

Learning Networks from Random Walk-Based Node Similarities

Jeremy G. Hoskins* Cameron Musco† Christopher Musco‡
Charalampos E. Tsourakakis§

January 22, 2018

Abstract

Digital presence in the world of online social media entails significant privacy risks [7, 11, 34, 58, 35]. In this work we consider a privacy threat to a social network in which an attacker has access to a subset of *random walk-based node similarities*, such as effective resistances (i.e., commute times) or personalized PageRank scores. Using these similarities, the attacker’s goal is to infer as much information as possible about the underlying network, including any remaining unknown pairwise node similarities and edges.

For the effective resistance metric, we show that with just a small subset of measurements, the attacker can learn a large fraction of edges in a social network (and in some cases all edges), even when the measurements are noisy. We also show that it is possible to learn a graph which accurately matches the underlying network on *all other effective resistances*. This second observation is interesting from a data mining perspective, since it can be expensive to accurately compute all effective resistances or other random walk-based similarities. As an alternative, our graphs learned from just a subset of approximate effective resistances can be used as surrogates in a wide range of applications that use effective resistances to probe graph structure, including for graph clustering, node centrality evaluation, and anomaly detection.

We obtain our results by formalizing the graph learning objective mathematically, using two optimization problems. One formulation is convex and can be solved provably in polynomial time. The other is not, but we solve it efficiently with projected gradient and coordinate descent. We demonstrate the effectiveness of these methods on a number of social networks obtained from Facebook. We also discuss how our methods can be generalized to other random walk-based similarities, such as personalized PageRank scores. Our code is available at <https://github.com/cnmusco/graph-similarity-learning>.

*Yale University. jeremy.hoskins@yale.edu

†Massachusetts Institute of Technology. cnmusco@mit.edu

‡Massachusetts Institute of Technology. cpmusco@mit.edu

§Boston University & Harvard University. ctsourak@bu.edu

1 Introduction

In graph mining and social network science, a variety of measures are used to quantify the similarity between nodes in a graph, including the shortest path distance, Jaccard’s coefficient between node neighborhoods, the Adamic-Adar coefficient [2], and hub-authority-based metrics [33, 12]. An important family of similarity measures are based on random walks, including SimRank [26], random walks with restarts [52], commute times [21], personalized PageRank [42, 27, 8], and DeepWalk embeddings [43]. These measures capture both local and global graph structure and hence are widely used in graph clustering and community detection [4, 46], anomaly detection [44], collaborative filtering [21, 47, 57], link prediction [39], and many other applications (including outside of network science, such as in computer vision [23]).

In this work we focus on these random walk-based similarity metrics. We initiate the study of a fundamental question:

How much information about a network can be learned given access to *a subset of potentially noisy estimates* of pairwise node similarities?

This question is important from a privacy perspective. A common privacy breach is *social link disclosure* [58], in which an attacker attempts to learn potentially sensitive links between nodes in a network. Such attacks are very common; fake accounts with engineered profiles are used to infiltrate and spy on social groups, potential employers may want to inspect the social network of a job candidate, and advertisers may wish to probe the demographic and interest information of a user to offer targeted ads. Thus, characterizing the ability of an attacker to reveal link information using pairwise node similarities is important in understanding the privacy implications of releasing such similarities, or information which can be used to compute them.

From a data mining perspective, computing all pairwise node similarities can be infeasible for large networks since the number of similarities grows quadratically in the number of nodes. Additionally, when the network cannot be accessed in full but can only be probed via crawling with random walks (e.g., a by third party studying a social network [31]), we may only have access to estimates of pairwise similarities rather than their exact values. Thus, understanding what information can still be learned from a partial, potentially noisy, set of node similarities is important when using these metrics in large scale graph mining.

1.1 Learning from Effective Resistances

In this paper, we focus on commute times, which are one of the most widely used random walk-based similarities. Commute times are a scaled version of *effective resistances*, they form a metric, and have major algorithmic applications, such as spectral graph sparsification [49]. Our ideas can be extended to related similarity measures, such as personalized PageRank, which we discuss in Section 4.4. It was shown in the seminal work of Liben-Nowell and Kleinberg [39] that effective resistances can be used to predict a significant fraction of future links appearing in networks from existing links, typically ranging from 5% up to 33%.

A difficulty associated with this task is that, in contrast to *local* similarity measures such as the number of common neighbors or the Adamic-Adar coefficient [2], node similarity under the effective resistance metric does not necessarily imply local connectivity. For example, two nodes connected by many long paths may be more similar than two nodes directly connected by a single edge.

Furthermore, in certain cases, the effective resistance between two nodes u, v tends to correlate well with a simple function of the degree sequence (specifically, $\frac{1}{d(u)} + \frac{1}{d(v)}$) [40, 54, 55], and it

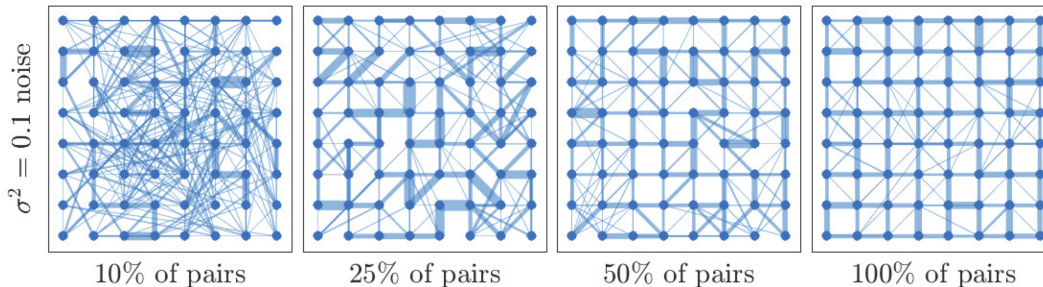


Figure 1: Grid graph learned from small, randomly selected fractions of effective resistance pairs using our proposed method. Edge thickness is proportional to edge weight.

is known that there are many graphs with the same degree sequence but very different global structures. Nevertheless, *considered in aggregate*, effective resistances encode global structure in a very strong way. For any graph, given all pairs effective resistances, it is possible to provably recover the full graph in polynomial time [48, 56]! This contrasts with purely local similarity metrics, which can be used heuristically for link prediction, but do not give network reconstruction in general. For instance, all-pairwise counts of common neighbors in any triangle free graph equal 0. Thus, these counts reveal no information about graph structure.

While the full information case is well understood, when all exact effective resistances are not available, little is known about what graph information can be learned. Some work considers reconstruction of trees based on a subset of effective resistances [18, 9, 50]. However outside of this special case, essentially nothing is known.

1.2 Our Contributions

We study in depth what can be learned about a graph given a subset of potentially noisy effective resistance estimates, from both a theoretical and empirical viewpoint. Our main contributions are:

Mathematical formulation. We provide an optimization-based formulation of the problem of learning a graph from effective resistances. Specifically, given a set of effective resistance measurements, we consider the problem of finding a graph whose effective resistances *match the given resistances as closely as possible*.

In general, there may be many different graphs which match any subset of all pairs effective resistances, and hence many minimizers to our optimization problem. If the resistances additionally have some noise, there may be no graph which matches them exactly but many which match them approximately. Nevertheless, as we show empirically, the graph obtained via our optimization approach typically recovers significant information about the underlying graph, including a large fraction of its edges, its global structure, and good approximations to *all* of its effective resistances.

Algorithms. We prove that, in some cases, the optimization problem we present can be solved exactly, in polynomial time. However, in general, the problem is non-convex and does not admit an obvious polynomial time solution. We give expressions for its gradient and Hessian, and show that it can be solved efficiently via iterative methods. In particular, we employ projected gradient and coordinate descent, as well as a powerful initialization strategy which allows us to find high quality solutions in most instances.

We also show that the problem can be relaxed to a convex formulation. Instead of searching for a

graph that matches all given effective resistance measurements, we just find a graph whose effective resistances are upper bounded by those given and which has minimum total edge weight. This modified problem is convex and can be solved in polynomial time via semidefinite programming.

Experimental Results. We evaluate our algorithms on several synthetic graphs and real Facebook ego networks, which contain all nodes in the social circle of a user. We demonstrate that given a small randomly selected fraction of all effective resistance pairs (10%-25%) we can learn a large fraction of a network – typically between 20% and 60% of edges. We show that this recovery is robust to adding Gaussian noise to the given effective resistances. A visual sample of our results is included in Figure 1 for a simple grid graph. We observe that using a small fraction of noisy pairwise similarities we can extract a lot of structural information about the graph. Here, noise is distributed for each edge independently as a Gaussian random variable with zero mean, and variance $\sigma^2 = 0.1$. We observe that as the size of the fraction of revealed similarities increases, an attacker can recover most of the graph.

We also demonstrate that by finding a graph which closely matches the given set of effective resistances (via our optimization approach), we in fact find a graph which closely matches the underlying network on *all* effective resistance pairs. This indicates that a significant amount of the information contained in all pairs effective resistances can be learned from just a small subset of these pairs, even when they are corrupted by noise.

2 Related Work

Link prediction and privacy in social networks. The link prediction problem was popularized by Liben-Nowell and Kleinberg [39]. The goal of *link prediction* is to infer which edges are likely to appear in the near-future, given a snapshot of the network. In [39] various node similarity measures are used to predict a non-trivial fraction of future interactions. Other works focus on predicting positive or negative links [36, 53]. For an extensive survey, see [3].

While closely related, link prediction differs from the problem we consider since, typically, one is given full access to a network snapshot (and in particular, could compute exact pairwise node similarities for this network) and seeks to make predictions about future evolutions of the network. In our setting, we are given partial information about a network (via a partial set of noisy similarity measures) and seek to learn existing links,

While link prediction is useful in applications ranging from understanding network evolution, to link recommendation, and predicting interactions between terrorists, it entails privacy risks. For instance, a malicious attacker can use link prediction to disclose information about sensitive connections in a sexual contact graph [34, 58]. Private link prediction is not adequately explored in the literature. Abebe and Nakos suggest a possible formalization [1].

Learning graphs. Learning graphs from data is central to many disciplines including machine learning [28], network tomography [14], bioinformatics [19], and phylogenetics [20]. The general framework is that there exists a hidden graph that we wish to discover by exploiting some kind of data, e.g., answers from a blackbox oracle returning certain (possibly noisy) graph measurements.

Theoretical work in this area has focused on worst case *query* complexity. Two representative examples include Angluin’s et al. work on learning graphs using edge detecting queries [5], and the recent work of Kannan, Mathieu, and Zhou using distance queries [29]. A number of works also consider learning graph parameters such as node count and mixing time by examining random walks

traces [31, 17, 10]. Experimental work has focused on recovering graphs from noisy measurements such as GPS traces [16], and distances between cell populations based on genetic differences [20, 19].

Learning trees. While a tree is special type of a graph, the case of learning trees deserves special mention since significant work on learning graphs from data has focused on trees. Distance-based reconstruction of trees aims to reconstruct a phylogenetic tree whose leaves correspond to n species, given their $\binom{n}{2}$ pairwise distances [20]. Note that on a tree, pairwise distances are identical to pairwise effective resistances. Batagelj et al. study tree realizability assuming access to fewer than $\binom{n}{2}$ leaves' pairwise distances [9]. A spectral method has been proposed by Stone and Griffing [50]. Culberson and Rudnicki [18] consider the problem of reconstructing a degree restricted tree given its distance matrix, see also [45].

3 Proposed Method

3.1 Notation and Preliminaries

For an undirected, weighted graph $G = (V, E, w)$ with n nodes, we let A be the $n \times n$ adjacency matrix. L denotes the (unnormalized) graph Laplacian: $L = D - A$, where D is a diagonal matrix with $D_{i,i}$ equal to the weighted degree of node i . For an integer $n > 0$, $[n]$ denotes the set $\{1, 2, \dots, n\}$. e_i denotes the i^{th} standard basis vector. For a matrix M , $M_{i,j}$ denotes the entry in its i^{th} row and j^{th} column.

Commute time and effective resistance. For two nodes $u, v \in V$, the hitting time $h_G(u, v)$ is the expected time it takes a random walk to travel from u to v . The *commute time* is its symmetrized version $c_G(u, v) = h_G(u, v) + h_G(v, u)$, i.e., the time to move from u to v and then back to u . For connected graphs, the effective resistance between u, v is a scaling of the commute time: $r_G(u, v) = \frac{c_G(u, v)}{\text{vol}(G)}$ where $\text{vol}(G) = 2 \sum_{e \in E} w_e$. Effective resistance has a natural electrical interpretation. When G is viewed as an electrical network on n nodes where each edge e corresponds to a link of conductance w_e (equivalently to a resistor of resistance $\frac{1}{w_e}$), the effective resistance is the voltage difference that appears across u, v when a unit current source is applied to them. Effective resistances (and hence commute times) always form a metric [32].

Letting $\chi_{u,v} = e_u - e_v$, the effective resistance between nodes u and v in a graph G with Laplacian matrix L can be computed as:

$$r_G(u, v) = \chi_{u,v}^T L^+ \chi_{u,v}. \quad (1)$$

Here L^+ denotes the Moore-Penrose pseudoinverse of L .

3.2 Problem Definition

We begin by providing a mathematic formulation of the problem introduced in Section 1 – that of learning the structure of a graph from partial and possibly noisy measurements of pairwise effective resistances. An analogous problem can be defined for other random walk-based similarities, such as personalized PageRank. We discuss initial results in this direction in Section 4.4.

Problem 1 (Graph Reconstruction From Effective Resistances). *Reconstruct an unknown graph G given a set of noisy effective resistance measurements,*

$$\bar{r}(u, v) = r_G(u, v) + n_{uv}$$

for each $(u, v) \in \mathcal{S}$, where $\mathcal{S} \subseteq [n] \times [n]$ is a set of node pairs and n_{uv} is a potentially random noise term.

We focus on three interesting cases of Problem 1:

Problem 1.1 $\mathcal{S} = [n] \times [n]$ and $n_{uv} = 0$ for all $(u, v) \in \mathcal{S}$. This is the *full information recovery problem*.

Problem 1.2 \mathcal{S} is a subset of $[n] \times [n]$ and $n_{uv} = 0$ for all $(u, v) \in \mathcal{S}$. In this setting we must learn G from a limited number of exact effective resistances.

Problem 1.3 \mathcal{S} is a subset of $[n] \times [n]$ and n_{uv} is a random term, e.g. a mean 0 normal random variable with variance σ^2 : $n_{uv} \sim \mathcal{N}(0, \sigma^2)$.

It is known that there is a unique graph consistent with any full set of effective resistance measurements (see e.g., [48] or the related problem in [56]). Additionally, this graph can be computed by solving a fully determined linear system. So, we can solve Problem 1.1 exactly in polynomial time. We illustrate this in Section 4.1.

From a privacy and data mining perspective, the limited information settings of Problems 1.2 and 1.3 are more interesting. In Section 4.1 we demonstrate that, when G is a tree, exact recovery is possible for Problem 1.2 whenever \mathcal{S} is a superset of G 's edges. However, in general, there is no simple closed form solution to these problems, and exact recovery of G is typically impossible. In particular, several graphs may be consistent with the measurements given. Thus, we address these cases by reposing Problem 1 as an optimization problem, in which we attempt to recover a graph matching the given effective resistances as best as possible.

3.3 Optimization Formulation

A natural formalization of Problem 1 is as a least squares problem.

Problem 2. *Given a set of vertex pairs $\mathcal{S} \subseteq [n] \times [n]$ and a target effective resistance $\bar{r}(u, v)$ for each $(u, v) \in \mathcal{S}$:*

$$\underset{\text{graph } H}{\text{minimize}} \quad F(H) \stackrel{\text{def}}{=} \sum_{(u,v) \in \mathcal{S}} [r_H(u, v) - \bar{r}(u, v)]^2. \quad (2)$$

Using formula (1) for effective resistances, Problem 2 can equivalently be viewed as an optimization problem over the set of graph Laplacians: we minimize $\sum_{(u,v) \in \mathcal{S}} [\chi_{u,v}^T L^+ \chi_{u,v} - \bar{r}(u, v)]^2$. While this set is convex, the objective function is not and thus it is unclear if it can be minimized provably in polynomial time. Nevertheless, we demonstrate that it is possible to solve the problem approximately using iterative methods. In Section 4.2 we derive projected gradient and coordinate

descent algorithms for the problem. Combined with a powerful initialization heuristic, our experiments show that these methods quickly converge to near global minimums of the objective function for many networks.

For Problem 1.2, where $\bar{r}(u, v)$ comprise a subset of the exact effective resistances for some graph G , $\min_H F(H) = 0$. This minimum may be achieved by multiple graphs (including G itself) if \mathcal{S} does not contain all effective resistance pairs. Nevertheless, we demonstrate experimentally in Section 5 that even when \mathcal{S} contains a small fraction of these pairs, an approximate solution to Problem 2 often recovers significant information about G , including a large fraction of its edges. Interestingly, we find that, while Problem 2 only minimizes over the subset \mathcal{S} , the recovered graph typically matches G on *all* effective resistances, explaining why it contains so much structural information about G . For Problem 1.3, if $\mathcal{S} = [n] \times [n]$ and the noise terms n_{uv} are distributed as i.i.d. Gaussians, it is not hard to see that Problem 2 gives the maximum likelihood estimator for G . We again show that an approximate solution can recover a large fraction of G 's edges.

We note that while we can solve Problem 2 quickly via iterative methods, we leave open provable polynomial time algorithms for solving this problem in the settings of both Problems 1.2 and 1.3.

Convex relaxation. As an alternative to Problem 2, we introduce an optimization formulation of Problem 1 that *is convex*. It is convenient here to optimize over the convex set of graph Laplacians.

Problem 3. Let \mathcal{L} be the convex set of $n \times n$ graph Laplacians. Given a set of vertex pairs $\mathcal{S} \subseteq [n] \times [n]$ and a target effective resistance $\bar{r}_{(u,v)}$ for every $(u, v) \in \mathcal{S}$,

$$\begin{aligned} & \underset{L \in \mathcal{L}}{\text{minimize}} && \text{Tr}(L) \\ & \text{subject to} && \chi_{u,v}^T L^+ \chi_{u,v} \leq \bar{r}(u, v) \quad \forall (u, v) \in \mathcal{S} \end{aligned}$$

Observe that we can trivially find a feasible solution for Problem 3 by setting L to be a large complete graph, in which case all effective resistances will be close to 0. By Rayleigh's monotonicity law, decreasing the weight on edges in L increases effective resistances. $\text{Tr}(L)$ is equal to the total degree of the graph corresponding to L , so the problem asks us to find a graph with as little total edge weight as possible that still satisfies the effective resistance constraints.

The disadvantage of this formulation is that it only encodes the target resistances $\bar{r}(u, v)$ as *upper bounds* on the resistances of L . The advantage is that we can solve Problem 3 provably in polynomial time via semidefinite programming (see Section 4.3). In practice, we find that it can sometimes effectively learn graph edges and structure from limited measurements.

Problem 3 is related to work on convex methods for minimizing total effective resistance or relatedly, mixing time in graphs [13, 51, 22, 22]. However, prior work does not consider pairwise resistance constraints and so is not suited to the graph learning problem.

4 Analytical Results and Algorithms

4.1 Full Graph Reconstruction – Problem 1

Problem 1 can be solved exactly in polynomial time when \mathcal{S} contains *all resistance pairs* of some graph G (i.e. Problem 1.1). In this case, there is a closed form solution for G 's Laplacian L and the solution is unique. This was pointed out in [48], however we include our own proof for completeness.

Theorem 1. *If there is a feasible solution to Problem 1.1 then it is unique and can be found in $O(n^3)$ time. Specifically, the Laplacian L of the unique solution G is given by*

$$-2 \cdot \left[\left(I - \frac{J}{n} \right) R \left(I - \frac{J}{n} \right) \right]^+ \quad (3)$$

where R is the matrix with $R_{u,v} = r_G(u, v)$ for all $u, v \in [n]$, I is the $n \times n$ identity matrix, and J is the $n \times n$ all ones matrix.

Proof. For an n node graph, the number of possible edges is $m = \binom{n}{2}$. Let $B \in \mathbb{R}^{m \times n}$ be the vertex edge incidence matrix of the complete graph with a row equal to $\chi_{u,v} = e_u - e_v$ for every (u, v) .

The Laplacian of any n vertex graph can be written as $L = B^T W B$, for some $W \in \mathbb{R}^{m \times m}$ which is a nonnegative diagonal matrix with entries corresponding to the edge weights.

We can rewrite the effective resistance formula in (1) as:

$$r_L(u, v) = \chi_{u,v}^\top L^+ \chi_{u,v} = (L^+)_{u,u} + (L^+)_{v,v} - 2(L^+)_{u,v}. \quad (4)$$

Since L^+ is symmetric we need only determine $\frac{n(n+1)}{2}$ unknown entries to determine the full matrix. Moreover, since the all ones vector is in the null space of L and therefore L^+ , we see that:

$$L^+_{u,u} = - \sum_{v \neq u} L^+_{u,v}, \quad (5)$$

and hence we can rewrite (4) as:

$$r_{u,v} = - \sum_{v' \neq v, u} (L^+)_{u,v'} - \sum_{u' \neq u, v} (L^+)_{u',v} - 4(L^+)_{u,v}. \quad (6)$$

Let M be the $m \times m$ matrix with rows and columns indexed by pairs $u, v \in [n]$ with $u \neq v$ and $(u_1, v_1), (u_2, v_2)$ entry given by:

$$M_{(u_1, v_1), (u_2, v_2)} = \begin{cases} -4 & \text{if } u_1 = u_2 \text{ and } v_1 = v_2 \\ -1 & \text{if } u_1 = u_2 \text{ or } v_1 = v_2 \\ 0 & \text{otherwise.} \end{cases}$$

Let $r \in \mathbb{R}^m$ contain each effective resistance $r_G(u, v)$. We can see from (6) that if we solve the linear system $Mx = -r$, as long as M is full rank and so the solution is unique, the entries of x will give us each $(L^+)_{u,v}$ with $u \neq v$. We can then use these entries to recover the remaining diagonal entries of L^+ using (5).

We can verify that M is in fact always full rank by writing $M = -|B||B|^T - 2I$, where $|B|$ denotes the matrix formed from B by taking the absolute value of each of its entries. We note that the non-zero eigenvalues of $|B||B|^T$ are equal to the non-zero eigenvalues of $|B|^T|B|$ which is the $n \times n$ matrix $\bar{M} = (n-2)I + J$. \bar{M} has eigenvalues $2n-2$ with multiplicity 1 and $n-2$ with multiplicity $n-1$. The remaining $m-n$ eigenvalues of $|B||B|^T$ are zero. Consequently, the eigenvalues of M are $-2n$ with multiplicity 1, $-n$ with multiplicity $n-1$ and -2 with multiplicity $m-n$. Thus M is full rank, proving that the effective resistances fully determine L^+ and thus L .

Solving for L via the linear system $Mx = -r$ would require $O(n^6)$ time, however, the closed form solution (3) given in Lemma 9.4.1 of [48] allows us to solve this problem in $O(n^3)$ time. \square

Reconstruction from hitting times. The above immediately generalizes to graph reconstruction from hitting times since, as discussed, for connected G , the effective resistance between u, v can be written as $r_G(u, v) = \frac{c_G(u, v)}{\text{vol}(G)} = \frac{h_G(u, v) + h_G(v, u)}{\text{vol}(G)}$. Thus, by Theorem 1, we can recover G up to a scaling from all pairs hitting times. This recovers a result in [56]. Note that if we scale all edge weights in G by a fixed factor, the hitting times do not change. Thus recovery up to a scaling is the best we can hope for in this setting.

Reconstruction from other similarity measures. An analogous result to Theorem 1 holds for graph recovery from all pairs personalized PageRank scores, and for related measures such as Katz similarity scores [30]. We discuss this direction in Section 4.4.

Are all pairs always necessary for perfect reconstruction? For general graphs, Problem 1 can only be solved exactly when \mathcal{S} contains all $\binom{n}{2}$ true effective resistances. However, given additional constraints on G , recovery is possible with much less information. In particular, when G is a tree, we can recover it (i.e., solve Problem 1.2) whenever \mathcal{S} is a superset of its edge set.

Roughly, since G is a tree, the effective resistance $r_G(u, v)$ is equal to the length of the unique path connecting u and v . As long as \mathcal{S} includes all edges in G , it fully determines all path lengths and hence the effective resistances for all pairs u, v . We can thus recover G via Theorem 1. Formally:

Theorem 2. *If G is a tree and a feasible solution to Problem 1.2 with edge set $E \subseteq \mathcal{S}$ then G is unique and can be found in $O(n^3)$ time.*

Proof. Let P_{uv} be the unique path between u, v in G . It is well known [15] that:

$$r_G(u, v) = \sum_{e \in P_{uv}} 1/w_e. \quad (7)$$

For $(u, v) \in \mathcal{S}$ set $\bar{r}(u, v) = r_G(u, v)$. Let \bar{G} be an undirected graph with an edge for each $(u, v) \in \mathcal{S}$ with length $r_G(u, v)$. For all $(u, v) \notin \mathcal{S}$, set $\bar{r}(u, v)$ to the shortest path distance between u and v in \bar{G} .

Claim. $\bar{r}(u, v) = r_G(u, v)$ for all $(u, v) \in [n] \times [n]$.

For *any pair* (u, v) , we have $\bar{r}(u, v) \leq r_G(u, v)$. The length of shortest path between u, v in \bar{G} is certainly at most the length of P_{uv} , which is contained in \bar{G} since $E \subseteq \mathcal{S}$. P_{uv} 's length in \bar{G} is: $\sum_{e \in P_{uv}} r_G(e) = \sum_{e \in P_{uv}} 1/w_e$ using (7). Thus, $\bar{r}(u, v) \leq \sum_{e \in P_{uv}} 1/w_e = r_G(u, v)$, again using (7).

Further, P_{uv} is in fact a shortest path between u, v in \bar{G} , giving that $\bar{r}(u, v) = r_G(u, v)$. This is because the length of every edge $(u, v) \in \mathcal{S}$ that is not in E just equals the length of path P_{uv} in \bar{G} (i.e., $r_G(u, v)$) and so removing this edge from \bar{G} does not change any shortest path distance. So we can assume that \bar{G} just contains the edges in E , and so P_{uv} is the unique path between u, v .

Given the above claim, the theorem follows since we can compute each $\bar{r}(u, v)$ from the effective resistances of the edges in \mathcal{S} and can then compute G from these resistances by Theorem 1. \square

The problem of recovering trees from pairwise distance measurements is a central problem in phylogenetics. There are other cases when just a subset of effective resistances is known to allow full recovery, for example when the effective resistances between any pair of leaves is known [50, 48]. Also related to our analysis for trees lies is the work of Mathieu and Zhou [41], which reconstructs graphs with bounded degree from pairwise distance measurements.

4.2 Graph Learning via Least Squares Minimization – Problem 2

When Problem 1 cannot be solved exactly, e.g. in the settings of Problems 1.2 and 1.3, an effective surrogate is to solve Problem 2 to find a graph with effective resistances close to the given target resistances. As we demonstrate experimentally in Section 5, this yields good solutions to Problems 1.2 and 1.3 in many cases. Problem 2 is non-convex, however we show that a good solution can often be found efficiently via projected gradient descent.

Optimizing over edge weights. Let $m = \binom{n}{2}$. We write the Laplacian of the graph H as $L(w) \stackrel{\text{def}}{=} B^T \text{diag}(w) B$, where $w \in \mathbb{R}^m$ is a non-negative vector whose entries correspond to the edge weights in H , $\text{diag}(w)$ is the $m \times m$ matrix with w as its diagonal, and $B \in \mathbb{R}^{m \times n}$ is the vertex edge incidence matrix with a row equal to $\chi_{u,v} = e_u - e_v$ for every possible edge $(u, v) \in [n] \times [n]$.

Optimizing the objective function $F(H)$ in Problem 2 is equivalent to optimizing $F(w)$ over the edge weight vector w , where we define $F(w) \stackrel{\text{def}}{=} F(H)$ for the unique H with Laplacian equal to $L(w)$.

We restrict $w_i \geq 0$ for all i and project to this constraint after each gradient step simply by setting $w_i := \max(w_i, 0)$. The gradient of $F(w)$ can be computed in closed form. We first define an auxiliary variable, $R(w) \in \mathbb{R}^{m \times m}$, whose diagonal contains all pairwise effective resistances of H with weight vector w :

Definition 1. For $w \in \mathbb{R}^m$ with $w_i \geq 0$ for all i , define

$$R(w) = BL(w)^+ B^T.$$

Using $R(w)$ we can compute the gradient of $F(w)$ by:

Proposition 1. Let \circ denote the Hadamard (entrywise) product for matrices. Define the error vector $\Delta(w) \in \mathbb{R}^m$ as having $\Delta(w)_i = \bar{r}(i) - [R(w)]_{i,i}$ for all $i \in \mathcal{S}$ and 0s elsewhere. We have:

$$\nabla F(w) = 2 (R \circ R) \Delta(w)$$

Proof. We begin by observing that, letting e_i be the i th standard basis vector in \mathbb{R}^m , for any weight vector w , the graph Laplacian corresponding to the weight vector $w + \epsilon e_i$, is $L(w) + \epsilon b_i b_i^T$. The Sherman-Morrison formula for the matrix pseudoinverse yields:

$$(L(w) + \epsilon b_i b_i^T)^+ = L(w)^+ - \epsilon \frac{L(w)^+ b_i b_i^T L(w)^+}{1 + \epsilon b_i^T L(w)^+ b_i},$$

and, hence thinking of $L(w)^+$ as a matrix-valued function of w ,

$$\begin{aligned} \frac{\partial L(w)^+}{\partial w_i} &= \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} \left[L(w)^+ - \epsilon \frac{L(w)^+ b_i b_i^T L(w)^+}{1 + \epsilon b_i^T L(w)^+ b_i} - L(w)^+ \right] \\ &= -L(w)^+ b_i b_i^T L(w)^+. \end{aligned}$$

Let R_i denote the i^{th} column of $R(w) = BL(w)^+ B^T$. By linearity:

$$\frac{\partial R}{\partial w_i} = -B (L(w)^+ b_i b_i^T L(w)^+) B^T = -R_i R_i^T.$$

Thus,
$$\frac{\partial F}{\partial w_i} = 2 \sum_{j \in \mathcal{S}} (\bar{r}(j) - [R(w)]_{j,j}) \cdot [R(w)]_{i,j}^2,$$

and so we obtain that the gradient equals $\nabla F(w) = 2(R \circ R) \Delta(w)$. \square

While gradient descent works well in our experiments, one may also apply second order methods, which require $F(w)$'s Hessian. Using similar computations to those in Proposition 1 we obtain:

Proposition 2. *Let $I_{\mathcal{S}} \in \mathbb{R}^{m \times m}$ be the diagonal matrix with a 1 at each entry corresponding to $i \in \mathcal{S}$ and 0s elsewhere and $\Delta(w)$ be as defined in Proposition 1. The Hessian matrix of $F(w)$ is:*

$$H_F(w) = -4[R \text{diag}(\Delta(w))R] \circ R + 2(R \circ R)I_{\mathcal{S}}(R \circ R).$$

Acceleration via coordinate descent. Naively computing the gradient $\nabla F(w)$ via Proposition 1 requires computing the full $m \times m$ matrix $R(w)$, which can be prohibitively expensive for large graphs – recall that $m = \binom{n}{2} = O(n^2)$. Note however, that the error vector $\Delta(w)$ only has nonzero entries at positions corresponding to the node pairs in \mathcal{S} . Thus, it suffices to compute just $|\mathcal{S}|$ columns of R corresponding to these pairs, which can give a significant savings.

We obtain further savings using block coordinate descent. At each step we restrict our updates to a subset of edges $\mathcal{B} \subseteq [n] \times [n]$. Let $I_{\mathcal{B}}$ be the matrix with a 1 at diagonal entries corresponding to elements of \mathcal{B} and 0's elsewhere. We step in the direction of $I_{\mathcal{B}} \nabla F(w)$. Computing this step only requires forming the rows of R corresponding to edges in \mathcal{B} . A typical way to choose \mathcal{B} is at random. See Section 5.1 for the actual implementation details.

Initialization. A good initialization for gradient descent can significantly accelerate the solution of Problem 2. We use a strategy based on the exact solution to Problem 1.1 in Theorem 1.

Since effective resistances form a metric, by triangle inequality, for any $u, v, w \in [n]$, $r_H(u, v) \leq r_H(u, w) + r_H(w, v)$. Guided by this fact, given target resistances $\bar{r}(u, v)$ for $(u, v) \in \mathcal{S}$, we first “fill in” the constraint set. For $(w, z) \notin \mathcal{S}$, we set $\bar{r}(w, z)$ equal to the shortest path distance in the graph \bar{G} which has an edge for each pair in \mathcal{S} with length $\bar{r}(u, v)$.

We thus obtain a full set of target effective resistances. We can form R with $R_{u,v} = \bar{r}(u, v)$ and initialize the Laplacian of H using the formula given in (3) in Theorem 1. However, this formula is quite unstable and generally yields an output which is far from a graph Laplacian even when R is corrupted by a small amount of noise. So we instead compute a regularized estimate,

$$\tilde{L} = -2 \cdot \left[\left(I - \frac{J}{n} \right) R \left(I - \frac{J}{n} \right) + \lambda I \right]^+,$$

where $\lambda > 0$ can be chosen e.g. by line search. Generally, \tilde{L} will not be a valid graph Laplacian, but by removing negative edge weights, we typically obtain a good initialization for Problem 2.

4.3 Graph Learning via Convex Optimization – Problem 3

We finally discuss how to efficiently solve our convex formulation, Problem 3. We express this problem as a semidefinite program (SDP) which can be solved via a number of available packages.

We can re-express our effective resistance constraint as a positive semidefinite constraint using the Schur complement condition:

$$\chi_{u,v}^T L^+ \chi_{u,v} \leq \bar{r}(u, v) \quad \text{iff} \quad \begin{bmatrix} L & \chi_{u,v} \\ \chi_{u,v}^T & \bar{r}(u, v) \end{bmatrix} \succeq 0 \text{ and } L \succeq 0.$$

Doing so yields the following program:

Problem 4 (SDP Form of Problem 3). *Given vertex pairs $\mathcal{S} \subseteq [n] \times [n]$, and target effective resistance $\bar{r}(u, v)$ for every $(u, v) \in \mathcal{S}$,*

$$\underset{L \in \mathcal{L}}{\text{minimize}} \text{Tr}(L)$$

subject to

$$L \succeq 0 \text{ and } \forall (u, v) \in \mathcal{S}, \begin{bmatrix} L & \chi_{u,v} \\ \chi_{u,v}^T & \bar{r}(u, v) \end{bmatrix} \succeq 0$$

We require L to be a valid graph Laplacian (i.e., constrain $L \in \mathcal{L}$) by adding linear constraints of the form:

$$\forall i, L_{i,i} = - \sum_{j \neq i} L_{i,j} \text{ and } \forall i \neq j, L_{i,j} \leq 0.$$

4.4 Extensions to Other Similarity Measures

As discussed, our results generalize to random walk-based node similarities beyond effective resistances, such as personalized PageRank scores. Given localization parameter $\alpha \geq 0$, the personalized PageRank score $p_G^\alpha(u, v)$ is the probability that a lazy random walk on graph G which jumps back to u with probability α in each step is at node v in its stationary distribution [42, 25, 4].

Letting $W = \frac{1}{2}(I + AD^{-1})$ be the lazy random walk matrix, $p_G^\alpha(u, v)$ is the v^{th} entry of the personalized PageRank vector:

$$p_G^\alpha(u) = \alpha(I - (1 - \alpha)W)^{-1}e_u. \quad (8)$$

This vector gives the stationary distribution for the random walk on G and thus the personalized PageRank $p_G^\alpha(u, v)$ is its v^{th} entry.

It is not hard to show an analogous result to Theorem 1, that given a full set of exact personalized PageRank scores, full recovery of G is possible. Roughly, if we let P be the matrix with $p_G^\alpha(u)$ as its u^{th} column, we have $P = \alpha(I - (1 - \alpha)W)^{-1}$ and can thus solve for the random walk matrix W , and the graph G . This gives:

Theorem 3. *For any connected graph G , given personalized PageRank score $p_G^\alpha(u, v)$ for each $(u, v) \in [n] \times [n]$, there is algorithm returning G (up to a scaling of its edge weights) in $O(n^3)$ time.*

Further, it is possible to formulate a problem analogous to Problem 2 and solve for a graph matching a subset of personalized PageRank measurements as closely as possible. As shown in Figure 2, personalized PageRank often gives a stronger signal of global graph structure than effective resistance. To create the plot, nodes are sorted by their value in the Laplacian Fielder vector, which corresponds roughly to residence in different clusters. In an extended version of our work, we will provide detailed empirical results with personalized PageRank, and other random walk measures.

5 Empirical results

In this section we present an in depth experimental study of how well our methods can learn a graph given a set of (noisy) effective resistance measurements. We seek to answer two key questions:




Figure 2: Personalized PageRank correlates better than commute times with the cluster structure in the FB SMALL C network (see Table 1). Heatmaps are shown in log scale.

1. Given a set of effective resistance measurements, can we find a graph matching these measurements via the optimization formulations posed in Problems 2 and 3 and the corresponding algorithms discussed in Sections 4.2 and 4.3?
2. What structure does the graph we learn via our optimization approach share with the underlying network that produced the resistance measurements? Can it be used to predict links in the network? Does it approximately match the network on effective resistances outside the given set of measurements, or share other global structure?

We address these questions by examining a variety of synthetic, and social network graphs.

5.1 Experimental Setup

Table 1 lists the networks analyzed in our experiments. These include two synthetic examples: an 8×8 two dimensional grid graph and a k -nearest neighbor graph constructed for vectors drawn from a Gaussian mixture model with two clusters. The other networks are Facebook ‘ego networks’ obtained from the Stanford Network Analysis Project (SNAP) collection [37, 38]. Each of these networks is formed by taking the largest connected component in the social circle of a specific user (whose nodeId is shown in Table 1).

For all experiments, we provide our algorithms with effective resistances that are uniformly sampled from the set of all $\binom{n}{2}$ effective resistances. We sample a fixed fraction $f \stackrel{\text{def}}{=} \frac{|S|}{\binom{n}{2}} \times 100\%$ of all possible measurements. We typically use $f \in \{10, 25, 50, 100\}\%$. In some cases, these resistances are corrupted with i.i.d. Gaussian noise $\eta \sim \mathcal{N}(0, \sigma^2)$. We experiment with different values of variance σ^2 .

For Problem 2 we implemented gradient decent based on the closed-form gradient calculation in Section 4.2. Line search was used to optimize step size at each iteration since it significantly outperformed implementations with fixed step sizes. For larger problems, block coordinate descent was used as described in Section 4.2, with the coordinate set chosen uniformly at random in each iteration. We set the block size $|\mathcal{B}| = 5000$. For Problem 3 we used MOSEK convex optimization software, accessed through the CVX interface [6, 24]. All experiments were run on a computer with a 2.6 GHz Intel Core i7 processor and 16 GB of main memory.

Name	# of nodes, n	# of edges, m
GRID (SYNTHETIC)	64	224
K-NN (SYNTHETIC)	80	560
FB SMALL A (NODEID 698)	40	220
FB SMALL B (NODEID 3980)	44	138
FB SMALL C (NODEID 414)	148	1692
FB SMALL D (NODEID 686)	168	1656
FB MEDIUM A (NODEID 348)	224	3192
FB MEDIUM B (NODEID 0)	324	2514
FB LARGE A (NODEID 3437)	532	4812
FB LARGE B (NODEID 1912)	795	30023

Table 1: Datasets for experiments. FB denotes “Facebook”.

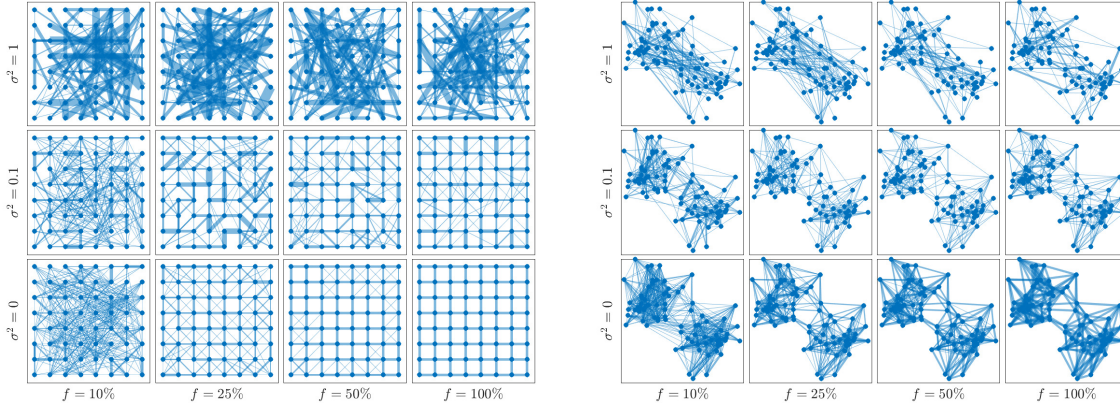


Figure 3: Graphs learned by solving Problem 2 with gradient descent run to convergence for uniformly sampled effective resistances with varying levels of Gaussian noise. Edge width is proportional to edge weight in the plots.

5.2 Learning Synthetic Graphs

We first evaluate our graph learning algorithms on GRID and K-NN, which are simple synthetic graphs with clear structure.

Least squares formulation. We first observe that gradient descent effectively minimizes the objective function of Problem 2 on the GRID and K-NN graphs. We consider the normalized objective for a constraint set \mathcal{S} and output graph H :

$$\hat{F}(H) = \frac{\sum_{(u,v) \in \mathcal{S}} [r_H(u,v) - \bar{r}(u,v)]^2}{\sum_{(u,v) \in \mathcal{S}} \bar{r}(u,v)^2}. \quad (9)$$

For noise variance 0, $\min_H \hat{F}(H) = 0$ and in Figure 4 we see that for GRID we in fact find H with $\hat{F}(H) \approx 0$ for varying sizes of \mathcal{S} . Convergence is notably faster when 100% of effective resistances are included in \mathcal{S} , but otherwise does not correlate strongly with the number of constraints.

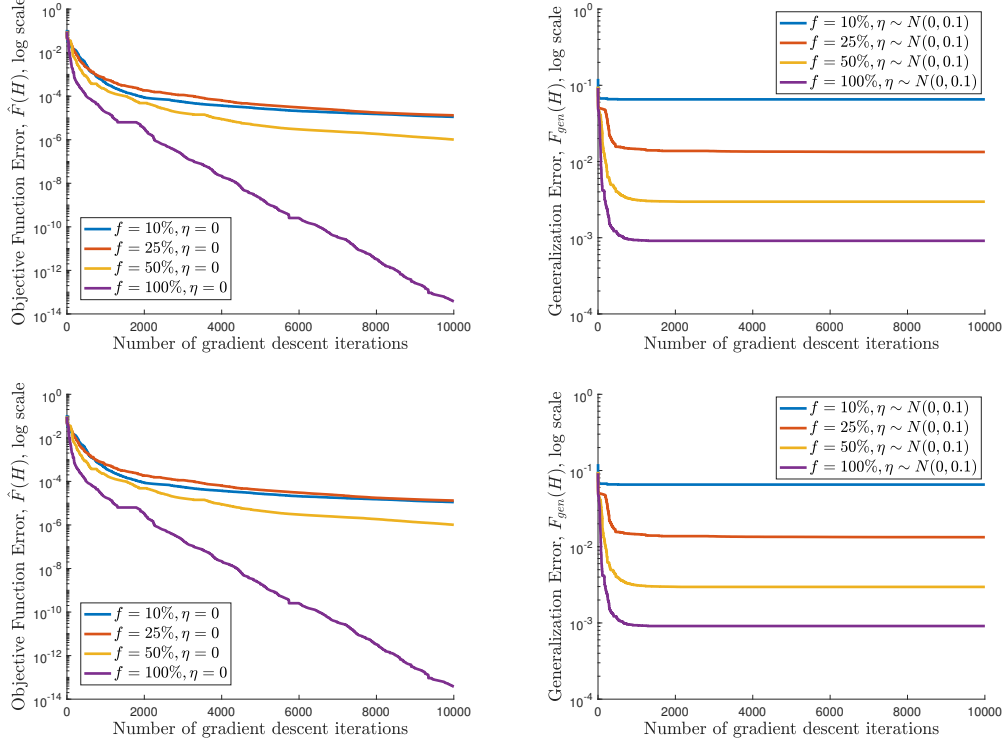


Figure 4: Objective error and generalization error for Problem 2, as defined in (10) for GRID. For details, see Section 5.2.

In Figure 4 we also plot the *generalization error*:

$$F_{gen}(H) = \frac{\sum_{(u,v) \in [n] \times [n]} [r_H(u,v) - r_G(u,v)]^2}{\sum_{(u,v) \in [n] \times [n]} r_G^2(u,v)}, \quad (10)$$

where $r_G(u,v)$ is the true effective resistance, uncorrupted by noise. $F_{gen}(H)$ measures how well the graph obtained by solving Problem 2 matches *all* effective resistances of the original network. We confirm that generalization decreases with improved objective function performance, indicating that optimizing Problem 2 effectively extracts network structure from a small set of effective resistances. We observe that the generalization error is small even when $f = 10\%$, and becomes negligible as we increase the fraction f of measurements, even in the presence of noise.

We repeat the same experiments with Gaussian noise added to each resistance measurement. The variance of the noise, σ^2 , is scaled relatively to the mean effective resistance in the graph, i.e., we set $\bar{r}(u,v) = r_G(u,v) + \mathcal{N}(0, \bar{\sigma}^2)$ where:

$$\bar{\sigma}^2 = \frac{\sigma^2}{\binom{n}{2}} \cdot \sum_{(u,v) \in [n] \times [n]} r_G(u,v). \quad (11)$$

While generally $\min_H \hat{F}(H) > 0$ when $\bar{r}(u,v)$ is noisy (it is likely that there is no graph consistent with these noisy measurements), we see that the objective value still decreases steadily with a larger number of iterations. Generalization error also decreases as desired.

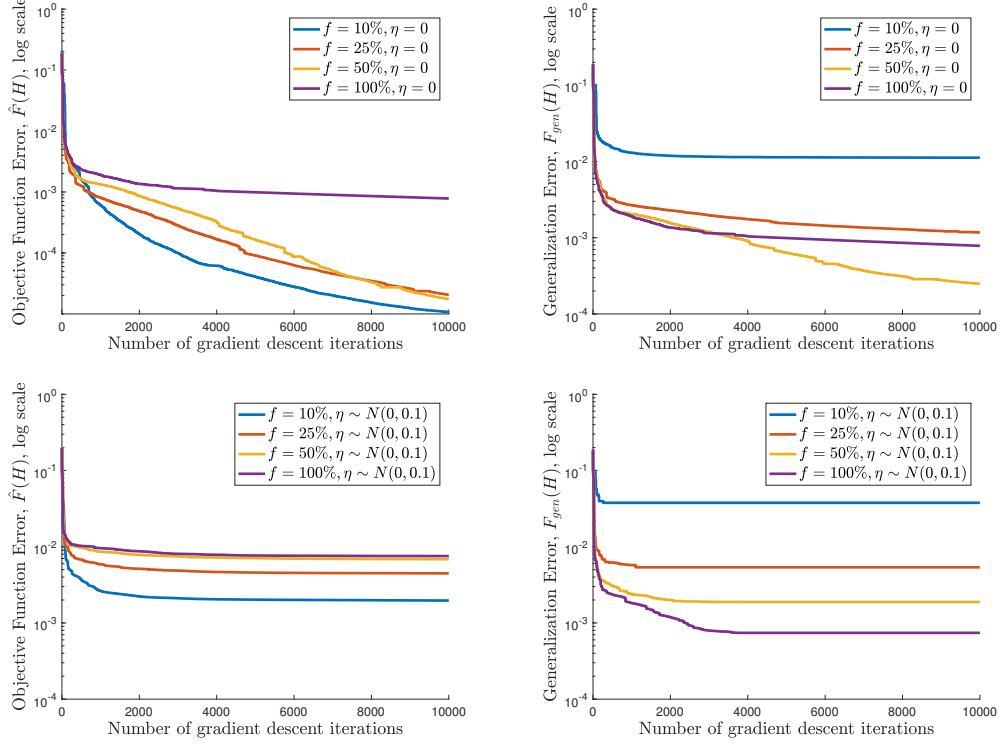


Figure 5: Objective error and generalization error for Problem 2, as defined in (10) for K-NN. Objective function error decreases steadily, leading to improved generalization error.

We obtain similar results by applying Problem 2 to a k -nearest neighbor graph with two clear clusters of data points. Again, it is apparent in Figure 5 that gradient descent converges easily for a variety of noise levels and constraint sets. Convergence on the least squares objective leads to improved generalization error.

Figure 3 shows the graphs obtained from solving the problem for varying σ^2 and f . For both graphs, when $\sigma^2 = 0$ and $f = 100\%$, the original network is recovered exactly. Reconstruction accuracy decreases with increasing noise and a decreasing number of constraints. For GRID, even with 25% of constraints, nearly full recovery is possible for $\sigma^2 = 0$ and recovery of approximately half of true edges is possible for $\sigma^2 = 0.1$. For K-NN, for $\sigma^2 = 0$ and $\sigma^2 = 0.1$ we observe that cluster structure is recovered. Detailed quantitative results for both networks are given in Table 3.

Convex formulation. We next evaluate the performance graph learning via the convex relaxation in Problem 3. In this case, we do not focus on convergence as we solve the problem directly using a semidefinite programming (SDP) routine. Unlike for Problem 2, solving Problem 3 does not recover the exact input graph, even in the noiseless all pairs effective resistance case. This is because the input graph does not necessarily minimize the objective of Problem 3 since there can be other graphs with smaller total edge weight and lower effective resistances.

However, the learned graphs *do capture* information about edges in the original: their heaviest edges typically align with true edges in the target graph. This property is captured in the quantitative results of Table 3. Qualitatively, it is very apparent for GRID: in Figure 6 we mark the 224

GRID, 64 nodes		K-NN, 80 nodes	
# of constraints	Time (min.)	# of constraints	Time (min.)
202	238.8	316	1216.0
504	761.6	790	3673.2
1008	1309.5	1580	8008.5
2016	2183.7	3160	16192.1

Table 2: Semidefinite program (SDP) optimization for Problem 3. Runtime is averaged over noise levels $\sigma^2 = 0, 0.1, 1$.

heaviest edges in the learned graph in red and note that this set converges exactly on the grid.

The convex formulation never significantly outperforms the least squares formulation, and significantly underperforms for small constraint sets. Additionally, the semidefinite program scales poorly with the number of nodes and effective resistance constraints. Sample runtimes are included in Table 2. Due to these considerations, we use the least squares formulation of Problem 2 instead of the convex formulation in our experiments on real social networks. However, we believe there is further opportunity for exploring Problem 3, especially given its provable runtime guarantees.

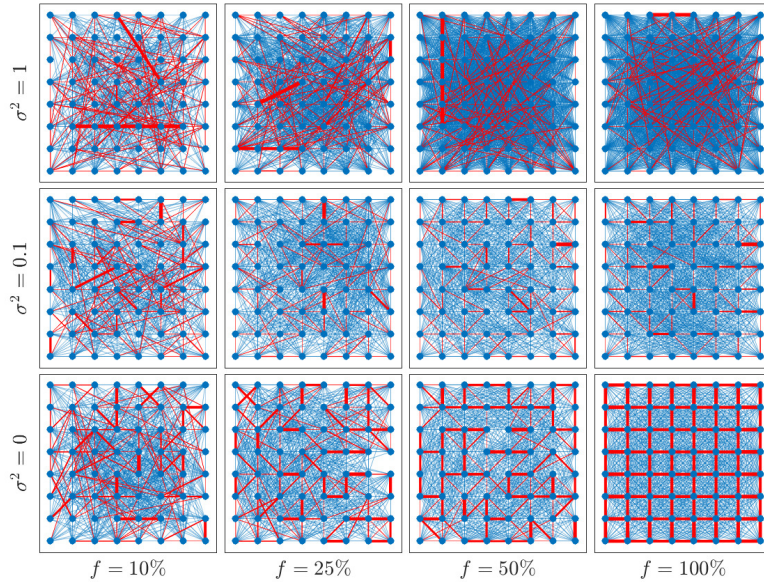


Figure 6: Graphs learned from solving the convex program in Problem 3 for uniformly sampled effective resistances from GRID with varying f, σ^2 . Heaviest edges marked in red.

5.3 Learning Social Network Graphs

We conclude by demonstrating the effectiveness of the least squares formulation of Problem 2 in learning Facebook ego networks from limited effective resistance measurements. We consider three metrics of performance, shown in Table 3 for a number of networks learned from randomly sampled subsets of effective resistances, corrupting with varying levels of noise.

Network	Algorithm	σ^2	Objective function error		Effect. resistance generalization error		% Edges learned baseline	% Edges learned	
			$f = 10\%$	$f = 25\%$	$f = 10\%$	$f = 25\%$		$f = 10\%$	$f = 25\%$
GRID	GD	0	.00001	.00001	.06559	.00099	5.56	20.54	88.39
	GD	.1	.00090	.00514	.08129	.01336		25.89	50.00
	SDP	0	na	na	.08758	.07422		16.07	25.00
	SDP	.1	na	na	.09549	.09343		12.50	26.79
K-NN	GD	0	.00001	.00002	.01122	.00117	11.58	44.54	72.68
	GD	.1	.00197	.00447	.05536	.00709		25.96	41.53
	SDP	0	na	na	.09314	.10399		27.05	48.36
	SDP	.1	na	na	.11899	.14097		24.32	39.89
FB SMALL A	GD	0	.01345	.00001	.21097	.00984	28.20	44.54	75.00
	GD	.1	.00017	.00204	.07964	.01687		53.64	60.00
FB SMALL B	GD	0	.00002	.00003	.01515	.00623	14.59	42.75	48.55
	GD	.1	.00032	.00206	.02229	.01291		36.23	43.48
FB SMALL C	GD	0	.00162	.00166	.00217	.00203	15.55	57.03	59.51
	GD	.1	.00532	.00644	.01542	.00218		52.66	57.51
FB SMALL D	GD	0	.00335	.00434	.00821	.00830	11.80	21.92	24.52
	GD	.1	.00610	.18384	.00923	.21426		21.38	21.20
FB MEDIUM A	GD	0	.00447	.00665	.02910	.01713	12.78	23.50	25.59
FB MEDIUM B	CD	0	.00224	.01255	.00862	.01471	4.80	18.97	22.15
	CD	.1	.01174	.03182	.01687	.03295		17.50	16.03
FB LARGE A	CD	0	.00516	.00796	.00682	.00862	3.41	10.52	12.45
FB LARGE B	CD	0	.00524	.00440	.00635	.00580	9.51	20.26	24.95
	CD	.1	.12745	.34646	.14532	.36095		19.43	16.97

Table 3: Graph recovery results. All results use a randomly sampled subset of $f = 10\%$ or $f = 25\%$ of all effective resistances. For “Algorithm”, GD denotes gradient descent and CD denotes block coordinate descent with random batches of size 5000. “Noise level, σ^2 ” indicates that the target resistances were set to $\bar{r}(u, v) = r_G(u, v) + \mathcal{N}(0, \sigma^2 \cdot \text{mean}_{u,v}(r_G(u, v)))$. “% Edges baseline”, is the edge density of the underlying network, equivalent to the expected edge prediction accuracy achieved with random guessing.

1. Objective Function Value: the value of the objective function of Problem 2, normalized as described in (9).
2. Generalization Error: the error in reconstructing the full set of effective resistances of the true graph, as defined in (10).
3. Edges Learned: the rate of recovery for edges in the true graph. We utilize a standard metric from the link prediction literature [33]: given underlying graph G with m edges and learned graph H , we consider the m heaviest edges of H and compute the percentage of G ’s edges contained in this set.

Results. We find that as for the synthetic GRID and K-NN graphs, we can effectively minimize the objective function of Problem 2 for the Facebook ego networks. Moreover, this minimization leads to very good generalization error in nearly all cases. i.e., we can effectively learn a graph matching our input on all effective resistances, even when we consider just a small subset.

For all graphs, we are able to recover a significant fraction of edges, even when just considering 10% or 25% of effective resistance pairs. We obtain the best recovery for small graphs, learning over half of the true edges in FB SMALL A and FB SMALL C. Even for larger graphs, we can typically recover over 20% of the true edges.

Typically, the number of edges learned increases as we increase the number of constraints and decrease the noise variance. However, occasionally, considering fewer effective resistances in fact improves learning, possibly because we more effectively solve the underlying optimization problem.

6 Conclusion

In this work, we show that a small subset of noisy effective resistances can be used to learn significant information about a network, including predicting a large percentage of edges and recovering global structure, such as accurate approximations to all pairwise effective resistances. From a privacy standpoint, our results raise major concerns about releasing random walk-based pairwise node similarity information as it entails significant privacy risk. From a data mining perspective, our methods can be used for graph mining, even when computing all effective resistances exactly is infeasible.

Our work leaves a number of future research directions open, including giving a provable polynomial time algorithm for the least squares formulation (Problem 2), extending our work to other similarity metrics, and scaling our methods to larger social networks.

References

- [1] R. Abebe and V. Nakos. Private link prediction in social networks. 2014.
- [2] L. A. Adamic and E. Adar. Friends and neighbors on the web. *Social Networks*, 25(3):211–230, 2003.
- [3] M. Al Hasan and M. J. Zaki. A survey of link prediction in social networks. In *Social Network Data Analytics*, pages 243–275. Springer, 2011.
- [4] R. Andersen, F. Chung, and K. Lang. Local graph partitioning using PageRank vectors. In *Foundations of Computer Science, 2006. FOCS’06. 47th Annual IEEE Symposium on*, pages 475–486. IEEE, 2006.
- [5] D. Angluin and J. Chen. Learning a hidden graph using $o(\log n)$ queries per edge. *Journal of Computer and System Sciences*, 74(4):546–556, 2008.
- [6] M. ApS. *The MOSEK Optimization Suite*, 2017.
- [7] L. Backstrom, C. Dwork, and J. Kleinberg. Wherefore art thou?: anonymized social networks, hidden patterns, and structural steganography. In *Proceedings of the 16th International Conference on World Wide Web*, pages 181–190. ACM, 2007.
- [8] B. Bahmani, A. Chowdhury, and A. Goel. Fast incremental and personalized PageRank. *Proceedings of the VLDB Endowment*, 4(3):173–184, 2010.
- [9] V. Batagelj, T. Pisanski, and J. M. S. S. os Pereira. An algorithm for tree-realizability of distance matrices. *International Journal of Computer Mathematics*, 34(3-4):171–176, 1990.
- [10] A. Ben-Hamou, R. I. Oliveira, and Y. Peres. Estimating graph parameters via random walks with restarts. *SODA 2018 (to appear)*, 2018.

- [11] S. Bhagat, G. Cormode, B. Krishnamurthy, and D. Srivastava. Class-based graph anonymization for social network data. *Proceedings of the VLDB Endowment*, 2(1):766–777, 2009.
- [12] V. D. Blondel, A. Gajardo, M. Heymans, P. Senellart, and P. Van Dooren. A measure of similarity between graph vertices: Applications to synonym extraction and web searching. *SIAM review*, 46(4):647–666, 2004.
- [13] S. Boyd, P. Diaconis, and L. Xiao. Fastest mixing Markov chain on a graph. *SIAM review*, 46(4):667–689, 2004.
- [14] R. Castro, M. Coates, G. Liang, R. Nowak, and B. Yu. Network tomography: Recent developments. *Statistical Science*, pages 499–517, 2004.
- [15] A. K. Chandra, P. Raghavan, W. L. Ruzzo, R. Smolensky, and P. Tiwari. The electrical resistance of a graph captures its commute and cover times. *Computational Complexity*, 6(4):312–340, 1996.
- [16] D. Chen, L. J. Guibas, J. Hershberger, and J. Sun. Road network reconstruction for organizing paths. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1309–1320. Society for Industrial and Applied Mathematics, 2010.
- [17] C. Cooper, T. Radzik, and Y. Siantos. Estimating network parameters using random walks. *Social Network Analysis and Mining*, 4(1):168, 2014.
- [18] J. C. Culberson and P. Rudnicki. A fast algorithm for constructing trees from distance matrices. *Information Processing Letters*, 30(4):215 – 220, 1989.
- [19] R. Desper, F. Jiang, O.-P. Kallioniemi, H. Moch, C. H. Papadimitriou, and A. A. Schäffer. Inferring tree models for oncogenesis from comparative genome hybridization data. *Journal of Computational Biology*, 6(1):37–51, 1999.
- [20] J. Felsenstein. Confidence limits on phylogenies: an approach using the bootstrap. *Evolution*, 39(4):783–791, 1985.
- [21] F. Fouss, A. Pirotte, J.-M. Renders, and M. Saerens. Random-walk computation of similarities between nodes of a graph with application to collaborative recommendation. *IEEE Transactions on Knowledge and Data Engineering*, 19(3):355–369, 2007.
- [22] A. Ghosh, S. Boyd, and A. Saberi. Minimizing effective resistance of a graph. *SIAM review*, 50(1):37–66, 2008.
- [23] L. Grady. Random walks for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(11):1768–1783, 2006.
- [24] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. <http://cvxr.com/cvx>, Mar. 2014.
- [25] T. H. Haveliwala. Topic-sensitive PageRank: A context-sensitive ranking algorithm for web search. *IEEE Transactions on Knowledge and Data Engineering*, 15(4):784–796, 2003.

- [26] G. Jeh and J. Widom. Simrank: a measure of structural-context similarity. In *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 538–543. ACM, 2002.
- [27] G. Jeh and J. Widom. Scaling personalized web search. In *Proceedings of the 12th International Conference on World Wide Web*, pages 271–279. ACM, 2003.
- [28] V. Kalofolias. How to learn a graph from smooth signals. In *The 19th International Conference on Artificial Intelligence and Statistics (AISTATS 2016)*. Journal of Machine Learning Research (JMLR), 2016.
- [29] S. Kannan, C. Mathieu, and H. Zhou. Near-linear query complexity for graph inference. In *International Colloquium on Automata, Languages, and Programming*, pages 773–784. Springer, 2015.
- [30] L. Katz. A new status index derived from sociometric analysis. *Psychometrika*, 18(1):39–43, 1953.
- [31] L. Katzir, E. Liberty, and O. Somekh. Estimating sizes of social networks via biased sampling. In *Proceedings of the 20th International Conference on World Wide Web*, pages 597–606. ACM, 2011.
- [32] D. J. Klein and M. Randić. Resistance distance. *Journal of Mathematical Chemistry*, 12(1):81–95, 1993.
- [33] J. Kleinberg, R. Kumar, P. Raghavan, S. Rajagopalan, and A. Tomkins. The web as a graph: measurements, models, and methods. *Computing and Combinatorics*, pages 1–17, 1999.
- [34] A. Korolova, R. Motwani, S. U. Nabar, and Y. Xu. Link privacy in social networks. In *Proceedings of the 17th ACM Conference on Information and Knowledge Management*, pages 289–298. ACM, 2008.
- [35] N. Korula and S. Lattanzi. An efficient reconciliation algorithm for social networks. *Proceedings of the VLDB Endowment*, 7(5):377–388, 2014.
- [36] J. Leskovec, D. Huttenlocher, and J. Kleinberg. Predicting positive and negative links in online social networks. In *Proceedings of the 19th International Conference on World Wide Web (WWW)*, pages 641–650. ACM, 2010.
- [37] J. Leskovec and A. Krevl. SNAP Datasets: Stanford large network dataset collection. <http://snap.stanford.edu/data>, June 2014.
- [38] J. Leskovec and J. J. McAuley. Learning to discover social circles in ego networks. In *Advances in neural information processing systems*, pages 539–547, 2012.
- [39] D. Liben-Nowell and J. Kleinberg. The link-prediction problem for social networks. *Journal of the Association for Information Science and Technology*, 58(7):1019–1031, 2007.
- [40] U. V. Luxburg, A. Radl, and M. Hein. Getting lost in space: Large sample analysis of the resistance distance. In *Advances in Neural Information Processing Systems*, pages 2622–2630, 2010.

- [41] C. Mathieu and H. Zhou. Graph reconstruction via distance oracles. In *International Colloquium on Automata, Languages, and Programming*, pages 733–744. Springer, 2013.
- [42] L. Page, S. Brin, R. Motwani, and T. Winograd. The PageRank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab, 1999.
- [43] B. Perozzi, R. Al-Rfou, and S. Skiena. DeepWalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 701–710. ACM, 2014.
- [44] M. J. Rattigan and D. Jensen. The case for anomalous link discovery. *ACM SIGKDD Explorations Newsletter*, 7(2):41–47, 2005.
- [45] L. Reyzin and N. Srivastava. On the longest path algorithm for reconstructing trees from distance matrices. *Information Processing Letters*, 101(3):98–100, 2007.
- [46] M. Saerens, F. Fouss, L. Yen, and P. Dupont. The principal components analysis of a graph, and its relationships to spectral clustering. In *ECML*, volume 3201, pages 371–383. Springer, 2004.
- [47] P. Sarkar and A. Moore. A tractable approach to finding closest truncated-commute-time neighbors in large graphs. In *Proceedings of the Twenty-Third Conference Annual Conference on Uncertainty in Artificial Intelligence (UAI-07)*, pages 335–343, 2007.
- [48] D. A. Spielman. Trees and distances. University Lecture, 2012.
- [49] D. A. Spielman and N. Srivastava. Graph sparsification by effective resistances. *SIAM Journal on Computing*, 40(6):1913–1926, 2011.
- [50] E. A. Stone and A. R. Griffing. On the Fiedler vectors of graphs that arise from trees by Schur complementation of the Laplacian. *Linear Algebra and its Applications*, 431(10):1869 – 1880, 2009.
- [51] J. Sun, S. Boyd, L. Xiao, and P. Diaconis. The fastest mixing Markov process on a graph and a connection to a maximum variance unfolding problem. *SIAM review*, 48(4):681–699, 2006.
- [52] H. Tong, C. Faloutsos, and J.-Y. Pan. Fast random walk with restart and its applications. 2006.
- [53] C. E. Tsourakakis, M. Mitzenmacher, J. Błasiok, B. Lawson, P. Nakkiran, and V. Nakos. Predicting positive and negative links with noisy queries: Theory & practice. *arXiv preprint arXiv:1709.07308*, 2017.
- [54] U. Von Luxburg, A. Radl, and M. Hein. Hitting and commute times in large graphs are often misleading. *arXiv preprint arXiv:1003.1266*, 2010.
- [55] U. Von Luxburg, A. Radl, and M. Hein. Hitting and commute times in large random neighborhood graphs. *Journal of Machine Learning Research*, 15(1):1751–1798, 2014.
- [56] D. M. Wittmann, D. Schmidl, F. Blochl, and F. J. Theis. Reconstruction of graphs based on random walks. *Theoretical Computer Science*, 410(38):3826 – 3838, 2009.

- [57] L. Yen, F. Fouss, C. Decaestecker, P. Francq, and M. Saerens. Graph nodes clustering based on the commute-time kernel. *Advances in Knowledge Discovery and Data Mining*, pages 1037–1045, 2007.
- [58] E. Zheleva, E. Terzi, and L. Getoor. Privacy in social networks. *Synthesis Lectures on Data Mining and Knowledge Discovery*, 3(1):1–85, 2012.