On the Optimal Recovery of Graph Signals

Simon Foucart, Chunyang Liao, and Nate Veldt — Texas A&M University

2023 International Conference on Sampling Theory and Applications (SampTA) | 979-8-3503-2885-1/23/$31.00 ©2023 IEEE | DOI: 10.1109/SampTA59647.2023.10301205

***Abstract*—Learning a smooth graph signal from partially observed data is a well-studied task in graph-based machine learning. We consider this task from the perspective of optimal recovery, a mathematical framework for learning a function from observational data that adopts a worst-case perspective tied to model assumptions on the function to be learned. Earlier work in the optimal recovery literature has shown that minimizing a regularized objective produces optimal solutions for a general class of problems, but did not fully identify the regularization parameter. Our main contribution provides a way to compute regularization parameters that are optimal or near-optimal (depending on the setting), specifically for graph signal processing problems. Our results offer a new interpretation for classical optimization techniques in graph-based learning and also come with new insights for hyperparameter selection. We illustrate the potential of our methods in numerical experiments on several semi-synthetic graph signal processing datasets.**

I. Introduction

In graph signal processing, one starts with a dataset defined over an irregular graph domain and the goal is to recover a signal on vertices of the graph (e.g. discrete labels or regression values) when given access to only one part of the signal [Dong et al., 2020, Ortega et al., 2018]. As a concrete example, the graph may encode US counties (nodes) and their physical adjacencies (edges) while the signals may represent voting patterns, birthrates, or any number of other attributes influenced by geographic region [Jia and Benson, 2020]. As other examples, in biology, the graph may represent a gene interaction network while signals may indicate expression levels of individual genes [Dong et al., 2020]; in neuroscience, brain activity signals coming from fMRI data may be analyzed over a graph representing physical connections or co-activations among regions of a brain [Huang et al., 2018], etc.

The task of recovering a graph signal from partial information about it is also known as graph-based semi-supervised learning [Zhou et al., 2003, Belkin et al., 2004, Zhu et al., 2003]. This task has been studied in depth by researchers from many related academic communities including machine learning, statistics, and of course signal processing. In all of these settings, a common assumption is that signals vary smoothly over the graph's edge structure, meaning that adjacent nodes often share similar labels [Zhou et al., 2003, Zhu et al., 2003, Belkin et al., 2004, Xu et al., 2010, Dong et al., 2020]. Many formal objective functions and theoretical results for graph signal processing and semi-supervised learning are justified by assuming that graph signals come from a certain well-behaved probability distributions [Zhu et al., 2003, Dong et al., 2019]. This often leads to objective functions that can be minimized using simple matrix-based methods [Zhu et al., 2003, Zhou et al., 2003, Belkin et al., 2004]. However, the performance is affected by the choice of a regularization parameter in the objective function and it is not always clear how to select such a parameter. In another direction, there has been a recent surge of interest in using graph-neural networks for learning over graphs. This is often successful in practice but typically comes with no mathematical guarantees.

In our work, we address the graph signal processing task from a novel perspective—that of optimal recovery. This perspective does not rely on the assumption that ground truth signals are drawn from a well-behaved distribution. Instead, the goal is to find optimal solutions under worst-case assumptions about graph smoothness and labeling error. This approach comes with several benefits. Primarily, we present new theoretical results on finding best solutions under the optimal recovery framework (locally and globally, see later sections for technical details). Along the way, we highlight the connections between the optimization problems stemming from this framework and the classical techniques encountered in graph signal processing. One significant contribution is to provide rigorous theoretical guarantees for selecting the regularization parameter in the objective function being minimized. Setting this parameter is not entirely free of challenges, as it actually depends on the parameters characterizing graph smoothness and labeling error. Nevertheless, our results offer fresh intuition on how to reasonably choose the regularization parameters intrinsic to objective functions common in graph signal processing. Finally, we provide a proof-of-concept implementation of our approach and illustrate its performance in several empirical graph signal processing experiments.

II. The Perspective from Optimal Recovery

Let $G = (V, E)$ be an undirected graph with $N = |V|$ vertices identified with $1, 2, \dots, N$. A signal f defined on V is thus identified with a vector $f \in \mathbb{R}^N$. The previously-mentioned common assumption that f varies smoothly over the graph's edge structure qualitatively translates into the fact that the values f_i and f_j do not differ much if the vertices i and j are strongly connected. Quantitatively, putting a weight $w_{i,j} \geq 0$ on the edge connecting i and j, thus defining a (weighted, symmetric) adjacency matrix $W \in \mathbb{R}^{N \times N}$, the assumption takes the form

$$\frac{1}{2} \sum_{i,j=1}^{N} w_{i,j} (f_i - f_j)^2 \leq \varepsilon^2$$

for a small $\varepsilon > 0$ standing for a graph smoothness parameter. Introducing the graph Laplacian $L = D - W \in \mathbb{R}^{N \times N}$, where D is the diagonal matrix with entries $D_{i,i} = \sum_{j=1}^{N} W_{i,j}$, the assumption succinctly reads

$$\langle Lf, f \rangle \leq \varepsilon^2, \qquad \text{or} \quad \|L^{1/2} f\|_2 \leq \varepsilon.$$

We recall that the square-root $L^{1/2}$ of L is well-defined because the graph Laplacian L is positive semidefinite. Note that it is not positive definite, since 0 is always an eigenvalue of L. In fact, its multiplicity equals the number of connected components C of G, with orthogonal eigenvectors provided by the indicator vectors $\mathbf{1}_C \in \{0,1\}^N$ of C. Throughout, we shall assume that the graph G is known, and hence that L is available to the user.

As for the unknown f, it is partially observed—or labeled[1]. In other words, there is a subset V_ℓ, with size $|V_\ell| = n_\ell$, of vertices for which the f_i, $i \in V_\ell$, are known. In reality, they are known up to additive errors, so that the user has access to

$$y_i = f_i + e_i, \qquad i \in V_\ell.$$

To abbreviate, we write $y = \Lambda f + e \in \mathbb{R}^{n_\ell}$, where the linear map $\Lambda : \mathbb{R}^N \to \mathbb{R}^{n_\ell}$ satisfies $\Lambda\Lambda^* = I_{n_\ell}$ here (since, up to a proper ordering of the vertices, (the matrix of) Λ takes the form $\begin{bmatrix} I_{n_\ell} \mid 0 \end{bmatrix}$). We shall assume that an ℓ_2-bound on the error vector e is available, namely

$$\|e\|_2 \le \eta$$

for a small $\eta > 0$ standing for a labeling error parameter.

Our objective is now to estimate the graph signal f on the set of unlabeled vertices, i.e., on $V_u = V \setminus V_\ell$, which has size $|V_u| = n_u = N - n_\ell$. In the framework of optimal recovery, we aim at doing so in a worst-case optimal way given the graph smoothness and labeling error assumptions, expressed as $f \in \mathcal{K}$ and $e \in \mathcal{E}$, where the model set $\mathcal{K}$ and the uncertainty set $\mathcal{E}$ are given as

$$\mathcal{K} = \{f \in \mathbb{R}^N : \|L^{1/2} f\|_2 \le \varepsilon\}, \tag{1}$$

$$\mathcal{E} = \{e \in \mathbb{R}^{n_\ell} : \|e\|_2 \le \eta\}. \tag{2}$$

A scheme to estimate $f_{|V_u} \in \mathbb{R}^{n_u}$ from $y \in \mathbb{R}^{n_\ell}$ is nothing but a map $\Delta : \mathbb{R}^{n_\ell} \to \mathbb{R}^{n_u}$, which we call a recovery map. We are interested in those recovery maps that are optimal

- *in the global setting,* i.e.,
$$\sup\{\|f_{|V_u} - \Delta(\Lambda f + e)\|_2 : f \in \mathcal{K}, e \in \mathcal{E}\}$$
is as small as possible;
- *in the local setting*, i.e., at any given $y \in \mathbb{R}^{n_\ell}$,
$$\sup\{\|f_{|V_u} - z\|_2 : f \in \mathcal{K}, e \in \mathcal{E}, \Lambda f + e = y\}$$
evaluated at $z = \Delta(y)$ is as small as possible. Such a $\Delta(y) \in \mathbb{R}^{n_u}$ is called a Chebyshev center for the set $\mathcal{S} = \{f_{|V_u} : f \in \mathcal{K}, e \in \mathcal{E}, \Lambda f + e = y\}$, as it is easily seen to be a center of a minimal-radius ball containing $\mathcal{S}$.

If we believe that the observed labels need to be adjusted, too, instead of estimating $f_{|V_u}$, we may want to estimate f in full. We may also want to estimate the average of f or its value at a particular vertex $i_0 \in V$. To deal with these situations all at once, we introduce a quantity of interest $Q : \mathbb{R}^N \to \mathbb{R}^n$, which in the examples above is given by, respectively,

$$Q(f) = f_{|V_u}, \;\; Q(f) = f, \;\; Q(f) = \frac{1}{N}\sum_{i \in V} f_i, \;\; Q(f) = f_{i_0}.$$

[1]Labels are real numbers here, not elements of a binary set such as $\{0,1\}$.

In this generality, the global and local worst-case errors for the estimation of Q are defined, for $\Delta : \mathbb{R}^{n_\ell} \to \mathbb{R}^n$ and for $y \in \mathbb{R}^{n_\ell}$, $z \in \mathbb{R}^n$, by

$$\mathrm{gwce}_Q(\Delta) = \sup_{f \in \mathcal{K}, e \in \mathcal{E}} \{\|Q(f) - \Delta(\Lambda f + e)\|_2\},$$

$$\mathrm{lwce}_Q(y,z) = \sup_{f \in \mathcal{K}, e \in \mathcal{E}} \{\|Q(f) - z\|_2 : \Lambda f + e = y\}.$$

We call a recovery map $\Delta : \mathbb{R}^{n_\ell} \to \mathbb{R}^n$ globally optimal if it minimizes $\mathrm{gwce}_Q(\Delta)$ and locally optimal if $\Delta(y)$ minimizes $\mathrm{lwce}_Q(y,z)$ for any given $y \in \mathbb{R}^{n_\ell}$. Of course, locally optimal recovery maps are automatically globally optimal, but they are typically harder to produce (as the current work will also illustrate). We may therefore relax the aspiration of genuine optimality to one of near optimality by merely requiring that $\mathrm{lwce}_Q(y, \Delta(y)) \le C \inf\{\mathrm{lwce}_Q(y,z), z \in \mathbb{R}^n\}$ for some absolute constant $C > 1$.

III. Selection of the Regularization Parameter

We now show that (near-)optimal recovery maps can be obtained through Tikhonov-style regularization and we uncover a principled way to choose the regularization parameter based on the graph smoothness and labeling error parameters. We start with some preparatory information about regularization before presenting our genuine optimality result in the global setting and our near optimality result in the local setting.

A. Rundown on regularization

When searching for the signal $f \in \mathbb{R}^N$ that produced the observation vector $y \in \mathbb{R}^{n_\ell}$, it is natural to try and make the data-fidelity term $\|\Lambda f - y\|_2^2$ small. Furthermore, to enforce the graph smoothness condition that $\|L^{1/2} f\|_2^2$ is small, one can incorporate this condition as a constraint in a miminization problem or add the regularization term $\gamma \|L^{1/2} f\|_2^2$ to the objective function, as done e.g. in [Belkin et al., 2004]. Instead of parametrizing by $\gamma > 0$, it will be more convenient for our purpose to parametrize by some $\tau \in (0,1)$, thus leading to the regularization map $\Delta_\tau : \mathbb{R}^{n_\ell} \to \mathbb{R}^N$ given by

$$\Delta_\tau : y \mapsto \operatorname*{argmin}_{f \in \mathbb{R}^N} \; (1-\tau)\|L^{1/2} f\|_2^2 + \tau \|\Lambda f - y\|_2^2.$$

This map is well defined under the assumption that at least one vertex is observed in each connected component of the graph, which translates into $\ker(L) \cap \ker(\Lambda) = \{0\}$ or equivalently into the invertibility of $(1-\tau)L + \tau\Lambda^*\Lambda$ (see §1 in the supplementary material). Indeed, as the minimizers f_τ of the above objective function are characterized by the normal equation $(1-\tau)Lf_\tau + \tau\Lambda^*(\Lambda f_\tau - y) = 0$, this invertibity shows that f_τ is unique and is equal to

$$\Delta_\tau(y) = \big((1-\tau)L + \tau\Lambda^*\Lambda\big)^{-1}(\tau\Lambda^* y). \tag{3}$$

This expression reveals in particular that Δ_τ is a linear map.

The extreme case $\tau \to 0$ is interpreted as the minimizer of $\|\Lambda f - y\|_2^2$ subject to $L^{1/2} f = 0$, which is not very interesting, see §2 for explanation. The extreme case $\tau \to 1$ is interpreted as the minimizer of $\|L^{1/2} f\|_2^2$ subject to $\Lambda f = y$, which appears commonly in graph signal processing under the names

of harmonic method [Zhu et al., 2003] or interpolatory method [Belkin et al., 2004].

B. Genuine optimality in the global setting

It was recognized already in [Melkman and Micchelli, 1979, Micchelli, 1993] that the regularization maps Δ_τ produce a globally optimal recovery map for *some* parameter $\tau \in (0,1)$, but the choice of this parameter was not made explicit. Theorem 1 below shows that this parameter can be obtained by solving a semidefinite program. Such a result was established in [Foucart and Liao, 2023] in a slightly more restrictive setting, namely the place of $L^{1/2}$ was taken by an orthogonal projector and only full recovery (i.e., $Q = I_N$) was considered.

Theorem 1. Given a linear quantity of interest $Q : \mathbb{R}^N \to \mathbb{R}^n$, let $\mathcal{K}$ and $\mathcal{E}$ be the model and uncertainty sets from (1)-(2). Defining $\tau_\flat = d_\flat/(c_\flat + d_\flat)$ where $c_\flat, d_\flat \geq 0$ are solutions to

$$\underset{c,d\geq 0}{\text{minimize}} \; c\varepsilon^2 + d\eta^2 \quad \text{s.to} \;\; cL + d\Lambda^*\Lambda \succeq Q^*Q, \qquad (4)$$

the linear map $Q \circ \Delta_{\tau_\flat} : \mathbb{R}^{n_\ell} \to \mathbb{R}^n$ is a globally optimal recovery map relative to $\mathcal{K}$ and $\mathcal{E}$, meaning that

$$\mathrm{gwce}_Q(Q \circ \Delta_{\tau_\flat}) = \inf_{\Delta : \mathbb{R}^{n_\ell} \to \mathbb{R}^n} \mathrm{gwce}_Q(\Delta).$$

The justification of this theorem relies on the two lemmas below, whose proofs appear in §3 and §4 of the supplementary material. Lemma 2 relies on a version of the S-procedure due to Polyak [1998] and follows [Foucart and Liao, 2023] closely, but the argument for Lemma 3 follows a different route to deal with a quantity of interest $Q \neq I_N$.

Lemma 2. For an arbitrary recovery map $\Delta : \mathbb{R}^{n_\ell} \to \mathbb{R}^n$, the squared global worst-case error is lower-bounded as

$$\begin{aligned}\mathrm{gwce}_Q(\Delta)^2 &\geq \sup_{h \in \mathbb{R}^N} \{\|Q(h)\|_2^2 : \|L^{1/2}h\|_2^2 \leq \varepsilon^2, \|\Lambda h\|_2^2 \leq \eta^2\} \\ &= \inf_{c,d\geq 0} \{c\varepsilon^2 + d\eta^2 : cL + d\Lambda^*\Lambda \succeq Q^*Q\}.\end{aligned}$$

Lemma 3. If $c, d \geq 0$ satisfy $cL + d\Lambda^*\Lambda \succeq Q^*Q$, then, setting $\tau = d/(c+d)$, one has, for all $f \in \mathbb{R}^N$ and all $e \in \mathbb{R}^{n_\ell}$,

$$\|Q(I - \Delta_\tau\Lambda)f - Q\Delta_\tau e\|_2^2 \leq c\|L^{1/2}f\|_2^2 + d\|e\|_2^2.$$

Proof of Theorem 1. Let $c_\flat, d_\flat \geq 0$ be minimizers of (4). On the one hand, according to Lemma 2, the squared global worst-case error of any recovery map $\Delta : \mathbb{R}^{n_\ell} \to \mathbb{R}^n$ satisfies

$$\mathrm{gwce}_Q(\Delta)^2 \geq c_\flat\varepsilon^2 + d_\flat\eta^2.$$

On the other hand, the linearity of $Q \circ \Delta_{\tau_\flat}$, $\tau_\flat := d_\flat/(c_\flat + d_\flat)$, implies that its squared global worst-case error becomes

$$\mathrm{gwce}_Q(Q \circ \Delta_{\tau_\flat})^2 = \sup_{f \in \mathcal{K}, e \in \mathcal{E}} \|Q(I - \Delta_\tau\Lambda)f - Q\Delta_\tau e\|_2^2,$$

which, according to Lemma 3, does not exceed

$$\sup_{f \in \mathcal{K}, e \in \mathcal{E}} c\|L^{1/2}f\|_2^2 + d\|e\|_2^2 = c_\flat\varepsilon^2 + d_\flat\eta^2.$$

All in all, we have shown that $\mathrm{gwce}_Q(Q \circ \Delta_{\tau_\flat}) \leq \mathrm{gwce}_Q(\Delta)$ for any map $\Delta : \mathbb{R}^{n_\ell} \to \mathbb{R}^n$, which is the desired result. □

Remark. To achieve genuine optimality, exact knowledge of the parameters ε and η is needed, but near optimality is achievable when these are overestimated by some $\bar{\varepsilon}$ and $\bar{\eta}$ satisfying $\bar{\varepsilon} \leq C\varepsilon$ and $\bar{\eta} \leq C\eta$, see §5.

Remark. The semidefinite program (4), featuring an $N \times N$ matrix, does not run when N is in the thousands. Nonetheless, it is expected that the computational burden could be alleviated if Q maps into a low-dimensional space $\mathbb{R}^n$. This is certainly the case in the extreme case $n = 1$, see §6.

C. Near optimality in the local setting

In contrast to the global setting, we are unaware of a general result stating that the regularization maps Δ_τ produce a locally optimal recovery map for *some* parameter $\tau \in (0,1)$. For full recovery ($Q = I_N$), such a statement is true in at least two situations, though. The first situation requires $\Lambda\Lambda^* = I_{n_\ell}$ (which is the case here) and an orthogonal projector P in place of $L^{1/2}$ (up to normalization, $L^{1/2}$ happens to be an orthogonal projector if G is an unweighted graph made of disconnected complete subgraphs of identical sizes, see §7): it was established in [Foucart and Liao, 2023] that $\Delta_{\tau_\sharp}$ is a locally optimal recovery map when $\tau_\sharp$ is the unique parameter τ between $1/2$ and $\varepsilon/(\varepsilon+\eta)$ satisfying the eigenvalue equation

$$\begin{aligned}&\lambda_{\min}((1-\tau)P + \tau\Lambda^*\Lambda) \\ &\qquad = \frac{(1-\tau)^2\varepsilon^2 - \tau^2\eta^2}{(1-\tau)\varepsilon^2 - \tau\eta^2 + (1-\tau)\tau(1-2\tau)\delta^2},\end{aligned}$$

where $\delta = \min\{\|Pf\|_2 : \Lambda f = y\} = \min\{\|\Lambda f - y\|_2 : Pf = 0\}$. The second situation requires working in the complex setting: it was established in [Beck and Eldar, 2007] that $\Delta_{\tau_\flat}$ is a locally optimal recovery map when $\tau_\flat = d_\flat/(c_\flat + d_\flat)$, with $c_\flat, d_\flat$ solving the semidefinite program

$$\begin{aligned}\underset{c,d,t\geq 0}{\text{minimize}} \; & c\varepsilon^2 + d(\eta^2 - \|y\|_2^2) + t \quad \text{s.to} \; cL + d\Lambda^*\Lambda \succeq I_N \\ & \text{and} \; \left[\begin{array}{c|c} cL + d\Lambda^*\Lambda & d\Lambda^*y \\ \hline dy^*\Lambda & t \end{array}\right] \succeq 0.\end{aligned}$$

In the real setting, although the value of $\mathrm{lwce}_{I_N}(y, \Delta_{\tau_\flat}(y))$ is only guaranteed to provide an upper bound for the minimal local worst-case error, it is not unlikely that $\Delta_{\tau_\flat}$ is genuinely a locally optimal recovery map—this is the case in the first situation (result not published yet).

Regardless of the above considerations, relaxing genuine optimality to near optimality, it can always be guaranteed that some regularization map Δ_τ produces a locally near optimal recovery map for a parameter $\tau \in (0,1)$ that can be computed. This is the gist of the following result, see §8 for a proof.

Theorem 4. Given a linear quantity of interest $Q : \mathbb{R}^N \to \mathbb{R}^n$, let $\mathcal{K}$ and $\mathcal{E}$ be the model and uncertainty sets from (1)-(2). For $y \in \mathbb{R}^{n_\ell}$, let $\widehat{f} \in \mathbb{R}^N$ be the solution to

$$\underset{f \in \mathbb{R}^N}{\text{minimize}} \; \max\left\{\|L^{1/2}f\|_2^2, \frac{\varepsilon^2}{\eta^2}\|\Lambda f - y\|_2^2\right\}.$$

Then $\widehat{f}$ agrees with $\Delta_{\tau_\natural}(y)$, where $\tau_\natural$ is the unique parameter $\tau \in (0,1)$ satisfying

$$\|L^{1/2}\Delta_\tau(y)\|_2 = \frac{\varepsilon}{\eta}\|\Lambda\Delta_\tau(y) - y\|_2.$$

Moreover, one has $\mathrm{lwce}_Q(y, Q(\widehat{f})) \leq 2\inf_{z \in \mathbb{R}^n} \mathrm{lwce}_Q(y, z)$.

IV. Numerical Validation

In this section, we illustrate the performance of optimal recovery methods on several semi-synthetic regression datasets and verify the near optimality of a regularization parameter in the local setting. We implement our algorithms in MATLAB and use CVX [Grant and Boyd, 2014] for solving semidefinite programs. All numerical experiments are available on the GitHub repository https://github.com/liaochunyang/ORofGraphSignals.

Let us recall that all our optimal recovery maps correspond to solving a regularized optimization program with a specific choice of hyperparameter. This type of regularized objective is already standard in graph-based machine learning [Belkin et al., 2004], but it is often unclear how this parameter should be chosen. The primary goal of our numerical experiments is to illustrate how our techniques for hyperparameter selection work in a controlled environment where we have access to a ground truth graph signal f and we can estimate the true smoothness and error parameters ϵ and η. In practical settings, we do not actually have access to f, nor can we estimate ε and η exactly, but we shall confirm that near optimality is achievable under mild overestimations of ϵ and η (see §5).

We consider several real-world graphs and, using a standard approach [Dong et al., 2019] (see §9 for more details), we generate synthetic signals f whose values at the nodes are normalizing to be between 0 and 1 for simplicity. In our first experiment, we show how the prediction error changes as the number of labeled vertices grows. We begin by running all methods for $n_\ell = 5$ and we keep adding 5 new labeled nodes at a time—by the end, we have run experiments for n_ℓ ranging from 5 to $N/2$ in increments of 5. For each choice of n_ℓ, our goal is to recover labels at unlabeled nodes, i.e., $Q(f) = f|_{V_u}$, where f is the vector of true node labels. The prediction error for any estimator $\hat{f}$ of $f|_{V_u}$ is defined as $\|f|_{V_u} - \hat{f}\|_2$. For the smoothness parameter, we set $\epsilon = 2\|L^{1/2}f\|_2^2$. Next, we introduce noise artificially by generating a uniform random vector e and subtract the mean before scaling so that $\|e\|_2 \leq \eta$, where η is chosen as $\eta = 2$ (see §9 for details on other types of noise). We then create the corrupted labels $y = f_{|V_\ell} + e$. In the supplementary material, we also consider a mild overestimation on η by setting $\bar{\eta} = 2\eta$.

In Figure 1, we display the prediction errors produced by locally/globally optimal recovery maps for the graph Adjnoun (see §9 for results on other graphs). The constructions of the globally optimal recovery map and the locally near optimal recovery map were presented in Theorem 1 and Theorem 4, respectively. We also use a grid-search approach to find the smallest prediction error over all regularization parameters in order to display a curve of the lowest possible prediction error. This is neither computationally efficient nor realistic, as it assumes that we can always check the prediction error for any estimator $\hat{f}$, but it provides a bound on the best case scenario for solving the regularized objective. Comparing the magenta and the red lines and the black and the green lines, we observe that an overestimation of η does not lead to large differences in prediction error for both local and global optimal recovery maps, which suggests that one can safely use a mild overestimation of η when we cannot access the true η. We also notice that the prediction errors produced by locally/globally optimal recovery maps are close to the prediction error produced by the best Tikhonov regularization method (blue line).

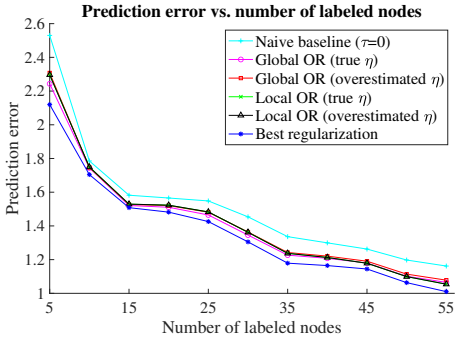

Fig. 1. Prediction error vs. number of labeled nodes on graph Adjnoun with added uniform noise.

In the second experiment (see Figure 2), we confirm the near optimality of the locally optimal recovery map described in Theorem 4. The setup of this experiment is similar to the first experiment. The parameter $\tau_\sharp$—the unique $\tau \in (0,1)$ such that $\|L^{1/2}\Delta_\tau(y)\|_2 = (\varepsilon/\eta)\|\Lambda\Delta_\tau(y) - y\|_2$— is found by the bisection method and is displayed in Figure 2 by the dashed vertical line. The blue curve represents an upper bound for the local worst-case error $\mathrm{lwce}_Q(y, Q \circ \Delta_\tau(y))$ as a function of the regularization parameter. For each τ in a grid of $[0,1]$, this upper bound was computed by solving a semidefinite relaxation for the local worst-case error, see §10 for details. Figure 2 not only supports the local near optimality of the recovery map $Q \circ \Delta_{\tau_\sharp}$, but it also hints that $\tau_\sharp$ is not far away the best regularization parameter.

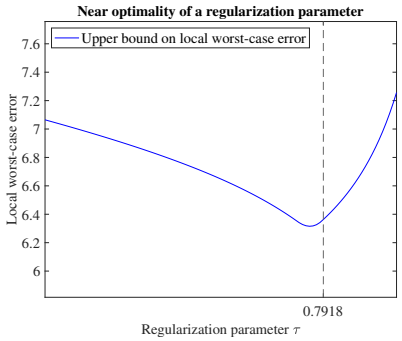

Fig. 2. Local worst-case error vs. regularization parameter.

Acknowledgment

S. F. is partially supported by grants from the NSF (DMS-2053172) and from the ONR (N00014-20-1-2787).

REFERENCES

A. Beck and Y. C. Eldar. Regularization in regression with bounded noise: a Chebyshev center approach. *SIAM Journal on Matrix Analysis and Applications*, 29(2):606–625, 2007.

M. Belkin, I. Matveeva, and P. Niyogi. Regularization and semi-supervised learning on large graphs. In *Learning Theory: 17th Annual Conference on Learning Theory*, pages 624–638. Springer, 2004.

T. A. Davis and Y. Hu. The University of Florida Sparse Matrix Collection. *ACM Trans. Math. Softw.*, 38(1), 2011. URL https://doi.org/10.1145/2049662.2049663.

X. Dong, D. Thanou, M. Rabbat, and P. Frossard. Learning graphs from data: A signal representation perspective. *IEEE Signal Processing Magazine*, 36(3):44–63, 2019. doi: 10.1109/MSP.2018.2887284.

X. Dong, D. Thanou, L. Toni, Mi. Bronstein, and P. Frossard. Graph signal processing for machine learning: A review and new perspectives. *IEEE Signal Processing Magazine*, 37(6): 117–127, 2020.

S. Foucart and C. Liao. Optimal recovery from inaccurate data in Hilbert spaces: Regularize, but what of the parameter? *Constructive Approximation*, 57:489–520, 2023.

A. L. Garkavi. On the optimal net and best cross-section of a set in a normed space. *Izvestiya Rossiiskoi Akademii Nauk. Seriya Matematicheskaya*, 26(1):87–106, 1962.

M Girvan and M. Newman. Community structure in social and biological networks. *Proceedings of the National Academy of Sciences*, 99(12):7821–7826, 2002.

M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. http://cvxr.com/cvx, March 2014.

W. Huang, T. Bolton, J. Medaglia, D. Bassett, A. Ribeiro, and D. Van De Ville. A graph signal processing perspective on functional brain imaging. *Proceedings of the IEEE*, 106(5): 868–885, 2018.

J. Jia and A. Benson. Residual correlation in graph neural network regression. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 588–598, 2020.

V. Krebs. Books about US Politics. URL http://networkdata.ics.uci.edu/data.php?d=polbooks.

D. Lusseau, K. Schneider, O. Boisseau, P. Haase, E. Slooten, and S. Dawson. The bottlenose dolphin community of doubtful sound features a large proportion of long-lasting associations: can geographic isolation explain this unique trait? *Behavioral Ecology and Sociobiology*, 54:396–405, 2003.

A. A. Melkman and C. A. Micchelli. Optimal estimation of linear operators in Hilbert spaces from inaccurate data. *SIAM Journal on Numerical Analysis*, 16(1):87–105, 1979.

C. A. Micchelli. Optimal estimation of linear operators from inaccurate data: a second look. *Numerical Algorithms*, 5(8): 375–390, 1993.

M. Newman. Finding community structure in networks using the eigenvectors of matrices. *Physical Review E*, 74(3): 036104, 2006.

A. Ortega, P. Frossard, J. Kovačević, J. Moura, and P. Vandergheynst. Graph signal processing: Overview, challenges, and applications. *Proc. IEEE*, 106(5):808–828, 2018.

B. T. Polyak. Convexity of quadratic transformations and its use in control and optimization. *Journal of Optimization Theory and Applications*, 99(3):553–583, 1998.

Y. Xu, J. Dyer, and A. Owen. Empirical stationary correlations for semi-supervised learning on graphs. *The Annals of Applied Statistics*, 4(2):589–614, 2010.

D. Zhou, O. Bousquet, T. Lal, J. Weston, and B. Schölkopf. Learning with local and global consistency. *Advances in Neural Information Processing Systems*, 16, 2003.

X. Zhu, Z. Ghahramani, and J. Lafferty. Semi-supervised learning using Gaussian fields and harmonic functions. In *Proceedings of the 20th International Conference on Machine Learning*, pages 912–919, 2003.