

Data-Adaptive Active Sampling for Efficient Graph-Cognizant Classification

Dimitris Berberidis , *Student Member, IEEE*, and Georgios B. Giannakis , *Fellow, IEEE*

Abstract—This paper deals with active sampling of graph nodes representing training data for binary classification. The graph may be given or constructed using similarity measures among nodal features. Leveraging the graph for classification builds on the premise that labels across neighboring nodes are correlated according to a categorical Markov random field (MRF). This model is further relaxed to a Gaussian (G)MRF with labels taking continuous values—an approximation that not only mitigates the combinatorial complexity of the categorical model, but also offers optimal unbiased soft predictors of the unlabeled nodes. The proposed sampling strategy is based on querying the node whose label disclosure is expected to inflict the largest change on the GMRF, and in this sense it is the most informative on average. Connections are established to other sampling methods including uncertainty sampling, variance minimization, and sampling based on the Σ -optimality criterion. A simple yet effective heuristic is also introduced for increasing the exploration capabilities of the sampler, and reducing bias of the resultant classifier, by adjusting the confidence on the model label predictions. The novel sampling strategies are based on quantities that are readily available without the need for model re-training, rendering them computationally efficient and scalable to large graphs. Numerical tests using synthetic and real data demonstrate that the proposed methods achieve accuracy that is comparable or superior to the state of the art even at reduced runtime.

Index Terms—Active learning, classification, expected change, graph-based.

I. INTRODUCTION

ACTIVE learning has recently gained popularity for various applications ranging from bioinformatics [1] to distributed signal classification and estimation [2]. In contrast to traditional passive supervised and semi-supervised learning methods, where classifiers are trained using given or randomly drawn labeled data, active learning allows for judiciously selecting which data are to be queried and added to the training set. Thus, a typical active learner begins with a small labeled set, and proceeds to select one or more informative query instances from a large unlabeled pool. Active learning yields markedly

Manuscript received February 17, 2018; revised May 20, 2018 and July 14, 2018; accepted August 13, 2018. Date of publication August 24, 2018; date of current version September 5, 2018. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Wee Peng Tay. This work was supported by National Science Foundation under Grant 1711471, under Grant 1500713, and under Grant 1442686. (Corresponding author: Georgios B. Giannakis.)

The authors are with the Department of Electronic and Communication Engineering and Digital Technology Center, University of Minnesota, Minneapolis, MN 55455 USA (e-mail: bermp001@umn.edu; georgios@umn.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSP.2018.2866812

improved classification accuracy over passive or random sampling when the number of training labels is fixed [3]–[6]. It can be particularly appealing when unlabeled data (instances) are readily available, but obtaining training labels is expensive. For instance, a classifier trained to predict the presence of cancer based on certain protein attributes requires labels that involve costly and time-consuming medical examinations; see e.g., [1].

Even intuitively, one expects active sampling to outperform random sampling when (un)labeled instances are correlated. Such a case emerges with graph-aware classification, where each instance is denoted by a node, while edges capture correlation among connected nodes. Although graphs may arise naturally in certain applications (e.g. social and citation networks), they can in general be constructed from any set of nodal feature vectors using proper similarity measures; see e.g., [7], [8]. Knowing the graph topology, *graph-cognizant* classification boils down to propagating the information from labeled nodes to unlabeled ones through edges of neighboring nodes; see e.g., [9]. As a result, classification on graphs is inherently semi-supervised and thus conducive to active learning.

Prior art in graph-based active learning can be divided in two categories. The first includes the *non-adaptive* design-of-experiments-type methods, where sampling strategies are designed *offline* depending only on the graph structure, based on ensemble optimality criteria. The non-adaptive category also includes the variance minimization sampling [10], as well as the error upper bound minimization in [11], and the data non-adaptive Σ -optimality approach in [12]. The second category includes methods that select samples adaptively and jointly with the classification process, taking into account both graph structure as well as previously obtained labels. Such *data-adaptive* methods give rise to sampling schemes that are not optimal on average, but adapt to a given realization of labels on the graph. Adaptive methods include the Bayesian risk minimization [13], the information gain maximization [14], as well as the manifold preserving method of [15]; see also [16]–[18]. Finally, related works deal with selective sampling of nodes that arrive sequentially in a gradually augmented graph [19]–[21], as well as active sampling to infer the graph structure [22], [23].

In this context, the present work develops data-adaptive pool based active samplers for graph-aware classification. The proposed sampling strategy relies on querying the node that is expected to inflict the largest change on the underlying label correlation model. Albeit in different context, a related criterion was adopted for semantic segmentation of images [24], and for regression of Gaussian processes [25]. The proposed

approach here advocates novel metrics of expected model change, establishing connections to existing methods such as uncertainty sampling, variance minimization and sampling based on the Σ –optimality criterion. A simple yet effective heuristic is also introduced for improving the exploration capabilities, and for reducing the bias of the resultant classifiers, by taking into account the confidence on the model label predictions.

The rest of the paper is organized as follows. Section II states the problem, and presents the GMRF model adopted to approximate the marginal distributions of the unknown categorical node labels. Section III develops active learning methods based on different measures of change, and establishes links to existing sampling schemes. Section IV points out the issue of sampling bias, and mitigates it by incorporating a confidence metric on the underlying model. Finally, Section V presents numerical experiments on real and synthetic datasets.

Notation: Lower- (upper-) case boldface letters denote column vectors (matrices). Calligraphic symbols are reserved for sets, while T stands for transposition. Vectors $\mathbf{0}$, $\mathbf{1}$, and \mathbf{e}_n denote the all-zeros, the all-ones, and the n -th canonical vector, respectively. Symbol $\mathbb{1}_E$ denotes the indicator for the event E . Notation $\mathcal{N}(\mathbf{m}, \mathbf{C})$ stands for the multivariate Gaussian distribution with mean \mathbf{m} and covariance matrix \mathbf{C} , while $\text{tr}(\mathbf{X})$, $\lambda_{\min}(\mathbf{X})$, and $\lambda_{\max}(\mathbf{X})$ are reserved for the trace, the minimum and maximum eigenvalues of matrix \mathbf{X} , respectively. Symbol $[\mathbf{x}]_i$ denotes the i -th entry of vector \mathbf{x} .

II. MODELING AND PROBLEM STATEMENT

Consider a connected undirected graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, where \mathcal{V} is the set of N nodes, and \mathcal{E} contains the edges that are also represented by the $N \times N$ weighted adjacency matrix \mathbf{W} whose (i, j) -th entry denotes the weight of the edge that connects nodes v_i and v_j . Let us further suppose that a binary label $y_i \in \{-1, 1\}$ is associated with each node v_i . The weighted binary labeled graph can either be given, or, it can be inferred from a set of N data points $\{\mathbf{x}_i, y_i\}_{i=1}^N$ such that each node of the graph corresponds to a data point. Matrix \mathbf{W} can be obtained from the feature vectors $\{\mathbf{x}_i\}_{i=1}^N$ using different similarity measures. For example, one may use the radial basis function $w_{i,j} = \exp(-\|\mathbf{x}_i - \mathbf{x}_j\|_2^2/\sigma^2)$ that assigns large edge weights to pairs of points that are neighbors in Euclidean space, or the Pearson correlation coefficients $w_{i,j} = \langle \mathbf{x}_i, \mathbf{x}_j \rangle / (\|\mathbf{x}_i\|_2 \|\mathbf{x}_j\|_2)$. If $w_{i,j} \neq 0 \forall i, j$, the resulting graph will be fully connected, but one may obtain a more structured graph by rounding small weights to 0.

Having embedded the data on a graph, semi-supervised learning amounts to propagating an observed subset of labels to the rest of the network. Thus, upon observing $\{y_i\}_{i \in \mathcal{L}}$ where $\mathcal{L} \subseteq \mathcal{V}$, henceforth collected in the $|\mathcal{L}| \times 1$ vector $\mathbf{y}_{\mathcal{L}}$, the goal is to infer the labels of the unlabeled nodes $\{y_i\}_{i \in \mathcal{U}}$ concatenated in the vector $\mathbf{y}_{\mathcal{U}}$, where $\mathcal{U} := \mathcal{V} \setminus \mathcal{L}$. Let us consider labels as random variables that follow an unknown joint distribution $(y_1, y_2, \dots, y_N) \sim p(y_1, y_2, \dots, y_N)$, or $\mathbf{y} \sim p(\mathbf{y})$ for brevity.

For the purpose of inferring unobserved from observed labels, it would suffice if the joint posterior distribution $p(\mathbf{y}_{\mathcal{U}}|\mathbf{y}_{\mathcal{L}})$ were available; then, $\mathbf{y}_{\mathcal{U}}$ could be obtained as a combination of labels

that maximizes $p(\mathbf{y}_{\mathcal{U}}|\mathbf{y}_{\mathcal{L}})$. Moreover, obtaining the marginal posterior distributions $p(y_i|\mathbf{y}_{\mathcal{L}})$ of each unlabeled node i is often of interest, especially in the present greedy active sampling approach. To this end, it is well documented that MRFs are suitable for modeling probability mass functions over undirected graphs using the generic form, see e.g., [13]

$$p(\mathbf{y}) := \frac{1}{Z_{\beta}} \exp\left(-\frac{\beta}{2} \Phi(\mathbf{y})\right) \quad (1a)$$

where the “partition function” Z_{β} ensures that (1a) integrates to 1, β is a scalar that controls the smoothness of $p(\mathbf{y})$, and $\Phi(\mathbf{y})$ is the so termed “energy” of a realization \mathbf{y} , given by

$$\Phi(\mathbf{y}) := \frac{1}{2} \sum_{i,j \in \mathcal{V}} w_{i,j} (y_i - y_j)^2 = \mathbf{y}^T \mathbf{L} \mathbf{y} \quad (1b)$$

that captures the graph-induced label dependencies through the graph Laplacian matrix $\mathbf{L} := \mathbf{D} - \mathbf{W}$ with $\mathbf{D} := \text{diag}(\mathbf{W}\mathbf{1})$. This categorical MRF model in (1a) naturally incorporates the known graph structure (through \mathbf{L}) in the label distribution by assuming label configurations where nearby labels (large edge weights) are similar, and have lower energy as well as higher likelihood. Still, finding the joint and marginal posteriors using (1a) and (1b) incurs *exponential complexity* since $\mathbf{y}_{\mathcal{U}} \in \{-1, 1\}^{|\mathcal{U}|}$. To deal with this challenge, less complex continuous-valued models are well motivated for a scalable approximation of the marginal posteriors. This prompts our next step to allow for continuous-valued label configurations $\psi_{\mathcal{U}} \in \mathbb{R}^{|\mathcal{U}|}$ that are modeled by a GMRF.

A. GMRF Relaxation

Consider approximating the binary field $\mathbf{y} \in \{-1, 1\}^{|\mathcal{U}|}$ that is distributed according to (1a) with the continuous-valued $\psi \sim \mathcal{N}(\mathbf{0}, \mathbf{C})$, where the covariance matrix satisfies $\mathbf{C}^{-1} = \mathbf{L}$. Label propagation under this relaxed GMRF model becomes readily available in closed form. Indeed, $\psi_{\mathcal{U}|\mathcal{L}}$ of unlabeled nodes conditioned on the labeled ones obeys

$$\psi_{\mathcal{U}|\mathcal{L}} \sim \mathcal{N}(\boldsymbol{\mu}_{\mathcal{U}|\mathcal{L}}, \mathbf{L}_{\mathcal{U}\mathcal{U}}^{-1}) \quad (2)$$

where $\mathbf{L}_{\mathcal{U}\mathcal{U}}$ is the part of the graph Laplacian that corresponds to unlabeled nodes in the partitioning

$$\mathbf{L} = \begin{bmatrix} \mathbf{L}_{\mathcal{U}\mathcal{U}} & \mathbf{L}_{\mathcal{U}\mathcal{L}} \\ \mathbf{L}_{\mathcal{L}\mathcal{U}} & \mathbf{L}_{\mathcal{L}\mathcal{L}} \end{bmatrix}. \quad (3)$$

Given the observed $\psi_{\mathcal{L}}$, the minimum mean-square error (MMSE) estimator of $\psi_{\mathcal{U}}$ is given by the conditional expectation

$$\begin{aligned} \boldsymbol{\mu}_{\mathcal{U}|\mathcal{L}} &= \mathbf{C}_{\mathcal{U}\mathcal{L}} \mathbf{C}_{\mathcal{L}\mathcal{L}}^{-1} \psi_{\mathcal{L}} \\ &= -\mathbf{L}_{\mathcal{U}\mathcal{U}}^{-1} \mathbf{L}_{\mathcal{U}\mathcal{L}} \psi_{\mathcal{L}} \end{aligned} \quad (4)$$

where the first equality holds because for jointly Gaussian zero-mean vectors the MMSE estimator coincides with the linear (L)MMSE one (see e.g., [26, p. 382]), while the second equality is derived in Appendix A1. When binary labels $\mathbf{y}_{\mathcal{L}}$ are obtained, they can be treated as measurements of the continuous field

$(\psi_{\mathcal{L}} := \mathbf{y}_{\mathcal{L}})$, and (4) reduces to

$$\boldsymbol{\mu}_{\mathcal{U}|\mathcal{L}} = -\mathbf{L}_{\mathcal{U}\mathcal{U}}^{-1} \mathbf{L}_{\mathcal{U}\mathcal{L}} \mathbf{y}_{\mathcal{L}}. \quad (5)$$

Interestingly, the conditional mean of the GMRF in (5) may serve as an approximation of the marginal posteriors of the unknown labels. Specifically, for the i -th node, we adopt the approximation

$$\begin{aligned} p(y_i = 1 | \mathbf{y}_{\mathcal{L}}) &= \frac{1}{2} (\mathbb{E} [[\mathbf{y}_{\mathcal{U}|\mathcal{L}}]_i] + 1) \\ &\approx \frac{1}{2} (\mathbb{E} [[\boldsymbol{\mu}_{\mathcal{U}|\mathcal{L}}]_i] + 1) \\ &= \frac{1}{2} ([\boldsymbol{\mu}_{\mathcal{U}|\mathcal{L}}]_i + 1) \end{aligned} \quad (6)$$

where the first equality follows from the fact that the expectation of a Bernoulli random variable equals its probability. Given the approximation of $p(y_i | \mathbf{y}_{\mathcal{L}})$ in (6), and the uninformative prior $p(y_i = 1) = 0.5 \forall i \in \mathcal{V}$, the maximum a posteriori (MAP) estimate of y_i , which in the Gaussian case here reduces to the minimum distance decision rule, is given as

$$\hat{y}_i = \begin{cases} 1 & [\boldsymbol{\mu}_{\mathcal{U}|\mathcal{L}}]_i > 0 \\ -1 & \text{else} \end{cases}, \quad \forall i \in \mathcal{U} \quad (7)$$

thus completing the propagation of the observed $\mathbf{y}_{\mathcal{L}}$ to the unlabeled nodes of the graph.

It is worth stressing at this point, that as the set of labeled samples changes, so does the dimensionality of the conditional mean in (5), along with the “auto-” and “cross-” Laplacian sub-matrices that enable soft label propagation via (5), and hard label propagation through (7). Two remarks are now in order.

Remark 1: It is well known that the Laplacian of a graph is not invertible, since $\mathbf{L}\mathbf{1} = \mathbf{0}$; see, e.g. [27]. To deal with this issue, we instead use $\mathbf{L} + \delta\mathbf{I}$, where $\delta \ll 1$ is selected arbitrarily small but large enough to guarantee the numerical stability of e.g., $\mathbf{L}_{\mathcal{U}\mathcal{U}}$ in (5). A closer look at the energy function $\Phi(\mathbf{y}) := \sum_{i,j \in \mathcal{V}} w_{i,j} (y_i - y_j)^2 + \delta \sum_{i \in \mathcal{V}} y_i^2$ reveals that this simple modification amounts to adding a “self-loop” of weight δ to each node of the graph. Alternatively, δ can be viewed as a regularizer that “pushes” the entries of the Gaussian field $\psi_{\mathcal{U}}$ closer to 0, which also causes the (approximated) marginal posteriors $p(y_i | \mathbf{y}_{\mathcal{L}})$ to be closer to 0.5 (cf. eq. (6)). In that sense, δ enforces the priors $p(y_i = 1) = p(y_i = -1) = 0.5$.

Remark 2: The method introduced here for label propagation (cf. (5)) is algorithmically similar to the one reported in [13]. The main differences are: i) we perform soft label propagation by minimizing the mean-square prediction error of unlabeled from labeled samples; and ii) our model approximates $\{-1, 1\}$ labels with a *zero-mean* Gaussian field, while the model in [13] approximates $\{0, 1\}$ labels also with a zero-mean Gaussian field (instead of one centered at 0.5). Apparently, [13] treats the two classes differently since it exhibits a bias towards class 0; thus, simply denoting class 0 as class 1 yields different marginal posteriors and classification results. In contrast, our model is bias-free and treats the two classes equally.

Algorithm 1: Active Graph Sampling Algorithm.

Input: Adjacency matrix \mathbf{W} , $\delta \ll 1$

Initialize: $\mathcal{U}^0 = \mathcal{V}$, $\mathcal{L}^0 = \emptyset$, $\boldsymbol{\mu} = \mathbf{0}$, $\mathbf{G}_0 = (\mathbf{L} + \delta\mathbf{I})^{-1}$

First query is chosen at random

for $t = 1 : T$ **do**

 Scan \mathcal{U}^{t-1} to find best query node v_{k_t} as in (8)

 Obtain label y_{k_t} of v_{k_t}

 Update the GMRF mean as in (9)

 Update \mathbf{G}_t as in (10)

$\mathcal{U}^t = \mathcal{U}^{t-1} / \{k_t\}$, $\mathcal{L}^t = \mathcal{L}^{t-1} \cup \{k_t\}$

end for

Predict remaining unlabeled nodes as in (7)

B. Active Sampling With GMRFs

In passive learning, \mathcal{L} is either chosen at random, or, it is determined a priori. In our pool based active learning setup, the learner can examine a set of instances (nodes in the context of graph-cognizant classification), and can choose which instances to label. Given its cardinality $|\mathcal{L}|$, one way to approximate the exponentially complex task of selecting \mathcal{L} is to *greedily* sample one node per iteration t with index

$$k_t = \arg \max_{i \in \mathcal{U}^{t-1}} U(v_i, \mathcal{L}^{t-1}) \quad (8)$$

where \mathcal{U}^{t-1} is the unlabeled set at time $t - 1$, and $U(v, \mathcal{L}^{t-1})$ is a utility function that evaluates how informative node v is while taking into account information already available in \mathcal{L}^{t-1} . Upon disclosing label y_{k_t} , it can be shown that instead of resolving (5), the GMRF mean can be updated recursively using the “dongle node” trick in [13] as

$$\boldsymbol{\mu}_{\mathcal{U}^{t-1}|\mathcal{L}^{t-1}}^{+y_{k_t}} = \boldsymbol{\mu}_{\mathcal{U}^{t-1}|\mathcal{L}^{t-1}} + \frac{1}{g_{k_t k_t}} (y_{k_t} - [\boldsymbol{\mu}_{\mathcal{U}^{t-1}|\mathcal{L}^{t-1}}]_{k_t}) \mathbf{g}_{k_t} \quad (9)$$

where $\boldsymbol{\mu}_{\mathcal{U}^{t-1}|\mathcal{L}^{t-1}}^{+y_{k_t}}$ is the conditional mean of the unlabeled nodes when node v_{k_t} is assigned label y_{k_t} (thus “gravitating” the GMRF mean $[\boldsymbol{\mu}_{\mathcal{U}^{t-1}|\mathcal{L}^{t-1}}]_{k_t}$ toward its replacement y_{k_t}); vector $\mathbf{g}_{k_t} := [\mathbf{L}_{\mathcal{U}^{t-1}\mathcal{U}^{t-1}}^{-1}]_{k_t}$ and scalar $g_{k_t k_t} := [\mathbf{L}_{\mathcal{U}^{t-1}\mathcal{U}^{t-1}}^{-1}]_{k_t k_t}$ are the k_t -th column and diagonal entry of the Laplacian inverse, respectively. Subsequently, the new conditional mean vector $\boldsymbol{\mu}_{\mathcal{U}^t|\mathcal{L}^t}$ defined over \mathcal{U}^t is given by removing the i -th entry of $\boldsymbol{\mu}_{\mathcal{U}^{t-1}|\mathcal{L}^{t-1}}^{+y_{k_t}}$. Using Shur’s lemma it can be shown that the inverse Laplacian $\mathbf{G}_t^{-k_t}$ when the k_t -th node is removed from the unlabeled sub-graph can be efficiently updated from $\mathbf{G}_t := \mathbf{L}_{\mathcal{U}^t\mathcal{U}^t}^{-1}$ as [12]

$$\begin{bmatrix} \mathbf{G}_t^{-k_t} & \mathbf{0} \\ \mathbf{0}^T & 0 \end{bmatrix} = \mathbf{G}_t - \frac{1}{g_{k_t k_t}} \mathbf{g}_{k_t} \mathbf{g}_{k_t}^T \quad (10)$$

which requires only $\mathcal{O}(|\mathcal{U}|^2)$ computations instead of $\mathcal{O}(|\mathcal{U}|^3)$ for matrix inversion. Alternatively, one may obtain $\mathbf{G}_t^{-k_t}$ by applying the matrix inversion lemma employed by the RLS-like solver in [13]. The resultant greedy active sampling scheme for graphs is summarized in Algorithm 1.

Remark 3: Existing data-adaptive sampling schemes, e.g., [13], [16], [14], often require *model-retraining* by examining

candidate labels per unlabeled node (cf. (8)). Thus, even when retraining is efficient, it still needs to be performed $|\mathcal{U}| \times \# \text{Classes}$ times per iteration of Algorithm 1, which in practice significantly increases runtime, especially for larger graphs.

In summary, different sampling strategies emerge by selecting distinct utilities $U(v, \mathcal{L}^{t-1})$ in (8). In this context, the goal of the present work is to develop novel active learning schemes within a maximum-expected change framework that achieve high accuracy with a small number of samples. A further desirable attribute of the sought approach is to bypass the need for GMRF retraining.

III. EXPECTED MODEL CHANGE

Judiciously selecting the utility function is of paramount importance in designing an efficient active sampling algorithm. In the present work, we introduce and investigate the relative merits of different choices under the prism of expected change (EC) measures that we advocate as information-revealing utility functions. From a high-level vantage point, the idea is to identify and sample nodes of the graph that are expected to have the greatest impact on the available GMRF model of the unknown labels. Thus, contrary to the expected error reduction and entropy minimization approaches that actively sample with the goal of increasing the “confidence” on the model, our focus is on effecting maximum perturbation of the model with each node sampled. The intuition behind our approach is that by sampling nodes with large impact, one may take faster steps towards an increasingly accurate model.

A. EC of Model Predictions

An intuitive measure of expected model change for a given node v_i is the expected number of unlabeled nodes whose label prediction will change after sampling the label of v_i . To start, consider per node i the measure

$$F(y_i, \boldsymbol{\mu}_{\mathcal{U}|\mathcal{L}}) := \sum_{j \in \mathcal{U} - \{i\}} \mathbb{1}_{\{\hat{y}_j^{+y_i} \neq \hat{y}_j\}} \quad (11)$$

where $\hat{y}_j^{+y_i}$ is the predicted label for the j -th node after the label of the i -th node is revealed, denoting the number of such “flips” in the predicted labels of (7). For notational brevity, we henceforth let $\mu_i = [\boldsymbol{\mu}_{\mathcal{U}|\mathcal{L}}]_i$. The corresponding utility function is

$$\begin{aligned} U_{FL}(v_i, \mathcal{L}) &:= \mathbb{E}_{y_i | \mathbf{y}_{\mathcal{L}}} [F(y_i, \boldsymbol{\mu}_{\mathcal{U}|\mathcal{L}})] \\ &= p(y_i = 1 | \mathbf{y}_{\mathcal{L}}) F(y_i = 1, \boldsymbol{\mu}_{\mathcal{U}|\mathcal{L}}) \\ &\quad + p(y_i = -1 | \mathbf{y}_{\mathcal{L}}) F(y_i = -1, \boldsymbol{\mu}_{\mathcal{U}|\mathcal{L}}) \\ &\approx \frac{1}{2} (\mu_i + 1) F(y_i = 1, \boldsymbol{\mu}_{\mathcal{U}|\mathcal{L}}) \\ &\quad + \left(1 - \frac{1}{2} (\mu_i + 1)\right) F(y_i = -1, \boldsymbol{\mu}_{\mathcal{U}|\mathcal{L}}) \quad (12) \end{aligned}$$

where the approximation is because (6) was used in place of $p(y_i = 1 | \mathbf{y}_{\mathcal{L}})$. Note that model retraining using (9) is required to be performed twice (in general, as many as the number of

classes) for each node in \mathcal{U} in order to obtain the labels $\{\hat{y}_j\}^{+y_i}$ in (11).

B. EC Using KL Divergence

The utility function in (12) depends on the hard label decisions of (7), but does not account for perturbations that do not lead to changes in label predictions. To obtain utility functions that are more sensitive to the soft GMRF model change, it is prudent to measure how much the continuous distribution of the unknown labels changes after sampling. Towards this goal, we consider first the KL divergence between two pdfs $p(\mathbf{x})$ and $q(\mathbf{x})$, which is defined as

$$\mathcal{D}_{KL}(p || q) := \int p(\mathbf{x}) \ln \frac{p(\mathbf{x})}{q(\mathbf{x})} d\mathbf{x} = \mathbb{E}_p \left[\ln \frac{p(\mathbf{x})}{q(\mathbf{x})} \right].$$

For the special case where $p(\mathbf{x})$ and $q(\mathbf{x})$ are Gaussian with identical covariance matrix \mathbf{C} and corresponding means \mathbf{m}_p and \mathbf{m}_q , their KL divergence is expressible in closed form as

$$\mathcal{D}_{KL}(p || q) = \frac{1}{2} (\mathbf{m}_p - \mathbf{m}_q)^T \mathbf{C}^{-1} (\mathbf{m}_p - \mathbf{m}_q) \quad (13)$$

Upon recalling that $\boldsymbol{\psi}_{\mathcal{U}}$ defined over the unlabeled nodes is Gaussian [cf. (2)], and since the Gaussian field obtained after node $v_i \in \mathcal{U}$ is designated label y_i is also Gaussian, we have

$$\boldsymbol{\psi}_{\mathcal{U}}^{+y_i} \sim \mathcal{N}(\boldsymbol{\mu}_{\mathcal{U}|\mathcal{L}}^{+y_i}, \mathbf{L}_{\mathcal{U}|\mathcal{L}}^{-1}). \quad (14)$$

It thus follows that the KL divergence induced on the GMRF after sampling y_i is (cf. (13))

$$\begin{aligned} \mathcal{D}_{KL}(\boldsymbol{\psi}_{\mathcal{U}}^{+y_i} || \boldsymbol{\psi}_{\mathcal{U}}) &= \frac{1}{2} \left[(\boldsymbol{\mu}_{\mathcal{U}|\mathcal{L}}^{+y_i} - \boldsymbol{\mu}_{\mathcal{U}|\mathcal{L}})^T \mathbf{L}_{\mathcal{U}|\mathcal{L}} (\boldsymbol{\mu}_{\mathcal{U}|\mathcal{L}}^{+y_i} - \boldsymbol{\mu}_{\mathcal{U}|\mathcal{L}}) \right] \\ &= \frac{1}{2g_{ii}^2} (y_i - \mu_i)^2 \mathbf{g}_i^T \mathbf{L}_{\mathcal{U}|\mathcal{L}} \mathbf{g}_i = \frac{1}{2g_{ii}} (y_i - \mu_i)^2 \end{aligned} \quad (15)$$

where the second equality relied on (9), and the last equality used the definition of g_{ii} . The divergence in (15) can be also interpreted as the normalized innovation of observation y_i . Averaging (15) over the candidate values of y_i yields the expected KL divergence of the GMRF utility as

$$\begin{aligned} U_{KLG}(v_i, \mathcal{L}) &:= \mathbb{E}_{y_i | \mathbf{y}_{\mathcal{L}}} [\mathcal{D}_{KL}(\boldsymbol{\psi}_{\mathcal{U}}^{+y_i} || \boldsymbol{\psi}_{\mathcal{U}})] \\ &= p(y_i = 1 | \mathbf{y}_{\mathcal{L}}) \mathcal{D}_{KL}(\boldsymbol{\psi}_{\mathcal{U}}^{+y_i=1} || \boldsymbol{\psi}_{\mathcal{U}}) \\ &\quad + p(y_i = -1 | \mathbf{y}_{\mathcal{L}}) \mathcal{D}_{KL}(\boldsymbol{\psi}_{\mathcal{U}}^{+y_i=-1} || \boldsymbol{\psi}_{\mathcal{U}}) \\ &\approx \frac{1}{2} \left[\frac{1}{2g_{ii}} (\mu_i + 1)(1 - \mu_i)^2 \right. \\ &\quad \left. + \left(1 - \frac{1}{2} (\mu_i + 1)\right) \frac{1}{g_{ii}} (-1 - \mu_i)^2 \right] \\ &= \frac{1}{2g_{ii}} (1 - \mu_i^2). \end{aligned} \quad (16)$$

Interestingly, the utility in (16) leads to a form of uncertainty sampling, since $(1 - \mu_i^2)$ is a measure of uncertainty of the model prediction for node v_i , further normalized by g_{ii} , which is the variance of the Gaussian field (cf. [10]). Note also that the

expected KL divergence in (16) also relates to the information gain between $\{\psi_j\}_{j \in \mathcal{U}/\{i\}}$ and y_i .

Albeit easy to compute since model retraining is not required, U_{KLG} quantifies the impact of disclosing y_i on the GMRF, but not the labels $\{y_j\}_{j \in \mathcal{U}/\{i\}}$ themselves. To account for the labels themselves, an alternative KL-based utility function could treat $\{y_j\}_{j \in \mathcal{U}/\{i\}}$ as Bernoulli variables [c.f. (6)]; that is

$$y_j \sim \text{Ber}((\mu_j + 1)/2). \quad (17)$$

In that case, one would ideally be interested in obtaining the expected KL divergence between the true posteriors, that is

$$\mathbb{E}_{y_i|\mathbf{y}_{\mathcal{L}}} [\mathcal{D}_{KL}(p(\mathbf{y}_{\mathcal{U}}|\mathbf{y}_{\mathcal{L}}, y_i) || p(\mathbf{y}_{\mathcal{U}}|\mathbf{y}_{\mathcal{L}}))]. \quad (18)$$

Nevertheless, the joint pdfs of the labels are not available by the GMRF model; in fact, any attempt at evaluating the joint posteriors incurs exponential complexity as mentioned in Section II. One way to bypass this limitation is by approximating the joint posterior $p(\mathbf{y}_{\mathcal{U}}|\mathbf{y}_{\mathcal{L}})$ with the product of marginal posteriors $\prod_{j \in \mathcal{U}} p(y_j|\mathbf{y}_{\mathcal{L}})$, since the later are readily given by the GMRF. Using this independence assumption causes the joint KL divergence in (18) to break down to the sum of marginal per-unlabeled-node KL divergences. The resulting utility function can be expressed as

$$U_{KL}(v_i, \mathcal{L}) := \sum_{j \in \mathcal{U}/\{i\}} I(y_j, y_i) \quad (19)$$

where

$$\begin{aligned} I(y_j, y_i) &:= \mathbb{E}_{y_i|\mathbf{y}_{\mathcal{L}}} [\mathcal{D}_{KL}(p(y_j|\mathbf{y}_{\mathcal{L}}, y_i) || p(y_j|\mathbf{y}_{\mathcal{L}}))] \\ &\approx \frac{1}{2}(\mu_i + 1) \mathcal{D}_{KL}(y_j^{+y_i=1} || y_j) \\ &\quad + \left(1 - \frac{1}{2}(\mu_i + 1)\right) \mathcal{D}_{KL}(y_j^{+y_i=-1} || y_j) \end{aligned} \quad (20)$$

since for univariate distributions the expected KL divergence between the prior and posterior is equivalent to the mutual information between the observed random variable and its unknown label. Note also that the KL divergence between univariate distributions is simply

$$\mathcal{D}_{KL}(y_j^{+y_i} || y_j) = H(y_j^{+y_i}, y_j) - H(y_j^{+y_i}) \quad (21)$$

where $H(y_j^{+y_i}, y_j)$ denotes the cross-entropy, which for Bernoulli variables is

$$\begin{aligned} H(y_j^{+y_i}, y_j) &= -\frac{1}{2}(\mu_j^{+y_i} + 1) \log \frac{1}{2}(\mu_j + 1) \\ &\quad - \left[1 - \frac{1}{2}(\mu_j^{+y_i} + 1)\right] \log \left[1 - \frac{1}{2}(\mu_j + 1)\right]. \end{aligned} \quad (22)$$

Combining (19)–(22) yields U_{KL} . Intuitively, this utility promotes the sampling of nodes that are expected to induce large change on the model (cross-entropy between old and new distributions), while at the same time increasing the “confidence” on the model (negative entropy of updated distribution). Furthermore, the mutual-information-based expressions (19) and (20) establish a connection to the information-based metrics in [14]

and [28], giving an expected-model-change interpretation of the entropy reduction method.

C. EC Without Model Retraining

In this section, we introduce two measures of model change that do not require model retraining (cf. Remark 3), and hence are attractive in their simplicity. Specifically, retraining (i.e., computing $\mu_{\mathcal{U}^{t-1}|\mathcal{L}^{t-1}}^{+y_i}$, $\forall i \in \mathcal{U}$ and $\forall y_i \in \mathcal{Y}$) is not required if per-node utility $U(v, \mathcal{L}^{t-1})$ can be given in *closed-form* as a function of \mathbf{G}_{t-1} and $\mu_{\mathcal{U}^{t-1}|\mathcal{L}^{t-1}}$. Two such measures are explored here: one based on the sum of total variations that a new sample inflicts on the (approximate) marginal distributions of the unknown labels, and one based on the mean-square deviation that a new sample is expected to inflict on the GMRF.

The *total variation* between two probability distributions $p(x)$ and $q(x)$ over a finite alphabet \mathcal{X} is

$$\delta(p, q) := \frac{1}{2} \sum_{x \in \mathcal{X}} |p(x) - q(x)|.$$

Using the approximation in (6), the total variation between the distribution of an unknown label y_j and the same label $y_j^{+y_i}$ after y_i becomes available is

$$\begin{aligned} \delta(y_j^{+y_i}, y_j) &= \frac{1}{2} (|\mu_j^{+y_i} - \mu_j| + |1 - \mu_j^{+y_i} - (1 - \mu_j)|) \\ &= |\mu_j^{+y_i} - \mu_j|. \end{aligned} \quad (23)$$

Consequently, the sum of total variations over all the unlabeled nodes $\{v_j\}_{j \in \mathcal{U}/\{i\}}$ is

$$\begin{aligned} \Delta(\mathbf{y}_{\mathcal{U}}^{+y_i}, \mathbf{y}_{\mathcal{U}}) &:= \sum_{j \in \mathcal{U}} \delta(y_j^{+y_i}, y_j) = \|\mu_{\mathcal{U}|\mathcal{L}}^{+y_i} - \mu_{\mathcal{U}|\mathcal{L}}\|_1 \\ &= \frac{1}{g_{ii}} |y_i - \mu_i| \|\mathbf{g}_i\|_1 \end{aligned}$$

where the second equality follows by concatenating all total variations (cf. (23)) in vector form, and the last one follows by the GMRF update rule in (9). Finally, the expected sum of total variations utility score-function is defined as

$$\begin{aligned} U_{TV}(v_i, \mathcal{L}) &:= \mathbb{E}_{y_i|\mathbf{y}_{\mathcal{L}}} [\Delta(\mathbf{y}_{\mathcal{U}}^{+y_i}, \mathbf{y}_{\mathcal{U}})] \\ &= \mathbb{E}_{y_i|\mathbf{y}_{\mathcal{L}}} [|y_i - \mu_i|] \frac{1}{g_{ii}} \|\mathbf{g}_i\|_1 \end{aligned}$$

and since

$$\begin{aligned} \mathbb{E}_{y_i|\mathbf{y}_{\mathcal{L}}} [|y_i - \mu_i|] &= p(y_i = 1|\mathbf{y}_{\mathcal{L}}) |1 - \mu_i| \\ &\quad + p(y_i = -1|\mathbf{y}_{\mathcal{L}}) |-1 - \mu_i| \\ &\approx 2(1 - \mu_i^2) \end{aligned}$$

it follows that the utility function based on total variation can be expressed as

$$U_{TV}(v_i, \mathcal{L}) = \frac{2}{g_{ii}} (1 - \mu_i^2) \|\mathbf{g}_i\|_1. \quad (24)$$

The second measure is based on the *mean-square deviation* (MSD) between two RV's X_1 and X_2

$$\begin{aligned} \text{MSD}(X_1, X_2) &:= \int (X_1 - X_2)^2 f(X_1, X_2) dX_1 dX_2 \\ &= \mathbb{E} \left[(X_1 - X_2)^2 \right]. \end{aligned}$$

Our next proposed utility score is the expected MSD between the Gaussian fields $\psi_{\mathcal{U}}$ and $\psi_{\mathcal{U}}^{+y_i}$ before and after obtaining y_i ; that is,

$$\begin{aligned} U_{MSD}(v_i, \mathcal{L}) &= \mathbb{E}_{y_i | \mathcal{Y}_{\mathcal{L}}} [\text{MSD}(\psi_{\mathcal{U}}^{+y_i}, \psi_{\mathcal{U}})] \\ &\approx \frac{1}{2}(\mu_i + 1) \text{MSD}(\psi_{\mathcal{U}}^{+y_i=1}, \psi_{\mathcal{U}}) \\ &\quad + \left(1 - \frac{1}{2}(\mu_i + 1)\right) \text{MSD}(\psi_{\mathcal{U}}^{+y_i=-1}, \psi_{\mathcal{U}}) \end{aligned} \quad (25)$$

where

$$\begin{aligned} \text{MSD}(\psi_{\mathcal{U}}^{+y_i}, \psi_{\mathcal{U}}) &:= \mathbb{E} [\|\psi_{\mathcal{U}}^{+y_i} - \psi_{\mathcal{U}}\|^2] \\ &= 2\text{tr}(\mathbf{L}_{\mathcal{U}\mathcal{U}}^{-1}) + \|\boldsymbol{\mu}_{\mathcal{U}|\mathcal{L}}^{+y_i} - \boldsymbol{\mu}_{\mathcal{U}|\mathcal{L}}\|^2 \\ &\propto \frac{1}{g_{ii}^2} (y_i - \mu_i)^2 \|\mathbf{g}_i\|_2^2. \end{aligned} \quad (26)$$

The second equality in (26) is derived in Appendix A2 under the assumption that $\psi_{\mathcal{U}}$ and $\psi_{\mathcal{U}}^{+y_i}$ are independent random vectors. Furthermore, the term $2\text{tr}(\mathbf{L}_{\mathcal{U}\mathcal{U}}^{-1})$ is ignored since it does not depend on y_i , and the final expression of (26) is obtained using (9). Finally, substituting (26) into (25) yields the following closed-form expression of the MSD-based utility score function

$$U_{MSD}(v_i, \mathcal{L}) \propto (1 - \mu_i^2) \frac{\|\mathbf{g}_i\|_2^2}{g_{ii}^2}. \quad (27)$$

Note that U_{TV} and U_{MSD} are proportional to the expected KL divergence of the Gaussian field U_{KLG} in the previous section since

$$U_{TV}(v_i, \mathcal{L}) \propto U_{KLG}(v_i, \mathcal{L}) \|\mathbf{g}_i\|_1 \quad (28)$$

and

$$U_{MSD}(v_i, \mathcal{L}) \propto U_{KLG}(v_i, \mathcal{L}) \|\mathbf{g}_i\|_2 \quad (29)$$

with the norms $\|\mathbf{g}_i\|_1$ and $\|\mathbf{g}_i\|_2$ quantifying the average influence of the i -th node over the rest of the unlabeled nodes.

It is worth mentioning that our TV- and MSD-based methods relate to the Σ -optimality-based active learning [12] and the variance minimization [10] correspondingly. This becomes apparent upon recalling that Σ -optimality and variance-minimization utility score functions are respectively given by

$$U_{\Sigma-opt}(v_i) = \frac{\|\mathbf{g}_i\|_1^2}{g_{ii}}$$

and

$$U_{VM}(v_i) := \frac{\|\mathbf{g}_i\|_2^2}{g_{ii}}.$$

Then, further inspection reveals that the metrics are related by

$$U_{TV}(v_i) \propto \frac{1}{g_{ii}} (1 - \mu_i^2) U_{\Sigma-opt}(v_i) \quad (30)$$

and correspondingly

$$U_{MSD}(v_i) \propto \frac{1}{g_{ii}} (1 - \mu_i^2) U_{VM}(v_i). \quad (31)$$

In fact, U_{TV} and U_{MSD} may be interpreted as *data-driven* versions of $U_{\Sigma-opt}$ and U_{VM} that are enhanced with the uncertainty term $g_{ii}^{-1}(1 - \mu_i^2)$. On the one hand, $U_{\Sigma-opt}$ and U_{VM} are design-of-experiments-type methods that rely on ensemble criteria and offer *offline* sampling schemes more suitable for applications where the set \mathcal{L} of nodes may *only* be labeled as a *batch*. On the other hand, U_{TV} and U_{MSD} are data-adaptive sampling schemes that adjust to the *specific realization* of labels, and are expected to outperform their batch counterparts in general. This connection is established due to $U_{VM}(v_i)$ and $U_{\Sigma-opt}(v_i)$ being l_2 and l_1 *ensemble* loss metrics on the GMRF (see equations 2.3 and 2.5 in [12]); similarly, MSD (mean square deviation) and TV (total variation) are also l_2 (on the GMRF distribution) and l_1 (on the binary labels pmf) metrics of *change*.

Remark 4: While the proposed methods were developed for binary classification, they can easily be modified to cope with multiple classes using the one-vs-the-rest trick. Specifically, for any set \mathcal{C} of possible classes, it suffices to solve $|\mathcal{C}|$ binary problems, each one focused on detecting the presence or absence of a class. Consequently, the maximum among the GMRF means $\mu_i^{(c)} \forall c \in \mathcal{C}$ reveals which class is the most likely for the i -th node. In addition, the marginal posteriors are readily given by normalizing the binary posteriors in (6), that is

$$p(y_i = c) = \bar{\mu}_i^{(c)} = \frac{\mu_i^{(c)} + 1}{\sum_{c \in \mathcal{C}} (\mu_i^{(c)} + 1)}.$$

Using this approximation, the TV-based scheme can be generalized to

$$U_{TV}(v_i, \mathcal{L}) \propto \sum_{c \in \mathcal{C}} \left[1 - (\bar{\mu}_i^{(c)})^2 \right] \frac{\|\mathbf{g}_i\|_1}{g_{ii}} \quad (32)$$

and similarly for the MSD-based scheme.

A summary of the five different methods that were considered in the context of the proposed EC-based active learning framework is given in Table I.

D. Computational Complexity Analysis

The present section analyzes the computational complexity of implementing the proposed adaptive sampling methods, as well as that of other common adaptive and non-adaptive active learning approaches on graphs. Complexity here refers to float-point multiplications and is given in $\mathcal{O}(\cdot)$ notation as function of the number of nodes N , number of edges $|\mathcal{E}|$ and number of classes $|\mathcal{C}|$. Three types of computational tasks are considered separately: computations that can be performed offline (e.g., initialization), computations required to update model after a new node is observed (only for adaptive methods), and the complexity of selecting a new node to sample (cf. eq. (8)).

TABLE I
SUMMARY OF EC METHODS BASED ON DIFFERENT METRICS OF CHANGE

Method	Change metric	Retraining	Utility function
FL	# of flipped labels	Yes	Eq. (11) and (12)
KLG	KL divergence of GMRF	No	$\propto \frac{1}{g_{ii}} (1 - \mu_i^2)$
KL	KL divergence of (Bernouli) discrete labels	Yes	Eq. (19) – (22)
MSD	Mean-square deviation of GMRF	No	$\propto (1 - \mu_i^2) \frac{\ \mathbf{g}_i\ _2^2}{g_{ii}^2}$
TV	Total variation of discrete labels	No	$\propto (1 - \mu_i^2) \frac{\ \mathbf{g}_i\ _1}{g_{ii}}$

Let us begin with the “plain-vanilla” label propagation scenario where nodes are randomly (or passively) sampled. In that case, the online framework described in Algorithm 1 and Section II.B is not necessary and the nodes can be classified offline after collecting $|\mathcal{L}|$ samples and obtaining (5) for each class in \mathcal{C} . Exploiting the sparsity of the \mathbf{L} , (5) can be approximated via a Power-like iteration (see, e.g., [29]) with $\mathcal{O}(|\mathcal{E}||\mathcal{C}|)$ complexity. Similarly to passive sampling, non-adaptive approaches such as the variance-minimization (VM) in [10] and Σ -opt design in [12] can also be implemented offline. However, unlike passive sampling, the non-adaptive sampling methods require computation of $\mathbf{G}_0 = (\mathbf{L} + \delta\mathbf{I})^{-1}$, which can be approximated with $\mathcal{O}(|\mathcal{E}|N)$ multiplications via the Jacobi method. The offline complexity of VM and Σ -opt is dominated by the complexity required to design the label set \mathcal{L} which is equivalent to $|\mathcal{L}|$ iterations of Algorithm 1 using $U_{VM}(v_i)$ and $U_{\Sigma\text{-}opt}(v_i)$ correspondingly. Thus, the total offline complexity of VM and Σ -opt is $\mathcal{O}(|\mathcal{L}|N^2)$, while $\mathcal{O}(N^2)$ memory is required to store and process \mathbf{G}_0 .

In the context of adaptive methods, computational efficiency largely depends on whether matrix \mathbf{G} is used for sampling and updating. Simple methods such as uncertainty sampling based on minimum margin do not require \mathbf{G} and have soft labels updated after each new sample using iterative label-propagation (see, e.g., [14]) with $\mathcal{O}(|\mathcal{E}||\mathcal{C}|)$ complexity. Uncertainty-sampling-based criteria are also typically very lightweight requiring for instance sorting class-wise the soft labels of each node ($\mathcal{O}(\log |\mathcal{C}|N)$ per sample). While uncertainty-based methods are faster and more scalable, their accuracy is typically significantly lower than that of more sophisticated methods that use \mathbf{G} . Methods that use \mathbf{G} such as the proposed EC algorithms in Section III, the expected-error minimization (EER) in [13], and the two-step approximation (TSA) algorithm in [16] all require $\mathcal{O}(N^2)$ to perform the update in (10). However, TSA and EER use retraining (cf. Remark 3) that incurs complexity $\mathcal{O}(|\mathcal{C}|^2N^2)$ to perform one sample selection; computing the “expected error” requires fictitiously labeling every unlabeled node, and re-computing the metric by treating the fictitious label as the true label. More specifically, consider the normalization of the binary posteriors that is required (similar to the one discussed in Remark 4) in order to define a posterior pmf over multiple classes ($|\mathcal{C}| > 2$). Normalization entails $|\mathcal{C}|$ divisions, and happens $|\mathcal{C}|$ times (once for every possible label of an unlabeled node). This gives rise to a nested loop where the outer loop repeats $|\mathcal{C}|$ times and the inner loop requires $|\mathcal{C}|N$ computations, yielding a total complexity $\mathcal{O}(|\mathcal{C}|^2N)$ for com-

puting the expected error score for one node. Since these scores have to be computed over all unlabeled nodes (in order to select the best one), the overall complexity to obtain a sample according to EER or TSA is $\mathcal{O}(|\mathcal{C}|^2N^2)$. In contrast, the proposed MSD and TV methods (cf. (24), (27)) only require $\mathcal{O}(|\mathcal{C}|N)$ for sampling. Note that the performance gap between EER and TSA on the one hand, and TV and MSD on the other grows as the number of classes $|\mathcal{C}|$ increases.

The complexity analysis is summarized in Table II and indicates that the proposed *retraining-free* adaptive methods exhibit lower overall complexity than EER and TSA. An important modification is proposed in the ensuing section in order to deal with the challenge of *bias* that is inherent to all data-adaptive sampling schemes.

IV. PROMOTING EXPLORATION BY ADJUSTING MODEL CONFIDENCE

It has been observed that active learning schemes may become “myopic” [30], meaning that they become overly focus on *exploiting* (focusing on) a small region of the sample space, and neglect *exploration*. Uncertainty sampling in particular can be prone to such behavior, due to the fact that it is more “myopic,” in the sense that it does not take into account the effect of a potential sample on the generalization capabilities of the classifier. Since the TV- and MSD-based utility score functions in (24) and (26) are influenced by the uncertainty factor $(1 - \mu_i^2)$, it is important to mitigate this effect before testing the performance of the proposed approaches.

Let us begin by observing that most active learning methods, including those based on EC we introduced here, are based on utility score functions that take the general form

$$U(v_i, \mathcal{L}) = \mathbb{E}_{y_i | \mathbf{y}_{\mathcal{L}}} [C(y_i, \mathcal{L})] \quad (33)$$

where $C(y_i, \mathcal{L})$ is any metric that evaluates the effect of node v_i on the model, *given that* its label is y_i . Using the existing probability model to predict how the model itself will change, induces “myopic” behavior especially in the early stages of the sampling process when the inferred model is most likely far from the true distribution.

One possible means of reducing bias is by complementing greedy active learning strategies with random sampling. That is, instead of selecting the index k_t of the node to be sampled at the t -th iteration according to (8), one can opt for a two-branch

TABLE II
COMPUTATIONAL AND MEMORY COMPLEXITY OF VARIOUS METHODS

	Offline	Sampling	Update	Memory
Random	$\mathcal{O}(\mathcal{E} \mathcal{C})$	*	*	$\mathcal{O}(\mathcal{E} + N \mathcal{C})$
VM [10], Σ -opt [12]	$\mathcal{O}(\mathcal{L} N^2)$	*	*	$\mathcal{O}(N^2)$
Uncertainty (min. margin)	*	$\mathcal{O}(\log \mathcal{C} N)$	$\mathcal{O}(\mathcal{E} \mathcal{C})$	$\mathcal{O}(\mathcal{E} + N \mathcal{C})$
EER [9], TSA [16]	$\mathcal{O}(\mathcal{E} N)$	$\mathcal{O}(\mathcal{C} ^2N^2)$	$\mathcal{O}(N^2)$	$\mathcal{O}(N^2)$
FL	$\mathcal{O}(\mathcal{E} N)$	$\mathcal{O}(\mathcal{C} N^2)$	$\mathcal{O}(N^2)$	$\mathcal{O}(N^2)$
TV, MSD	$\mathcal{O}(\mathcal{E} N)$	$\mathcal{O}(\mathcal{C} N)$	$\mathcal{O}(N^2)$	$\mathcal{O}(N^2)$

hybrid rule

$$k_t = \begin{cases} \arg \max_{i \in \mathcal{U}^{t-1}} U(v_i, \mathcal{L}^{t-1}), & \text{w.p. } (1 - \pi^t) \\ \text{Unif}\{1, \dots, |\mathcal{L}^{t-1}|\}, & \text{w.p. } \pi^t \end{cases}. \quad (34)$$

where π^t is the probability that at iteration t the sampling strategy switches to uniform random sampling over the unlabeled nodes. Naturally, one should generally select a sequence $\{\pi^t\}$ such that $\pi^t \rightarrow 0$ as t increases the model becomes more accurate. Upon testing the simple heuristic in (34) we observed that it can significantly improve the performance of the more “myopic” active sampling strategies. Specifically, uncertainty sampling which relies purely on exploitation can be greatly enhanced by completing it with the exploration queries introduced by (34).

Another option is to sample nodes that maximize the minimum over all possible labels change. That is, instead of (33) one can adopt utility scores of the general form

$$U(v_i, \mathcal{L}) = \min_{y_i \in \{-1, 1\}} C(y_i, \mathcal{L}). \quad (35)$$

Albeit intuitive, (34) is not as appropriate for promoting exploration in more sophisticated strategies such as the ones presented in this work, because it does not account for the graph structure, and it is somewhat aggressive in assuming that with probability π^t the model is completely uninformative. For similar reasons, (35) also does not produce satisfactory results.

In the present section, we introduce a “softer” heuristic that is better tailored to the sampling strategies at hand. The main idea is to implement $U(v_i, \mathcal{L})$ in (33) using a different set of probabilities than the ones provided by the model (cf. (6)). Specifically, we suggest to use label predictions that are closer to a “non-informative” prior early on, and gradually converge to the ones provided by the trained model as our confidence on the latter increases. Thus, instead of taking the expectation in (33) over $p(y_i|\mathbf{y}_{\mathcal{L}})$, one may instead use a convex combination of the latter and a node prior $\pi(y_i)$, that is

$$\check{p}(y_i|\mathbf{y}_{\mathcal{L}}; \alpha_t) = \alpha_t \pi(y_i) + (1 - \alpha_t) p(y_i|\mathbf{y}_{\mathcal{L}}) \quad (36)$$

where $0 \leq \alpha_t \leq 1$ is a constant that quantifies the confidence on the current estimate of the posterior. If no prior is available, one may simply use $\pi(y_i = 1) = \pi(y_i = -1) = 1/2$. Intuitively pleasing results were obtained when combining (36) with EC methods. For instance, combining (36) with our proposed TV method yields the following modified MSD utility

score function

$$U_{MSD}(v_i, \mathcal{L}, a_t) \propto [a_t + (1 - a_t)(1 - \mu_i^2)] \frac{\|\mathbf{g}_i\|_2^2}{g_i^2} \quad (37)$$

where a_t tunes the sensitivity of the sampling process to the uncertainty metric $(1 - \mu_i^2)$. As more samples become available, the confidence that the current estimate of the posterior is close to the true distribution may increase. Thus, instead of using a constant α throughout the sampling process, one may use a sequence $\{\alpha_t\}_{t=1}^T$, where t is the iteration index, T the total number of samples, and a_t is *inversely proportional* to t . Finally, note that by setting $\alpha_t = 1 \forall t$ the uncertainty terms vanish with MSD and TV becoming non-adaptive.

V. EXPERIMENTAL RESULTS

The present section includes numerical experiments carried to assess the performance of the proposed methods in terms of prediction accuracy. Specifically, the ensuing subsections present plots of accuracy

$$\text{Accuracy} = \frac{1}{|\mathcal{U}|} \sum_{i \in \mathcal{U}} \mathbf{1}_{\{\hat{y}_i = y_i\}}$$

as a function of the number of nodes sampled by the GMRF-based active learning algorithms (cf. Algorithm 1). We compare the proposed methods (number of flips (FL), KL divergence, MSD, sum of TVs) with the variance minimization (VM) [10], Σ -optimality [12], expected error minimization (EER) [13], and two-step approximation method (TSA) [16]. Furthermore, we compare with the minimum-margin uncertainty sampling (UNC) scheme that samples the node with smallest difference between the largest soft labels, which is equivalent to using the utility function $U_{UNC}(v_i, \mathcal{L}) := -|\mu_i^{(c_1)} - \mu_i^{(c_2)}|$, where c_1 and c_2 is the most-probable and second-most probable class for node v_i correspondingly. Finally, all methods are compared to the predictions that are given by the GMRF method used here (cf. (2)–(7)) with nodes sampled randomly (passive learning). For all graph tested the prediction accuracy remained high for a large range of $\delta \in [0.1, 0.001]$ with the exact value tuned for every graph in order to maximize accuracy for passive (random) sampling.

A. Synthetic Graphs

Following [16], we first considered a 10×10 rectangular grid similar to the one in Fig. 1, where each node is connected to

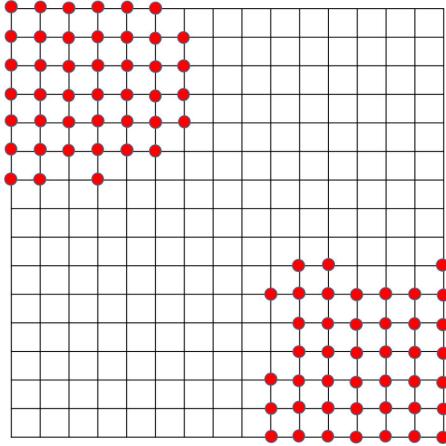


Fig. 1. Rectangular grid synthetic graph with two separate class 1 regions.

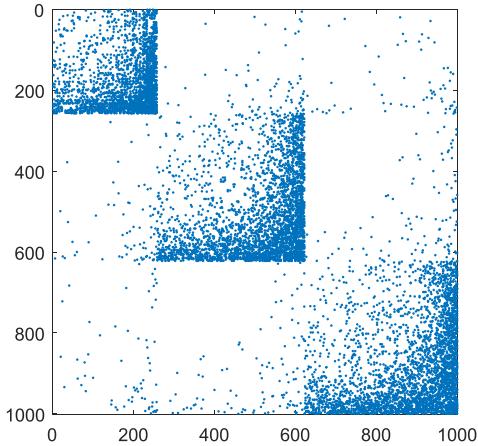


Fig. 2. Adjacency matrix of LFR graph with 1,000 nodes and 3 classes.

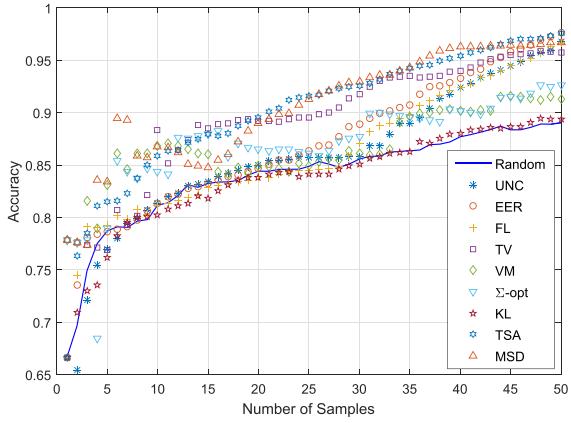


Fig. 3. Test results for synthetic grid in Fig. 1.

four neighboring nodes. Red dots correspond to nodes belonging to class 1, and uncolored intersections correspond to nodes belonging to class -1 . To make the classification task more challenging, the class 1 region was separated into two 3×3 squares (upper left and lower right) and additional class 1 nodes were added w.p. 0.5 along the dividing lines. Plotted in Fig. 3 is the accuracy-vs-number of samples performance averaged over 50

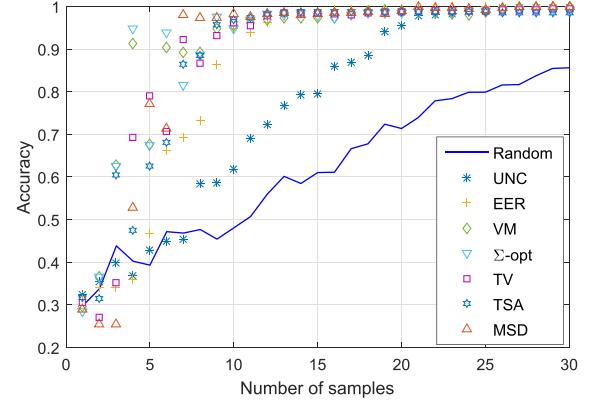


Fig. 4. Test results for synthetic LFR graph in Fig. 2.

Monte Carlo runs. As expected, most algorithms outperform random sampling. In addition, one observes that purely exploratory non-adaptive methods (VM and Σ -optimality) enjoy relatively high accuracy for a small number of samples, but are eventually surpassed by adaptive methods. It can also be observed that the novel TV method with $a_t = t^{-1/2}$ performs equally well to the state-of-the-art TSA method. Interestingly, it does so while using a much simpler criterion that avoids model retraining, and therefore requires significantly shorter runtime. Note finally that the performance of EER is poor because the sampler easily becomes “trapped” in one of the two class 1 regions, and does not explore the graph.

The purpose of the experiment in Fig. 1 was to simulate problems where a complex label distribution appears on a simple uniform graph (e.g., image segmentation). To simulate more structured graphs, we generated a 1000-node network using the Lancichinetti–Fortunato–Radicchi (LFR) method [31]. The LFR algorithm is widely used to generate benchmark graphs that resemble real world networks by exhibiting community structure and degree distributions that follow the power law. Figure 2 reveals the sparsity pattern of the adjacency matrix of the LFR graph that was used, while the 3 clearly visible clusters correspond to groups of nodes in the same class, that is

$$y_i = \begin{cases} 1, & i \in [1, 250] \\ 2, & i \in [251, 600] \\ 3, & i \in [601, 1000] \end{cases}$$

Note that, unlike the one in Fig. 1, the graph used here is characterized by a community structure that matches the nodes labels. This is a highly favorable scenario for the non-adaptive VM and Σ -opt approaches that rely solely on the graph structure. Indeed, as seen in Fig. 4, VM and Σ -opt quickly reach 90% accuracy by selecting 5 most influential samples. Nevertheless, between 5 and 10 samples our proposed MSD and TV adaptive methods enjoy superior accuracy before converging to 100% accuracy.

B. Similarity Graphs From Real Datasets

Real binary classification datasets taken from the UC Irvine Machine Learning Repository [32] and the LibSVM webpage [33] were used for further testing of the proposed methods. First,

TABLE III
DATASET LIST

Dataset	# of nodes	Baseline Accuracy
Coloncancer	62	0.64
Ionosphere	351	0.64
Leukemia	70	0.65
Australian	690	0.55
Parkinsons	191	0.75
Ecoli	326	0.57

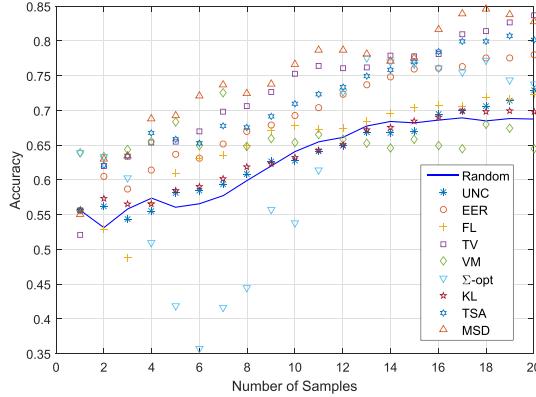


Fig. 5. Coloncancer dataset.

each entry of the feature vectors was normalized to lie between -1 and 1 . Then, a graph was constructed using the Pearson correlations among pairs of normalized feature vectors as weights of the adjacency matrix \mathbf{W} ; thresholding was also applied to negative and small weights leading to sparse adjacency matrices. It was observed that sparsification generally improves the prediction accuracy, while also reducing the computational burden. In the presented experiments, thresholds were tuned until one of the methods achieved the highest possible classification accuracy.

Having constructed the graphs, the proposed expected model change sampling schemes were compared with UNC, TSA, EER, VM and Σ -optimality on seven real datasets listed in Table III; in the latter, “baseline accuracy” refers to the proportion of the largest class in each dataset, and thus the highest accuracy that can be achieved by naively assuming that all labels belong to the majority class. Plotted in Figs. 5 to 11 are the results of the numerical tests, where it is seen that the performance of the proposed low-complexity TV- and MSD-based scheme is comparable or superior to that of competing alternatives. The confidence parameter was set to $a_t = 1/\sqrt{t}$ for the smaller datasets, where only few data were sampled, and the model was expected to be less accurate, whereas for the larger ones it was set to $a_t = 0$.

C. Real Graphs

Experiments were also performed on real labeled graphs. Specifically, the CORA and CITESEER [34] citation networks with 2708 and 3312 nodes correspondingly were used; similarly to [12], we isolated the largest connected components. In citation networks, each node corresponds to a scientific pub-

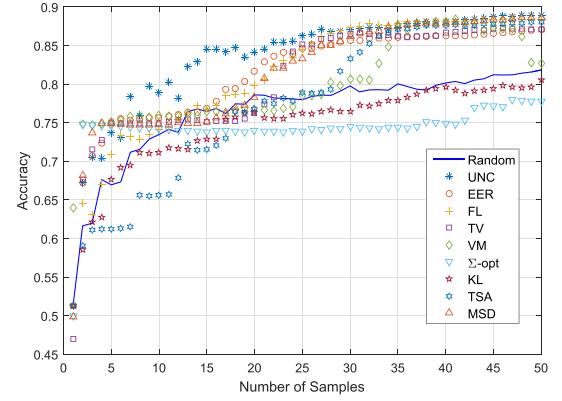


Fig. 6. Ionosphere dataset.

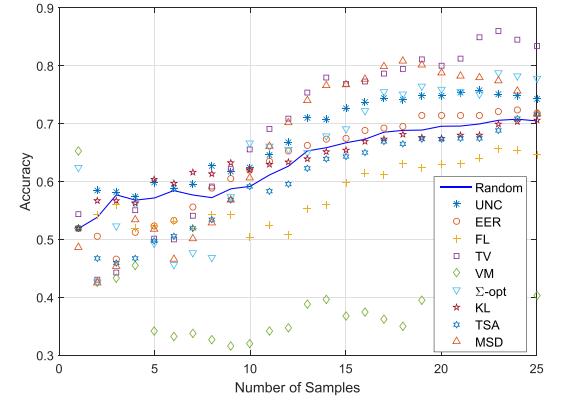


Fig. 7. Leukemia dataset.

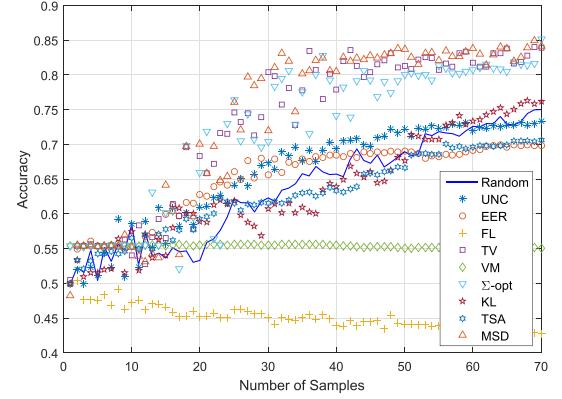


Fig. 8. Australian dataset.

lication and is linked only with cited or citing papers. Nodal labels correspond to the scientific field that each paper belongs to (6 classes for CITESEER and 7 for CORA). Lately, CORA and CITESEER have been used as benchmarks for graph convolutional neural networks (GCNs) [35], as well as for classification based on node embeddings (Planetoid-G) [36]. For this reason, together with the GMRF-based passive (random) sampling benchmark, we also use GCNs and Planetoid-G as passive benchmarks. Note that the latter two methods require validation samples for early stopping during training, while GMRF does not. In order to proceed with comparisons, a set of validation

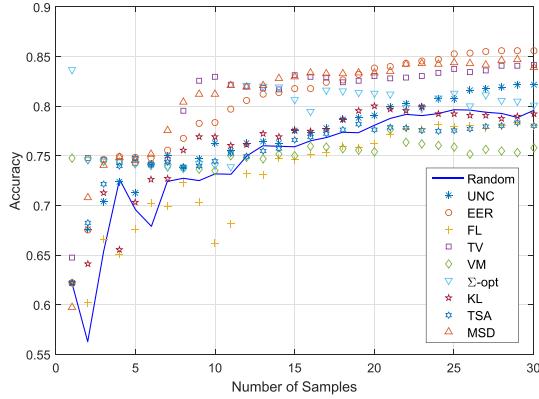


Fig. 9. Parkinsons dataset.

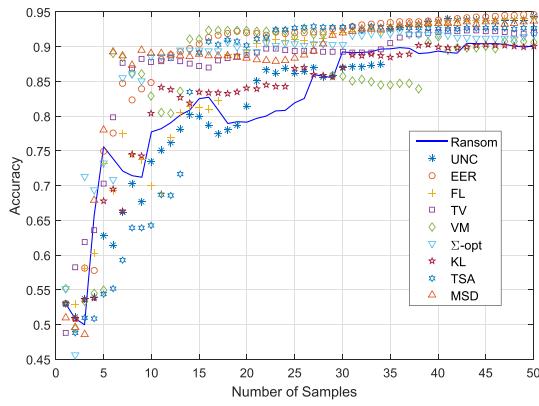


Fig. 10. Ecoli dataset.

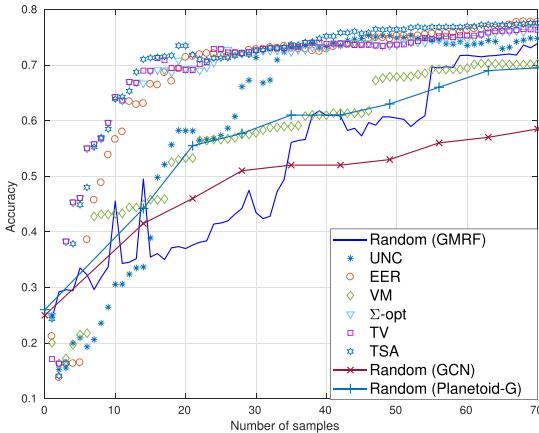


Fig. 11. CORA citation network.

samples was given to GCN and Planetoid-G equal in size to the training set. The benchmark political-blog network [37] with 1490 nodes and two classes was also used. The confidence sequence $\alpha_t = t^{-1/2}$ was used for all graphs, with $\delta = 0.005$ similarly to [11]. The results of the experiments are depicted in Figs. 11–13 and demonstrate the effectiveness of the proposed MSD and TV algorithms on these social graphs. For the CORA network, TV achieves state of the art performance equal to EER, TSA and Σ -opt, while for the CITESEER network its accuracy slightly surpasses that of competing methods. Note that for both

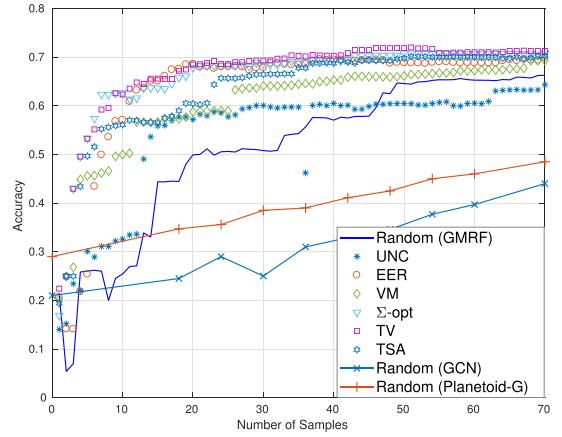


Fig. 12. CITESEER citation network.

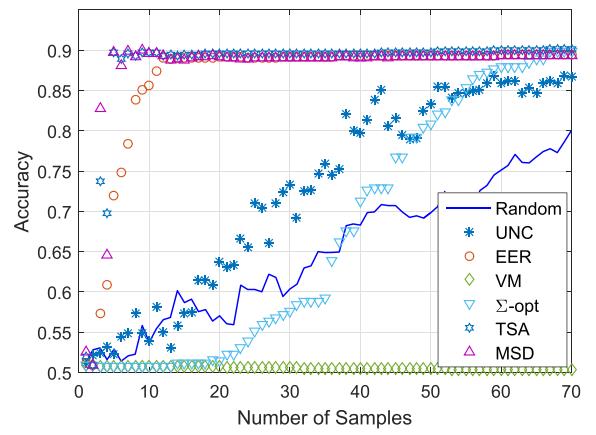


Fig. 13. Political blogs network.

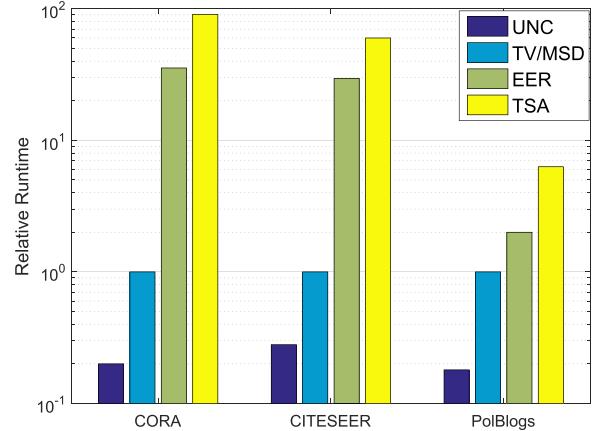


Fig. 14. Relative runtime of different adaptive methods for experiments on real social graphs.

citation networks and given randomly selected samples, the performance of GCNs and Planetoid-G barely reaches that of the simple GMRF classifier. This is mostly attributed to the latter being more suitable for the *graph-only* setting that we are dealing with here, whereas the former mostly exploit node features that are available for citation networks (bag-of-words description

of abstracts). For the political-blogs network, non-adaptive TV and Σ -opt methods perform poorly, while the proposed MSD method performs at least as good as the significantly more complex TSA. The bar plot in Fig. 14 depicts the relative runtimes of different adaptive methods. Observe that MSD and TV are two orders of magnitude faster than EER and TSA for the larger multilabel citation graphs, and one order of magnitude faster for the smallest binary-labeled political blogs network.

VI. CONCLUSIONS AND FUTURE DIRECTIONS

This paper unified existing and developed novel utility functions for data-adaptive graph-cognizant active classification using GMRFs. These utility functions rely on metrics that capture expected changes in GMRF models. Specifically, the proposed samplers query the node that is expected to inflict the largest change on the model. Towards this direction, several measures of expected model change were introduced, sharing similarities and connections with existing methods such as uncertainty sampling, variance minimization, and sampling based on the Σ -optimality criterion. A simple yet effective heuristic was also introduced for increasing the exploration capabilities and reducing bias of the proposed methods, by taking into account the confidence on the model label predictions. Numerical tests using synthetic and real data confirm that the proposed methods achieve accuracy that is comparable or superior to state of the art at smaller runtime.

Future research directions will focus on developing even more efficient adaptive sampling schemes for graphs by finding the sweet spot of how a given graph structure attains the desirable exploration versus exploitation trade-off. Furthermore, our research agenda includes developing adaptive sampling methods tailored for Markov-chain-Monte-Carlo-based and random-walk-based inference on graphs.

APPENDIX

A1. Since $\mathbf{C}^{-1} = \mathbf{L}$ and upon partitioning the two matrices according to labeled and unlabeled nodes, we have

$$\begin{bmatrix} \mathbf{L}_{\mathcal{U}\mathcal{U}} & \mathbf{L}_{\mathcal{U}\mathcal{L}} \\ \mathbf{L}_{\mathcal{L}\mathcal{U}} & \mathbf{L}_{\mathcal{L}\mathcal{L}} \end{bmatrix} \begin{bmatrix} \mathbf{C}_{\mathcal{U}\mathcal{U}} & \mathbf{C}_{\mathcal{U}\mathcal{L}} \\ \mathbf{C}_{\mathcal{L}\mathcal{U}} & \mathbf{C}_{\mathcal{L}\mathcal{L}} \end{bmatrix} = \begin{bmatrix} \mathbf{I}_{|\mathcal{U}|} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{|\mathcal{L}|} \end{bmatrix} \quad (38)$$

which gives rise to four matrix equations. Specifically, the equation that corresponds to the upper right part of (38) is

$$\mathbf{L}_{\mathcal{U}\mathcal{U}} \mathbf{C}_{\mathcal{U}\mathcal{L}} + \mathbf{L}_{\mathcal{U}\mathcal{L}} \mathbf{C}_{\mathcal{L}\mathcal{L}} = \mathbf{0}. \quad (39)$$

Multiplying (39) from the left by $\mathbf{L}_{\mathcal{U}\mathcal{U}}^{-1}$ and from the right by $\mathbf{C}_{\mathcal{L}\mathcal{L}}^{-1}$ yields $\mathbf{C}_{\mathcal{U}\mathcal{L}} \mathbf{C}_{\mathcal{L}\mathcal{L}}^{-1} = -\mathbf{L}_{\mathcal{U}\mathcal{U}}^{-1} \mathbf{L}_{\mathcal{U}\mathcal{L}}$, which verifies (4).

A2. Let $\mathbf{x}_1 \sim \mathcal{N}(\mathbf{m}_1, \mathbf{C})$ and $\mathbf{x}_2 \sim \mathcal{N}(\mathbf{m}_2, \mathbf{C})$, and assume that \mathbf{x}_1 and \mathbf{x}_2 are uncorrelated. Then,

$$\begin{aligned} \text{MSD}(\mathbf{x}_1, \mathbf{x}_2) &:= \mathbb{E} [\|\mathbf{x}_1 - \mathbf{x}_2\|_2^2] \\ &= \mathbb{E} [\|\mathbf{x}_1\|_2^2 + \|\mathbf{x}_2\|_2^2 - 2\mathbf{x}_1^T \mathbf{x}_2] \end{aligned} \quad (40)$$

where

$$\begin{aligned} \mathbb{E} [\|\mathbf{x}_1\|_2^2] &= \mathbb{E} [\|(\mathbf{x}_1 - \mathbf{m}_1) + \mathbf{m}_1\|_2^2] \\ &= \mathbb{E} [\|\mathbf{x}_1 - \mathbf{m}_1\|_2^2] + 2\mathbb{E} [(\mathbf{x}_1 - \mathbf{m}_1)^T \mathbf{m}_1] \\ &\quad + \|\mathbf{m}_1\|_2^2 \\ &= \text{tr}(\mathbf{C}) + \|\mathbf{m}_1\|_2^2 \end{aligned} \quad (41)$$

and similarly for $\mathbb{E} [\|\mathbf{x}_2\|_2^2]$. Finally, note that

$$\begin{aligned} \mathbb{E} [(\mathbf{x}_1 - \mathbf{m}_1)^T (\mathbf{x}_2 - \mathbf{m}_2)] &= \mathbb{E} [\mathbf{x}_1^T \mathbf{x}_2 - \mathbf{x}_1^T \mathbf{m}_2 \\ &\quad - \mathbf{m}_1^T \mathbf{x}_2 + \mathbf{m}_1^T \mathbf{m}_2] \\ &= \mathbb{E} [\mathbf{x}_1^T \mathbf{x}_2] - \mathbf{m}_1^T \mathbf{m}_2 \end{aligned} \quad (42)$$

and since \mathbf{x}_1 and \mathbf{x}_2 are uncorrelated it follows that (42) equals to 0; hence,

$$\mathbb{E} [\mathbf{x}_1^T \mathbf{x}_2] = \mathbf{m}_1^T \mathbf{m}_2. \quad (43)$$

Substituting (41) and (43) into (40) yields

$$\begin{aligned} \text{MSD}(\mathbf{x}_1, \mathbf{x}_2) &= 2\text{tr}(\mathbf{C}) + \|\mathbf{m}_1\|_2^2 + \|\mathbf{m}_2\|_2^2 - 2\mathbf{m}_1^T \mathbf{m}_2 \\ &= 2\text{tr}(\mathbf{C}) + \|\mathbf{m}_1 - \mathbf{m}_2\|_2^2 \end{aligned}$$

which implies that the MSD between two Gaussian fields with the same covariance matrix is proportional to the Euclidean norm of the difference of their means.

REFERENCES

- [1] S. A. Danziger *et al.*, “Predicting positive p53 cancer rescue regions using most informative positive (MIP) active learning,” *PLOS Comput. Biol.*, vol. 5, no. 9, Sep. 2009, Art. no. e1000498.
- [2] J. Haupt, R. M. Castro, and R. Nowak, “Distilled sensing: Adaptive sampling for sparse detection and estimation,” *IEEE Trans. Inf. Theory*, vol. 57, no. 9, pp. 6222–6235, Sep. 2011.
- [3] Y. Fu, X. Zhu, and B. Li, “A survey on instance selection for active learning,” *J. Knowl. Inf. Syst.*, vol. 35, no. 2, pp. 249–283, 2013.
- [4] B. Settles and M. Craven, “An analysis of active learning strategies for sequence labeling tasks,” in *Proc. Conf. Empirical Methods Natural Lang. Process.*, Waikiki, Honolulu, HI, USA, Oct. 2008.
- [5] D. A. Cohn, Z. Ghahramani, and M. I. Jordan, “Active learning with statistical models,” *J. Artif. Intell. Res.*, vol. 4, no. 1, pp. 129–145, 1996.
- [6] K. Chaloner and I. Verdinelli, “Bayesian experimental design: A review,” *Statist. Sci.*, vol. 10, pp. 273–304, 1995.
- [7] G. V. Karanikolas, G. B. Giannakis, K. Slavakis, and R. M. Leahy, “Multi-kernel based nonlinear models for connectivity identification of brain networks,” in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, Shanghai, China, Mar. 2016.
- [8] B. Baingana Y. Shen, and G. B. Giannakis, “Kernel-based structural equation models for topology identification of directed networks,” *IEEE Trans. Signal Process.*, vol. 65, no. 10, pp. 2503–2516, May 2017.
- [9] X. Zhu, Z. Ghahramani, and J. Lafferty, “Semi-supervised learning using Gaussian fields and harmonic functions,” in *Proc. Int. Conf. Mach. Learn.*, Washington, DC, USA, Aug. 2003.
- [10] M. Ji and J. Han, “A variance minimization criterion to active learning on graphs,” in *Proc. Int. Conf. Artif. Intell. Statist.*, La Palma, Canary Islands, Apr. 2012.
- [11] Q. Gu and J. Han, “Towards active learning on graphs: An error bound minimization approach,” in *Proc. IEEE 12th Int. Conf. Data Mining*, Brussels, Belgium, Dec. 2012, pp. 882–887.
- [12] Y. Ma, R. Garnett, and J. Schneider, “ σ -optimality for active learning on Gaussian random fields,” in *Proc. 26th Int. Conf. Adv. Neural Inf. Process. Syst.*, Dec. 2013, pp. 2751–2759.
- [13] X. Zhu, J. Lafferty, and Z. Ghahramani, “Combining active learning and semi-supervised learning using Gaussian fields and harmonic functions,” in *Proc. Int. Conf. Mach. Learn.*, Washington, DC, USA, Aug. 2003, pp. 215–222.

- [14] J. Long, J. Yin, W. Zhao, and E. Zhu, "Graph-based active learning based on label propagation," in *Proc. Int. Conf. Model. Decisions Artif. Intell.*, Catalonia, Spain, Oct. 2008.
- [15] J. Zhou and S. Sun, "Active learning of Gaussian processes with manifold-preserving graph reduction," *J. Neural Comput. Appl.*, vol. 25, no. 7/8, pp. 1615–1625, 2014.
- [16] K.-S. Jun and R. Nowak, "Graph-based active learning: A new look at expected error minimization," in *Proc. Global Conf. Signal Inf. Process.*, Washington DC, Apr. 2017.
- [17] E. E. Gad, A. Gadde, A. S. Avestimehr, and A. Ortega, "Active learning on weighted graphs using adaptive and non-adaptive approaches," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, Shanghai, China, Mar. 2016.
- [18] N. Cesa-Bianchi, C. Gentile, F. Vitale, and G. Zappella, "Active learning on trees and graphs," arXiv:1301.5112, 2013.
- [19] Q. Gu, C. Aggarwal, J. Liu, and J. Han, "Selective sampling on graphs for classification," in *Proc. 19th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Chicago, IL, USA, Aug. 2013, pp. 131–139.
- [20] K. Fujii and H. Kashima, "Budgeted stream-based active learning via adaptive submodular maximization," in *Proc. 30th Int. Conf. Neural Inf. Process. Syst.*, Barcelona, Spain, Dec. 2016, pp. 514–522.
- [21] H. Su, Z. Yin, T. Kanade, and S. Huh, "Active sample selection and correction propagation on a gradually-augmented graph," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Boston, MA, USA, Jun. 2015, pp. 1975–1983.
- [22] J. J. Pfeiffer III, J. Neville, and P. N. Bennett, "Active sampling of networks," in *Proc. Int. Workshop Mining Learn. Graphs*, Edinburgh, Scotland, Jul. 2012.
- [23] A. Hauser and P. Bühlmann, "Two optimal strategies for active learning of causal models from interventions," in *Proc. Eur. Workshop Probability Graph. Models*, Granada, Spain, Sep. 2012, pp. 123–130.
- [24] A. Vezhnevets, J. M. Buhmann, and V. Ferrari, "Active learning for semantic segmentation with expected change," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Portland, OR, USA, 2013, pp. 3162–3169.
- [25] A. Freytag, E. Rodner, and J. Denzler, "Selecting influential examples: Active learning with expected model output changes," in *Proc. Eur. Conf. Comput. Vis.*, Zurich, Switzerland, Sep. 2014, pp. 562–577.
- [26] S. M. Kay, *Fundamentals of Statistical Signal Processing, Vol. I: Estimation Theory*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1993.
- [27] E. D. Kolaczyk, *Statistical Analysis of Network Data*. New York, NY, USA: Springer, 2009.
- [28] A. Krause, A. Singh, and C. Guestrin, "Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies," *J. Mach. Learn. Res.*, vol. 9, pp. 235–284, 2008.
- [29] Z. Ghahramani and X. Zhu, "Learning from labeled and unlabeled data with label propagation," CMU CALD Tech Rep. CMU-CALD-02-107, 2002.
- [30] B. Settles, "Active learning," *Synthesis Lectures Artif. Intell. Mach. Learn.*, vol. 6, no. 1, pp. 1–114, 2012.
- [31] A. Lancichinetti, S. Fortunato, and F. Radicchi, "Benchmark graphs for testing community detection algorithms," *Phys. Rev. E*, vol. 78, no. 4, pp. 046–110, 2008.
- [32] [Online]. Available: <https://archive.ics.uci.edu/ml/index.html>
- [33] [Online]. Available: <https://www.csie.ntu.edu.tw/cjlin/libsvmtools/datasets/binary.html>
- [34] [Online]. Available: <https://linqs.soe.ucsc.edu/data>
- [35] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *Proc. Int. Conf. Learn. Represent.*, Toulon, France, Apr. 2017.
- [36] W. Yang, Z. Cohen, and R. Salakhutdinov, "Revisiting semi-supervised learning with graph embeddings," in *Proc. 33rd Int. Conf. Int. Conf. Mach. Learn.*, Feb. 2016, pp. 40–48.
- [37] [Online]. Available: <http://www-personal.umich.edu/mejn/netdata/>



Dimitris Berberidis (S'15) received the Diploma degree in electrical and computer engineering (ECE) from the University of Patras, Patras, Greece, in 2012 and the M.Sc. degree in ECE from the University of Minnesota, Minneapolis, MN, USA, where he is currently working toward the Ph.D. degree. His research interests include the areas of statistical signal processing, focusing on sketching and tracking of large-scale processes, and in machine learning, focusing on the development of algorithms for scalable learning over (and embedding of) graphs.



Georgios B. Giannakis (F'97) received the Diploma degree in electrical Engineering from the National Technical University of Athens, Zografou, Greece, 1981. From 1982 to 1986, he was with the University of Southern California where he received the M.Sc. degree in electrical engineering in 1983, the M.Sc. degree in mathematics in 1986, and the Ph.D. degree in electrical engineering in 1986. He was with the University of Virginia from 1987 to 1998, and since 1999 he has been a Professor with the University of Minnesota, where he holds an Endowed Chair in

Wireless Telecommunications, a University of Minnesota McKnight Presidential Chair in ECE, and serves as the Director of the Digital Technology Center.

His general interests span the areas of communications, networking and statistical learning - subjects on which he has published more than 430 journal papers, 720 conference papers, 25 book chapters, two edited books, and two research monographs (h-index 133). Current research focuses on learning from big data, wireless cognitive radios, and network science with applications to social, brain, and power networks with renewables. He is the (co-) inventor of 32 patents issued, and the (co-) recipient of 9 best journal paper awards from the IEEE Signal Processing (SP) and Communications Societies, including the G. Marconi Prize Paper Award in Wireless Communications. He also received Technical Achievement Awards from the SP Society (2000), from EURASIP (2005), a Young Faculty Teaching Award, the G. W. Taylor Award for Distinguished Research from the University of Minnesota, and the IEEE Fourier Technical Field Award (inaugural recipient in 2015). He is a Fellow of EURASIP, and has served the IEEE in a number of posts, including that of a Distinguished Lecturer for the IEEE-SP Society.