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- Lovro Šubelj, Dalibor Fiala, Tadej Ciglarič, and Luka Kronegger. Convexity in scientific collaboration networks. *Journal of Informetrics*, 13(1):10–31, 2019.
- Michel Talagrand. Sharper bounds for gaussian and empirical processes. *The Annals of Probability*, pages 28–76, 1994.
- Maximilian Thiessen and Thomas Gärtner. Active learning of convex halfspaces on graphs. *NeurIPS*, 2021.
- Maximilian Thiessen and Thomas Gärtner. Online learning of convex sets on graphs. In *ECMLP-KDD*, 2022.
- Leslie G. Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.
- Marcel L. J. van de Vel. *Theory of convex structures*, volume 50 of *North-Holland Mathematical Library*. North Holland, 1993.
- Vladimir Vapnik and Alexey Ya. Chervonenkis. *Theory of pattern recognition*. Nauka, Moscow, 1974.
- Xianghong Zhou, Ming-Chih J Kao, and Wing Hung Wong. Transitive functional annotation by shortest-path analysis of gene expression data. *Proceedings of the National Academy of Sciences*, 99(20):12783–12788, 2002.



## Appendix A. Proofs of Auxiliary Lemmas

In this section, we provide the proofs of some structural properties that are stated in the main body, and are crucial in proving the main results of this work. We restate each lemma for convenience.

**Lemma 17** *Let  $H \in \mathcal{H}_{\text{MP}}(G)$ , let  $uv \in \delta(H)$ , and let  $\Delta_{uv} = (N(u) \cap N(v)) \cup \{u, v\}$ . Then,  $H$  consists precisely of those vertices whose label can be inferred from  $u, v$ , and a vertex in  $H \cap \Delta_{uv}$ . Formally:*

$$H = \bigcup_{z \in H \cap \Delta_{uv}} z/v .$$

**Proof of Lemma 17.** We first show that  $\bigcup_{z \in H \cap \Delta_{uv}} z/v \subseteq H$ . This follows from the fact that for  $z \in H$  and  $v \in \overline{H}$  it holds  $z/v \subseteq H$ , as  $H$  is a monophonic halfspace.

For the other direction, we have to show that for all  $x \in H$  there exists a  $z \in H \cap \Delta_{uv}$  such that  $x \in z/v$ . If  $x \in u/v$  we are done. Also,  $x \notin v/u$ , as  $v/u \subseteq \overline{H}$ . Hence, we can assume that  $x \notin u/v \cup v/u$ . Take an induced path  $\pi_u$  from  $u$  to  $x$  and an induced path  $\pi_v$  from  $v$  to  $x$ . Denote the closest vertex to  $u$  where  $\pi_u$  and  $\pi_v$  meet as  $z$  ( $z = x$  is possible). Let  $C_{uvz}$  denote the cycle given by the subpaths going from  $u$  and  $v$  to  $z$  and the edge  $\{u, v\}$ . Assume that  $C_{uvz}$  is not an induced cycle. Then there is a chord  $ab$  in  $C_{uvz}$  going from  $\pi_u$  to  $\pi_v$  (or vice versa) and we could modify  $\pi_u$  (respectively  $\pi_v$ ) to meet  $\pi_v$  already in  $b$ . We can iterate this process until we are left with an induced cycle  $C_{uvz}$ . Assume that  $C_{uvz}$  consists of at least 5 vertices. Then either  $u, v \in V(C_{uvz}) \subseteq H$  or  $u, v \in V(C_{uvz}) \subseteq \overline{H}$  a contradiction to  $u \in H$  and  $v \in \overline{H}$ . Also,  $C_{uvz}$  cannot consist of exactly 4 vertices as in this case  $x \in u/v \cup v/u$ . Thus,  $C_{uvz}$  is an induced triangle and  $z \in (N(u) \cap N(v)) \setminus \{u, v\}$ . By construction  $z$  is along an induced path from  $u$  to  $x$  and from  $v$  to  $x$  and hence  $x \in z/v$  as required.  $\blacksquare$

**Lemma 6** *Let  $\{z, v\} \in E$ . Then,  $z/v = \{x \in V \mid N(v) \setminus \{z\}$  is not a  $(z, x)$ -separator in  $G\}$ . As a consequence one can compute  $z/v$  in time  $\text{poly}(n)$ .*

**Proof of Lemma 6.** Let  $L = N(v) \setminus \{z\}$  and  $S = \{x \in V \mid L \text{ is not a } (z, x)\text{-separator}\}$ . It is straightforward that  $S$  can be computed in polynomial time, hence we only need to show that  $z/v = S$ .

Suppose first that  $x \in S$ . By definition of  $S$  there exists a path  $\pi$  between  $x$  and  $z$  such that  $V(\pi) \cap L = \emptyset$ . Thus  $v$  is not adjacent to  $V(\pi)$ , and this holds obviously for any subpath  $\pi'$  of  $\pi$ . Take in particular an induced subpath  $\pi'$  of  $\pi$  between  $x$  and  $z$ , and observe that extending it to  $v$  yields an induced path between  $x$  and  $v$  that contains  $z$ . Hence  $x \in z/v$ .

Suppose now that  $x \in z/v$ . By definition of  $z/v$  there exists an induced path  $\pi$  between  $x$  and  $v$  that contains  $z$ . Clearly  $\pi$  does not contain any  $y \in N(v) \setminus \{z, v\}$ , and thus its prefix  $\pi'$  between  $x$  and  $z$  does not contain any  $y \in N(v) \setminus \{z\} = L$ . It follows that  $L$  is not a  $(z, x)$ -separator.  $\blacksquare$

**Lemma 16**  *$H \in \mathcal{H}_{\text{MP}}(G)$  if and only if  $G[\Gamma(H)]$  and  $G[\Gamma(\overline{H})]$  are cliques.*

**Proof of Lemma 16.** Let  $H \in \mathcal{H}_{\text{MP}}(G)$ . Suppose  $\{x, x'\} \notin E$  for some distinct  $x, x' \in \Gamma(H)$ . Choose  $y \in N(x) \setminus H$  and  $y' \in N(x') \setminus H$  that have smallest distance in  $G[\overline{H}]$ . (Note that  $N(x) \setminus H \subseteq \overline{H}$  and  $N(y') \setminus H \subseteq \overline{H}$ , hence  $y, y'$  exist). Let  $\pi$  be a shortest path between  $y$  and  $y'$  in

$G[\overline{H}]$ . Observe that the concatenation of  $x, \pi, x'$  is an induced path between  $x$  and  $x'$  that intersects  $\overline{H}$ . This contradicts  $H \in \mathcal{H}_{\text{MP}}(G)$ . We conclude that  $\{x, x'\} \in E$  for all distinct  $x, x' \in \Gamma(H)$ , hence  $G[\Gamma(H)]$  is a clique. The same holds with  $\overline{H}$  in place of  $H$ .

Now let  $H \subseteq V$  and suppose  $G[\Gamma(H)]$  and  $G[\Gamma(\overline{H})]$  are cliques; we show that both  $H$  and  $\overline{H}$  are monophonically convex. Suppose for contradiction that an induced path between distinct  $x, x' \in H$  intersects  $\overline{H}$ . Without loss of generality we may assume  $x, x'$  are the endpoints of the path, and thus by definition in  $\Gamma(H)$  by definition. This implies  $\{x, x'\} \in E$ , contradicting the fact that the path is induced. The same holds with  $\overline{H}$  in place of  $H$ .  $\blacksquare$

**Lemma 19** *Let  $H \in \mathcal{H}_{\text{MP}}$  and  $uv$  be a cut-edge of  $H$ . Then, the number  $N_{uv}$  of connected components in  $\neg G[\Delta_{uv}]$  satisfies  $N_{uv} \leq \omega(G)$ .*

**Proof of Lemma 19.** Observe that  $N_{uv} \leq \alpha(\neg G[\Delta_{uv}]) = \omega(G[\Delta_{uv}]) \leq \omega(G)$ .  $\blacksquare$

## Appendix B. Missing Details for Section 8

### B.1. Full list of constraints for $\varphi_{uv}$

$$\varphi_{\Delta_{uv}^-} = \bigwedge_{\{x,y\} \in E(\neg G[\Delta_{uv}^-])} (x \in H \vee y \in H) \wedge (x \notin H \vee y \notin H) \quad (6)$$

$$\varphi_{\text{conv}} = \bigwedge_{x \in \text{conv}(A^u)} (x \in H) \quad \wedge \quad \bigwedge_{x \in \text{conv}(A^v)} (x \notin H) \quad (7)$$

$$\varphi_T = \bigwedge_{T_i \in \text{cc}(T)} \bigwedge_{x,y \in V(T_i)} (x \in H \vee y \notin H) \wedge (x \notin H \vee y \in H) \quad (8)$$

$$\varphi_{u,3} = \bigwedge_{\substack{\pi=(x,y,z) \\ \pi \in \Pi_3(A^u, \Delta_{uv}^-)}} (y \in H \vee z \notin H), \quad \varphi_{v,3} = \bigwedge_{\substack{\pi=(x,y,z) \\ \pi \in \Pi_3(A^v, \Delta_{uv}^-)}} (y \notin H \vee z \in H) \quad (9)$$

$$\varphi_{u,4} = \bigwedge_{\substack{\pi=(x,w,y,z) \\ \pi \in \Pi_4(A^u, \Delta_{uv}^-)}} (y \in H \vee z \notin H), \quad \varphi_{v,4} = \bigwedge_{\substack{\pi=(x,w,y,z) \\ \pi \in \Pi_4(A^v, \Delta_{uv}^-)}} (y \notin H \vee z \in H) \quad (10)$$

$$\varphi_{A^u} = \bigwedge_{\substack{\{x,y\} \in E \\ x \in \text{conv}(A^u)}} \bigwedge_{z \in \Delta_{uv}^- \setminus N(x)} (y \in H \vee z \notin H), \quad \varphi_{A^v} = \bigwedge_{\substack{\{x,y\} \in E \\ x \in \text{conv}(A^v)}} \bigwedge_{z \in \Delta_{uv}^- \setminus N(x)} (y \notin H \vee z \in H) \quad (11)$$

### B.2. Proof of Theorem 5

Before proving Theorem 5 we need another technical lemma.

**Lemma 21** *Let  $H \in \mathcal{H}_{\text{MP}}(G)$  and  $uv \in \delta(H)$ . Then, (i)  $A^u \subseteq H$  and  $A^v \subseteq \overline{H}$ , and (ii)  $T_i \subseteq H$  or  $T_i \subseteq \overline{H}$ , for all  $T_i \in \text{cc}(T)$ .*

**Proof** (i) Let  $x \in A^u$  and suppose by contradiction  $x \notin H$ . In particular  $x \neq u$ , so  $d(x, u) \geq 1$ . Now  $(x, u, v)$  is an induced path with endpoints in  $H$  containing a vertex of  $\overline{H}$ , contradicting  $H \in \mathcal{H}_{\text{MP}}(G)$ . A symmetric argument applies to  $A^v$ .

(ii) First, let us prove that  $\Gamma(H) \cup \Gamma(\overline{H}) \subseteq \square_{uv} \cup \Delta_{uv}^-$ . Let  $x \in \Gamma(H)$ . If  $x = u$  then  $x \in \square_{uv}$ , so we may assume  $x \in \Gamma(H) \setminus \{u\}$ , which by Lemma 16 implies  $x \in N(u)$ . Consider then any edge  $xy \in \delta(H)$ . If  $y = v$  then  $u \in \Delta_{uv}^-$ . If instead  $y \neq v$ , then  $y \in N(v)$  by Lemma 16; so  $(u, x, y, v)$  is a 4-cycle, and  $x, y \in \square_{uv}$ . We conclude that  $\Gamma(H) \cup \Gamma(\overline{H}) \subseteq \square_{uv} \cup \Delta_{uv}^-$ . This implies that  $T$  is an edge-subgraph of  $R = G \setminus E(G[\Gamma(H) \cup \Gamma(\overline{H})])$ . Therefore, every connected component  $T_i \in \text{cc}(T)$  is entirely contained in some connected component  $R_i \in \text{cc}(R)$ . Thus, it is sufficient to prove the claim for  $R_i$ . Suppose for contradiction that  $R_i$  intersects both  $H$  and  $\overline{H}$ . Since  $R_i$  is connected there is an edge  $\{x, y\} \in E(R_i)$  with  $x \in H$  and  $y \in \overline{H}$ , so  $\{x, y\} \in E(R_i) \cap \delta(H)$ . But  $\delta(H) \subseteq E(G[\Gamma(H) \cup \Gamma(\overline{H})])$  and thus  $E(R) \cap \delta(H) = \emptyset$ , a contradiction. ■

**Theorem 5 (Consistency check is in polynomial time)** *Let  $G = (V, E)$  and  $(S, y)$  a labeled sample. Then,  $\text{MH-CHECK}(G, (S, y))$  (Algorithm 1) runs in time  $\text{poly}(n)$  and returns an  $H \in \mathcal{H}_{\text{MP}}(G)$  consistent with  $(S, y)$ , or reports that no such  $H$  exists.*

**Proof of Theorem 5.** First of all, note that  $\varphi_{uv}$  is a 2-SAT instance, and that it can be computed in polynomial time (see above). The same is then obviously true for  $\varphi_{uv}(S_-, S_+)$ , too. It remains to prove that  $H \in \mathcal{H}_{\text{MP}}(G)$  consistent with  $(S_-, S_+)$  exists if and only if  $\varphi_{uv}(S_-, S_+)$  is satisfied for some  $uv \in E$ . We prove the two directions separately.

**Forward direction.** Suppose  $H \in \mathcal{H}_{\text{MP}}(G)$  is consistent with  $(S_-, S_+)$ . We check (7)–(11) and show  $H$  satisfies  $\varphi_{uv}$  for every  $uv \in \delta(H)$ . This implies that  $H$  satisfies  $\varphi_{uv}(S_-, S_+)$  for every  $uv \in \delta(H)$  and, thus, for some  $uv \in E$ .

- $\varphi_{\text{conv}}$ . First,  $A^u \subseteq H$  and  $A^v \subseteq \overline{H}$  by Lemma 21. Since  $H$  is closed under  $\text{conv}$  this implies  $\text{conv}(A^u) \subseteq H$  and  $\text{conv}(A^v) \subseteq \overline{H}$ , too.
- $\varphi_{\Delta_{uv}^-}$ . Note that  $\Delta_{uv}^- \subseteq \Gamma(H) \cup \Gamma(\overline{H})$  and apply Lemma 22.
- $\varphi_T$ . Apply Lemma 21.
- $\varphi_{u,3}$  and  $\varphi_{v,3}$ . Suppose by contradiction  $y \notin H$  and  $z \in H$ . Since  $(x, y, z)$  is induced, this violates the convexity of  $H$ . Thus,  $\varphi_{u,3}$  is satisfied. A symmetric argument applies to  $\varphi_{v,3}$ .
- $\varphi_{u,4}$ . The arguments are similar to those for  $\varphi_{u,3}$  and  $\varphi_{v,3}$ .
- $\varphi_{A^u}$  and  $\varphi_{A^v}$ . Suppose the formula fails, so there exists  $\{x, y\} \in E$  with  $x \in \text{conv}(A^u)$  and  $y \notin H$ , as well as  $z \in \Delta_{uv}^- \setminus N(x)$  with  $z \in H$ . As  $x \in H, y \notin H$ , and  $\{x, y\} \in E$ , then  $\{x, y\} \in \delta(H)$ . As  $z \in H, v \notin H$ , and  $z \in \Delta_{uv}^-$ , then  $z \in \delta(H)$ , too. But then by Lemma 16 we have  $\{x, y\} \in E$ , a contradiction. A symmetric argument applies to  $\varphi_{A^v}$ .

**Backward direction.** Suppose  $H \subseteq V$  satisfies  $\varphi_{uv}(S_-, S_+)$  for some  $uv \in E$ . This implies that  $H$  satisfies  $\varphi_{uv}$ , which in turn implies  $uv \in \delta(H)$  by the constraints of (7). We now show that  $H \in \mathcal{H}_{\text{MP}}(G)$ . Suppose indeed by contradiction  $H \notin \mathcal{H}_{\text{MP}}(G)$ . Note how this implies  $\text{conv}(H) \neq H$  or  $\text{conv}(\overline{H}) \neq \overline{H}$ . Without loss of generality we assume  $\text{conv}(H) \neq H$ . Therefore there exists an induced path  $P = G[\{x_1, \dots, x_k\}]$  in  $G$  such that  $x_1, x_k \in H$  and  $x_i \in \overline{H}$  for all  $i = 2, \dots, k-1$ . We will use the path  $P$  to prove a contradiction.

First, since  $\varphi_T$  is satisfied, then  $\{x_1, x_2\} \in E(G[\square_{uv} \cup \Delta_{uv}^-]) = E(G[A \cup \Delta_{uv}^-])$ . By the same argument  $\{x_{k-1}, x_k\} \in E(G[A \cup \Delta_{uv}^-])$ . Moreover, since  $\varphi_{\text{conv}}$  is satisfied then  $\{x_1, x_k\} \cap A^v = \emptyset$  and  $\{x_2, x_{k-1}\} \cap A^u = \emptyset$ ; hence,  $x_1, x_k \in A^u \cup \Delta_{uv}^-$  and  $x_2, x_{k-1} \in A^v \cup \Delta_{uv}^-$ .

Now, if  $x_1, x_k \in A^u$ , then  $V(P) \subseteq H$  since  $\varphi_{\text{conv}}$  is satisfied; hence  $|\{x_1, x_k\} \cap A^u| \leq 1$ . If instead  $x_1, x_k \in \Delta_{uv}^-$ , then since  $\varphi_{\Delta_{uv}^-}$  is satisfied  $\{x_1, x_k\} \notin E(\neg G[\Delta_{uv}^-])$ , implying the absurd  $\{x_1, x_k\} \in E(G)$ ; hence  $|\{x_1, x_k\} \cap \Delta_{uv}^-| \leq 1$ . We conclude that precisely one of  $x_1, x_k$  is in  $A^u$  and the other one in  $\Delta_{uv}^-$ . Without loss of generality say  $x_1 \in A^u$  and  $x_k \in \Delta_{uv}^-$ . Then  $k \geq 5$ , for otherwise  $\pi = P$  would falsify  $\varphi_{u,3}$  or  $\varphi_{u,4}$ . As  $k \geq 5$  and  $P$  is induced we deduce  $x_2, x_{k-1}$  are distinct and  $\{x_1, x_{k-1}\}, \{x_2, x_{k-1}\} \notin E$ .

Next we show that  $x_2 \in \Delta_{uv}^-$  and  $x_{k-1} \in A^v$ . Suppose by contradiction  $x_2 \in A^v$ . Letting  $x, y = x_1, x_2$  and  $z = x_k$ , then,  $\varphi_{A^v}$  fails. Thus  $x_2 \in \Delta_{uv}^-$ . Now suppose by contradiction  $x_{k-1} \in \Delta_{uv}^-$ . Then  $\{x_2, x_{k-1}\} \in E(\neg G[\Delta_{uv}^-])$ , and since  $\varphi_{\Delta_{uv}^-}$  is satisfied then  $x_2 \in H$  or  $x_{k-1} \in H$ , a contradiction. Thus  $x_{k-1} \in A^v$ . Now let  $x, y = x_{k-1}, x_k$  and  $z = x_2$ . Then  $\varphi_{A^v}$  fails. We conclude that  $P$  does not exist. Hence  $H = \text{conv}(H)$ , and  $H \in \mathcal{H}_{\text{MP}}(G)$ .  $\blacksquare$

### B.3. Proof of Theorem 8

**Theorem 8 (Listing the version space with polynomial delay)** *There exists an algorithm MH-LIST that, given a graph  $G$  and a labeled sample  $(S, y)$ , lists  $\text{VS} = \text{VS}(S, \mathcal{H}_{\text{MP}})$  with delay  $\text{poly}(n)$  and total time  $|\text{VS}| \text{poly}(n)$ .*

The pseudocode of MH-LIST is given in Algorithm 2. It is straightforward to see that the algorithm constructs a binary search tree whose leaves are precisely the elements of the version space  $\text{VS}((S, y), \mathcal{H}_{\text{MP}}(G))$ , which are the only ones that are output. Moreover, by Theorem 5 MH-LIST spends time  $\text{poly}$  for every node in the search tree. The claim of Theorem 8 follows.

---

#### Algorithm 2: MH-LIST

---

**Input:**  $(S, y), G = (V, E)$

**Output:** the version space  $\text{VS}((S, y), \mathcal{H}_{\text{MP}}(G))$

let  $H^+ = y^{-1}(1)$  and  $H^- = y^{-1}(0)$

**if**  $H^+ \cup H^- = V$  **then**

    | **print**  $H^+$  and **return**

pick any  $x \in V \setminus (H^+ \cup H^-)$

**if** MH-CHECK( $G, H^+ \cup \{x\}, H^-$ ) **then** MH-LIST( $G, H^+ \cup \{x\}, H^-$ )

**if** MH-CHECK( $G, H^+, H^- \cup \{x\}$ ) **then** MH-LIST( $G, H^+, H^- \cup \{x\}$ )

---

### Appendix C. MH-FPT-LIST, and a proof of Theorem 20

Before proving the theorem we need one auxiliary lemma.

**Lemma 22** *Let  $H \in \mathcal{H}_{\text{MP}}(G)$  and let  $F = \neg G[\Gamma(H) \cup \Gamma(\overline{H})]$ . Then, every connected component  $\hat{F}$  of  $F$  is bipartite with sides  $V(\hat{F}) \cap H$  and  $V(\hat{F}) \cap \overline{H}$ .*

---

**Algorithm 3:** MH-FPT-LIST
 

---

**Input:**  $G = (V, E)$   
**Output:**  $\mathcal{H}_{\text{MP}}(G)$   
 let  $B = \emptyset$   
**print**  $\emptyset, V$   
**for**  $\{u, v\} \in E$  **do**  
     let  $\Delta_{uv} = (N(u) \cap N(v)) \cup \{u, v\}$   
     **if**  $\neg G[\Delta_{uv}]$  *is bipartite* **then**  
         let  $R \subseteq \Delta_{uv}$  contain one vertex  $x_C$  from each connected component  $C$  of  $\neg G[\Delta_{uv}]$   
         let  $S = \emptyset$   
         **for each subset**  $U \subseteq R$  **do**  
             **for each connected component**  $C$  of  $\neg G[\Delta_{uv}]$  **do**  
                 **if**  $x_C \in U$  **then**  $S = S \cup C(x_C)$  **else**  $S = S \cup V(C) \setminus C(x_C)$   
             **end**  
             let  $X = \bigcup_{z \in S} z/v$   
             **if**  $X$  and  $V \setminus X$  are  $m$ -convex and  $\Gamma(X) \cap B = \emptyset$  **then print**  $X$   
         **end**  
          $B = B \cup \{\{u, v\}\}$   
     **end**  
**end**

---

**Proof** The first claim follows immediately from Lemma 16. For the second claim observe that if  $\{x, y\} \in E(F)$  then  $\{x, y\} \notin E(G)$  and thus, again by Lemma 16,  $\{x, y\} \not\subseteq H$  and  $\{x, y\} \not\subseteq \bar{H}$ ; therefore, precisely one of  $x$  and  $y$  is in  $H$ . ■

**Proof of Theorem 20.** For the correctness, the condition at line 3 ensures MH-FPT-LIST lists only elements of  $\mathcal{H}_{\text{MP}}(G)$ ; thus we need just to prove that every  $H \in \mathcal{H}_{\text{MP}}(G)$  is listed exactly once. Let  $\{u, v\}$  be the first edge of  $\delta(H)$  considered at line 3 (recall that we assume  $G$  to be connected, hence such an edge exists). Note that every  $x \in N(u) \cap N(v)$  is incident with an edge in  $\delta(H)$  (the other endpoint being  $u$  if  $x \in \bar{H}$  and  $v$  if  $x \in H$ ), and the same holds obviously for  $x \in \{u, v\}$ . Therefore  $\Delta_{uv} \subseteq \Gamma(H) \cup \Gamma(\bar{H})$ . Lemma 22 then implies  $\neg G[\Delta_{uv}]$  is bipartite, so the condition at line 3 holds. Next, consider the loop at line 3. Clearly, at some iteration  $U = R \cap H$ . At that point, again by Lemma 22 the loop at line 3 ensures  $S = \Delta_{uv} \cap H$ . By Lemma 17, then,  $X = \bigcup_{z \in S} z/v = H$ . The condition at line 3 is then satisfied, thus  $H$  is printed; moreover in every subsequent iteration of the loop of line 3  $\{u, v\} \in B$  and thus  $H$  will not be printed.

For the running time, note that the loop at line 3 is executed only if  $\neg G[\Delta_{uv}]$  is bipartite, in which case the loop performs  $2^{N_{uv}}$  iterations where  $N_{uv}$  is the number of connected components of  $\neg G[\Delta_{uv}]$ . By Lemma 19  $N_{uv} \leq \omega(G)$ . Hence the loop at line 3 makes at most  $2^{\omega(G)}$  iterations. To conclude the proof note that every other step performed by MH-FPT-LIST (including computing  $X$ , see Lemma 6) takes time  $\text{poly}(n)$ .

### C.1. Proof of Theorem 9

**Theorem 9**  $|\mathcal{H}_{\text{MP}}(G)| \leq \frac{4m2^{\omega(G)}}{\omega(G)} + 2$  for all graphs  $G$ .

**Proof** For any  $e \in E$ , let  $\mathcal{H}_{\text{MP}}(e) = \{H \in \mathcal{H}_{\text{MP}}(G) \mid e \in \delta(H)\}$ . Recall that each  $x \in \Gamma(H) \cup \Gamma(\overline{H})$  is incident with some edge in  $\delta(H)$  by definition, and therefore

$$|\delta(H)| \geq \max\{|\Gamma(H)|, |\Gamma(\overline{H})|\} \geq \frac{|\Gamma(H) \cup \Gamma(\overline{H})|}{2}. \quad (12)$$

Moreover, for any  $e \in E$  let  $\omega(e) = \omega(G[\Delta_{uv}])$  where  $e = \{u, v\}$ , see Algorithm 3 and the proof of Theorem 20. Note that, since  $\omega(G[\Delta_{uv}]) \geq \frac{|\Delta_{uv}|}{2}$ , then

$$\frac{2^{\omega(e)}}{|\Delta_{uv}|} = \frac{2^{\omega(G[\Delta_{uv}])}}{|\Delta_{uv}|} \leq \frac{2^{1+\omega(G[\Delta_{uv}])}}{\omega(G[\Delta_{uv}])} \leq \frac{2^{1+\omega(G)}}{\omega(G)}, \quad (13)$$

where the last inequality follows by monotonicity of  $2^x/x$  for  $x \geq 2$  and since  $\omega(G[\Delta_{uv}]) \geq 2$  as  $\{u, v\} \subseteq \Delta_{uv}$ . We then obtain (not counting the two trivial halfspaces  $\emptyset, V$ ):

$$|\mathcal{H}_{\text{MP}}(G)| - 2 = \sum_{H \in \mathcal{H}_{\text{MP}}(G)} \sum_{e \in \delta(H)} \frac{1}{|\delta(H)|} \quad (14)$$

$$= \sum_{e \in E} \sum_{H \in \mathcal{H}_{\text{MP}}(e)} \frac{1}{|\delta(H)|} \quad \text{by rearranging terms} \quad (15)$$

$$\leq \sum_{e \in E} \sum_{H \in \mathcal{H}_{\text{MP}}(e)} \frac{2}{|\Gamma(H) \cup \Gamma(\overline{H})|} \quad \text{by (12)} \quad (16)$$

$$\leq \sum_{e \in E} |\mathcal{H}_{\text{MP}}(e)| \frac{2}{|\Delta_e|} \quad \text{as } \Delta_e \subseteq \Gamma(H) \cup \Gamma(\overline{H}) \quad (17)$$

$$\leq \sum_{e \in E} 2^{\omega(e)} \frac{2}{|\Delta_e|} \quad \text{see proof of Theorem 20} \quad (18)$$

$$\leq \sum_{e \in E} \frac{2^{2+\omega(G)}}{\omega(G)} \quad \text{by (13)} \quad (19)$$

$$= 4m \frac{2^{\omega(G)}}{\omega(G)}. \quad (20)$$

The proof is complete. ■

## Appendix D. Supervised Learning

We study the standard PAC (probably approximately correct) setting with a given finite instance space  $V$  and a known hypothesis space  $\mathcal{H} \subseteq 2^V$ . In our case  $\mathcal{H} = \mathcal{H}_{\text{MP}}(G)$  is defined implicitly by the input graph  $G$  given to the algorithm. The standard approach to PAC learning is to return a hypothesis, a m-halfspace in our case, that is as close as possible to being consistent with the given labeled sample  $(S, y)$ . In particular, we say that it is possible to  $(\varepsilon, \delta)$ -PAC learn  $(V, \mathcal{H}_{\text{MP}})$  if we can provide a m-halfspace  $H = H(S, y)$  that has error bounded by  $\varepsilon$  with probability at least  $1 - \delta$ , for any  $\varepsilon, \delta \in (0, 1)$ , under some unknown joint distribution over vertices and their labels. The aim of PAC learning is to guarantee this property with the smallest sample complexity possible.

While standard PAC learning formulations typically measure the running time as a function of the sample size (see, e.g., Valiant (1984); Shalev-Shwartz and Ben-David (2014)), we measure running time as a function of  $n$ . Note that a dependence on  $n$  is unavoidable: even just to check whether two vertices  $u, v \in V$  are separated by some  $H \in \mathcal{H}_{\text{MP}}(G)$ , one needs to distinguish between  $G$  being a path or a cycle, which takes time  $\Omega(n)$ .

We start with a bound on the VC dimension of  $\mathcal{H}_{\text{MP}}(G)$ , thus determining learnability in the supervised setting. We write  $\text{VC}(V, \mathcal{H})$ , or simply  $\text{VC}(\mathcal{H})$ , for the VC dimension of the hypothesis space  $\mathcal{H} \subseteq 2^V$ .

**Proposition 4** *If  $G \in S_4$  then  $\text{VC}(\mathcal{H}_{\text{MP}}(G)) \geq \omega(G)$ . Moreover,  $\text{VC}(\mathcal{H}_{\text{MP}}(G)) \leq \tilde{\omega}(G)$  for all  $G$ .*

**Proof** We first provide a proof for the upper bound by showing that the chain of inequalities  $\text{VC}(V, \mathcal{H}_{\text{MP}}(G)) \leq r(G) \leq \tilde{\omega}(G)$  holds. The first inequality holds for halfspaces in arbitrary convexity spaces, as any set shattered by halfspaces must be Radon independent (Moran and Yehudayoff, 2020). The second inequality is a classic result in monophonic graph convexity theory (Duchet, 1988). Now, we show that the lower bound holds too. Indeed, if  $G \in S_4$  then every clique in  $G$  is shattered by  $\mathcal{H}_{\text{MP}}(G)$ , since (i) it is Radon independent, and (ii) any of its partitions can be extended to a halfspace (Thiessen and Gärtner, 2021). Hence,  $\text{VC}(V, \mathcal{H}_{\text{MP}}(G)) \geq \omega(G)$ . ■

We now move on to the study of the sample complexity for PAC learning  $(V, \mathcal{H}_{\text{MP}})$  in both the realizable and the agnostic settings.

**Realizable case.** Since we are in the realizable setting, we know there exists some m-halfspace  $H^* \in \mathcal{H}_{\text{MP}}(G)$  that is consistent with  $(S, y)$ . While it is known that computing m-convex  $p$ -partitions of a graph for  $p \geq 3$  is NP-hard (González et al., 2020), the complexity of computing m-halfspaces ( $p = 2$ ) is not known. A first step towards closing this gap is to use our listing algorithm to obtain an FPT algorithm for this problem. We go beyond that and fully close this gap by providing, in Theorem 5, a polynomial-time algorithm relying on a non-trivial reduction to the 2-SAT problem.

By standard PAC results (Blumer et al., 1989) we obtain the following near-optimal sample complexity bound, save possibly for the  $\log(1/\varepsilon)$  factor.

**Theorem 2 (Realizable PAC is poly-time)** *There exists an algorithm that, for any  $\varepsilon, \delta \in (0, 1)$ , can  $(\varepsilon, \delta)$ -PAC learn  $(V, \mathcal{H}_{\text{MP}})$  in the realizable setting using  $\mathcal{O}\left(\frac{\omega(G) \log(1/\varepsilon) + \log(1/\delta)}{\varepsilon}\right)$  labeled samples and time  $\text{poly}(n)$ .*

**Proof** Blumer et al. (1989) give a sample complexity of  $\mathcal{O}\left(\frac{\text{VC}(V, \mathcal{H}_{\text{MP}}) \log(1/\varepsilon) + \log(1/\delta)}{\varepsilon}\right)$  and the stated bound follows from the fact that  $\text{VC}(V, \mathcal{H}_{\text{MP}}) \leq \tilde{\omega}(G)$  (Proposition 4). The runtime follows by Theorem 5. ■

**Agnostic case.** In the agnostic case we are given a sample  $(S, y)$  that is not necessarily realizable; that is, there may be no  $H^* \in \mathcal{H}_{\text{MP}}(G)$  that is consistent with  $(S, y)$ . The classic goal is to compute an ERM (empirical risk minimizer), i.e., some  $H_{\text{ERM}} \in \mathcal{H}_{\text{MP}}$  that minimizes the empirical risk  $L(H, (S, y)) = \frac{1}{|S|} \sum_{s \in S} \mathbb{1}\{y(s) \neq \mathbb{1}_{s \in H}\}$  over all  $H \in \mathcal{H}_{\text{MP}}$ . As a further corollary of our results above we obtain:

**Corollary 10 (ERM is FPT)** *Given  $G = (V, E)$  and a labeled sample  $(S, y)$ , one can compute  $H_{\text{ERM}} \in \arg \min_{H \in \mathcal{H}_{\text{MP}}} \frac{1}{|S|} \sum_{s \in S} \mathbb{1}\{y(s) \neq \mathbb{1}_{s \in H}\}$  in time  $|\mathcal{H}_{\text{MP}}| \text{poly}(n) \leq 2^{\omega(G)} \text{poly}(n)$ .*

**Proof** We list all m-halfspaces using MH-LIST and return any of the listed m-halfspaces that minimizes the number of mistakes on  $S$ . ■

Similarly to the realizable case, we get the following agnostic PAC bound.

**Theorem 3 (Agnostic PAC is FPT)** *There exists an algorithm that, for any  $\varepsilon, \delta \in (0, 1)$ , can  $(\varepsilon, \delta)$ -PAC learn  $(V, \mathcal{H}_{\text{MP}})$  in the agnostic setting using  $\mathcal{O}\left(\frac{\omega(G) + \log(1/\delta)}{\varepsilon^2}\right)$  labeled samples and time  $|\mathcal{H}_{\text{MP}}| \text{poly}(n) \leq 2^{\omega(G)} \text{poly}(n)$ .*

**Proof** It is well known that an ERM yields the optimal agnostic sample complexity of the order  $\mathcal{O}\left(\frac{\text{VC}(V, \mathcal{H}_{\text{MP}}) + \log(1/\delta)}{\varepsilon^2}\right)$  (Vapnik and Chervonenkis, 1974; Talagrand, 1994). The sample complexity follows by  $\text{VC}(V, \mathcal{H}_{\text{MP}}) \leq \tilde{\omega}(G)$  (Proposition 4). We can compute an ERM in time  $|\mathcal{H}_{\text{MP}}| \text{poly}(n) \leq 2^{\omega(G)} \text{poly}(n)$  by Corollary 10. ■

## Appendix E. Active Learning

This section presents an efficient algorithm for the problem of actively learning monophonic halfspaces, defined as follows. The algorithm is given a graph  $G = (V, E)$ , and nature selects a concept  $H \in \mathcal{H}_{\text{MP}}(G)$ . The algorithm can query any vertex  $x \in V$  for its label  $\mathbb{1}_{x \in H}$ . The goal of the algorithm is to output  $H$  by making as few queries as possible. This problem is a special case of realizable transductive active learning on a graph (Afshani et al., 2007; Guillory and Bilmes, 2009; Cesa-Bianchi et al., 2010; Dasarathy et al., 2015), and can be seen as a variant of query learning (Angluin, 1988; Hegedűs, 1995) with an additional fixed graph given where vertices correspond to the instance space. Now let  $\mathcal{H} \subseteq 2^V$  be a concept class. The *query complexity* of  $(V, \mathcal{H})$  is the maximum number of queries an optimal algorithm makes over  $H \in \mathcal{H}$ . More precisely, for any algorithm  $A$  and any  $H \in \mathcal{H}$ , let  $\text{qc}(A, H)$  denote the number of queries  $A$  make on  $G$  when nature chooses  $H$ . The query complexity of  $A$  on  $\mathcal{H}$  is  $\text{qc}(A, \mathcal{H}) = \max_{H \in \mathcal{H}} \text{qc}(A, H)$ . The query complexity of  $\mathcal{H}$  is  $\text{qc}(\mathcal{H}) = \min_A \text{qc}(A, \mathcal{H})$ .

Relying on Lemma 17 we get the main result in this section: a polynomial-time algorithm achieving a graph-dependent near-optimal query complexity.

**Theorem 11 (Poly-time active learning)** *It is possible to actively learn  $(V, \mathcal{H}_{\text{MP}})$  in time  $\text{poly}(n)$  with query complexity  $\mathcal{O}(h(G) + \log \text{diam}_g(G) + \omega(G))$ .*

**Proof of Theorem 11.** We compute a minimum monophonic hull set  $S$  of size  $h(G)$  in polynomial time using the algorithm of Dourado et al. (2010). We query all vertices of  $S$ . If they have all the same label, all vertices in the graph have the same label, by the definition of hull sets. Otherwise, we take two vertices with different labels and find a shortest path, which is induced, in particular, between them. As each label class corresponds to a halfspace, the path has exactly one cut edge  $e = \{u, v\}$ . We identify  $e$  using  $\mathcal{O}(\log \text{diam}_g(G))$  queries through binary search on the path. Using the labels of  $u$  and  $v$  we can determine the labels of  $u/v$  (same as  $u$ ) and  $v/u$  (same as  $v$ ). We compute these sets efficiently using Lemma 6. We only have to query one vertex in each connected



component of  $\neg G[\Delta_{uv}]$ , which are at most  $\omega(G)$  by Lemma 19. By Lemma 17 the labels of the remaining vertices can be determined through the  $m$ -shadows of vertices in  $\Delta_{uv}$  and  $u$  or  $v$ . ■

Along the previously mentioned separation axioms from Definition 1 we achieve increasingly tighter lower bounds on the query complexity, eventually matching our algorithmic upper bound from Theorem 11 for all  $S_4$  graphs, the strongest separability assumption.

**Theorem 12** *Let  $G$  be a graph. The following holds for the query complexity  $\text{qc}(\mathcal{H}_{\text{MP}}(G))$ :*

- if  $G \in S_2$ , then  $\text{qc}(\mathcal{H}_{\text{MP}}(G)) \geq \Omega(\log \text{diam}_m(G))$ ,
- if  $G \in S_3$ , then  $\text{qc}(\mathcal{H}_{\text{MP}}(G)) \geq \Omega(\log \text{diam}_m(G) + h(G))$ , and
- if  $G \in S_4$ , then  $\text{qc}(\mathcal{H}_{\text{MP}}(G)) \geq \Omega(\log \text{diam}_m(G) + h(G) + \omega(G))$ .

**Proof of Theorem 12.** For  $S_2$  graphs note that any edge on a fixed induced path  $P$  can be a cut edge of a  $m$ -halfspace. Locating this cut edge takes  $\Omega(\log |V(P)|)$  queries in the worst case. For  $S_3$  graphs take any minimum hull set  $S$  (of size  $h(G)$ ). Having queried only a proper subset  $S' \subsetneq S$ , the algorithm cannot decide between the halfspace  $V$  and a halfspace separating  $\text{conv}(S')$  and some  $s \in S \setminus S'$ . Note that by minimality of  $S$ , the vertex  $s \notin \text{conv}(S')$ . For  $S_4$  graphs any clique can be shattered by halfspaces. Hence, we fix a clique of maximum size and force the learner to query it completely. ■

## Appendix F. Online Learning

In this section, we devise near-optimal efficient algorithms for the problem of online learning of monophonic halfspaces. The classical realizable online learning problem of Littlestone (1988) (while the agnostic/unrealizable setting has been studied by, e.g., Ben-David et al. (2009)) can be modelled as an iterative game between our learner and the environment over a finite number  $T$  of rounds. The instance space  $V$  and a hypothesis space  $\mathcal{H} \subseteq 2^V$  is known and fixed. First, the environment chooses a hypothesis  $h$  from a hypothesis space  $\mathcal{H}$ . Then, in each round  $t = 1, \dots, T$ :

1. the environment chooses a point  $v_t \in V$ ,
2. the learner predicts the label  $\hat{y}_t \in \{0, 1\}$  of  $v_t$ ,
3. the environment reveals the true label  $y_t = \mathbb{1}_{v_t \in h}$ ,
4. the learner made a mistake if  $\hat{y}_t \neq y_t$ .

The goal of the learner is to minimize the total number of mistakes. More precisely let  $A$  be an algorithm for this online learning problem. Then, let  $M(A, h)$  for  $h \in \mathcal{H}$  denote the worst-case number of mistakes  $A$  would make on any sequence labeled by  $h$ . The mistake bound of  $A$  on  $\mathcal{H}$  is thus defined as  $M(A, \mathcal{H}) = \max_{h \in \mathcal{H}} M(A, h)$  and we are interested in the optimal mistake bound  $M(\mathcal{H}) = \min_A M(A, \mathcal{H})$ , also known as the Littlestone dimension of  $\mathcal{H}$ .

The node classification (or node labeling) variant of this problem is well studied (Herbster et al., 2005; Cesa-Bianchi et al., 2013; Herbster et al., 2015). As in the active learning variant the main parameter is the (potentially effective resistance weighted) cutsize, linearly determining the mistake bounds. In this section, we mainly study the variant of the above realizable online learning problem over our hypothesis class  $\mathcal{H}_{\text{MP}}$ .

We prove the main result in this section (Theorem 13) below, discussing the strategy using Halving and Winnow individually.

### F.1. Halving over Monophonic Halfspaces

The Halving algorithm is a classical algorithm for online binary classification over a finite hypothesis class, used by [Littlestone \(1988\)](#) for providing mistake bounds, which implements an “iterative halving” procedure. The main idea behind Halving is to maintain a set of hypotheses from the original class  $\mathcal{H}$  that are consistent with the labels observed so far; these consistent hypotheses form the so called *version space*. The version space of Halving begins with the entire class  $\mathcal{H}$  and it gradually shrinks while guaranteeing that the target hypothesis fixed by the environment belongs to the version space (by consistency). The prediction of Halving at each round consists of a majority vote for the current point over the entire version space.

Our realizable online learning problem perfectly fits the requirements for running Halving, especially considering the fact that we can construct the nontrivial hypothesis class  $\mathcal{H}_{\text{MP}}(G)$  by running MH-FPT-LIST (Algorithm 3). Therefore, we can derive the following guarantees as a consequence of well-known results on Halving (e.g., [Littlestone \(1988, Theorem 1\)](#)).

**Proposition 23** *Halving applied to  $(V, \mathcal{H}_{\text{MP}}(G))$  achieves a mistake bound of  $\mathcal{O}\left(\omega(G) + \log \frac{n}{\omega(G)}\right)$  and can be run in time  $2^{\omega(G)} \text{poly}(n)$  per round.*

A straightforward implementation of Halving requires the enumeration of the version space in each round. By Theorem 9, we know that  $|\mathcal{H}_{\text{MP}}| \leq \frac{4m2^{\omega(G)}}{\omega(G)} + 2$  and that this inequality can be almost tight. Hence, if we run Halving by enumerating the version space we cannot hope for a better runtime. Halving’s mistake bound is given by the fact that it corresponds to  $\mathcal{O}(\log |\mathcal{H}_{\text{MP}}|)$ .

### F.2. Winnow over Monophonic Halfspaces

The main downside of running the Halving algorithm is its running time. In general we cannot hope to improve upon it unless we can manage to further restrict the initial size of the hypothesis class. Nonetheless, if the target hypothesis can be represented as a sparse disjunction, the Winnow algorithm ([Littlestone, 1988](#)) can achieve a similar mistake bound as Halving without the need to enumerate the version space. A similar usage of Winnow for online node classification problems was discussed by [Gentile et al. \(2013\)](#).

**Proposition 24 ([Littlestone \(1988\)](#))** *Winnow achieves a mistake bound of  $\mathcal{O}(k \log d)$  in  $\mathcal{O}(\text{poly}(kd))$  time to online learn monotone  $k$ -literal disjunctions on  $d$  variables.*

By leveraging on this property of the Winnow algorithm, we can show the following result.

**Theorem 25** *Online learning of monophonic halfspaces is possible in time  $\text{poly}(n)$  with a mistake bound of  $\mathcal{O}(\omega(G) \log n)$ .*

**Proof** Let  $E = \{\{u_1, v_1\}, \dots, \{u_m, v_m\}\}$  and  $S = (u_1/v_1, \dots, u_m/v_m, v_1/u_1, \dots, v_m/u_m)$ . Each  $S_i = u_i/v_i$  can be computed in polynomial time by computing the distance matrix once (all pairs shortest path). We represent each vertex  $v \in V$  as a binary vector  $\phi(v) \in \{0, 1\}^{2m}$  (where, we recall,  $m = |E(G)|$ ) such that  $\phi(v)_i = \mathbb{1}_{\{v \in S_i\}}$ . Using Lemma 18, we know that any monophonic halfspace can be represented as the union of at most  $k \leq \omega(G)$  sets in the sequence  $S$ . Hence, any monophonic halfspace can be represented as a disjunction

$$\phi(\cdot)_{i_1} \vee \dots \vee \phi(\cdot)_{i_k} \tag{21}$$

of at most  $k \leq \omega(G)$  of the  $2m$  variables in  $\phi(\cdot)$ . Thus, we can run Winnow on  $\phi(V) = \{\phi(v) \mid v \in V\}$  and achieve a mistake bound of  $\mathcal{O}(\omega(G) \log n)$  where we note that  $\log(2m) = \mathcal{O}(\log n)$ . As each step of Winnow takes time  $\mathcal{O}(m) = \mathcal{O}(n^2)$ , the overall runtime is dominated by the construction of  $\phi(V)$ . ■

Winnow nearly achieves the mistake bound as Halving in Theorem 23 (worse by only a logarithmic factor), yet in polynomial time without the need to enumerate the version space. The achieved mistakes is near-optimal. In particular we have the following simple result.

**Proposition 15** *For any  $S_4$  graph  $G$ , it holds that  $M(\mathcal{H}_{\text{MP}}(G)) \geq \omega(G)$ .*

**Proof of Proposition 15.** For  $S_4$  graphs  $\omega(G) \leq \text{VC}(V, \mathcal{H}_{\text{MP}}(G))$  by Proposition 4, which is a lower bound on the optimal number of mistakes (Littlestone, 1988). ■

### E.3. Agnostic Online Learning

Compared to the realizable setting studied so far, in the agnostic case we have no guarantee that there exists a hypothesis  $H \in \mathcal{H}_{\text{MP}}$  that makes no mistakes with respect to the true labels. In other words, the true labels revealed by the environment during the online learning process can be determined according to any (fixed) arbitrary  $f \in 2^V$ . Even though we now cannot show generally good mistake bounds as the ones in the realizable setting, we can at least provide guarantees on the performance with respect to the best fixed hypothesis, i.e.,  $H^* \in \arg \min_{H \in \mathcal{H}_{\text{MP}}(G)} M(H, f)$ . Indeed, if we let  $M^* = M(H^*, f)$  be the minimum number of mistakes achieved by some hypothesis in  $\mathcal{H}_{\text{MP}}(G)$ , we can design algorithms  $A$  whose mistake bound  $M(A, f)$  is a function of  $M^*$ . We can show, in particular, that  $M(A, f) - M^*$  can be upper bounded to provide guarantees for online learning in the agnostic case.<sup>3</sup>

In order to achieve such guarantees, though, we can easily incur in large running times if we adopt one of the common strategies for agnostic online learning. For example, adopting a learning-with-expert-advice algorithm  $A$  such as Hedge (Cesa-Bianchi et al., 1997)—also see, e.g., Cesa-Bianchi and Lugosi (2006, Chapter 2)—where the experts are the hypotheses  $\mathcal{H}_{\text{MP}}(G)$  listed by MH-FPT-LIST, would provide an expected mistake bound of

$$\mathbb{E}[M(A, f)] = M^* + \mathcal{O}\left(\sqrt{T \log |\mathcal{H}_{\text{MP}}(G)|}\right) = M^* + \mathcal{O}\left(\sqrt{T\left(\omega(G) + \log \frac{m}{\omega(G)}\right)}\right) \quad (22)$$

for any  $f \in 2^V$ , without prior knowledge on  $\omega(G)$  or  $M^*$ . Moreover, one could provide an improved guarantee with a finer tuning for the learning rate (Ben-David et al., 2009) of  $A$  to further show that

$$\mathbb{E}[M(A, f)] \leq M^* + \sqrt{2M^* \ln |\mathcal{H}_{\text{MP}}(G)|} + \ln |\mathcal{H}_{\text{MP}}(G)| \quad (23)$$

$$= M^* + \mathcal{O}\left(\sqrt{M^*\left(\omega(G) + \log \frac{m}{\omega(G)}\right)} + \omega(G) + \log \frac{m}{\omega(G)}\right), \quad (24)$$

which can similarly be achieved by lifting prior knowledge on  $M^*$  via a doubling trick. However, approaches that implement learning with expert advice in a straightforward way would require to

3. Note that  $M(A, f) - M^*$ , or better  $\mathbb{E}[M(A, f)] - M^*$  if  $A$  is randomized, takes the name of (expected) *regret*.

maintain weights over all  $|\mathcal{H}_{\text{MP}}(G)| \leq \frac{4m2^{\omega(G)}}{\omega(G)} + 2$  hypotheses, thus incurring in  $\mathcal{O}(2^{\omega(G)} \text{poly}(n))$  running time per round.

If we instead run the Winnow algorithm over the same task, it is possible to achieve similar rates (only losing a factor of roughly  $\omega(G)$ ) for the mistake bound while keeping computational efficiency even in the agnostic case. As shown by [Blum \(1996\)](#), Winnow has the following mistake bound in the agnostic case.

**Proposition 26 ([Blum \(1996, Theorem 6\)](#))** *For any sequence of examples consistent with any  $f \in 2^{[n]}$  and any  $r$ -literal disjunction  $h$  over  $n$  variables, the number of mistakes made by Winnow is  $\mathcal{O}(r(M(h, f) + \log n))$ .*

This immediately implies [Theorem 14](#) for agnostic online learning with  $\mathcal{H}_{\text{MP}}(G)$ .