# Active Learning on Trees and Graphs 

Nicolò Cesa-Bianchi<br>Università degli Studi di Milano Italy<br>nicolo.cesa-bianchi@unimi.it

Claudio Gentile<br>Università dell'Insubria<br>Italy<br>claudio.gentile@uninsubria.it

Fabio Vitale Giovanni Zappella<br>Università degli Studi di Milano<br>Italy<br>fabio.vitale@unimi.it<br>giovanni.zappella@studenti.unimi.it


#### Abstract

We investigate the problem of active learning on a given tree whose nodes are assigned binary labels in an adversarial way. Inspired by recent results by Guillory and Bilmes, we characterize (up to constant factors) the optimal placement of queries so to minimize the mistakes made on the non-queried nodes. Our query selection algorithm is extremely efficient, and the optimal number of mistakes on the non-queried nodes is achieved by a simple and efficient mincut classifier. Through a simple modification of the query selection algorithm we also show optimality (up to constant factors) with respect to the trade-off between number of queries and number of mistakes on nonqueried nodes. By using spanning trees, our algorithms can be efficiently applied to general graphs, although the problem of finding optimal and efficient active learning algorithms for general graphs remains open. Towards this end, we provide a lower bound on the number of mistakes made on arbitrary graphs by any active learning algorithm using a number of queries which is up to a constant fraction of the graph size.


## 1 Introduction

The abundance of networked data in various application domains (web, social networks, bioinformatics, etc.) motivates the development of scalable and accurate graph-based prediction algorithms. An important topic in this area is the graph binary classification problem: Given a graph with unknown binary labels on its nodes, the learner receives the labels on a subset of the nodes (the training set) and must predict the labels on the remaining vertices. This is typically done by relying on some notion of label regularity depending on the graph topology, such as that nearby nodes are likely to be labeled similarly. Standard approaches to this problem predict with the assignment of labels minimizing the induced cutsize (e.g., [4, 5]), or by binarizing the assignment that minimizes certain real-valued extensions of the cutsize function (e.g., $[10,2,3]$ and references therein).

In the active learning version of this problem the learner is allowed to choose the subset of training nodes. Similarly to standard feature-based learning, one expects active methods to provide a significant boost of predictive ability compared to a noninformed (e.g., random) draw of the training set. The following simple example provides some intuition of why this could happen when the labels are chosen by an adversary, which is the setting considered in this paper. Consider a "binary star system" of two star-shaped graphs whose centers are connected by a bridge, where one star is a constant fraction bigger than the other. The adversary draws two random binary labels and assigns the first label to all nodes of the first star graph, and the second label to all nodes of the second star graph. Assume that the training set size is two. If we choose the centers of the two stars and predict with a mincut strategy, ${ }^{1}$ we are guaranteed to make zero mistakes on all unseen vertices. On the other hand, if we query two nodes at random, then with constant probability both of them will belong to the bigger star, and all the unseen labels of the smaller star will be mistaken. This simple example shows that the gap between the performance of passive and active learning on graphs can be made arbitrarily big.

In general, one would like to devise a strategy for placing a certain budget of queries on the vertices of a given graph. This should be done so as to minimize the number of mistakes made on the non-queried nodes by some reasonable classifier like mincut. This question has been investigated from a theoretical viewpoint

[^0]by Guillory and Bilmes [6], and by Afshani et al. [1]. Our work is related to an elegant result from [6] which bounds the number of mistakes made by the mincut classifier on the worst-case assignment of labels in terms of $\Phi / \Psi(L)$. Here $\Phi$ is the cutsize induced by the unknown labeling, and $\Psi(L)$ is a function of the query (or training) set $L$, which depends on the structural properties of the (unlabeled) graph. For instance, in the above example of the binary system, the value of $\Psi(L)$ when the query set $L$ includes just the two centers is 1 . This implies that for the binary system graph, Guillory and Bilmes' bound on the mincut strategy is $\Phi$ mistakes in the worst case (note that in the above example $\Phi \leq 1$ ). Since $\Psi(L)$ can be efficiently computed on any given graph and query set $L$, the learner's task might be reduced to finding a query set $L$ that maximizes $\Psi(L)$ given a certain query budget (size of $L$ ). Unfortunately, no feasible general algorithm for solving this maximization problem is known, and so one must resort to heuristic methods - see [6].

In this work we investigate the active learning problem on graphs in the important special case of trees. We exhibit a simple iterative algorithm which, combined with a mincut classifier, is optimal (up to constant factors) on any given labeled tree. This holds even if the algorithm is not given information on the actual cutsize $\Phi$. Our method is extremely efficient, requiring $\mathcal{O}(n \ln Q)$ time for placing $Q$ queries in an $n$-node tree, and space linear in $n$. As a byproduct of our analysis, we show that $\Psi$ can be efficiently maximized over trees to within constant factors. Hence the bound $\min _{L} \Phi / \Psi(L)$ can be achieved efficiently.

Another interesting question is what kind of trade-off between queries and mistakes can be achieved if the learner is not constrained by a given query budget. We show that a simple modification of our selection algorithm is able to trade-off queries and mistakes in an optimal way up to constant factors.

Finally, we prove a general lower bound for predicting the labels of any given graph (not necessarily a tree) when the query set is up to a constant fraction of the number of vertices. Our lower bound establishes that the number of mistakes must then be at least a constant fraction of the cutsize weighted by the effective resistances. This lower bound apparently yields a contradiction to the results of Afshani et al. [1], who constructs the query set adaptively. This apparent contradiction is also obtained via a simple counterexample that we detail in Section 5.

## 2 Preliminaries and basic notation

A labeled tree $(T, \boldsymbol{y})$ is a tree $T=(V, E)$ whose nodes $V=\{1, \ldots, n\}$ are assigned binary labels $\boldsymbol{y}=$ $\left(y_{1}, \ldots, y_{n}\right) \in\{-1,+1\}^{n}$. We measure the label regularity of $(T, \boldsymbol{y})$ by the cutsize $\Phi_{T}(\boldsymbol{y})$ induced by $\boldsymbol{y}$ on $T$, i.e., $\Phi_{T}(\boldsymbol{y})=\left|\left\{(i, j) \in E: y_{i} \neq y_{j}\right\}\right|$. We consider the following active learning protocol: given a tree $T$ with unknown labeling $\boldsymbol{y}$, the learner obtains all labels in a query set $L \subseteq V$, and is then required to predict the labels of the remaining nodes $V \backslash L$. Active learning algorithms work in two-phases: a selection phase, where a query set of given size is constructed, and a prediction phase, where the algorithm receives the labels of the query set and predicts the labels of the remaining nodes. Note that the only labels ever observed by the algorithm are those in the query set. In particular, no labels are revealed during the prediction phase.

We measure the ability of the algorithm by the number of prediction mistakes made on $V \backslash L$, where it is reasonable to expect this number to depend on both the uknown cutsize $\Phi_{T}(\boldsymbol{y})$ and the number $|L|$ of requested labels. A slightly different prediction measure is considered in Section 4.3.

Given a tree $T$ and a query set $L \subseteq V$, a node $i \in V \backslash L$ is a fork node generated by $L$ if and only if there exist three distinct nodes $i_{1}, i_{2}, i_{3} \in L$ that are connected to $i$ through edge disjoint paths. Let $\operatorname{FORK}(L)$ be the set of all fork nodes generated by $L$. Then $L^{+}$is the query set obtained by adding to $L$ all the generated fork nodes, i.e., $L^{+} \triangleq L \cup \operatorname{FORK}(L)$. We say that $L \subseteq V$ is 0 -forked iff $L^{+} \equiv L$. Note that $L^{+}$is 0 -forked. That is, $\operatorname{FORK}\left(L^{+}\right) \equiv \emptyset$ for all $L \subseteq V$.

Given a node subset $S \subseteq V$, we use $T \backslash S$ to denote the forest obtained by removing from the tree $T$ all nodes in $S$ and all edges incident to them. Moreover, given a second tree $T^{\prime}$, we denote by $T \backslash T^{\prime}$ the forest $T \backslash V^{\prime}$, where $V^{\prime}$ is the set of nodes of $T^{\prime}$. Given a query set $L \subseteq V$, a hinge-tree is any connected component of $T \backslash L^{+}$. We call connection node of a hinge-tree a node of $L$ adjacent to any node of the hinge tree. We distinguish between 1-hinge and 2-hinge trees. A 1-hinge-tree has one connection node only, whereas a 2-hinge-tree has two (note that a hinge tree cannot have more than two connection nodes because $L^{+}$is zero-forked, see Figure 1).

## 3 The active learning algorithm

We now describe the two phases of our active learning algorithm. For the sake of exposition, we call SEL the selection phase and PRED the prediction phase. SEL returns a 0 -forked query set $L_{\text {SEL }}^{+} \subseteq V$ of desired size. PRED takes in input the query set $L_{\text {SEL }}^{+}$and the set of labels $y_{i}$ for all $i \in L_{\text {SEL }}^{+}$. Then PRED returns a prediction for the labels of all remaining nodes $V \backslash L_{\text {SEL }}^{+}$.

In order to see the way SEL operates, we formally introduce the function $\Psi^{*}$. This is the reciprocal of the $\Psi$ function introduced in [6] and mentioned in Section 1.


Figure 1: A tree $T=(V, E)$ whose nodes are shaded (the query set $L$ ) or white (the set $V \backslash L$ ). The shaded nodes are also the connection nodes of the depicted hinge trees (not all hinge trees are contoured). The fork nodes generated by $L$ are denoted by double circles. The thick black edges connect the nodes in $L$.

Definition 1 Given a tree $T=(V, E)$ and a set of nodes $L \subseteq V$,

$$
\Psi^{*}(L) \triangleq \max _{\emptyset \neq V^{\prime} \subseteq V \backslash L} \frac{\left|V^{\prime}\right|}{\left|\left\{(i, j) \in E: i \in V^{\prime}, j \in V \backslash V^{\prime}\right\}\right|}
$$

In words, $\Psi^{*}(L)$ measures the largest set of nodes not in $L$ that share the least number of edges with nodes in $L$. From the adversary's viewpoint, $\Psi^{*}(L)$ can be described as the largest return in mistakes per unit of cutsize invested. We now move on to the description of the algorithms SEL and PRED.

The selection algoritm SEL greedily computes a query set that minimizes $\Psi^{*}$ to within constant factors. To this end, sEL exploits Lemma 10 (a) (see Section 4.2) stating that, for any fixed query set $L$, the subset $V^{\prime} \subseteq V$ maximizing $\frac{\left|V^{\prime}\right|}{\left|\left\{(i, j) \in E: i \in V^{\prime}, j \in V \backslash V^{\prime}\right\}\right|}$ is always included in a connected component of $T \backslash L$. Thus SEL places its queries in order to end up with a query set $L_{\text {SEL }}^{+}$such that the largest component of $T \backslash L_{\text {SEL }}^{+}$is as small as possible.

SEL operates as follows. Let $L_{t} \subseteq L$ be the set including the first $t$ nodes chosen by SEL, $T_{\max }^{t}$ be the largest connected component of $T \backslash L_{t-1}$, and $\sigma\left(T^{\prime}, i\right)$ be the size (number of nodes) of the largest component of the forest $T^{\prime} \backslash\{i\}$, where $T^{\prime}$ is any tree. At each step $t=1,2, \ldots$, SEL simply picks the node $i_{t} \in T_{\max }^{t}$ that minimizes $\sigma\left(T_{\max }^{t}, i\right)$ over $i$ and sets $L_{t}=L_{t-1} \cup\left\{i_{t}\right\}$. During this iterative construction, SEL also maintains a set containing all fork nodes generated in each step by adding nodes $i_{t}$ to the sets $L_{t-1} .{ }^{2}$ After the desired number of queries is reached (also counting the queries that would be caused by the stored fork nodes), SEL has terminated the construction of the query set $L_{\text {SEL }}$. The final query set $L_{\text {SEL }}^{+}$, obtained by adding all stored fork nodes to $L_{\text {SEL }}$, is then returned.

The Prediction Algorithm PRED receives in input the labeled nodes of the 0-forked query set $L_{\text {SEL }}^{+}$and computes a mincut assignment. Since each component of $T \backslash L_{\text {SEL }}^{+}$is either a 1-hinge-tree or a 2-hinge-tree, PRED is simple to describe and is also very efficient. The algorithm predicts all the nodes of hinge-tree $\mathcal{T}$ using the same label $\widehat{y}_{\mathcal{T}}$. This label is chosen according to the following two cases:

1. If $\mathcal{T}$ is a 1-hinge-tree, then $\widehat{y}_{\mathcal{T}}$ is set to the label of its unique connection node;
2. If $\mathcal{T}$ is a 2-hinge-tree and the labels of its two connection nodes are equal, then $\widehat{y}_{\mathcal{T}}$ is set to the label of its connection nodes, otherwise $\widehat{y}_{\mathcal{T}}$ is set as the label of the closer connection node (ties are broken arbitrarily).
In Section 6 we show that SEL requires overall $\mathcal{O}(|V| \log Q)$ time and $\mathcal{O}(|V|)$ memory space for selecting $Q$ query nodes. Also, we will see that the total running time taken by PRED for predicting all nodes in $V \backslash L$ is linear in $|V|$.

## 4 Analysis

For a given tree $T$, we denote by $m_{A}(L, \boldsymbol{y})$ the number of prediction mistakes that algorithm $A$ makes on the labeled tree $(T, \boldsymbol{y})$ when given the query set $L$. Introduce the function

$$
m_{A}(L, K)=\max _{\boldsymbol{y}: \Phi_{T}(\boldsymbol{y}) \leq K} m_{A}(L, \boldsymbol{y})
$$

[^1]



Figure 2: The SEL algorithm at work. The upper pane shows the initial tree $T=T_{\max }^{1}$ (in the box tagged with " 1 "), and the subsequent subtrees $T_{\max }^{2}, T_{\max }^{3}, T_{\max }^{4}$, and $T_{\max }^{5}$. The left pane also shows the nodes selected by SEL in chronological order. The four lower panes show the connected components of $T \backslash L_{t}$ resulting from this selection. Observe that at the end of round 3, SEL detects the generation of fork node $3^{\prime}$. This node gets stored, and is added to $L_{\text {SEL }}$ at the end of the selection process.
denoting the number of prediction mistakes made by $A$ with query set $L$ on all labeled trees with cutsize bounded by $K$. We will also find it useful to deal with the "lower bound" function $\operatorname{LB}(L, K)$. This is the maximum expected number of mistakes that any prediction algorithm $A$ can be forced to make on the labeled tree $(T, \boldsymbol{y})$ when the query set is $L$ and the cutsize is not larger than $K$.

We show that the number of mistakes made by PRED on any labeled tree when using the query set $L_{\text {SEL }}^{+}$ satisfies

$$
m_{\mathrm{PRED}}\left(L_{\mathrm{SEL}}^{+}, K\right) \leq 10 \mathrm{LB}(L, K)
$$

for all query sets $L \subseteq V$ of size up to $\frac{1}{8}\left|L_{\text {SEL }}^{+}\right|$. Though neither SEL nor PRED do know the actual cutsize of the labeled tree $(T, \boldsymbol{y})$, the combined use of these procedures is competitive against any algorithm that knows the cutsize budget $K$ beforehand.

While this result implies the optimality (up to constant factors) of our algorithm, it does not relate the mistake bound to the cutsize, which is a clearly interpretable measure of the label regularity. In order to address this issue, we show that our algorithm also satisfies the bound

$$
m_{\text {PRED }}\left(L_{\mathrm{SEL}}^{+}, \boldsymbol{y}\right) \leq 4 \Psi^{*}(L) \Phi_{T}(\boldsymbol{y})
$$

for all query sets $L \subseteq V$ of size up to $\frac{1}{8}\left|L_{\text {SEL }}^{+}\right|$. The proof of these results needs a number of preliminary lemmas.

Lemma 2 For any tree $T=(V, E)$ it holds that $\min _{v \in V} \sigma(T, v) \leq \frac{1}{2}|V|$.
Proof: Let $i \in \operatorname{argmin}_{v \in V} \sigma(T, v)$. For the sake of contradiction, assume there exists a component $T_{i}=$ $\left(V_{i}, E_{i}\right)$ of $T \backslash\{i\}$ such that $\left|V_{i}\right|>|V| / 2$. Let $s$ be the sum of the sizes all other components. Since
$\left|V_{i}\right|+s=|V|-1$, we know that $s \leq|V| / 2-1$. Now let $j$ be the node adjacent to $i$ which belongs to $V_{i}$ and $T_{j}=\left(V_{j}, E_{j}\right)$ be the largest component of $T \backslash\{j\}$. There are only two cases to consider: either $V_{j} \subset V_{i}$ or $V_{j} \cap V_{i} \equiv \emptyset$. In the first case, $\left|V_{j}\right|<\left|V_{i}\right|$. In the second case, $V_{j} \subseteq\{i\} \cup\left(T \backslash V_{i}\right)$, which implies $\left|V_{j}\right| \leq 1+s \leq|V| / 2<\left|V_{i}\right|$. In both cases, $i \notin \operatorname{argmin}_{v \in V} \sigma(T, v)$, which provides the desired contradiction.

Lemma 3 For all subsets $L \subset V$ of the nodes of a tree $T=(V, E)$ we have $\left|L^{+}\right| \leq 2|L|$.
Proof: Pick an arbitrary node of $T$ and perform a depth-first visit of all nodes in $T$. This visit induces an ordering $\mathcal{T}_{1}, \mathcal{I}_{2}, \ldots$ of the connected components in $T \backslash L$ based on the order of the nodes visited first in each component. Now let $\mathcal{T}_{1}^{\prime}, \mathcal{T}_{2}^{\prime}, \ldots$ be such that each $\mathcal{T}_{i}^{\prime}$ is a component of $\mathcal{T}_{i}$ extended to include all nodes of $L$ adjacent to nodes in $\mathcal{T}_{i}$. Then the ordering implies that, for $i \geq 2, \mathcal{T}_{i}^{\prime}$ shares exactly one node (which must be a leaf) with all previously visited trees. Since in any tree the number of nodes of degree larger than two must be strictly smaller than the number of leaves, we have $\left|\operatorname{FORK}\left(\mathcal{T}_{i}^{\prime}\right)\right|<\left|\Lambda_{i}\right|$ where, with slight abuse of notation, we denote by $\operatorname{FORK}\left(\mathcal{T}_{i}^{\prime}\right)$ the set of all fork nodes in subtree $\mathcal{T}_{i}^{\prime}$. Also, we let $\Lambda_{i}$ be the set of leaves of $\mathcal{T}_{i}^{\prime}$. This implies that, for $i=1,2, \ldots$, each fork node in $\operatorname{FORK}\left(\mathcal{T}_{i}^{\prime}\right)$ can be injectively associated with one of the $\left|\Lambda_{i}\right|-1$ leaves of $\mathcal{T}_{i}^{\prime}$ that are not shared with any of the previously visited trees. Since $|\operatorname{FORK}(L)|$ is equal to the sum of $\left|\operatorname{FORK}\left(\mathcal{T}_{i}\right)\right|$ over all indices $i$, this implies that $|\operatorname{FORK}(L)| \leq|L|$.

Lemma 4 Let $L_{t-1} \subseteq L_{\text {SEL }}$ be the set of the first $t-1$ nodes chosen by SEL. Given any tree $T=(V, E)$, the largest subtree of $T \backslash L_{t-1}$ contains no more than $\frac{2}{t}|V|$ nodes.

Proof: Recall that $i_{s}$ denotes the $s$-th node selected by SEL during the incremental construction of the query set $L_{\text {SEL }}$, and that $T_{\max }^{s}$ is the largest component of $T \backslash L_{s-1}$. The first $t$ steps of the recursive splitting procedure performed by SEL can be associated with a splitting tree $T^{\prime}$ defined in the following way. The internal nodes of $T^{\prime}$ are $T_{\max }^{s}$, for $s \geq 1$. The children of $T_{\max }^{s}$ are the connected components of $T_{\max }^{s} \backslash\left\{i_{s}\right\}$, i.e., the subtrees of $T_{\max }^{s}$ created by the selection of $i_{s}$. Hence, each leaf of $T^{\prime}$ is bijectively associated with a tree in $T \backslash L_{t}$.

Let $T_{\text {nol }}^{\prime}$ be the tree obtained from $T^{\prime}$ by deleting all leaves. Each node of $T_{\text {nol }}^{\prime}$ is one of the $t$ subtrees split by SEL during the construction of $L_{t}$. As $T_{\max }^{t}$ is split by $i_{t}$, it is a leaf in $T_{\text {nol }}^{\prime}$. We now add a second child to each internal node $s$ of $T_{\text {nol }}^{\prime}$ having a single child. This second child of $s$ is obtained by merging all the subtrees belonging to leaves of $T^{\prime}$ that are also children of $s$. Let $T^{\prime \prime}$ be the resulting tree.

We now compare the cardinality of $T_{\max }^{t}$ to that of the subtrees associated with the leaves of $T^{\prime \prime}$. Let $\Lambda$ be the set of all leaves of $T^{\prime \prime}$ and $\Lambda_{\text {add }}=T^{\prime \prime} \backslash T_{\text {nol }}^{\prime} \subset \Lambda$ be the set of all leaves added to $T_{\text {nol }}^{\prime}$ to obtain $T^{\prime \prime}$. First of all, note that $\left|T_{\max }^{t}\right|$ is not larger than the number of nodes in any leaf of $T_{\text {nol }}^{\prime}$. This is because the selection rule of SEL ensures that $T_{\max }^{t}$ cannot be larger than any subtree associated with a leaf in $T_{\text {nol }}^{\prime}$, since it contains no node selected before time $t$. In what follows, we write $|s|$ to denote the size of the forest or subtree associated with a node $s$ of $T^{\prime \prime}$. We now prove the following claim:
Claim. For all $\ell \in \Lambda,\left|T_{\max }^{t}\right| \leq|\ell|$, and for all $\ell \in \Lambda_{\text {add }},\left|T_{\max }^{t}\right|-1 \leq|\ell|$.
Proof of Claim. The first part just follows from the observation that any $\ell \in \Lambda$ was split by sel before time $t$. In order to prove the second part, pick a leaf $\ell \in \Lambda_{\text {add }}$. Let $\ell^{\prime}$ be its unique sibling in $T^{\prime \prime}$ and let $p$ be the parent of $\ell$ and $\ell^{\prime}$, also in $T^{\prime \prime}$. Lemma 2 applied to the subtree $p$ implies $\left|\ell^{\prime}\right| \leq \frac{1}{2}|p|$. Moreover, since $|\ell|+\left|\ell^{\prime}\right|=|p|-1$, we obtain $|\ell|+1 \geq \frac{1}{2}|p| \geq\left|\ell^{\prime}\right| \geq\left|T_{\max }^{t}\right|$, the last inequality using the first part of the claim. This implies $\left|T_{\max }^{t}\right|-1 \leq|\ell|$, and the claim is proven.
Let now $N(\Lambda)$ be the number of nodes in subtrees and forests associated with the leaves of $T^{\prime \prime}$. With each internal node of $T^{\prime \prime}$ we can associate a node of $L_{\text {SEL }}$ which does not belong to any leaf in $\Lambda$. Moreover, the number $\left|T^{\prime \prime} \backslash \Lambda\right|$ of internal nodes in $T^{\prime \prime}$ is bigger than the number $\left|\Lambda_{\text {add }}\right|$ of internal nodes of $T_{\text {nol }}^{\prime}$ to which a child has been added. Since these subtrees and forests are all distinct, we obtain $N(\Lambda)+\left|T^{\prime \prime} \backslash \Lambda\right|<N(\Lambda)+$ $\left|\Lambda_{\text {add }}\right| \leq|V|$. Hence, using the above claim we can write $N(\Lambda) \geq\left(|\Lambda|-\left|\Lambda_{\text {add }}\right|\right)\left|T_{\max }^{t}\right|+\left|\Lambda_{\text {add }}\right|\left(\left|T_{\max }^{t}\right|-1\right)$, which implies $\left|T_{\max }^{t}\right| \leq\left(N(\Lambda)+\left|\Lambda_{\text {add }}\right|\right) /|\Lambda| \leq|V| /|\Lambda|$. Since each internal node of $T^{\prime \prime}$ has at least two children, we have that $|\Lambda| \geq\left|T^{\prime \prime}\right| / 2 \geq\left|T_{\text {nol }}^{\prime}\right| / 2=t / 2$. Hence, we can conclude that $\left|T_{\max }^{t}\right| \leq 2|V| / t$.

### 4.1 Lower bounds

We now state and prove a lower bound on the number of mistakes that any prediction algorithm (even knowing the cutsize budget $K$ ) makes on any given tree, when the query set $L$ is 0 -forked. The bound depends on the following quantity: Given a tree $T(V, E)$, a node subset $L \subseteq V$ and an integer $K$, the component function $\Upsilon(L, K)$ is the sum of the sizes of the $K$ largest components of $T \backslash L$, or $|V \backslash L|$ if $T \backslash L$ has less than $K$ components.

Theorem 5 For all trees $T=(V, E)$, for all 0 -forked subsets $L^{+} \subseteq V$, and for all cutsize budgets $K=$ $0,1, \ldots,|V|-1$, we have that $\operatorname{LB}\left(L^{+}, K\right) \geq \frac{1}{2} \Upsilon\left(L^{+}, K\right)$.

Proof: We describe an adversarial strategy causing any algorithm to make at least $\Upsilon\left(L^{+}, K\right) / 2$ mistakes even when the cutsize budget $K$ is known beforehand. Since $L^{+}$is 0 -forked, each component of $T \backslash L^{+}$is a hinge-tree. Let $F_{\max }$ be the set of the $K$ largest hinge-trees of $T \backslash L^{+}$, and $E(\mathcal{T})$ be the set of all edges in $E$ incident to at least one node of a hinge-tree $\mathcal{T}$. The adversary creates at most one $\phi$-edge ${ }^{3}$ in each edge set $E\left(\mathcal{T}_{1}\right)$ for all 1-hinge-trees $\mathcal{T}_{1} \in F_{\max }$, exactly one $\phi$-edge in each edge set $E\left(\mathcal{T}_{2}\right)$ for all 2-hinge-trees $\mathcal{T}_{2} \in F_{\text {max }}$, and no $\phi$-edges in the edge set $E(\mathcal{T})$ of any remaining hinge-tree $\mathcal{T} \notin F_{\max }$. This is done as follows. By performing a depth-first visit of $T$, the adversary can always assign disagreeing labels to the two connection nodes of each 2-hinge-tree in $F_{\max }$, and agreeing labels to the two connection nodes of each 2-hinge-tree not in $F_{\max }$. Then, for each hinge-tree $\mathcal{T} \in F_{\max }$, the adversary assigns a unique random label to all nodes of $\mathcal{T}$, forcing $|\mathcal{T}| / 2$ mistakes in expectation. The labels of the remaining hinge-trees not in $F_{\max }$ are chosen in agreement with their connection nodes.

Remark 1 Note that Theorem 5 holds for all query sets, not only those that are 0-forked, since any adversarial strategy for a query set $L^{+}$can force at least the same mistakes on the subset $L \subseteq L^{+}$. Note also that it is not difficult to modify the adversarial strategy described in the proof of Theorem 5 in order to deal with algorithms that are allowed to adaptively choose the query nodes in $L$ depending on the labels of the previously selected nodes. The adversary simply assigns the same label to each node in the query set and then forces, with the same method described in the proof, $\frac{1}{2} \Upsilon\left(L^{+}, \frac{K}{2}\right)$ mistakes in expectation on the $\frac{K}{2}$ largest hinge-trees. Thus there are at most two $\phi$-edges in each edge set $E(\mathcal{T})$ for all hinge-trees $\mathcal{T}$, yielding at most $K \phi$-edges in total. The resulting (slightly weaker) bound is $\operatorname{LB}\left(L^{+}, K\right) \geq \frac{1}{2} \Upsilon\left(L^{+}, \frac{K}{2}\right)$. Theorem 8 and Corollary 9 can also be easily rewritten in order to extend the results in this direction.

### 4.2 Upper bounds

We now bound the total number of mistakes that PRED makes on any labeled tree when the queries are decided by SEL. We use Lemma 2 and 3, together with the two lemmas below, to prove that $m_{\text {PRED }}\left(L_{\text {SEL }}^{+}, K\right) \leq$ $10 \mathrm{LB}(L, K)$ for all cutsize budgets $K$ and for all node subset $L \subseteq V$ such that $|L| \leq \frac{1}{8}\left|L_{\text {SEL }}^{+}\right|$.

Lemma 6 For all labeled trees $(T, \boldsymbol{y})$ and for all 0 -forked query sets $L^{+} \subseteq V$, the number of mistakes made by PRED satisfies $m_{\text {PRED }}\left(L^{+}, \boldsymbol{y}\right) \leq \Upsilon\left(L^{+}, \Phi_{T}(\boldsymbol{y})\right)$.

Proof: As in the proof of Theorem 5, we first observe that each component of $T \backslash L^{+}$is a hinge-tree. Let $E(\mathcal{T})$ be the set of all edges in $E$ incident to nodes of a hinge-tree $\mathcal{T}$, and $F_{\phi}$ be the set of hinge-trees such that, for all $\mathcal{T} \in F_{\phi}$, at least one edge of $E(\mathcal{T})$ is a $\phi$-edge. Since $E(\mathcal{T}) \cap E\left(\mathcal{T}^{\prime}\right) \equiv \emptyset$ for all $\mathcal{T}, \mathcal{T}^{\prime} \in T \backslash L^{+}$, we have that $\left|F_{\phi}\right| \leq \Phi_{T}(\boldsymbol{y})$. Moreover, since for any $\mathcal{T} \notin F_{\phi}$ there are no $\phi$-edges in $E(\mathcal{T})$, the nodes of $\mathcal{T}$ must be labeled as its connections nodes. This, together with the prediction rule of PRED, implies that PRED makes no mistakes over any of the hinge-trees $\mathcal{T} \notin F_{\phi}$. Hence, the number of mistakes made by PRED is bounded by the sum of the sizes of all hinge-trees $\mathcal{T} \in F_{\phi}$, which (by definition of $\Upsilon$ ) is bounded by $\Upsilon\left(L^{+}, \Phi_{T}(\boldsymbol{y})\right)$.

The next lemma, whose proof is a bit involved, provides the relevant properties of the component function $\Upsilon(\cdot, \cdot)$. Figure 3 helps visualizing the main ingredients of the proof.

Lemma 7 Given a tree $T=(V, E)$, for all node subsets $L \subseteq V$ such that $|L| \leq \frac{1}{2}\left|L_{\text {SEL }}\right|$ and for all integers $k$, we have: (a) $\Upsilon\left(L_{\mathrm{SEL}}, k\right) \leq 5 \Upsilon(L, k)$; (b) $\Upsilon\left(L_{\mathrm{SEL}}, 1\right) \leq \Upsilon(L, 1)$.

Proof: We prove part (a) by constructing, via SEL, three bijective mappings $\mu_{1}, \mu_{2}, \mu_{3}: \mathcal{P}_{\text {sEL }} \rightarrow \mathcal{P}_{L}$, where $\mathcal{P}_{\text {SEL }}$ is a suitable partition of $T \backslash L_{\text {SEL }}, \mathcal{P}_{L}$ is a subset of $2^{V}$ such that any $S \in \mathcal{P}_{L}$ is all contained in a single connected component of $T \backslash L$, and the union of the domains of the three mappings covers the whole set $T \backslash L_{\text {SEL }}$. The mappings $\mu_{1}, \mu_{2}$ and $\mu_{3}$ are shown to satisfy, for all forests ${ }^{4} F \in \mathcal{P}_{\text {SEL }}$,

$$
|F| \leq\left|\mu_{1}(F)\right|, \quad|F| \leq 2\left|\mu_{2}(F)\right|, \quad|F| \leq 2\left|\mu_{3}(F)\right|
$$

Since each $S \in \mathcal{P}_{L}$ is all contained in a connected component of $T \backslash L$, this we will enable us to conclude that, for each tree $T^{\prime} \in T \backslash L$, the forest of all trees $T \backslash L_{\text {SEL }}$ mapped (via any of these mappings) to any node subset of $T^{\prime}$ has at most five times the number of nodes of $T^{\prime}$. This would prove the statement in (a).

[^2]



Figure 3: The upper pane illustrates the different kinds of nodes chosen by SEL. Numbers in the square tags indicate the first six subtrees $T_{\text {max }}^{t}$, and their associated nodes $i_{t}$, selected by SEL. Node $i_{1}$ is a $[\geq 1 ; \geq 1]$ node, $i_{2}$ is an initial $[0 ; \geq 1]$-node, $i_{3}$ is a (noninitial) $[0 ; \geq 1]$-node, $i_{4}$ is an initial collision node, $i_{5}$ is a (noninitial) collision node, and $i_{6}$ is a $[0 ; 0]$-node. As in Figure 2, we denote by $3^{\prime}$ the fork node generated by the inclusion of $i_{3}$ into $L_{\text {SEL }}$. Note that node $i_{6}$ may be chosen arbitrarily among the four nodes in $T_{\max }^{4} \backslash i_{4}$. The two black nodes are the set of nodes we are competing against, i.e., the nodes in the query set $L$. Forest $T \backslash L$ is made up of one large subtree and two small subtrees. In the lower panes we illustrate some steps of the proof of Lemma 7, with reference to the upper pane. Time $t=2$ : Trees $T_{\max }^{2}$ and $T_{i_{2}}$ are shown. As explained in the proof, $\left|T_{i_{2}}\right| \leq\left|T_{\max }^{2} \backslash T_{i_{2}}\right|$. The circled black node is captured by $i_{2}$. The nodes of tree $T_{\max }^{2} \backslash T_{i_{2}}$ are shaded, and can be used for mapping any $\zeta$-component through $\mu_{2}$. Time $t=3$ : Trees $T_{\text {max }}^{3}$ and $T_{i_{3}}$ are shown. Again, one can easily verify that $\left|T_{i_{3}}\right| \leq\left|T_{\max }^{3} \backslash T_{i_{3}}\right|$. As before, the nodes of $T_{\max }^{3} \backslash T_{i_{3}}$ are shaded, and can be used for mapping any $\zeta$-component via $\mu_{2}$. The reader can see that, according to the injectivity of $\mu_{2}$, these grey nodes are well separated from the ones in $T_{\max }^{2} \backslash T_{i_{2}}$. Time $t=4$ : $T_{\max }^{4}$ and the initial collision node $i_{4}$ are depicted. The latter is enclosed in a circled black node since it captures itself. Time $t=5,6$ : We depicted trees $T_{\max }^{5}$ and $T_{\max }^{6}$, together with nodes $i_{5}$ and $i_{6}$. Node $i_{5}$ is a collision node, which is not initial since it was already captured by the $[0 ; \geq 1]$-node $i_{2}$. Node $i_{6}$ is a $[0 ; 0]$ node, so that the whole tree $T_{\max }^{6}$ is completely included in a component (the largest, in this case) of $T \backslash L$. Tree $T_{\max }^{6}$ can be used for mapping via $\mu_{3}$ any $\zeta$-component. The resulting forest $T \backslash L_{6}$ includes several single-node trees and one two-node tree. If $i_{6}$ is the last node selected by $L_{\text {SEL }}$, then each component of $T \backslash L_{6}$ can be exploited by mapping $\mu_{1}$, since in this specific case none of these components contains nodes of $L$, i.e., there are no $\zeta$-components left.

The construction of these mappings requires some auxiliary definitions. We call $\zeta$-component each connected component of $T \backslash L_{\text {SEL }}$ containing at least one node of $L$. Let $i_{t}$ be the $t$-th node selected by SEL during the incremental construction of the query set $L_{\text {SEL }}$. We distinguish between four kinds of nodes chosen by SEL-see Figure 3 for an example.

Node $i_{t}$ is:

1. A collision node if it belongs to $L_{\text {SEL }} \cap L$;
2. a $[0 ; 0]$-node if, at time $t$, the tree $T_{\max }^{t}$ does not contain any node of $L$;
3. a $[0 ; \geq 1]$-node if, at time $t$, the tree $T_{\text {max }}^{t}$ contains $k \geq 1$ nodes $j_{1}, \ldots, j_{k} \in L$ all belonging to the same connected component of $T_{\text {max }}^{t} \backslash\left\{i_{t}\right\}$;
4. a $[\geq 1 ; \geq 1]$-node if $i_{t} \notin L$ and, at time $t$, the tree $T_{\text {max }}^{t}$ contains $k \geq 2$ nodes $j_{1}, \ldots, j_{k} \in L$, which do not belong to the same connected component of $T_{\max }^{t} \backslash\left\{i_{t}\right\}$.

We now turn to building the three mappings.
$\mu_{1}$ simply maps each tree $T^{\prime} \in T \backslash L_{\text {SEL }}$ that is not a $\zeta$-component to the node set of $T^{\prime}$ itself. This immediately implies $|F| \leq\left|\mu_{1}(F)\right|$ for all forests $F$ (which are actually single trees) in the domain of $\mu_{1}$. Mappings $\mu_{2}$ and $\mu_{3}$ deal with the $\zeta$-components of $T \backslash L_{\text {SEL }}$. Let $Z$ be the set of all such $\zeta$-components, and denote by $V_{0 ; 0}, V_{0 ; 1}$, and $V_{1 ; 1}$ the set of all $[0 ; 0]$-nodes, $[0 ; \geq 1]$-nodes, and $[\geq 1 ; \geq 1]$-nodes, respectively. Observe that $\left|V_{1 ; 1}\right|<|L|$. Combined with the assumption $\left|L_{\text {SEL }}\right| \geq 2|L|$, this implies that $\left|V_{0 ; 0}\right|+\left|V_{0 ; 1}\right|$ plus the total number of collision nodes must be larger than $|L|$; as a consequence, $\left|V_{0 ; 0}\right|+\left|V_{0 ; 1}\right|>|Z|$. Each node $i_{t} \in V_{0 ; 1}$ chosen by SEL splits the tree $T_{\max }^{t}$ into one component $T_{i_{t}}$ containing at least one node of $L$ and one or more components all contained in a single tree $T_{i_{t}}^{\prime}$ of $T \backslash L$. Now mapping $\mu_{2}$ can be constructed incrementally in the following way. For each $[0 ; \geq 1]$-node selected by SEL at time $t, \mu_{2}$ sequentially maps any $\zeta$-component generated to the set of nodes in $T_{\max }^{t} \backslash T_{i_{t}}$, the latter being just a subset of a component of $T \backslash L$. A future time step $t^{\prime}>t$ might feature the selection of a new $[0 ; \geq 1]$-node within $T_{i_{t}}$, but mapping $\mu_{2}$ would cover a different subset of such component of $T \backslash L$. Now, applying Lemma 2 to tree $T_{\max }^{t}$, we can see that $\left|T_{\max }^{t} \backslash T_{i_{t}}\right| \geq\left|T_{\max }^{t}\right| / 2$. Since the selection rule of SEL guarantees that the number of nodes in $T_{\max }^{t}$ is larger than the number of nodes of any $\zeta$-component, we have $|F| \leq 2\left|\mu_{2}(F)\right|$, for any $\zeta$-component $F$ considered in the construction of $\mu_{2}$.

Mapping $\mu_{3}$ maps all the remaining $\zeta$-components that are not mapped through $\mu_{2}$. Let $\sim$ be an equivalence relation over $V_{0 ; 0}$ defined as follows: $i \sim j$ iff $i$ is connected to $j$ by a path containing only $[0 ; 0]$ nodes and nodes in $V \backslash\left(L_{\text {SEL }} \cup L\right)$. Let $i_{t_{1}}, i_{t_{2}}, \ldots, i_{t_{k}}$ be the sequence of nodes of any given equivalence class $[C]_{\sim}$, sorted according to SEL's chronological selection. Lemma 4 applied to tree $T_{\text {max }}^{t_{1}}$ shows that $\left|T_{\max }^{t_{k}}\right| \leq 2\left|T_{\max }^{t_{1}}\right| / k$. Moreover, the selection rule of SEL guarantees that the number of nodes of $T_{\max }^{t_{k}}$ cannot be smaller than the number of nodes of any $\zeta$-component. Hence, for each equivalence class $\left[{ }^{\max }\right]_{\sim}^{\sim}$ containing $k$ nodes of type $[0 ; 0]$, we map through $\mu_{3}$ a set $F_{\zeta}$ of $k$ arbitrarily chosen $\zeta$-components to $T_{\max }^{t_{1}}$. Since the size of each $\zeta$-component is $\leq\left|T_{\max }^{t_{k}}\right|$, we can write $\left|F_{\zeta}\right| \leq k\left|T_{\max }^{t_{k}}\right| \leq 2\left|T_{\max }^{t_{1}}\right|$, which implies $\left|F_{\zeta}\right| \leq 2\left|\mu_{3}\left(F_{\zeta}\right)\right|$ for all $F_{\zeta}$ in the domain of $\mu_{3}$. Finally, observe that the number of $\zeta$-components that are not mapped through $\mu_{2}$ cannot be larger than $\left|V_{0 ; 0}\right|$, thus the union of mappings $\mu_{2}$ and $\mu_{3}$ do actually map all $\zeta$-components. This, in turn, implies that the union of the domains of the three mappings covers the whole set $T \backslash L_{\text {SEL }}$, thereby concluding the proof of part (a).

The proof of (b) is built on the definition of collision nodes, $[0 ; 0]$-nodes, $[0 ; \geq 1]$-nodes and $[\geq 1 ; \geq 1]$ nodes given in part (a). Let $L_{t} \subseteq L_{\text {SEL }}$ be the set of the first $t$ nodes chosed by SEL. Here, we make a further distinction within the collision and $[0 ; \geq 1]$-nodes. We say that during the selection of node $i_{t} \in V_{0 ; 1}$, the nodes in $L \cap T_{\max }^{t}$ are captured by $i_{t}$. This notion of capture extends to collision nodes by saying that a collision node $i_{t} \in L \cap L_{\text {SEL }}$ just captures itself. We say that $i_{t}$ is an initial $[0 ; \geq 1]$-node (resp., initial collision node) if $i_{t}$ is a $[0 ; \geq 1]$-node (resp., collision node) such that the whole set of nodes in $L$ captured by $i_{t}$ contains no nodes captured so far. See Figure 3 for reference. The simple observation leading to the proof of part (b) is the following. If $i_{t}$ is a $[0 ; 0]$-node, then $T_{\max }^{t}$ cannot be larger than the component of $T \backslash L$ that contains $T_{\max }^{t}$, which in turn cannot be larger than $\Upsilon(L, 1)$. This would already imply $\Upsilon\left(L_{t-1}, 1\right) \leq \Upsilon(L, 1)$. Let now $i_{t}$ be an initial $[0 ; \geq 1]$-node and $T_{i_{t}}$ be the unique component of $T_{\max }^{t} \backslash\left\{i_{t}\right\}$ containing one or more nodes of $L$. Applying Lemma 2 to tree $T_{\max }^{t}$ we can see that $\left|T_{i_{t}}\right|$ cannot be larger than $\left|T_{\max }^{t} \backslash T_{i_{t}}\right|$, which in turn cannot be larger than $\Upsilon(L, 1)$. If at time $t^{\prime}>t$ the procedure SEL selects $i_{t^{\prime}} \in T_{i_{t}}$ then $\left|T_{\max }^{t^{\prime}}\right| \leq\left|T_{i_{t}}\right| \leq \Upsilon(L, 1)$. Hence, the maximum integer $q$ such that $\Upsilon\left(L_{q}, 1\right)>\Upsilon(L, 1)$ is bounded by the number of $[\geq 1 ; \geq 1]$-nodes plus the number of initial $[0 ; \geq 1]$-nodes plus the number of initial collision nodes. We now bound this sum as follows. The number of $[\geq 1 ; \geq 1]$-nodes is clearly bounded by $|L|-1$. Also, any initial $[0 ; \geq 1]$-node or initial collision node selected by SEL captures at least a new node in $L$, thereby implying that the total number of initial $[0 ; \geq 1]$-node or initial collision node must be $\leq|L|$. After $q=2|L|-1$ rounds, we are sure that the size of the largest tree of $T_{\text {max }}^{q}$ is not larger than the size of the largest component of $T \backslash L$, i.e., $\Upsilon(L, 1)$.

We now put the above lemmas together to prove our main result concerning the number of mistakes made by PRED on the query set chosen by SEL.

Theorem 8 For all trees $T$ and all cutsize budgets $K$, the number of mistakes made by PRED on the query set $L_{\text {SEL }}^{+}$satisfies

$$
m_{\text {PRED }}\left(L_{\text {SEL }}^{+}, K\right) \leq \min _{L \subseteq V:|L| \leq \frac{1}{8}\left|L_{\text {SEL }}^{+}\right|} 10 \mathrm{LB}(L, K)
$$

Proof: Pick any $L \subseteq V$ such that $|L| \leq \frac{1}{8}\left|L_{\text {SEL }}^{+}\right|$. Then

$$
m_{\text {PRED }}\left(L_{\mathrm{SEL}}^{+}, K\right) \stackrel{(\text { Lem. 6) }}{\leq} \Upsilon\left(L_{\mathrm{SEL}}^{+}, K\right) \stackrel{(\mathrm{A})}{\leq} \Upsilon\left(L_{\mathrm{SEL}}, K\right) \stackrel{(\text { Lem. } 7 \text { (a)) }}{\leq} 5 \Upsilon\left(L^{+}, K\right) \stackrel{(\mathrm{Thm} .5)}{\leq} 10 \mathrm{LB}\left(L^{+}, K\right) \stackrel{(\mathrm{B})}{\leq} 10 \mathrm{LB}(L, K)
$$

Inequality (A) holds because $L_{\mathrm{SEL}} \subseteq L_{\mathrm{SEL}}^{+}$, and thus $T \backslash L_{\mathrm{SEL}}^{+}$has connected components of smaller size than $L_{\text {SEL }}$. In order to apply Lemma 7 (a), we need the condition $\left|L^{+}\right| \leq \frac{1}{2}\left|L_{\text {SEL }}\right|$. This condition is seen to hold after combining Lemma 3 with our assumptions: $\left|L^{+}\right| \leq 2|L| \leq \frac{1}{4}\left|L_{\text {SEL }}^{+}\right| \leq \frac{1}{2}\left|L_{\text {SEL }}\right|$. Finally, inequality (B) holds because any adversarial strategy using query set $L$ can also be used with the larger query set $L^{+} \supseteq L$.

Note also that Theorem 5 and Lemma 6 imply the following statement about the optimality of PRED over 0 -forked query sets.
Corollary 9 For all trees $T$, for all cutsize budgets $K$, and for all 0 -forked query sets $L^{+} \subseteq V$, the number of mistakes made by PRED satisfies $m_{\text {PRED }}\left(L^{+}, K\right) \leq 2 \operatorname{LB}\left(L^{+}, K\right)$.
In the rest of this section we derive a more intepretable bound on $m_{\text {PRED }}\left(L^{+}, \boldsymbol{y}\right)$ based on the function $\Psi^{*}$ introduced in [6]. To this end, we prove that $L_{\text {SEL }}$ minimizes $\Psi^{*}$ up to constant factors, and thus is an optimal query set according to the analysis of [6].

For any subset $V^{\prime} \subseteq V$, let $\Gamma\left(V^{\prime}, V \backslash V^{\prime}\right)$ be the number of edges between nodes of $V^{\prime}$ and nodes of $V \backslash V^{\prime}$. Using this notation, we can write

$$
\Psi^{*}(L)=\max _{\emptyset \not \equiv V^{\prime} \subseteq V \backslash L} \frac{\left|V^{\prime}\right|}{\Gamma\left(V^{\prime}, V \backslash V^{\prime}\right)} .
$$

Lemma 10 For any tree $T=(V, E)$ and any $L \subseteq V$ the following holds.
(a) A maximizer of $\frac{\left|V^{\prime}\right|}{\Gamma\left(V^{\prime}, V \backslash V^{\prime}\right)}$ exists which is included in the node set of a single component of $T \backslash L$;
(b) $\Psi^{*}(L) \leq \Upsilon(L, 1)$.

Proof: Let $V_{\max }^{\prime}$ be any maximizer of $\frac{\left|V^{\prime}\right|}{\Gamma\left(V^{\prime}, V \backslash V^{\prime}\right)}$. For the sake of contradiction, assume that the nodes of $V_{\text {max }}^{\prime}$ belong to $k \geq 2$ components $\mathcal{T}_{1}, \mathcal{T}_{2}, \ldots, \mathcal{T}_{k} \in T \backslash L$. Let $V_{i}^{\prime} \subset V_{\text {max }}^{\prime}$ be the subset of nodes included in the node set of $\mathcal{T}_{i}$, for $i=1, \ldots, k$. Then $\left|V^{\prime}\right|=\sum_{i \leq k}\left|V_{i}^{\prime}\right|$ and $\Gamma\left(V^{\prime}, V \backslash V^{\prime}\right)=\sum_{i \leq k} \Gamma\left(V_{i}^{\prime}, V \backslash V_{i}^{\prime}\right)$. Now let $i^{*}=\operatorname{argmax}_{i \leq k}\left|V_{i}^{\prime}\right| / \Gamma\left(V_{i}^{\prime}, V \backslash V_{i}^{\prime}\right)$. Since $\left(\sum_{i} a_{i}\right) /\left(\sum_{i} b_{i}\right) \leq \max _{i} a_{i} / b_{i}$ for all $a_{i}, b_{i} \geq 0$, we immediately obtain $\bar{\Psi}\left(V_{i^{*}}^{\prime}\right) \geq \Psi\left(V_{\max }^{\prime}\right)$, contradicting our assumption. This proves (a). Part (b) is an immediate consequence of (a).

Lemma 11 For any tree $T=(V, E)$ and any 0-forked subset $L^{+} \subseteq V$ we have $\Upsilon\left(L^{+}, 1\right) \leq 2 \Psi^{*}\left(L^{+}\right)$.
Proof: Let $\mathcal{T}_{\text {max }}$ be the largest component of $T \backslash L^{+}$and $V_{\max }$ be its node set. Since $L^{+}$is a 0 -forked query set, $\mathcal{T}_{\max }$ must be either a 1-hinge-tree or a 2 -hinge-tree. Since the only edges that connect a hinge-tree to external nodes are the edges leading to connection nodes, we find that $\Gamma\left(V_{\max }, V \backslash V_{\max }\right) \leq 2$. We can now write

$$
\Psi^{*}\left(L^{+}\right)=\max _{\emptyset \not \equiv V^{\prime} \subseteq V \backslash L^{+}} \frac{\left|V^{\prime}\right|}{\Gamma\left(V^{\prime}, V \backslash V^{\prime}\right)} \geq \frac{\left|V_{\max }\right|}{\Gamma\left(V_{\max }, V \backslash V_{\max }\right)} \geq \frac{\left|V_{\max }\right|}{2}=\frac{\Upsilon\left(L^{+}, 1\right)}{2}
$$

thereby concluding the proof.
Lemma 12 For any tree $T=(V, E)$ and any subset $L \subseteq V$ we have $\Psi^{*}\left(L^{+}\right) \leq \Psi^{*}(L)$.
Proof: Let $V_{\max }^{\prime}$ be any set maximizing $\Psi^{*}\left(L^{+}\right)$. Since $V_{\max }^{\prime} \in V \backslash L^{+}, V_{\max }^{\prime}$ cannot contain any node of $L \subseteq L^{+}$. Hence

$$
\Psi^{*}(L)=\max _{\emptyset \neq V^{\prime} \subseteq V \backslash L} \frac{\left|V^{\prime}\right|}{\Gamma\left(V^{\prime}, V \backslash V^{\prime}\right)} \geq \frac{\left|V_{\max }^{\prime}\right|}{\Gamma\left(V_{\max }^{\prime}, V \backslash V_{\max }^{\prime}\right)}=\Psi^{*}\left(L^{+}\right)
$$

which concludes the proof.
We now put together the previous lemmas to show that the query set $L_{\text {SEL }}$ minimizes $\Psi^{*}$ up to constant factors.

Theorem 13 For any tree $T=(V, E)$ we have $\Psi^{*}\left(L_{\mathrm{SEL}}\right) \leq \min _{L \subseteq V:|L| \leq \frac{1}{4}\left|L_{\mathrm{SEL}}\right|} 2 \Psi^{*}(L)$.
Proof: Let $L$ be a query set such that $|L| \leq\left|L_{\text {SEL }}\right| / 4$. Then we have the following chain of inequalities:

$$
\Psi^{*}\left(L_{\mathrm{SEL}}\right) \stackrel{(\text { Lemma } 10(\mathrm{~b}))}{\leq} \Upsilon\left(L_{\mathrm{SEL}}, 1\right) \stackrel{(\text { Lemma } 7(\mathrm{~b}))}{\leq} \Upsilon\left(L^{+}, 1\right) \stackrel{(\text { Lemma 11) }}{\leq} 2 \Psi^{*}\left(L^{+}\right) \stackrel{(\text { Lemma 12) }}{\leq} 2 \Psi^{*}(L)
$$

In order to apply Lemma 7 (b), we need the condition $\left|L^{+}\right| \leq \frac{1}{2}\left|L_{\text {sEL }}\right|$. This condition holds because, by Lemma $3,\left|L^{+}\right| \leq 2|L| \leq \frac{1}{2}\left|L_{\text {SEL }}\right|$.

Finally, as promised, the following corollary contains an interpretable mistake bound for PRED run with a query set returned by SEL.

Corollary 14 For any labeled tree $(T, \boldsymbol{y})$, the number of mistakes made by PRED when run with query set $L_{\text {SEL }}^{+}$satisfies

$$
m_{\mathrm{PRED}}\left(L_{\mathrm{SEL}}^{+}, \boldsymbol{y}\right) \leq 4 \min _{L \subseteq V:|L| \leq \frac{1}{8}\left|L_{\mathrm{SEL}}^{+}\right|} \Psi^{*}(L) \Phi_{T}(\boldsymbol{y})
$$

Proof: Observe that PRED assigns labels to nodes in $V \backslash L_{\text {SEL }}^{+}$so as to minimize the resulting cutsize given the labels in the query set $L_{\text {SEL }}^{+}$. We can then invoke [6, Lemma 1], which bounds the number of mistakes made by the mincut strategy in terms of the functions $\Psi^{*}$ and the cutsize. This yields

$$
m_{\mathrm{PRED}}\left(L_{\mathrm{SEL}}^{+}, \boldsymbol{y}\right) \stackrel{[6, \text { Lemma 1] }}{\leq} 2 \Psi^{*}\left(L_{\mathrm{SEL}}^{+}\right) \Phi_{T}(\boldsymbol{y}) \stackrel{(\mathrm{A})}{\leq} 2 \Psi^{*}\left(L_{\mathrm{SEL}}\right) \Phi_{T}(\boldsymbol{y}) \stackrel{(\text { Theorem 13) }}{\leq} 4 \Psi^{*}(L) \Phi_{T}(\boldsymbol{y})
$$

Inequality (A) holds because $L_{\text {SEL }} \subseteq L_{\text {SEL }}^{+}$, and thus $T \backslash L_{\mathrm{SEL}}^{+}$has connected components of smaller size than $L_{\text {SEL }}$. In order to apply Theorem 13 , we need the conditon $|L| \leq \frac{1}{4}\left|L_{\text {SEL }}\right|$, which follows from a simple combination of Lemma 3 and our assumptions: $|L| \leq \frac{1}{8}\left|L_{\text {SEL }}^{+}\right| \leq \frac{1}{4}\left|L_{\text {SEL }}\right|$.

Remark 2 A mincut algorithm exists which efficiently predicts even when the query set $L$ is not 0-forked (thereby gaining a factor of 2 in the cardinality of the competing query sets $L$ - see Theorem 8 and Corollary 14). This algorithm is a "batch" variant of the TreeOpt algorithm analyzed in [7]. The algorithm can be implemented in such a way that the total time for predicting $|V|-|L|$ labels is $\mathcal{O}(|V|)$.

### 4.3 Automatic calibration of the number of queries

A key aspect to the query selection task is deciding when to stop asking queries. Since the more queries are asked the less mistakes are made afterwards, a reasonable way to deal with this trade-off is to minimize the number of queries issued during the selection phase plus the number of mistakes made during the prediction phase. For a given pair $A=\langle S, P\rangle$ of prediction and selection algorithms, we denote by $[q+m]_{A}$ the sum of queries made by $S$ and prediction mistakes made by $P$. Similarly to $m_{A}$ introduced in Section $4,[q+m]_{A}$ has to scale with the cutsize $\Phi_{T}(\boldsymbol{y})$ of the labeled tree $(T, \boldsymbol{y})$ under consideration.

As a simple example of computing $[q+m]_{A}$, consider a line graph $T=(V, E)$. Since each query set on $T$ is 0 -forked, Theorem 5 and Corollary 9 ensure that an optimal strategy for selecting the queries in $T$ is choosing a sequence of nodes such that the distance between any pair of neighbor nodes in $L$ is equal. The total number of mistakes that can be forced on $V \backslash L$ is, up to a constant factor, $(|V| /|L|) \Phi_{T}(\boldsymbol{y})$. Hence, the optimal value of $[q+m]_{A}$ is about

$$
\begin{equation*}
|L|+\frac{|V|}{|L|} \Phi_{T}(\boldsymbol{y}) \tag{1}
\end{equation*}
$$

Minimizing the above expression over $|L|$ clearly requires knowledge of $\Phi_{T}(\boldsymbol{y})$, which is typically unavailable. In this section we investigate a method for choosing the number of queries when the labeling is known to be sufficiently regular, that is when a bound $K$ is known on the cutsize $\Phi_{T}(\boldsymbol{y})$ induced by the adversarial labeling. ${ }^{5}$

We now show that when a bound $K$ on the cutsize is known, a simple modification of SEL(we call it SEL $\star$ ) exists which optimizes the $[q+m]_{A}$ criterion. This means that the combination of SEL $\star$ and PRED can trade-off optimally (up to constant factors) queries against mistakes.

[^3]$$
\text { Given a selection algorithm } S \text { and a prediction algorithm } P \text {, define }[q+m]_{\langle S, P\rangle} \text { by }
$$
$$
[q+m]_{\langle S, P\rangle}=\min _{Q \geq 1}\left(Q+m_{P}\left(L_{S(Q)}, K\right)\right)
$$
where $L_{S(Q)}$ is the query set output by $S$ given query budget $Q$, and $m_{P}\left(L_{S(Q)}, K\right)$ is the maximum number of mistakes made by $P$ with query set $L_{S(Q)}$ on any labeling $\boldsymbol{y}$ with $\Phi_{T}(\boldsymbol{y}) \leq K-$ see definition in Section 4 . Define also $[q+m]_{\mathrm{OPT}}=\inf _{S, P}[q+m]_{\langle S, P\rangle}$, where OPT $=\left\langle S^{*}, P^{*}\right\rangle$ is an optimal pair of selection and prediction algorithms. If SEL knows the size of the query set $L^{*}$ selected by $S^{*}$, so that SEL can choose a query budget $Q=8\left|L^{*}\right|$, then a direct application of Theorem 8 guarantees that $\left|L_{\text {SEL }}^{+}\right|+m_{\text {PRED }}\left(L_{\mathrm{SEL}}^{+}, K\right) \leq$ $10[q+m]_{\text {opt }}$. We now show that SEL $\star$, the announced modification of SEL, can efficiently search for a query set size $Q$ such that $Q+m_{\text {PRED }}\left(L_{\operatorname{SEL}(Q)}^{+}, K\right)=\mathcal{O}\left([q+m]_{\mathrm{OPT}}\right)$ when only $K$, rather than $\left|L^{*}\right|$, is known. In fact, Theorem 5 and Corollary 9 ensure that $m_{\text {PRED }}\left(L_{\text {SEL }}^{+}, K\right)=\Theta\left(\Upsilon\left(L_{\text {SEL }}^{+}, K\right)\right)$. When $K$ is given as side information, SEL $\star$ can operate as follows. For each $t \leq|V|$, the algorithm builds the query set $L_{t}^{+}$and computes $\Upsilon\left(L_{t}^{+}, K\right)$. Then it finds the smallest value $t^{*}$ minimizing $t+\Upsilon\left(L_{t}^{+}, K\right)$ over all $t \leq|V|$, and selects $L_{\text {SEL* }} \equiv L_{t^{*}}$. We stress that the above is only possible because the algorithm can estimate within constant factors its own future mistake bound (Theorem 5 and Corollary 9), and because the combination of SEL and PRED is competitive against all query sets whose size is a constant fraction of $\left|L_{\text {SEL }}^{+}\right|$ -see Theorem 8. Putting together, we have shown the following result.

Theorem 15 For all trees $(T, \boldsymbol{y})$, for all cutsize budgets $K$, and for all labelings $\boldsymbol{y}$ such that $\Phi_{T}(\boldsymbol{y}) \leq K$, the combination of SEL $\star$ and PRED achieves $\left|L_{\text {SEL }}\right|+m_{\text {PRED }}\left(L_{\mathrm{SEL} \star}^{+}, K\right)=\mathcal{O}\left([q+m]_{\mathrm{OPT}}\right)$ when $K$ is given to SEL $\star$ as input.

Just to give a few simple examples of how SELぇ works, consider a star graph. It is not difficult to see that in this case $t^{*}=1$ independent of $K$, i.e., SEL $\star$ always selects the center of the star, which is intuitively the optimal choice. If $T$ is the "binary system" mentioned in the introduction, then $t^{*}=2$ and SEL $\star$ always selects the centers of the two stars, again independent of $K$. At the other extreme, if $T$ is a line graph, then SEL $\star$ picks the query nodes in such a way that the distance between two consecutive nodes of $L$ in $T$ is (up to a constant factor) equal to $\sqrt{|V| / K}$. Hence $|L|=\Theta(\sqrt{|V| K})$, which is the minimum of (1) over $|L|$ when $\Phi_{T}(\boldsymbol{y}) \leq K$.

## 5 On the prediction of general graphs

In this section we provide a general lower bound for prediction on arbitrary labeled graphs $(G, y)$. We then contrast this lower bound to some results contained in Afshani et al. [1].

Let $\Phi_{G}^{R}(\boldsymbol{y})$ be the sum of the effective resistances (see, e.g., [9]) on the $\phi$-edges of $G=(V, E)$. The theorem below shows that any prediction algorithm using any query set $L$ such that $|L| \leq \frac{1}{4}|V|$ makes at least order of $\Phi_{G}^{R}(\boldsymbol{y})$ mistakes. This lower bound holds even if the algorithm is allowed to use a randomized adaptive strategy for choosing the query set $L$, that is, a randomized strategy where the next node of the query set is chosen after receiving the labels of all previously chosen nodes.

Theorem 16 Given a labeled graph $(G, y)$, for all $K \leq|V| / 2$, there exists a randomized labeling strategy such that for all prediction algorithms $A$ choosing a query set of size $|L| \leq \frac{1}{4}|V|$ via a possibly randomized adaptive strategy, the expected number of mistakes made by $A$ on the remaining nodes $V \backslash L$ is at least $K / 4$, while $\Phi_{G}^{R}(\boldsymbol{y}) \leq K$.
The above lower bound (whose proof is omitted) appears to contradict an argument by Afshani et al. [1, Section 5]. This argument establishes that for any $\varepsilon>0$ there exists a randomized algorithm using at most $K \ln (3 / \varepsilon)+K \ln (|V| / K)+\mathcal{O}(K)$ queries on any given graph $G=(V, E)$ with cutsize $K$, and making at most $\varepsilon|V|$ mistakes on the remaining vertices. This contradiction is easily obtained through the following simple counterexample: assume $G$ is a line graph where all node labels are +1 but for $K=o(|V| / \ln |V|)$ randomly chosen nodes, which are also given random labels. For all $\varepsilon=o\left(\frac{K}{|V|}\right)$, the above argument implies that order of $K \ln |V|=o(|V|)$ queries are sufficient to make at most $\varepsilon|V|=o(K)$ mistakes on the remaining nodes, among which $\Omega(K)$ have random labels - which is clearly impossible.

## 6 Efficient Implementation

In this section we describe an efficient implementation of SEL and PRED. We will show that the total time needed for selecting $Q$ queries is $\mathcal{O}(|V| \log Q)$, the total time for predicting $|V|-Q$ nodes is $\mathcal{O}(|V|)$, and that the overall memory space is again $\mathcal{O}(|V|)$.

In order to locate the largest subtree of $T \backslash L_{t-1}$, the algorithm maintains a priority deque [8] $D$ containing at most $Q$ items. This data-structure enables to find and eliminate the item with the smallest (resp., largest) key in time $\mathcal{O}(1)$ (resp., time $\mathcal{O}(\log Q)$ ). In addition, the insertion of a new element takes time $\mathcal{O}(\log Q)$.

Each item in $D$ has two records: a reference to a node in $T$ and the priority key associated with that node. Just before the selection of the ${ }^{6} t$-th query node $i_{t}$, the $Q$ references point to nodes contained in the $Q$ largest subtrees in $T \backslash L_{t-1}$, while the corresponding keys are the sizes of such subtrees. Hence at time $t$ the item top of $D$ having the largest key points to a node in $T_{\text {max }}^{t}$.

First, during an initialization step, sel creates, for each edge $(i, j) \in E$, a directed edge $[i, j]$ from $i$ to $j$ and the twin directed edge $[j, i]$ from $j$ to $i$. During the construction of $L_{\text {SEL }}$ the algorithm also stores and maintains the current size $\sigma(D)$ of $D$, i.e., the total number of items contained in $D$. We first describe the way SEL finds node $i_{t}$ in $T_{\max }^{t}$. Then we will see how SEL can efficiently augment the query set $L_{\text {SEL }}$ to obtain $L_{\text {SEL }}^{+}$.

Starting from the node $r$ of $T_{\max }^{t}$ referred to by ${ }^{7} D$, sEL performs a depth-first visit of $T_{\max }^{t}$, followed by the elimination of the item with the largest key in $D$. For the sake of simplicity, consider $T_{\max }^{t}$ as rooted at node $r$. Given any edge $(i, j)$, we let $T_{i}$ and $T_{j}$ be the two subtrees obtained from $T_{\max }^{t}$ after removing edge $(i, j)$, where $T_{i}$ contains node $i$, and $T_{j}$ contains node $j$. During each backtracking step of the depth-first visit from a node $i$ to a node $j$, sEL stores the number of nodes $\left|T_{i}\right|$ contained in $T_{i}$. This number gets associated with $[j, i]$. Observe that this task can be accomplished very efficiently, since $\left|T_{i}\right|$ is equal to 1 plus the number of nodes of the union of $T_{c(i)}$ over all children $c(i)$ of $i$. These numbers can be recursively calculated by summing the size values that SEL associates with all direct edges $[i, c(i)]$ in the previous backtracking steps. Just after storing the value $\left|T_{i}\right|$, the algorithm also stores $\left|T_{j}\right|=\left|T_{\max }^{t}\right|-\left|T_{i}\right|$ and associates this value with the twin directed edge $[i, j]$. The size of $T_{\max }^{t}$ is then stored in $D$ as the key record of the pointer to node $r$.

It is now important to observe that the quantity $\sigma\left(T_{\max }^{t}, i\right)$ used by SEL (see Section 3 ) is simply the largest key associated with the directed edges $[i, j]$ over all $j$ such that $(i, j)$ is an edge of $T_{\text {max }}^{t}$. Hence, a new depth-first visit is enough to find in time $\mathcal{O}\left(\left|T_{\max }^{t}\right|\right)$ the $t$-th node $i_{t}=\arg \min _{i \in T_{\max }^{t}} \sigma\left(T_{\text {max }}^{t}, i\right)$ selected by SEL. Let $N\left(i_{t}\right)$ be the set of all nodes adjacent to node $i_{t}$ in $T_{\max }^{t}$. For all nodes $i^{\prime} \in N\left(i_{t}\right)$, SEL compares $\left|T_{i^{\prime}}\right|$ to the smallest key bottom stored in $D$. We have three cases:

1. If $\left|T_{i^{\prime}}\right| \leq$ bottom and $\sigma(D) \geq Q-t$ then the algorithm does nothing, since $T_{i^{\prime}}$ (or subtrees thereof) will never be largest in the subsequent steps of the construction of $L_{\text {SEL }}$, i.e., there will not exist any node $i_{t^{\prime}}$ with $t^{\prime}>t$ such that $i_{t^{\prime}} \in T_{i^{\prime}}$.
2. If $\left|T_{i^{\prime}}\right| \leq$ bottom and $\sigma(D)<Q-t$, or if $\left|T_{i^{\prime}}\right|>$ bottom and $\sigma(D)<Q$ then SEL inserts a pointer to $i^{\prime}$ together with the associated key $\left|T_{i^{\prime}}\right|$. Note that, since $D$ is not full (i.e., $\sigma(D)<Q$ ), the algorithm need not eliminate any item in $D$.
3. If $\left|T_{i^{\prime}}\right|>$ bottom and $\sigma(D)=Q$ then SEL eliminates from $D$ the item having the smallest key, and inserts a pointer to $i^{\prime}$, together with the associated key $\left|T_{i^{\prime}}\right|$.
Finally, sel eliminates node $i_{t}$ and all edges (both undirected and directed) incident to it. Note that this elimination implies that we can easily perform a depth-first visit within $T_{\max }^{s}$ for each $s \leq Q$, since $T_{\max }^{s}$ is always completely disconnected from the rest of the tree $T$.

In order to turn $L_{\text {SEL }}$ into $L_{\mathrm{SEL}}^{+}$, the algorithm proceeds incrementally, using a technique borrowed from [7]. Just after the selection of the first node $i_{1}$, a depth-first visit starting from $i_{1}$ is performed. During each backtracking step of this visit, the algorithm associates with each edge $(i, j)$, the closer node to $i_{1}$ between the two nodes $i$ and $j$. In other words, SEL assigns a direction to each undirected edge $(i, j)$ so as to be able to efficiently find the path connecting each given node $i$ to $i_{1}$. When the $t$-th node $i_{t}$ is selected, SEL follows these edge directions from $i_{t}$ towards $i_{1}$. Let us denote by $\pi(i, j)$ the path connecting node $i$ to node $j$. During the traversal of $\pi\left(i_{1}, i_{t}\right)$, the algorithm assigns a special mark to each visited node, until the algorithm reaches the first node $j \in \pi\left(i_{1}, i_{t}\right)$ which has already been marked. Let $\eta(i, L)$ be the maximum number of edge disjoint paths connecting $i$ to nodes in the query set $L$. Observe that all nodes $i$ for which $\eta\left(i, L_{t}\right)>\eta\left(i, L_{t-1}\right)$ must necessarily belong to $\pi\left(i_{t}, j\right)$. We have $\eta\left(i_{t}, L_{t}\right)=1$, and $\eta\left(i, L_{t}\right)=2$, for all internal nodes $i$ in the path $\pi\left(i_{t}, j\right)$. Hence, $j$ is the unique node that we may need to add as a new fork node (if $j \notin \operatorname{FORK}\left(L_{t-1}\right)$ ). In fact, $j$ is the unique node such that the number of edge-disjoint paths connecting it to query nodes may increase, and be actually larger than 2 .

Therefore if $j \in L_{t-1}^{+}$we need not add any fork node during the incremental construction of $L_{\mathrm{SEL}}^{+}$. On the other hand, if $j \notin L_{t-1}^{+}$then $\eta\left(i, L_{t-1}\right)=2$, which implies $\eta\left(i, L_{t}\right)=3$. This is the case when SEL views $j$ as new fork node to be added to the query set $L_{\text {SEL }}$ under consideration.

In order to bound the total time required by SEL for selecting $Q$ nodes, we rely on Lemma 4, showing that $\left|T_{\max }^{t}\right| \leq 2|V| / t$. The two depth-first visits performed for each node $i_{t}$ take $\mathcal{O}\left(\left|T_{\max }^{t}\right|\right)$ steps. Hence the overall running time spent on the depth-first visits is $\mathcal{O}\left(\sum_{t \leq Q} 2|V| / t\right)=\mathcal{O}(|V| \log Q)$. The total time

[^4]spent for incrementally finding the fork nodes of $L_{\text {SEL }}$ is linear in the number of nodes marked by the algorithm, which is equal to $|V|$. Finally, handling the priority deque $D$ takes $|V|$ times the worst-case time for eliminating an item with the smallest (or largest) key or adding a new item. This is again $\mathcal{O}(|V| \log Q)$.

We now turn to the implementation of the prediction phase. PRED operates in two phases. In the first phase, the algorithm performs a depth-first visit of each hinge-tree $\mathcal{T}$, starting from each connection node (thereby visiting the nodes of all 1-hinge-tree once, and the nodes of all 2-hinge-tree twice). During these visits, we add to the nodes a tag containing (i) the label of node $i_{\mathcal{T}}$ from which the depth-first visit started, and (ii) the distance between $i_{\mathcal{T}}$ and the currently visited node. In the second phase, we perform a second depth-first visit, this time on the whole tree $T$. During this visit, we predict each node $i \in V \backslash L$ with the label coupled with smaller distance stored in the tags of ${ }^{8} i$. The total time of these visits is linear in $|V|$ since each node of $T$ gets visited at most 3 times.

## 7 Conclusions and ongoing work

The results proven in this paper characterize, up to constant factors, the optimal algorithms for adversarial active learning on trees in two main settings. In the first setting the goal is to minimize the number of mistakes on the non-queried vertices under a certain query budget. In the second setting the goal is to minimize the sum of queries and mistakes under no restriction on the number of queries.

An important open question is the extension of our results to the general case of active learning on graphs. While a direct characterization of optimality on general graphs is likely to require new analytical tools, an alternative line of attack is reducing the graph learning problem to the tree learning problem via the use of spanning trees. Certain types of spanning trees, such as random spanning trees, are known to summarize well the graph structure relevant to passive learning - see, e.g., [7]. In the case of active learning, however, we want good query sets on the graph to correspond to good query sets on the spanning tree, and random spanning trees may fail to do so in simple cases. For example, consider a set of $m$ cliques connected through bridges, so that each clique is connected to, say, $k$ other cliques. The breadth-first spanning tree of this graph is a set of connected stars. This tree clearly reveals a query set (the star centers) which is good for regular labelings (cfr., the binary system example of Section 1). On the other hand, for certain choices of $m$ and $k$ a random spanning tree has a good probability of hiding the clustered nature of the original graph, thus leading to the selection of bad query sets.

In order to gain intuition about this phenomenon, we are currently running experiments on various realworld graphs using different types of spanning trees, where we measure the number of mistakes made by our algorithm (for various choices of the budget size) against common baselines.

We also believe that an extension to general graphs of our algorithm does actually exist. However, the complexity of the methods employed in [6] suggests that techniques based on minimizing $\Psi^{*}$ on general graphs are computationally very expensive.
Acknowledgments. This work was supported in part by Google Inc. through a Google Research Award and by the PASCAL2 Network of Excellence under EC grant no. 216886. This publication only reflects the authors' views.

## References

[1] P. Afshani, E. Chiniforooshan, R. Dorrigiv, A. Farzan, M. Mirzazadeh, N. Simjour, H. Zarrabi-Zadeh. On the complexity of finding an unknown cut via vertex queries. COCOON 2007, pages 459-469.
[2] Belkin, M., Matveeva, I., and Niyogi, P. Regularization and semi-supervised learning on large graphs. COLT 2004, pages 624-638.
[3] Bengio, Y., Delalleau, O., and Le Roux, N. Label propagation and quadratic criterion. In SemiSupervised Learning, pages 193-216. MIT Press, 2006.
[4] Blum, A. and Chawla, S. Learning from labeled and unlabeled data using graph mincuts. ICML 2001, pages 19-26.
[5] A. Blum, J. Lafferty, R. Reddy, and M.R. Rwebangira. Semi-supervised learning using randomized mincuts. ICML 2004.
[6] A. Guillory and J. Bilmes. Label Selection on Graphs. NIPS 2009.
[7] N. Cesa-Bianchi, C. Gentile, F. Vitale. Fast and optimal prediction of a labeled tree. COLT 2009.
[8] J. Katajainen, F. Vitale. Navigation piles with applications to sorting, priority queues, and priority deques. Nordic Journal of Computing, 10(3):238-262, 2003.
[9] R. Lyons and Y. Peres. Probability on Trees and Networks. Manuscript, 2009.
[10] Zhu, X., Ghahramani, Z., and Lafferty, J. Semi-supervised learning using Gaussian fields and harmonic functions. In ICML Workshop on the Continuum from Labeled to Unlabeled Data in Machine Learning and Data Mining, 2003.

[^5]
[^0]:    ${ }^{1}$ A mincut strategy considers all labelings consistent with the labels observed so far, and chooses among them one that minimizes the resulting cutsize over the whole graph.

[^1]:    ${ }^{2}$ In Section 6 we will see that during each step $L_{t-1} \rightarrow L_{t}$ at most a single new fork node may be generated.

[^2]:    ${ }^{3} \mathrm{~A} \phi$-edge $(i, j)$ is one where $y_{i} \neq y_{j}$.
    ${ }^{4}$ In this proof, $|\mu(A)|$ denotes the number of nodes in the set (of nodes) $\mu(A)$. Also, with a slight abuse of notation, for all forests $F \in \mathcal{P}_{\text {SEL }}$, we denote by $|F|$ the sum of the number of nodes in all trees of $F$. Finally, whenever $F \in \mathcal{P}_{\text {sEL }}$ contains a single tree, we refer to $F$ as it were a tree, rather than a (singleton) forest containing only one tree.

[^3]:    ${ }^{5}$ In [1] a labeling $\boldsymbol{y}$ of a graph $G$ is said to be $\alpha$-balanced if, after the elimination of all $\phi$-edges, each connected component of $G$ is not smaller than $\alpha|V|$ for some known constant $\alpha \in(0,1)$. In the case of labeled trees, the $\alpha$-balancing condition is stronger than our regularity assumption. This is because any $\alpha$-balanced labeling $\boldsymbol{y}$ implies $\Phi_{T}(\boldsymbol{y}) \leq$ $1 / \alpha-1$. In fact, getting back to the line graph example, we immediately see that, if $\boldsymbol{y}$ is $\alpha$-balanced, then the optimal number of queries $|L|$ is order of $\sqrt{|V|(1 / \alpha-1)}$, which is also $\inf _{A}[q+m]_{A}$.

[^4]:    ${ }^{6}$ If $t=1$ the priority deque $D$ is empty.
    ${ }^{7}$ In the initial step $t=1$ (i.e., when $T_{\max }^{t} \equiv T$ ) node $r$ can be chosen arbitrarily .

[^5]:    ${ }^{8}$ If $i$ belongs to a 1-hinge-tree, we simply predict $y_{i}$ with the unique label stored in the tag.

