

Factorized Graph Representations for Semi-Supervised Learning from Sparse Data

Krishna Kumar P.
IIT Madras

Paul Langton
Northeastern University

Wolfgang Gatterbauer
Northeastern University

ABSTRACT

Node classification is an important problem in graph data management. It is commonly solved by various *label propagation* methods that iteratively pass messages along edges, starting from a few labeled seed nodes. For graphs with arbitrary compatibilities between classes, these methods crucially depend on knowing the *compatibility matrix*, which must thus be provided by either domain experts or heuristics. We instead suggest a principled and scalable method for directly estimating the compatibilities from a sparsely labeled graph. This method, which we call *distant compatibility estimation*, works even on extremely sparsely labeled graphs (e.g., 1 in 10,000 nodes is labeled) in a fraction of the time it later takes to label the remaining nodes. Our approach first creates multiple *factorized graph representations* (with size independent of the graph) and then performs estimation on these smaller graph sketches. We refer to *algebraic amplification* as the underlying idea of leveraging algebraic properties of an algorithm's update equations to amplify sparse signals in data. We show that our estimator is by orders of magnitude faster than alternative approaches and that the end-to-end classification accuracy is comparable to using gold standard compatibilities. This makes it a cheap pre-processing step for any existing label propagation method and removes the current dependence on heuristics.

ACM Reference Format:

Krishna Kumar P., Paul Langton, and Wolfgang Gatterbauer. 2020. Factorized Graph Representations for Semi-Supervised Learning from Sparse Data. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data (SIGMOD'20)*, June 14–19, 2020, Portland, OR, USA. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3318464.3380577>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. SIGMOD'20, June 14–19, 2020, Portland, OR, USA

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6735-6/20/06...\$15.00

<https://doi.org/10.1145/3318464.3380577>

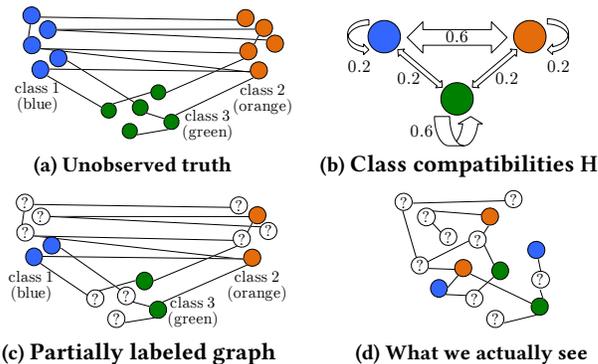


Figure 1: (a, b): Graphs are formed based on relative compatibilities between classes of nodes. (c, d): We have access to only a few labels $n_l \ll n$ and want to classify the remaining nodes *without* knowing the compatibilities between classes.

1 INTRODUCTION

Node classification (or label prediction) [7] is an important component of graph data management. In a broadly applicable scenario, we are given a large graph with edges that reflect affinities between their adjoining nodes and a small fraction of labeled nodes. Most graph-based semi-supervised learning (SSL) methods attempt to infer the labels of the remaining nodes by assuming similarity of neighboring labels. For example, people with similar political affiliations are more likely to follow each other on social networks. This problem is well-studied, and solutions are often variations of random walks that are fast and sufficiently accurate.

However, at other times opposites attract or complement each other (also called *heterophily* or disassortative mixing) [28]. For example, predators might form a functional group in a biological food web, not because they interact with each other, but because they eat similar prey [42], groups of proteins that serve a certain purpose often don't interact with each other but rather with complementary protein [9], and in some social networks pairs of nodes are more likely connected if they are from different classes (e.g., members on the social network studied by [57] being more likely to interact with the opposite gender than the same one).

In more complicated scenarios, such as online auction fraud, fraudsters are more likely linked to accomplices, and we have a mix of homophily and heterophily between multiple classes of nodes [48].

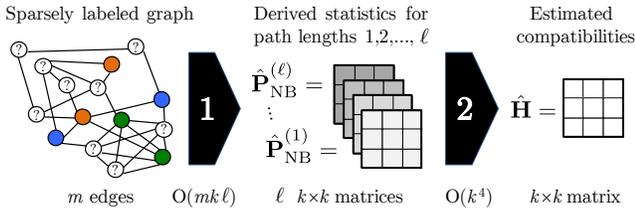


Figure 2: Our approach for compatibility estimation proceeds in two steps: (1) an efficient graph summarization that creates sketches in linear time of the number of edges m and classes k , see Section 4.6; and (2) an optimization step which is independent of the size of the graph, see Section 4.4.

Example 1.1 (Email). Consider a corporate email network with three different classes of users. Class 1, the marketing people, often email class 2, the engineers (and v.v.), whereas users of class 3, the C-Level Executives, tend to email amongst themselves (Fig. 1b). Assume we are given the labels (classes) of very few nodes (Fig. 1c). How can we infer the labels of the remaining nodes?

For these scenarios, standard random walks do not work as they cannot capture such arbitrary compatibilities. Early works addressing this problem propose belief propagation (BP) for labeling graphs, since BP can express arbitrary compatibilities between labels. However, the update equations of BP are more complicated than standard label propagation algorithms. They have well-known convergence problems [44, Sec. 22], and are difficult to use in practice [53]. A number of recent papers found ways to circumvent the convergence problems of BP by linearizing the update equations [13, 15, 17, 18, 29, 31], and thus transforming the update equations of BP into an efficient matrix formulation. The resulting updates are similar to random walks but propagate messages “modulated” with *relative class compatibilities*.

A big challenge for deploying this family of algorithms is knowing the appropriate compatibility matrix H , where each entry H_{ij} captures the relative affinity between neighboring nodes of labels i and j . Finding appropriate compatibilities was identified as a challenging open problem [38], and the current state of the art is to have them given by domain experts or by ad-hoc and rarely justified heuristics.

Our contribution. We propose an approach that *does not need any prior domain knowledge of compatibilities*. Instead, we estimate the compatibilities on the *same* graph for which we later infer the labels of unlabeled nodes (Fig. 1d). We achieve this by deriving an estimation method that (i) can handle extreme label scarcity, (ii) is orders of magnitude faster than textbook estimation methods, and (iii) results in labeling accuracy that is nearly indistinguishable from the actual gold standard (GS) compatibilities. In other words, we suggest an end-to-end solution for a difficult *within-network classification*, where compatibilities are not given to us.

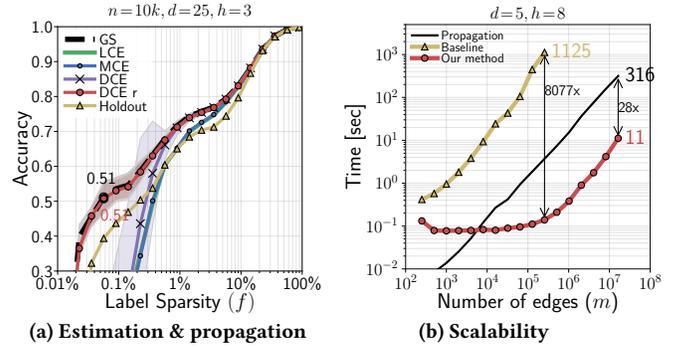


Figure 3: (a): Our methods infer labels with similar accuracy as if we were given the gold standard compatibilities (GS): e.g., labeling accuracy of 0.51 in a graph with 10k nodes and only 8 labeled nodes with our best method *distance compatibility estimation with restarts* (DCEr) in red as compared to the same accuracy with GS. (b): The additional step of estimating compatibilities is fast: DCEr learns the compatibilities on a graph with 16.4m edges in 11 sec, which is 28 times faster than node labeling (316 sec) and 3-4 orders of magnitude faster than a baseline holdout method.

PROBLEM 1.2 (AUTOMATIC NODE CLASSIFICATION). Given an undirected graph $G(V, E)$ with a set of labeled nodes $V_\ell \subset V$ from k classes and unknown compatibilities between classes. Classify the remaining nodes, $v \in V \setminus V_\ell$.

Summary of approach. We develop a novel, consistent, and scalable graph summarization that allows us to split compatibility estimation into two steps (Fig. 2): (1) First calculate the number of paths of various lengths ℓ between nodes for all pairs of classes. While the number of paths is exponential in the path’s length, we develop efficient factorization and sparse linear algebra methods that calculate them in time linear of the graph size and path length.¹ (2) Second use a combination of these compact graph statistics to estimate H . We derive an explicit formula for the gradient of the loss function that allows us to find the global optimum quickly. Importantly, this second optimization step takes time *independent of the graph size* (!). In other words, we *reduce compatibility estimation over a sparsely labeled graph into an optimization problem over a set of small factorized graph representations* with an explicit gradient. Our approach has only one, relatively insensitive, hyperparameter.

Our approach is orders of magnitude faster than common parameter estimation methods that rely on log-likelihood estimations and variants of expectation maximization. For example, recent work [42] develops methods that can learn compatibilities on graphs with hundreds of nodes in minutes time. In contrast, we learn compatibilities in graphs with 16.4 million edges in 11 sec using an off-the-shelf optimizer

¹Example 4.6 will illustrate evaluating 10^{14} such paths in less than 0.1 sec.

and running on a single CPU (see Fig. 3b). In a graph with $10k$ nodes and only 8 labeled nodes, we estimate \mathbf{H} such that the subsequent labeling has equivalent accuracy (0.51) to a labeling using the actual compatibilities (GS in Fig. 3a). We are not aware of any reasonably fast approach that can learn the compatibilities from the sparsely labeled graph. All recent work in the area uses simple heuristics to specify the compatibilities: e.g., [15, 17, 18, 29].

Outline. We start by giving a precise meaning to *compatibility matrices* by showing that prior label propagation methods based on *linearized belief propagation* essentially propagate *frequency distributions* of labels between neighbors (Section 3.1) and deriving the corresponding energy minimization framework (Section 3.2). Based on this formulation, we derive two convex optimization methods for parameter estimation (Section 4): *linear compatibility estimation* (LCE) and *myopic compatibility estimation* (MCE). We then develop a novel consistent estimator which counts ℓ -distance non-backtracking paths: *distant compatibility estimation* (DCE). Its objective function is not convex anymore, but well-behaved enough so we can find the global optimum in practice with a few repeated restarts, an approach we call *DCE with restarts* (DCER). Section 5 gives an extensive comparative study on synthetic and real-world data. Proofs and more experimental results are available in our extended version on arXiv [30].

2 FORMAL SETUP AND RELATED WORK

We first define essential concepts and review related work on semi-supervised node labeling. We denote vectors (\mathbf{x}) and matrices (\mathbf{X}) in bold. We use row-wise ($\mathbf{X}_{i\cdot}$), column-wise ($\mathbf{X}_{\cdot j}$), and element-wise (X_{ij}) matrix indexing, e.g., \mathbf{X}_i is the i -th row vector of \mathbf{X} (and thus bold), whereas X_{ij} is a single number (and thus not bold).

2.1 Semi-Supervised Learning (SSL)

Traditional graph-based Semi-Supervised Learning (SSL) predict the labels of unlabeled nodes under the assumption of *homophily* or *smoothness*. Intuitively, a label distribution is “smooth” if a label “ x ” on a node makes the same label on a neighboring node more likely, i.e. nodes of the same class tend to link to each other. The various methods differ mainly in their definitions of “smoothness” between classes of neighboring nodes [6, 35, 56, 63, 65, 67].²

Common to all approaches, we are given a graph $G = (V, E)$ with $n = |V|$, $m = |E|$, and real edge weights given by $w : E \rightarrow \mathbb{R}$. The weight $w(e)$ of an edge e indicates the

²Notice a possible naming ambiguity: “learning” in SSL stands for *classifying* unlabeled nodes (usually assuming homophily). In our setup, we first need to “learn” (or *estimate*) the compatibility parameters, before we can classify the remaining nodes with a variant of label propagation.

similarity of the incident nodes, and a missing edge corresponds to zero similarity. These weights are captured in the symmetric weighted adjacency matrix $\mathbf{W} \in \mathbb{R}^{n \times n}$ defined by $W_{ij} \triangleq w(e)$ if $e = (i, j) \in E$, and 0 otherwise. Each node is a member of exactly one of k classes which have increased edge incidence between members of the same class. Given a set of labeled nodes $V_L \subset V$ with labels in $[k]$, predict the labels of the remaining unlabeled nodes $V \setminus V_L$.

Most binary SSL algorithms [60, 64, 66] specify the existing labels by a vector $\mathbf{x} = [x_1, \dots, x_n]^T$ with $x_i \in L = \{+1, -1\}$ for $i \leq n_L$ and $x_i = 0$ for $n_L + 1 \leq i \leq n$. Then a real-valued “labeling function” assigns a value f_i with $1 \leq i \leq n$ to each data point i . The final classification is performed as $\text{sign}(f_i)$ for all unlabeled nodes. This binary approach can be extended to multi-class classification [60] by assigning a vector to each node. Each entry represents the belief that a node is in the corresponding class. Each of the classes is propagated *separately* and, at convergence, compared at each node with a “one-versus-all” approach [10]. SSL methods differ in how they compute f_i for each node i and commonly justify their formalism from a “regularization framework”; i.e., by motivating a different energy function and proving that the derived labeling function f is the solution to the objective of minimizing the energy function.

Contrast to our work. The labeling problem we are interested in this work is a generalization of standard SSL. In contrast to the commonly used smoothness assumption (i.e. labels of the same class tend to connect more often), we are interested in the more general scenario of arbitrary compatibilities between classes.

2.2 Belief Propagation (BP)

Belief Propagation (BP) [53] is a widely used method for reasoning in networked data. In contrast to typical semi-supervised label propagation, BP handles the case of arbitrary compatibilities. By using the symbol \odot for the component-wise multiplication and writing \mathbf{m}_{ji} for the k -dimensional “message” that node j sends to node i , the BP update equations [44, 61] can be written as:

$$\mathbf{f}_i \leftarrow Z_i^{-1} \mathbf{x}_i \odot \left(\bigodot_{j \in N(i)} \mathbf{m}_{ji} \right) \quad \mathbf{m}_{ij} \leftarrow \mathbf{H} \left(\mathbf{x}_i \odot \left(\bigodot_{v \in N(i) \setminus j} \mathbf{m}_{vi} \right) \right)$$

Here, Z_i is a normalizer that makes the elements of \mathbf{f}_i sum to 1, and each entry H_{ce} in \mathbf{H} is a proportional “compatibility” that indicates the relative influence of a node of class c on its neighbor of class e . Thus, an outgoing message from a node is computed by multiplying all incoming messages (except the one sent previously by the recipient) and then multiplying the outgoing message by the edge potential \mathbf{H} .

Unlike other SSL methods, BP has no simple linear algebra formulation and has well-known convergence problems. Despite extensive research on the convergence of BP [14, 41]

exact criteria for convergence are not known [44, Sec. 22] and practical use of BP is non-trivial [53].

Contrast to our work. Parameter estimation in graphical models quickly becomes intractable for even moderately-sized datasets [42]. We transform the original problem into a linear algebra formulation that allows us to leverage existing highly optimized tools and that can learn compatibilities often *faster* than the time needed to label the graph.

2.3 Linearized Belief Propagation

Recent work [18, 29] suggested to “linearize” BP and showed that the original update equations of BP can be reasonably approximated by linearized equations

$$\tilde{\mathbf{f}}_i \leftarrow \tilde{\mathbf{x}}_i + \frac{1}{k} \cdot \sum_{j \in N(i)} \tilde{\mathbf{m}}_{ji} \quad \tilde{\mathbf{m}}_{ij} \leftarrow \tilde{\mathbf{H}} \left(\tilde{\mathbf{f}}_i - \underbrace{\frac{1}{k} \tilde{\mathbf{m}}_{ji}}_{\text{EC}} \right)$$

by “centering” the belief vectors \mathbf{x} , \mathbf{f} and the potential matrix around $\frac{1}{k}$. If a vector \mathbf{x} is centered around c , then the residual vector around c is defined as $\tilde{\mathbf{x}} = [x_1 - c, x_2 - c, \dots]$ and centered around 0. This centering allowed the authors to rewrite BP in terms of the residuals. The “echo cancellation” (EC) term is a result of the condition “ $v \in N(i) \setminus j$ ” in the original BP equations.

While the EC term has a strong theoretical justification for BP and appears to have been kept for the correspondence between BP and LinBP, in our extensive simulations, we have not identified any parameter regime where including the EC term for propagation consistently gives better results. It rather slows down evaluation and complicates determining the convergence threshold (the top eigenvalue becomes negative slightly above the convergence threshold). We will thus explicitly ignore the EC term in the remainder of this paper. The update equations of LinBP then become:

$$\tilde{\mathbf{F}} \leftarrow \tilde{\mathbf{X}} + \mathbf{W}\tilde{\mathbf{H}} \quad (\text{LinBP}) \quad (1)$$

The advantage of LinBP over standard BP is that the linearized formulation allows provable convergence guarantees. The process was shown to converge iff the following condition holds on the spectral radii³ ρ of $\tilde{\mathbf{H}}$ and \mathbf{W} :

$$\rho(\tilde{\mathbf{H}}) < 1/\rho(\mathbf{W}) \quad (2)$$

Follow-up work [17] generalizes LinBP to the most general case of arbitrary pairwise Markov networks which include heterogeneous graphs with fixed number of node and edge types. Independently, ZooBP [15] follows a similar motivation, yet restricts itself to the mathematically less challenging special case of constant row-sum symmetric potentials.

Contrast to our work. Our work focuses on homogeneous graphs and makes a complementary contribution to

³The spectral radius of a matrix is the largest absolute value among its eigenvalues.

that of label propagation: that of learning compatibilities from a sparsely labeled graph *in a fraction of the time* it takes to propagate the labels (Section 4). This avoids the reliance on domain experts or heuristics and results in an end-to-end estimation and propagation method. An earlier version of the ideas in our paper was made available on arXiv as [16].

2.4 Iterative Classification Methods

Random walks with Restarts (RWR). Random walk-based methods make the assumption that the graph is homophilous; i.e., that instances belonging to the same class tend to link to each other or have higher edge weight between them [34]. In general, given a graph $G = (V, E)$, random walk algorithms return as output a ranking vector \mathbf{f} that results from iterating following equation until convergence:

$$\mathbf{f} \leftarrow \bar{\alpha} \mathbf{u} + \alpha \mathbf{W}^{\text{col}} \mathbf{f} \quad (3)$$

Here, \mathbf{u} is a normalized teleportation vector with $\|\mathbf{u}\| = |V|$ and $\|\mathbf{u}\|_1 = 1$, and \mathbf{W}^{col} is column-normalized. Notice that above Eq. (3) can be interpreted as the probability of a random walk on G arriving at node i , with teleportation probability $\bar{\alpha}$ at every step to a node with distribution \mathbf{u} [34]. Variants of this formulation are used by PageRank [46], Personalized PageRank [11, 24], Topic-sensitive PageRank [23], Random Walks with Restarts [47], and MultiRankWalk [34] which runs k random walks in parallel (one for each class c).

To compare it with our setting, MultiRankWalk [34] and other forms of random walks can be stated as special cases of the more general formulation: (1) For each class $c \in [k]$: (a) set $\mathbf{u}_i \leftarrow 1$ if node i is labeled c , (b) normalize \mathbf{u} s.t. $\|\mathbf{u}\|_1 = 1$. (2) Let \mathbf{U} be the $n \times k$ matrix with column i equal \mathbf{u}_i . (3) Then iterate until convergence:

$$\mathbf{F} \leftarrow \bar{\alpha} \mathbf{U} + \alpha \mathbf{W}^{\text{col}} \mathbf{F} \mathbf{I}_k$$

(4) After convergence, label each node i with the class c with maximum value: $c = \arg \max_j F_{ij}$.

Other Iterative Classification Methods. Goldberg et al. [20] consider a concept of similarity and dissimilarity between nodes. This method only applies to classification tasks with 2 labels and cannot generalize to arbitrary compatibilities. Bhagat et al. [8] look at commonalities across the direct neighbors of nodes in order to classify them. The paper calls this method leveraging “co-citation regularity” which is indeed equally expressive as heterophily. The experiments in that paper require at least 2% labeled data (Figure 6e in [8]), which is similar to the regimes up to which MCE works. Similarly, Peel [50] suggests an interesting method that skips compatibility matrices by propagating information across nodes with common neighbors. The method was tested on networks with 10% labeled nodes and it will be interesting to investigate its performance in the sparse label regime.

2.5 Recent Neural Network Approaches

Several recent papers propose neural network (NN) architectures for node labeling (e.g., [21, 27, 43]). In contrast to our work (and all other work discussed in this section), those NN-based approaches *require additional features* from the nodes. For example, in the case of Cora, [27] also has access to node content (i.e. which words co-occur in a paper). Having access to the actual text of a paper allows better classification than the network structure alone. As a result, [27] can learn and use a large number of parameters in their trained NN.

Contrast to our work. We classify the nodes based on *the graph structure alone*, without access to additional features. The result is that while [27] achieves an accuracy of 81.5% for 5.2% labeled nodes in Cora (see Section 5.1 and Section 6.1 of [27]), we still achieve 66% accuracy based on the network alone and only 21 estimated parameters.

2.6 Non-Backtracking Paths (NB)

Section 4.5 derives estimators for the powers of \mathbf{H} by counting labels over all “non-backtracking” (NB) paths in a partially labeled graph. We prove our estimator to be consistent and thus with negligible bias for increasing n . Prior work already points to the advantages of NB paths for various different graph-related problems, such as graph sampling [32], calculating eigenvector centrality [36], increasing the detectability threshold for community detection [31], improving estimation of graphlet statistics [12], or measuring the distance between graphs [58]. To make this work, all these papers replace the $n \times n$ adjacency matrix with a $2m \times 2m$ “Hashimoto matrix” [22] which represents the link structure of a graph in an augmented state space with $2m$ states (one state for each directed pair of nodes) and in the order of $O(m(d-1))$ non-zero entries, and then perform random walks. The only work we know that uses NB paths without Hashimoto is [2], which calculates the mixing rate of a NB random walk on a regular expanders (thus graphs with identical degree across all nodes). That work does not generalize to graphs with varying degree distribution and does not allow an efficient path summarization.

Contrast to our work. Our approach does not perform random walks, does not require an augmented state space (see Proposition 4.3), and still allows an efficient path summarization (see Proposition 4.5). To the best of our knowledge, ours is the first proposal to (i) estimate compatibilities from NB paths and (ii) propose an efficient calculation.

2.7 Distant Supervision

The idea of *distant supervision* is to adapt existing ground truth data from a related yet different task for providing *additional lower quality labels* (also called *weak labels*) to

sparsely labeled data [25, 39, 51]. The methods are thus also often referred to as *weak supervision*.

Contrast to our work. In our setting, we are given no other outside ground truth data nor heuristic rules to label more data. Instead, we leverage certain algebraic properties of an algorithm’s update equations to *amplify sparse signals in the available data*. We thus refer to the more general idea of our approach as *algebraic amplification*.

3 PROPERTIES OF LABEL PROPAGATION

This section makes novel observations about linearized versions of BP that help us later find efficient ways to learn the compatibility matrix \mathbf{H} from sparsely labeled graphs.

3.1 Propagating Frequency Distributions

Our first observation is that centering of prior beliefs \mathbf{X} and compatibility matrix \mathbf{H} in LinBP Eq. (1) is *not necessary* and that the final labels are identical whether we use $\tilde{\mathbf{X}}$ or \mathbf{X} , and $\tilde{\mathbf{H}}$ or \mathbf{H} . We state this result in a slightly more general form: Let $\mathbf{F} = \text{LinBP}(\mathbf{W}, \mathbf{X}, \mathbf{H}, \epsilon, r)$ stand for the label distribution after iterating the LinBP update equations r times, starting from \mathbf{X} and using scaling factor ϵ . Let $\mathbf{l} = \text{label}(\mathbf{F})$ stand for the operation of assigning each node the class with the maximum belief: $l_i = \arg \max_j F_{ij}$. Then:

THEOREM 3.1 (CENTERING IN LINBP IS UNNECESSARY). *Given constants c_1 and c_2 s.t. $\mathbf{H}_2 = \mathbf{H}_1 + c_1$ and $\mathbf{X}_2 = \mathbf{X}_1 + c_2$.⁴ Then, $\forall \mathbf{W}, \epsilon, r$: $\text{label}(\text{LinBP}(\mathbf{W}, \mathbf{X}_2, \mathbf{H}_2, \epsilon, r)) = \text{label}(\text{LinBP}(\mathbf{W}, \mathbf{X}_1, \mathbf{H}_1, \epsilon, r))$.*

Modulating beliefs of a node with \mathbf{H} instead of $\tilde{\mathbf{H}}$ allows a natural interpretation of label propagation as “propagating frequency distributions” and thus imposing an *expected frequency distribution* on the labels of neighbors of a node. This observation gives us an intuitive interpretation of our later derived approaches for learning \mathbf{H} from observed frequency distributions (Section 4.3). For the rest of this paper, we will thus replace Eq. (1) with the “uncentered” version:

$$\mathbf{F} \leftarrow \mathbf{X} + \mathbf{WFH} \quad (4)$$

A consequence is that compatibility propagation works identically whether the compatibility matrix \mathbf{H} is centered or kept as doubly-stochastic. In other words, if the relative frequencies by which different node classes connect to each other is known, then this matrix can be used *without centering* for compatibility propagation and will lead to identical results and thus node labels.

⁴We use here “broadcasting notation:” adding a number to a vector or matrix is a short notation for adding the number to each entry in the vector.

3.2 Labeling as Energy Minimization

Our next goal is to formulate the solution to the update equations of LinBP as the solution to an optimization problem; i.e., as an *energy minimization framework*. While LinBP was derived from probabilistic principles (as approximation of the update equations of belief propagation [18]), it is currently not known whether there is a simple objective function that a solution minimizes. Knowledge of such an objective is helpful as it allows principled extensions to the core algorithm. We will next give the objective function for LinBP and will use it later in Section 4 to solve the problem of parameter learning; i.e., estimating the compatibility matrix from a partially labeled graph.

PROPOSITION 3.2 (LINBP OBJECTIVE FUNCTION). *The energy function minimized by the LinBP update equations Eq. (1) is given by:*

$$E(\mathbf{F}) = \|\mathbf{F} - \mathbf{X} - \mathbf{WFH}\|^2 \quad (5)$$

4 COMPATIBILITY ESTIMATION

In this section we develop a scalable algorithm to learn compatibilities from partially labeled graph. We proceed step-by-step, starting from a baseline until we finally arrive at our suggested consistent and scalable method called “*Distant Compatibility Estimation with restarts*” (DCER).

The compatibility matrix we wish to estimate is a $k \times k$ -dimensional doubly stochastic matrix \mathbf{H} . Because any symmetric doubly-stochastic matrix has $k^* \triangleq \frac{k(k-1)}{2}$ degrees of freedom, we parameterize all k^2 entries as a function of k^* appropriately chosen parameters. In all following approaches, we parameterize \mathbf{H} as a function of the k^* entries of H_{ij} with $i \leq j, j \neq k$. We can calculate the remaining matrix entries from symmetry and stochasticity conditions as follows:

$$H_{ij} = \begin{cases} H_{ji}, & \text{if } i < j, j \neq k \\ 1 - \sum_{\ell=1}^{k-1} H_{i\ell}, & \text{if } i \neq k, j = k \\ 1 - \sum_{\ell=1}^{k-1} H_{\ell j}, & \text{if } i = k, j \neq k \\ 2 - k + \sum_{\ell,r < k} H_{\ell r}, & \text{if } i = j = k \end{cases} \quad (6)$$

For example, for $k = 3$, \mathbf{H} can be reconstructed from a $k^* = 3$ -dimensional vector $\mathbf{h} = [H_{11}, H_{21}, H_{22}]^\top$ as follows:

$$\mathbf{H}(\mathbf{h}) = \begin{bmatrix} H_{11} & H_{12} & 1-H_{11}-H_{12} \\ H_{21} & H_{22} & 1-H_{21}-H_{22} \\ 1-H_{11}-H_{21} & 1-H_{12}-H_{22} & H_{11}+2H_{21}+H_{22}-1 \end{bmatrix}$$

More generally, let $\mathbf{h} \in \mathbb{R}^{k^*}$ and define \mathbf{H} as function of the $k^* \triangleq \frac{k(k-1)}{2}$ entries of \mathbf{h} as follows:

$$\mathbf{H} = \begin{bmatrix} h_1 & \cdot & \cdot & \dots & \cdot & \cdot \\ h_2 & h_3 & \cdot & \dots & \cdot & \cdot \\ h_4 & h_5 & h_6 & \dots & \cdot & \cdot \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ h_{\dots} & h_{\dots} & h_{\dots} & \dots & h_{k^*} & \cdot \end{bmatrix}$$

The remaining matrix entries can be calculated from Eq. (6).

4.1 Baseline: Holdout Method

Our first approach for estimating \mathbf{H} is a variant of a standard textbook method [28, 40, 62] and serves as baseline against which we compare all later approaches: we split the labeled data into two sets and learn the compatibilities that fit best when propagating labels from one set to the other.

Formally, let \mathcal{Q} be a partition of the available labels into a Seed and a Holdout set. For a fixed partition \mathcal{Q} and given compatibility matrix \mathbf{H} , the “holdout method” runs label propagation Eq. (1) with Seed as seed labels and evaluates accuracy over Holdout. Denote $\text{Acc}_{\mathcal{Q}}(\mathbf{H})$ the resulting accuracy. Its goal is then to find the matrix \mathbf{H} that maximizes the accuracy. In other words, the energy function that holdout minimizes is the negative accuracy:

$$E(\mathbf{H}) = -\text{Acc}_{\mathcal{Q}}(\mathbf{H})$$

The optimization itself is then a search over the parameter space given by the k^* free parameters of \mathbf{H} :

$$\hat{\mathbf{H}} = \arg \min_{\mathbf{H}} E(\mathbf{H}), \text{ s.t. Eq. (6)}$$

The result may depend on the choice of partition \mathcal{Q} . We could thus use b different partitions $\mathcal{Q}_i, i \in [b]$: For a fixed \mathbf{H} we run label propagation b times, each starting from a different Seed $_i$, and each evaluated over their corresponding test set Holdout $_i$. The energy function to minimize is then the negative compound accuracy:

$$E(\mathbf{H}) = - \sum_i \text{Acc}_{\mathcal{Q}_i}(\mathbf{H}) \quad (\text{Holdout}) \quad (7)$$

We suggest this method as reasonable baseline as it mimics parameter estimation methods in probabilistic graphical models that optimize over a parameter space by *using multiple executions of inference as a subroutine* [28]. Similarly, our holdout method maximizes the accuracy by using inference as a “black box” subroutine. The downside of the holdout method is that each step in this iterative algorithm performs inference over the whole graph which makes parameter estimation considerably more expensive than inference (label propagation). The number of splits b has an obvious trade-off: higher b smoothens the energy function and avoids overfitting to one partition, but increases runtime.

In the following sections, we introduce novel path summarizations that avoid running estimation over the whole graph. Instead we use a few concise graph summaries of size $O(k^2)$, independent of the graph size. In other words, the expensive iterative estimation steps can now be performed on a reduced size summary of the partially labeled graph. This conceptually simple idea allows us to perform *estimation faster than inference* (recall Fig. 3b).

4.2 Linear Compatibility Estimation (LCE)

We obtain our first novel approach from energy minimization objective of LinBP in [Proposition 3.2](#):

$$E(\mathbf{F}) = \|\mathbf{F} - \mathbf{X} - \mathbf{WFH}\|^2$$

Note that for an unlabeled node i , the final label distribution is the weighted average of its neighbors: $\mathbf{F}_i = (\mathbf{WFH})_i$. To see this, consider a single row for a node i :

$$\|(\mathbf{F} - \mathbf{X} - \mathbf{WFH})_i\|^2$$

If i is unlabeled then its corresponding entries in \mathbf{X}_i are 0, and the minimization objective is equivalent to

$$\|(\mathbf{F} - \mathbf{WFH})_i\|^2$$

which leads to $\mathbf{F}_i = (\mathbf{WFH})_i$ for an unlabeled node. Next notice that if we knew \mathbf{F} and ignored the few explicit labels, then \mathbf{H} could be learned from minimizing

$$E(\mathbf{H}) = \|\mathbf{F} - \mathbf{WFH}\|^2$$

In our case, we only have few labels in the form of \mathbf{X} instead of \mathbf{F} . Our first novel proposal for learning the compatibility matrix \mathbf{H} is to thus use the available labels \mathbf{X} and to minimize the following energy function:

$$E(\mathbf{H}) = \|\mathbf{X} - \mathbf{WXH}\|^2 \quad (\text{LCE}) \quad (8)$$

Notice that [Eq. \(8\)](#) defines a *convex* optimization problem. Thus any standard optimizer can solve it in considerably faster time than the Holdout method and it is no longer necessary to use inference as subroutine. We call this approach “linear compatibility estimation” as the optimization criterion stems directly from the optimization objective of linearized belief propagation.

4.3 Myopic Compatibility Estimation: MCE

We next introduce a powerful yet simple idea that allows our next approaches to truly scale: we first (1) summarize the partially labeled graph into a small summary, and then (2) use this summary to perform the optimization. This idea was motivated by the observation that [Eq. \(8\)](#) requires an iterative gradient descent algorithm and has to multiply large adjacency matrix \mathbf{W} in each iteration. We try to derive an approach that can “factor out” this calculation into small but sufficient *factorized graph representation*, which can then be repeatedly used during optimization.

Our first method is called *myopic compatibility estimation* (MCE). It is “myopic” in the sense that it tries to summarize *the relative frequencies of classes between observed neighbors*. We describe below the three variants to transform this summary into a symmetric, doubly-stochastic matrix.

Consider a partially labeled $n \times k$ -matrix \mathbf{X} with $X_{ic} = 1$. If node i has label c (recall that unlabeled nodes have a corresponding null row vector in \mathbf{X}), then the $n \times k$ -matrix $\mathbf{N} \triangleq \mathbf{WX}$ has entries N_{ic} representing the number of labeled

neighbors of node i with label c . Furthermore, the $k \times k$ -matrix $\mathbf{M} \triangleq \mathbf{X}^T \mathbf{N} = \mathbf{X}^T \mathbf{W} \mathbf{X}$ has entries M_{cd} representing the number of nodes with label c that are neighbors of nodes with label d . This symmetric matrix represents the observed number of labels among labeled nodes. Intuitively, we are trying to find a compatibility matrix which is “similar” to \mathbf{M} . We normalize \mathbf{M} into an observed *neighbor statistics* matrix $\hat{\mathbf{P}}$ and then find the closest doubly-stochastic matrix \mathbf{H} :

We consider three variants for normalizing \mathbf{M} . The first one appears most natural (creating a stochastic matrix representing label frequency distributions between neighbors, then finding the closest doubly-stochastic matrix). We conceived of two other approaches, just to see if the choice of normalization has an impact on the final labeling accuracy.

- (1) Make \mathbf{M} row-stochastic by dividing each row by its sum. The vector of row-sums can be expressed in matrix notation as $\mathbf{M}\mathbf{1}$. We thus define the first variant of neighbor statistics matrix as:

$$\hat{\mathbf{P}} = \mathbf{M}^{\text{row}} \triangleq \text{diag}(\mathbf{M}\mathbf{1})^{-1} \mathbf{M} \quad (\text{Variant 1}) \quad (9)$$

Note, we use \mathbf{M}^{row} as short notation for row-normalizing the matrix \mathbf{M} . For each class c , the entry \hat{P}_{ce} gives the relative frequency of a node being connected to class e . While the matrix is row-stochastic, it is not yet doubly-stochastic.

- (2) The second variant uses the symmetric normalization method LGC [64] from [Section 2](#):

$$\hat{\mathbf{P}} = \text{diag}(\mathbf{M}\mathbf{1})^{-\frac{1}{2}} \mathbf{M} \text{diag}(\mathbf{M}\mathbf{1})^{-\frac{1}{2}} \quad (\text{Variant 2}) \quad (10)$$

The resulting $\hat{\mathbf{P}}$ is symmetric but not stochastic.

- (3) The third variant scaled \mathbf{M} s.t. the average matrix entry is $\frac{1}{k}$. This divisor is the sum of all entries divided by k (in vector notation written as $\mathbf{1}^T \mathbf{M} \mathbf{1}$):

$$\hat{\mathbf{P}} = k(\mathbf{1}^T \mathbf{M} \mathbf{1})^{-1} \mathbf{M} \quad (\text{Variant 3}) \quad (11)$$

This scaled matrix is neither symmetric nor stochastic.

We then find the “closest” symmetric, doubly stochastic matrix \mathbf{H} (i.e., it fulfills the $k^* \triangleq \frac{k(k-1)}{2}$ conditions implied by symmetry $\mathbf{H} = \mathbf{H}^T$ and stochasticity $\mathbf{H}\mathbf{1} = \mathbf{1}$). We use the Frobenius norm because of its favorable computational properties and thus minimize the following function:

$$E(\mathbf{H}) = \|\mathbf{H} - \hat{\mathbf{P}}\|^2 \quad (\text{MCE}) \quad (12)$$

Notice that all three normalization variants above have an alternative, simple justification: on a fully labeled graph, each variant will learn the same compatibility matrix; i.e., the matrix that captures the relative label frequencies between neighbors in a graph. On a graph with sampled nodes, however, \mathbf{M} will not necessarily be constant row-sum anymore. The three normalizations and the subsequent optimization are alternative approaches for finding a “smoothened” matrix \mathbf{H} that is close to the observations. Our experiments

have shown that the “most natural” normalization variant 1 consistently performs best among the three methods. Unless otherwise stated, we thus imply using variant 1.

Notice that MCE and all following approaches estimate \mathbf{H} *without performing label propagation*; and only because we avoid propagation, our method turns out to be faster than label propagation on large graphs and moderate k .

4.4 Distant Compatibility Estimation: DCE

While MCE addresses the scalability issue, it still requires a sufficient number of neighbors that are both labeled. For very small fractions f of labeled nodes, this may not be enough. Our next method, “*distant compatibility estimation*” (DCE), takes into account longer distances between labeled nodes.

In a graph with m edges and a small fraction f of labeled nodes, the number of neighbors that are *both* labeled can be quite small ($\sim mf^2$). Yet the number of “distance-2-neighbors” (i.e., nodes which are connected via a path of length 2) is higher in proportion to the average node degree d ($\sim dmf^2$). Similarly for distance- ℓ -neighbors ($\sim d^{\ell-1}mf^2$). As information travels via a path of length ℓ , it gets modulated ℓ times; i.e., via a power of the compatibility matrix: \mathbf{H}^ℓ .⁵ We propose to use powers of the matrix \mathbf{H} to be estimated by comparing them against an “*observed length- ℓ statistics matrix*.”

Powers of the adjacency matrix \mathbf{W}^ℓ with entries W_{ij}^ℓ count the number of paths of length ℓ between any nodes i and j . Extending the ideas in Section 4.3, let $\mathbf{N}^{(\ell)} \triangleq \mathbf{W}^\ell \mathbf{X}$ and $\mathbf{M}^{(\ell)} \triangleq \mathbf{X}^\top \mathbf{N}^{(\ell)} = \mathbf{X}^\top \mathbf{W}^\ell \mathbf{X}$. Then entries $M_{ce}^{(\ell)}$ represent the number of labeled nodes of class c that are connected to nodes of class e by an ℓ -distance path. Normalize this matrix (in any of the previous 3 variants) to get the observed length- ℓ statistics matrix $\hat{\mathbf{P}}^{(\ell)}$. Calculate these length- ℓ statistics for several path lengths ℓ , and then find the compatibility matrix that best fits these multiple statistics.

Concretely, minimize a “distance-smoothed” energy

$$E(\mathbf{H}) = \sum_{\ell=1}^{\ell_{\max}} w_\ell \|\mathbf{H}^\ell - \hat{\mathbf{P}}^{(\ell)}\|^2 \quad (\text{DCE}) \quad (13)$$

where ℓ_{\max} is the maximal distance considered, and the weights w_ℓ balance having more (but weaker) data points for bigger ℓ the more reliable (but sparser) signal from smaller ℓ .

To parameterize the weight vector \mathbf{w} , we use a “scaling factor” λ defined by $w_{\ell+1} = \lambda w_\ell$. For example, a distance-3 weight vector is then $[1, \lambda, \lambda^2]^\top$. The intuition is that in a typical graph, the fraction of number of paths of length ℓ to the number of paths of length $\ell - 1$ is *largely independent* of ℓ (but proportional to the average degree). Thus, λ determines the relative weight of paths of one more hop. As consequence our framework has only one single hyperparameter λ .

⁵Notice that this is strictly correct only in graphs with balanced labels. Our experiments verify the quality of estimation also on unbalanced graphs.

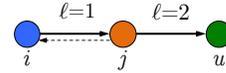


Figure 4: Illustration for non-backtracking paths

In our experiments (Section 5), we initialize the optimization with a k^* -dimensional vector with all entries equal to $\frac{1}{k}$ and discuss our choice of ℓ_{\max} and λ .

4.5 Non-Backtracking Paths (NB)

In our previous approach of learning from more distant neighbors, we made a slight but consistent mistake. We illustrate this mistake with Fig. 4. Consider the blue node i which has one orange neighbor j , which has two neighbors, one of which is green node u . Then the blue node i has one distance-2 neighbor u that is green. However, our previous approach will consider all length-2 paths, one of which leads back to node i . Thus, the row entry for node i in \mathbf{N} is $\mathbf{N}_{i \cdot}^{(2)} = [1, 0, 1]$ (assuming blue, orange, and green represent classes 1, 2, and 3, respectively). In other words, $\mathbf{M}^{(2)}$ will consistently *overestimate the diagonal entries*.

To address this issue, we consider only *non-backtracking paths* (NB) in the powers of the adjacency matrix. A NB path on an undirected graph G is a path which does not traverse the same edge twice in a row. In other words, a path $(u_1, u_2, \dots, u_{\ell+1})$ of length ℓ is non-backtracking iff $\forall j \leq \ell - 1 : u_j \neq u_{j+2}$. In our notation, we replace \mathbf{W}^ℓ with $\mathbf{W}_{\text{NB}}^{(\ell)}$. For example, $\mathbf{W}_{\text{NB}}^{(2)} = \mathbf{W}^2 - \mathbf{D}$ (a node i with degree d_i has $D_{ii} = d_i$ as diagonal entry in \mathbf{D}). A more general calculation of $\mathbf{W}_{\text{NB}}^{(\ell)}$ for any length ℓ is presented in Section 4.6. We now calculate new graph statistics $\hat{\mathbf{P}}_{\text{NB}}^{(\ell)}$ from $\mathbf{M}_{\text{NB}}^{(\ell)} \triangleq \mathbf{X}^\top \mathbf{W}_{\text{NB}}^{(\ell)} \mathbf{X}$ instead of $\mathbf{M}^{(\ell)}$, and replace $\hat{\mathbf{P}}^{(\ell)}$ with $\hat{\mathbf{P}}_{\text{NB}}^{(\ell)}$ in Eq. (13):

$$E(\mathbf{H}) = \sum_{\ell=1}^{\ell_{\max}} w_\ell \|\mathbf{H}^\ell - \hat{\mathbf{P}}_{\text{NB}}^{(\ell)}\|^2 \quad (\text{DCE NB}) \quad (14)$$

We next show that—assuming a label-balanced graph—this change gives us a *consistent estimator* with bias in the order of $O(1/m)$, in contrast to the prior bias in the order of $O(1/d)$:

THEOREM 4.1 (CONSISTENCY OF STATISTICS $\hat{\mathbf{P}}_{\text{NB}}^{(\ell)}$). *Under mild assumptions for the degree distributions, $\hat{\mathbf{P}}_{\text{NB}}^{(\ell)}$ is a consistent estimator for \mathbf{H}^ℓ , whereas $\hat{\mathbf{P}}^{(\ell)}$ is not:*

$$\lim_{n \rightarrow \infty} \hat{\mathbf{P}}_{\text{NB}}^{(\ell)} = \mathbf{H}^\ell \quad \text{whereas,} \quad \lim_{n \rightarrow \infty} \hat{\mathbf{P}}^{(\ell)} \neq \mathbf{H}^\ell$$

Example 4.2 (Non-backtracking paths). Consider the compatibility matrix $\mathbf{H} = \begin{bmatrix} 0.2 & 0.6 & 0.2 \\ 0.6 & 0.2 & 0.2 \\ 0.2 & 0.2 & 0.6 \end{bmatrix}$. Then $\mathbf{H}^2 = \begin{bmatrix} 0.44 & 0.28 & 0.28 \\ 0.28 & 0.44 & 0.28 \\ 0.28 & 0.28 & 0.44 \end{bmatrix}$, and the maximum entry (permuting between first and second position in the first row) follows the series

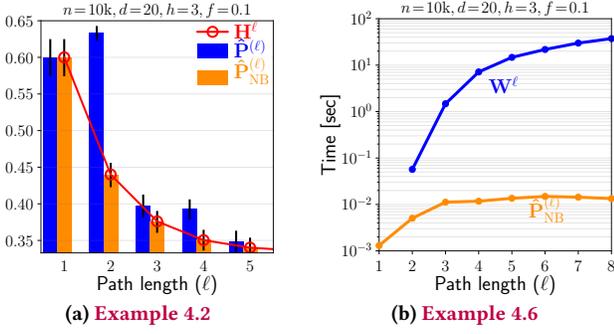


Figure 5: (a): **Example 4.2:** $\hat{P}_{NB}^{(\ell)}$ uses non-backtracking paths only and is a consistent estimator, in contrast to $\hat{P}^{(\ell)}$. (b) **Example 4.6:** Calculating W^ℓ for increasing ℓ is costly, while our factorized calculation of $\hat{P}_{NB}^{(\ell)}$ avoids evaluating W^ℓ explicitly and thus scales to arbitrary path lengths ℓ .

0.6, 0.44, 0.376, 0.3504, ... for increasing ℓ (shown as continuous green line H^ℓ in Fig. 5a). We create synthetic graphs with $n = 10k$ nodes, average node degree $d = 20$, uniform degree distribution, and compatibility matrix H . We remove the labels from $1 - f = 90\%$ nodes, then calculate the top entry in both $\hat{P}^{(\ell)}$ and $\hat{P}_{NB}^{(\ell)}$. The two bars in Fig. 5a show the mean and standard deviation of the corresponding matrix entries, illustrating that the approach based on non-backtracking paths leads to an unbiased estimator (height of orange bars are identical to the red circles), in contrast to the full paths (blue bars are higher than the red circles). \square

4.6 Scalable, Factorized Path Summation

Calculating longer NB paths is more involved. For example: $W_{NB}^{(3)} = W^3 - (DW + WD - W)$. However, we can calculate them recursively as follows:

PROPOSITION 4.3 (NON-BACKTRACKING PATHS). Let $W_{NB}^{(\ell)}$ be the matrix with $W_{NB}^{(\ell)}_{ij}$ being the number of non-backtracking paths of length ℓ from node i to j . Then $W_{NB}^{(\ell)}$ for $\ell \geq 3$ can be calculated via following recurrence relation:

$$W_{NB}^{(\ell)} = WW_{NB}^{(\ell-1)} - (D - I)W_{NB}^{(\ell-2)} \quad (15)$$

with starting values $W_{NB}^{(1)} = W$ and $W_{NB}^{(2)} = W^2 - D$. \square

Calculating $\hat{P}_{NB}^{(\ell)}$ requires multiple matrix multiplications. While matrix multiplication is associative, the order in which we perform the multiplications considerably affects the time to evaluate a product. A straight-forward evaluation strategy quickly becomes infeasible for increasing ℓ .

We illustrate with $M^{(3)}$: a default strategy is to first calculate $W^{(3)} = W(WW)$ and then $M^{(3)} = X^T(W^{(3)}X)$. The problem is that the intermediate result $W^{(\ell)}$ becomes dense. Concretely, if W is sparse with m entries and average node

degree d , then W^2 has in the order of d more entries ($\sim dm$), and W^ℓ exponential more entries ($\sim d^{\ell-1}m$). Thus intuitively, we like to choose the evaluation order so that intermediate results are as sparse as possible.⁶ The ideal way to calculate the expressions is to keep $n \times k$ intermediate matrices as in $M^{(3)} = X^T(W(WX))$.

Our solution is thus to re-structure the calculation in a way that minimizes the result sizes of intermediate results and caches results used across estimators with different ℓ . The reason of the scalability of our approach is that we can calculate all ℓ_{\max} graph summaries very efficiently.

ALGORITHM 4.4 (FACTORIZED PATH SUMMATION). Iteratively calculate the graph summaries $\hat{P}_{NB}^{(\ell)}$, for $\ell \in [\ell_{\max}]$ as follows:

- (1) Starting from $N_{NB}^{(1)} = WX$ and $N_{NB}^{(2)} = WN_{NB}^{(1)} - DX$, iteratively calculate $N_{NB}^{(\ell)} = WN_{NB}^{(\ell-1)} - (D - I)N_{NB}^{(\ell-2)}$.
- (2) Calculate $M_{NB}^{(\ell)} = X^T N_{NB}^{(\ell)}$.
- (3) Calculate $\hat{P}_{NB}^{(\ell)}$ from normalizing $M^{(\ell)}$ with Eq. (9).

PROPOSITION 4.5 (FACTORIZED PATH SUMMATION). Algorithm 4.4 calculates all $\hat{P}_{NB}^{(\ell)}$ for $\ell \in [\ell_{\max}]$ in $O(mk\ell_{\max})$.

Example 4.6 (Factorized path summation). Using the setup from Example 4.2, Fig. 5b shows the times for evaluating W^ℓ against our more efficient evaluation strategy for $\hat{P}_{NB}^{(\ell)}$. Notice the three orders of magnitude speed-up for $\ell = 5$. Also notice that $\hat{P}_{NB}^{(8)}$ summarizes statistics over more than 10^{14} paths in a graph with 100k edges in less than 0.02 sec.

4.7 Gradient-based Optimization

Our objective to find a symmetric, doubly stochastic matrix that minimizes Eq. (14) can be represented as

$$\hat{H} = \underset{H}{\operatorname{argmin}} E(H) \text{ s.t. } H\mathbf{1} = \mathbf{1}, H^T = H \quad (16)$$

For $\ell_{\max} > 1$, the function is non-convex and unlikely to have a closed-form solution. We thus minimize the function with gradient descent. However, we would require to calculate the gradient with regard to the free parameters.

PROPOSITION 4.7 (GRADIENT). The gradient for Eq. (16) and energy function Eq. (14) with regard to the free parameters

⁶This is well known in linear algebra and is analogous to query optimization in relational algebra: The two query plans $\pi_y(R(x) \bowtie S(x, y))$ and $R(x) \bowtie (\pi_y S(x, y))$ return the same values, but the latter can be considerably faster. Similarly, the ‘‘evaluation plans’’ $(WW)X$ and $W(WX)$ are algebraically equivalent, but the latter can be considerably faster for $n \gg k$. Efficient factorized representations are also the focus of factorized databases [45].

$H_{ij}, i \leq j, j \neq k$ is the dot product SG calculated from

$$\mathbf{G} = 2 \sum_{\ell=1}^{\ell_{\max}} w_{\ell} \left(\ell \mathbf{H}^{2\ell-1} - \sum_{r=0}^{\ell-1} \mathbf{H}^r \hat{\mathbf{H}}^{(\ell)} \mathbf{H}^{\ell-r-1} \right)$$

$$S^{ij} = \begin{cases} \mathbf{J}^{ij} + \mathbf{J}^{ji} - \mathbf{J}^{ik} - \mathbf{J}^{kj} - \mathbf{J}^{jk} - \mathbf{J}^{ki} + 2\mathbf{J}^{kk}, & \text{if } i < j, j \neq k \\ \mathbf{J}^{ij} - \mathbf{J}^{ik} - \mathbf{J}^{kj} + \mathbf{J}^{kk}, & \text{if } i = j, j \neq k \end{cases}$$

where \mathbf{J}^{ij} is single-entry matrix with 1 at (i, j) and 0 elsewhere.

4.8 DCE with Restarts (DCer)

Whereas MCE solves a convex optimization problem, the objective function for DCE becomes non-convex for a sparsely labeled graph (i.e. $f \ll 1$). Given a number of classes k , DCE optimizes over $k^* = \Theta(k^2)$ free parameters. Since the parameter space have several local minimas, the optimization should be restarted from multiple points in the k^* -dimensional parameter space, in order to find the global minimum. Thus DCer optimizes the same energy function Eq. (13) as DCE, but with *multiple restarts* from different initial values.

Here our two-step approach of separating the estimation into two steps (recall Fig. 2) becomes an asset: Because the optimizations run on small sketches of the graph that are independent of the graph size, starting multiple optimizations is actually cheap. For small k , restarting from within each of the 2^{k^*} possible hyper-quadrants of parameter space (each free parameter being $\frac{1}{k} \pm \delta$ for some small $\delta < \frac{1}{k^2}$) is negligible as compared to the graph summarization: This is so as for increasing m (large graphs), calculation of the graph statistics dominates the cost for optimization (see Fig. 6k, where DCE and DCer are effectively equal for large graphs). For higher k , our extensive experiments show that in practice Eq. (13) has nice enough properties that just restarting from a limited number of restarts usually leads to a compatibility matrix that achieves the optimal labeling accuracy (see Fig. 6h and discussion in Section 5.2).

4.9 Complexity Analysis

Proposition 4.5 shows that our factorized approach for calculating all ℓ_{\max} graph estimators $\hat{\mathbf{P}}_{\text{NB}}^{(\ell)}$ is $O(mk\ell_{\max})$ and thus is linear in the size of the graph (number of edges). The second step of estimating the compatibility matrix \mathbf{H} is then *independent of the graph size* and depends only on k and the number of restarts r . The number of free parameters in the optimization is $k^* = O(k^2)$, and calculating the Hessian is quadratic in this number. Thus, the second step is $O(k^4r)$.

5 EXPERIMENTS

We designed the experiments to answer two key questions: (1) How *accurate* is our approach in predicting the remaining nodes (and how sensitive is it with respect to our single hyperparameter)? (2) How *fast* is it and how does it scale?

We use two types of datasets: we first perform *carefully controlled experiments* on synthetic datasets that allow us to change various graph parameters. We then use *8 real-world datasets* with high levels of class imbalance, various mixes between homophily and heterophily, and extreme skews of compatibilities. There, we verify that our methods also work well on a variety of datasets which we did not generate.

Using both datasets, we show that: (1) Our method “Distant Compatibility Estimation with restarts” (DCer) is largely insensitive to the choice of its hyperparameter and consistently competes with the labeling accuracy of using the “true” compatibilities (gold standard). (2) DCer is faster than label propagation with LinBP [18] for large graphs, which makes it a simple and cheap pre-processing step (and thereby again rendering heuristics and domain experts obsolete).

Synthetic graph generator. We first use a completely controlled simulation environment. This setup allows us to systematically change parameters of the planted compatibility matrix and see the effect on the accuracy of the techniques as result of such changes. We can thus make observations from many repeated experiments that would not be feasible otherwise. Our synthetic graph generator is a variant of the widely studied stochastic block-model described in [52], but with two crucial generalizations: (1) we actively control the degree distributions in the resulting graph (which allows us to plant more realistic power-law degree distributions); and (2) we “*plant exact graph properties*” instead of fixing a property only in expectation. In other words, our generator creates a desired degree distribution and compatibility matrix during graph generation, which allows us to control all important parameters. The input to the generator is a tuple of parameters $(n, m, \alpha, \mathbf{H}, \text{dist})$ where n is the number of nodes, m the number of edges, α the node label distribution with $\alpha(i)$ being the fraction of nodes of class i ($i \in [k]$), \mathbf{H} any symmetric doubly-stochastic compatibility matrix, and “dist” a family of degree distributions. Notice that α allows us to simulate arbitrary *node imbalances*. In some of our synthetic experiments, we parameterize the compatibility matrix by a value h representing the ratio between min and max entries. Thus parameter h models the “skew” of the potential: For $k = 3$, $\mathbf{H} = \begin{bmatrix} 1 & h & 1 \\ h & 1 & 1 \\ 1 & 1 & h \end{bmatrix} / (2 + h)$. For example, $\mathbf{H} = \begin{bmatrix} 0.1 & 0.8 & 0.1 \\ 0.8 & 0.1 & 0.1 \\ 0.1 & 0.1 & 0.8 \end{bmatrix}$ for $h = 8$, and $\mathbf{H} = \begin{bmatrix} 0.2 & 0.6 & 0.2 \\ 0.6 & 0.2 & 0.2 \\ 0.2 & 0.2 & 0.6 \end{bmatrix}$ for $h = 3$ (see Example 4.2). We create graphs with n nodes and, assuming class balance, assign equal fractions of nodes to one of the k classes, e.g. $\alpha = [\frac{1}{3}, \frac{1}{3}, \frac{1}{3}]$. We also vary the average degree of the graph $d = 2\frac{m}{n}$ and perform experiments assuming power law (coefficient 0.3) distributions.

Quality assessment. We randomly sample a stratified fraction f of nodes as seed labels (i.e. classes are sampled in proportion to their frequencies) and evaluate end-to-end

accuracy as the fraction of the remaining nodes that receive correct labels. Random sampling of seed nodes mimic real-world setting, like social networks, where people who choose to disclose their data, like gender label or political affiliation, are randomly scattered. Notice that decreasing f represents increasing *label sparsity*. To account for class imbalance, we macro-average the accuracy, i.e. we take the mean of the partial accuracies for each class.

Computational setup and code. We implement our algorithms in Python using optimized libraries for sparse matrix operations (NumPy [59] and Scipy [26]). Timing data was taken on a 2.5 Ghz Intel Core i5 with 16G of main memory and a 1TB SSD hard drive. Holdout method uses `scipy.optimize` with the Nelder-Mead Simplex algorithm [1], which is specifically suited for discrete non-contiguous functions.⁷ All other estimation methods use Sequential Least Squares Programming (SLSQP). The spectral radius of a matrix is calculated with an approximate method from the PyAMG library [5] that implements a technique described in [4]. Our code (including the data generator) is inspired by Scikit-learn [49] and will be made publicly available to encourage reproducible research.⁸

5.1 Accuracy of Compatibility Estimation

We show accuracy of parameter estimation by DCER and compare it with “holdout” baseline and linear, myopic and simple distant variants. We consider propagation using ‘true compatibility’ matrix as our gold standard (GS).

Result 1 (Parameter choice of DCER) Normalization variant 1 and longer paths $\ell_{\max} = 5$ are optimal for DCE. Choosing the hyperparameter $\lambda = 10$ performs robustly for a wide range of average degrees d and label sparsity f .

Figure 6a shows DCE used with our three normalization variants and different maximal path lengths ℓ_{\max} . The vertical axis shows the L2-norm between estimation and GS for \mathbf{H} . Variant 3 generally performs worse, and variant 2 generally has higher variance. Our explanation is that finding the L2-norm closest symmetric doubly-stochastic matrix to a stochastic one is a well-behaved optimization problem.

Figure 6b shows DCER for various values of λ and ℓ_{\max} . Notice that DCER for $\ell_{\max} = 1$ is identical to MCE, and that DCER works better for longer paths $\ell_{\max} = 5$, as those can overcome sparsity of seed labels. This observation holds over a wide range of parameters and becomes stronger for small f . Also notice that even numbers $\ell_{\max} = 2$ do not work as well as the objective has multiple minima with identical value.

⁷We tried alternative optimizers, such as the Broyden-Fletcher-Goldfarb-Shanno (“BFGS”) algorithm [3]. Nelder-Mead performed best for the baseline holdout method due to its gradient-free nature.

⁸<https://github.com/northeastern-datalab/factorized-graphs/>

Figures 6c and 6d show the optimal choices of hyperparameter λ (giving the smallest L2 norm from GS) for a wide range of d and f . Each red dot shows an optimal choice of λ . Each gray dot shows a choice with L2-norm that is within 10% of the optimal choice. The red line shows a moving trend line of averaged choices. We see that choosing $\lambda = 10$ is a general robust choice for good estimation, unless we have enough labels (high f): then we don’t need longer paths and can best just learn from immediate neighbors (small λ).

Figure 6e shows the advantage of using $\ell_{\max} = 5$, $\lambda = 10$ and random restarts for estimating \mathbf{H} as compared to just MCE or DCE: for small f , DCE may get trapped in local optima (see Section 4.8); randomly restarting the optimization a few times overcomes this issue.

Result 2 (Accuracy performance of DCER) Label accuracy with DCER is within ± 0.01 of GS performance and is quasi indistinguishable from GS.

Figure 6f show results from first estimating \mathbf{H} on a partially labeled graph and then labeling the remaining nodes with LinBP. We see that *more accurate estimation of \mathbf{H} also translates into more accurate labeling*, which provides strong evidence that state-of-the-art approaches that use simple heuristics are not optimal. GS runs LinBP with gold standard parameters and is the best LinBP can do. The holdout method was varied with $b \in \{1, 2, 4, 8\}$ in Fig. 6f and $b = 1$ else. Increasing the number of splits moderately increases the accuracy for the holdout method, but comes at proportionate cost in time. DCER is faster and more accurate throughout all parameters. In all plots, estimating with DCER gives identical or similar labeling accuracy as knowing the GS. DCE is as good as DCER for $f > 1\%$ for 10k and $f > 0.1\%$ for 100k nodes. MCE and LCE both rely on labeled neighbors and have similar accuracy.

Figures 6b and 6j show that neighbor frequency distributions alone do not work with sparse labels, and that our ℓ -distance trick successfully overcomes its shortcomings.

Result 3 (Restarts required for DCER) With $r = 10$ restarts, DCER obtains the performance levels of GS.

Figure 6h shows propagation accuracy of DCER for different number of restarts r compared against the global minimum baseline, which is calculated by initializing DCE optimization with GS. Notice, the optimal baseline is the best any estimation based method can perform. Averaged over 35 runs, Fig. 6h shows that DCER approaches the global minima with just 10 restarts; we thus use $r = 10$ in our experiments.

Figure 6i serves as sanity check and demonstrates what happens if we use standard random walks (here the harmonic functions method [66]) to label nodes in graphs with

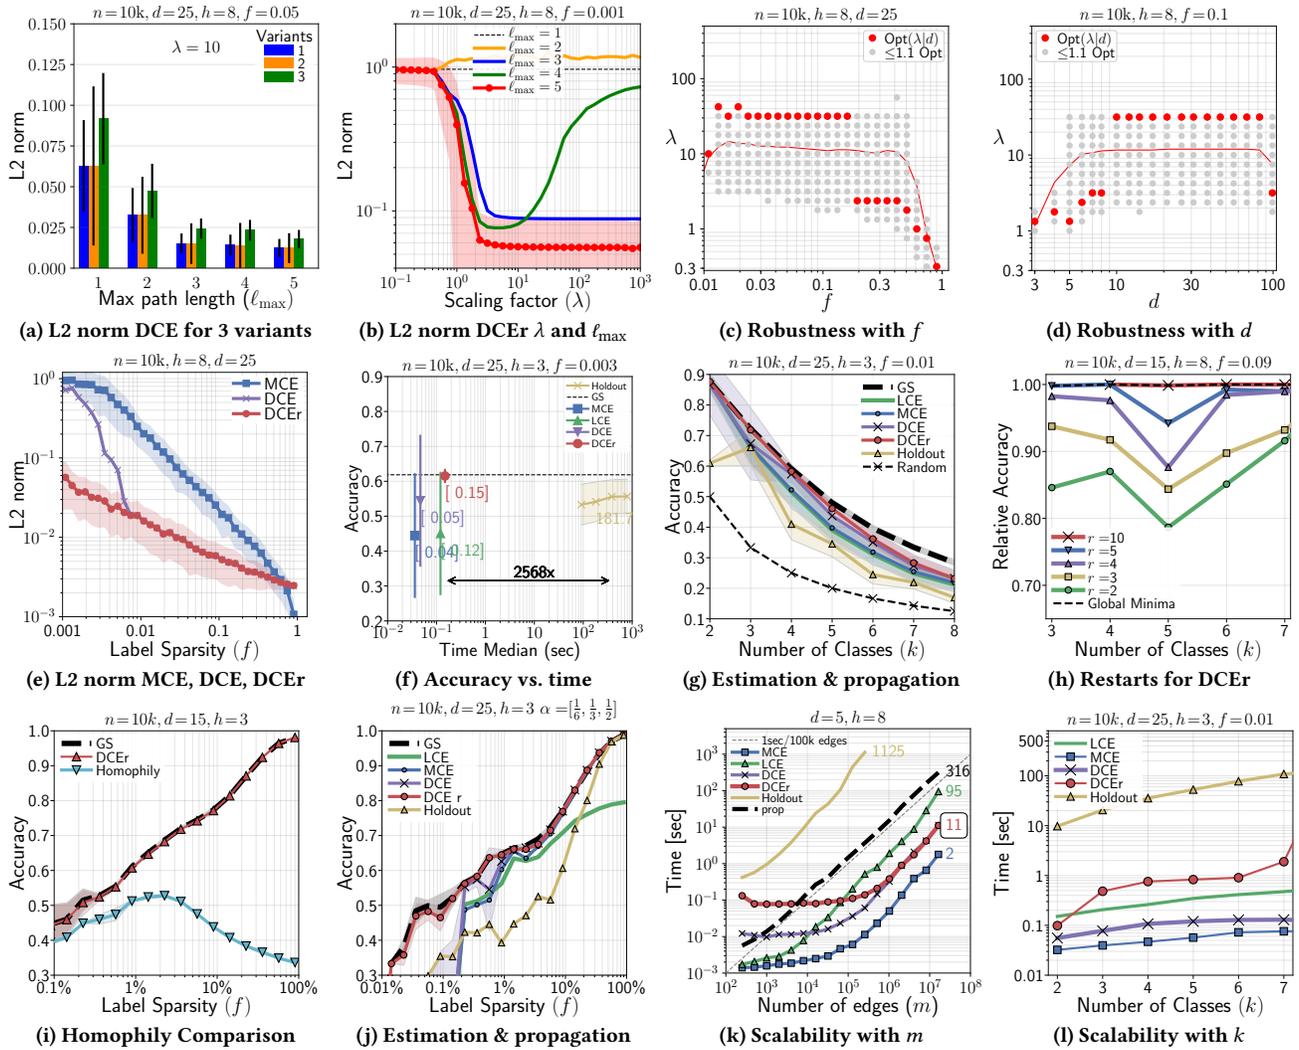


Figure 6: Experimental results for (a)-(j) accuracy (Section 5.1), and scalability of our methods (Section 5.2)

arbitrary compatibilities: Baselines with a homophily assumption fall behind tremendously on graphs that do not follow assortative mixing.

Result 4 (Robustness of DCER) Performance of DCER remains consistently above other baselines for skewed label distributions and large number of classes, whereas other SSL methods deteriorate for $k > 3$.

To illustrate the approaches for class imbalance and more general H , we include an experiment with $\alpha = [\frac{1}{6}, \frac{1}{3}, \frac{1}{2}]$ and $H = \begin{bmatrix} 0.2 & 0.6 & 0.2 \\ 0.6 & 0.1 & 0.3 \\ 0.2 & 0.3 & 0.5 \end{bmatrix}$. Figure 6j (contrast to Fig. 3a) shows that DCER works robustly better than alternatives, can deal with label imbalance, and can learn the more general H .

Figure 6g compares accuracy against random labeling for fixed n, m, h, f , and increasing k . DCER restarts up to 10 times

and works robustly better than alternatives. Recall that the number of compatibilities to learn is $O(k^2)$.

5.2 Scalability of Compatibility Estimation

Figure 6k shows the scalability of our combined methods. On our largest graph with 6.6m nodes and 16.4m edges, propagation takes 316 sec for 10 iterations, estimation with LCE 95 sec, DCE or DCER 11 sec, and MCE 2 sec.

Result 5 (Scalability with increasing graph size) DCER scales linearly in m and experimentally is more than 3 orders of magnitude faster than Holdout method.

Our estimation method DCER is more than 25 times faster than inference (used by Holdout as a subroutine) and thus comes “for free” as m scales. Also notice that DCE and DCER need the same time for large graphs because of our two-step

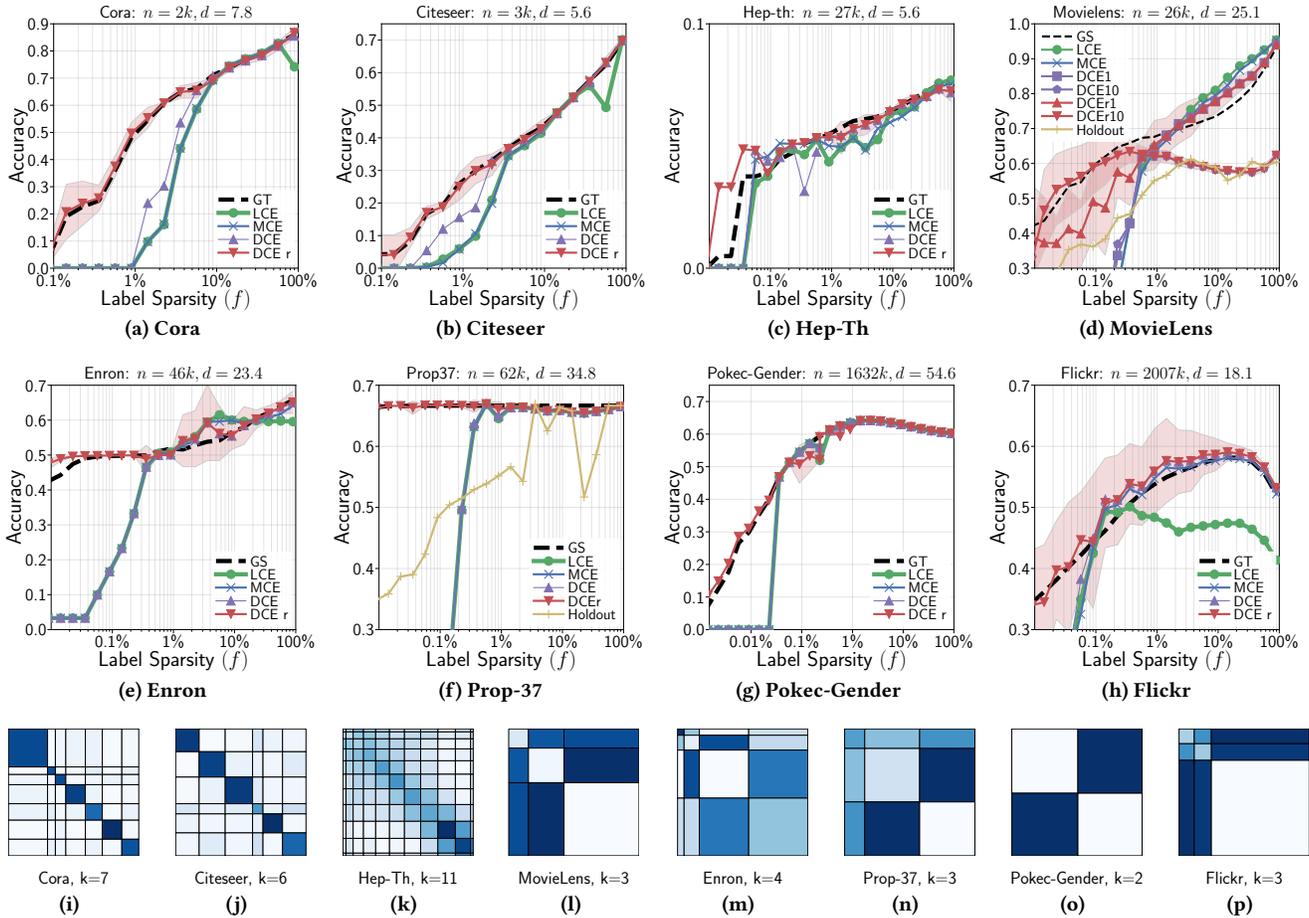


Figure 7: Experiments over 8 real-world datasets (Section 5.3): (a)-(h): Accuracy of end-to-end estimation and propagation. (i)-(p): Illustration of class imbalance and heterophily in their gold standard compatibility matrices (darker colors represent higher number of edges): the first 3 show homophily, the latter 5 arbitrary heterophily.

calculation: the time needed to calculate the graph statistics $\hat{\mathbf{P}}_{\text{NB}}^{(\ell)}$, $\ell \in [5]$ becomes dominant; each of the 8 optimizations of Eq. (13) then takes less than 0.1 sec for any size of the graph. The Holdout method for $b = 1$ takes 1125 sec for a graph with 256k edges. Thus the extrapolated difference in scalability between DCE_r and Holdout is 3-4 orders of magnitude if $b = 1$. Figure 6f shows that increasing the number of splits b for the holdout method can slightly increase accuracy at even higher runtime cost.

Figure 6l uses a setup identical to Figure 6g and shows that our methods also scale nicely with respect to number of classes k , as long as the graph is large and thus the graph summarization take most time. Here, DCE_r uses 10 restarts.

5.3 Performance on Real-World Datasets

We next evaluate our approach over 8 real-world graphs, described in Figure 8, that have a variety of complexities : (i)

Dataset	n	m	d	k	DCE _r
Cora [53]	2,708	10,858	8.0	7	3.33
Citeseer [53]	3,312	9,428	5.7	6	1.13
Hep-Th [19]	27,770	352,807	25.4	11	10.61
MovieLens [54]	26,850	336,742	25.0	3	0.07
Enron [33]	46,463	613,838	26.4	4	0.20
Prop-37 [55]	62,383	2,167,809	69.4	3	0.09
Pokec-Gender [57]	1,632,803	30,622,564	37.5	2	5.12
Flickr [37]	2,007,369	18,147,504	18.1	3	2.39

Figure 8: Real-world dataset statistics, Section 5.3. The last column shows runtime of DCE_r (in sec).

Graphs are formed by very different processes, (ii) class distributions are often highly imbalanced, (iii) compatibilities are often skewed by orders of magnitude, and (iv) graphs are so large that it is infeasible to even run Holdout. Our 8 datasets are as follows:

- (1) Cora [53] is also a citation graph containing publications from 7 categories in the area of ML (*neural nets, rule*

learning, reinforcement learning, probabilistic methods, theory, genetic algorithms, case based).

- (2) **Citeseer** [53] contains 3264 publications from 6 categories in the area of Computer Science. The citation graph connects papers from six different classes (*agents, IR, DB, AI, HCI, ML*).
- (3) **Hep-Th** [37] is the High Energy Physics Theory publication network is from arXiv e-print and covers their citations. The node are labeled based on one of 11 years of publication (from 1993 to 2003).
- (4) **MovieLens** [54] is from a movie recommender system that connects 3 classes: *users, movies*, and assigned *tags*.
- (5) **Enron** [33] has 4 types of nodes: *person, email address, message* and *topic*. Messages are connected to topics and email addresses; people are connected to email addresses.
- (6) **Prop-37** [55] comes from a California ballot initiative on labeling genetically engineered foods. It is based on Twitter data and has 3 classes: *users, tweets*, and *words*.
- (7) **Pokec-Gender** [57] is a social network graph connecting people (*male* or *female*) to their friends. More interaction edges exist between people of opposite gender.
- (8) **Flickr** [37] connects *users* to *pictures* they upload and other users pictures in the same *group*. Pictures are also connected to groups they are posted in.

Result 6 (Accuracy on real datasets) DCEr consistently labels nodes with accuracy ± 0.01 compared to true compatibility matrix for $f < 10\%$ and ± 0.03 for $f > 10\%$ averaged across datasets, basically indistinguishable from GS.

Our experimental setup is similar to before: we estimate \mathbf{H} on a random fraction f of labeled seed sets and repeat many times. We retrieve the Gold-Standard (GS) compatibilities from the relative label distribution on the fully labeled graph (if we know all labels in a graph, then we can simply “measure” the relative frequencies of classes between neighboring nodes). The remaining nodes are then labeled with LinBP, 10 iterations, $s = 0.5$, as suggested by [18].

The relative accuracy for varying f can look quite differently for different networks. Notice that our real-world graphs show a wide variety of (i) label imbalance, (ii) number of classes ($2 \leq k \leq 12$), (iii) mixes of homophily and heterophily, and (iv) a wide variety of skews in compatibilities. This variety can be seen in our illustrations of the gold-standard compatibility matrices (Figs. 7i to 7p): we clustered all nodes by their respective classes, and shades of blue represent the relative frequency of edges between classes. It appears that the accuracy of LinBP for varying f can be widely affected by combinations of graph parameters. Also, for MovieLens Fig. 7d, choosing $\lambda = 10$ worked better for DCEr in the sparse regime ($f < 1\%$), while $\lambda = 1$ worked better for $f > 1\%$. This observation appears consistent with our

understanding of λ where larger values can amplify weaker distant signals, yet smaller λ is enough to propagate stronger signals. Fine-tuning of λ on real datasets remains interesting future work. However, all our experiments consistently show that our methods for learning compatibilities robustly compete with the gold-standard and beat all other methods, especially for sparsely labeled graphs. This suggests that our approach renders any prior heuristics obsolete.

Scalability with increasing number of classes k . We know from Section 4.9 that joint complexity of our two-step compatibility estimation is $O(mk + k^4r)$. For very large graphs with small k , the first factor stemming from graph summarization is dominant and estimation runs in $O(mk)$, which is faster than propagation. However, for moderate graphs with high k (e.g. Hep-Th with $k = 11$), the second factor can dominate since calculating the Hessian matrix has $O(k^4)$ complexity. Thus for high k and small graphs, our approach will not remain faster than propagation, but in most practical settings estimation will act as a computationally cheap *pre-processing* step.

6 CONCLUSIONS

Label propagation methods that can propagate arbitrary class-to-class compatibilities between nodes require those compatibilities as input. Prior work assumes these compatibilities as given. We instead propose methods to accurately learn compatibilities from sparsely labeled graphs in a fraction of the time it takes to later propagate the labels, resolving the long open question of where the compatibilities come from. The take-away for practitioners is that *prior knowledge of compatibilities is no longer necessary* and estimation can become a cheap pre-processing step before labeling.

Our approach *amplifies signals from sparse data* by leveraging *algebraic properties* (notably the distributive law) in the update equations of linear label propagation algorithms, an idea we call *algebraic amplification*. We use linearized belief propagation [18], which is derived from the widely used inference method belief propagation by applying linearization as an algebraic simplification, and compliment it with a method for parameter estimation. Thus our estimation method *uses the same approximations used for inference, for learning the parameters as well*. Such linear methods have well-known benefits: they have computational advantages (especially during learning), have fewer parameters (which helps with label sparsity), have fewer hyperparameters (which simplifies tuning), are easier to interpret, and have favorable algebraic properties that are crucial to our approach. It remains to be seen how graph neural networks (GNNs) can be adapted to similarly sparse data.

Acknowledgements. This work was supported in part by NSF under award CAREER IIS-1762268.

REFERENCES

- [1] James W. Akitt. 1977. Function Minimisation Using the Nelder and Mead Simplex Method with Limited Arithmetic Precision: The Self Regenerative Simplex. *Comput. J.* 20, 1 (1977), 84–85. <https://doi.org/10.1093/comjnl/20.1.84>
- [2] Noga Alon, Itai Benjamini, Eyal Lubetzky, and Sasha Sodin. 2007. Non-backtracking random walks mix faster. *Communications in Contemporary Mathematics* 09, 04 (2007), 585–603. <https://doi.org/10.1142/S0219199707002551>
- [3] Paul Armand, Jean Charles Gilbert, and Sophie Jan-Jégou. 2000. A Feasible BFGS Interior Point Algorithm for Solving Convex Minimization Problems. *SIAM Journal on Optimization* 11, 1 (2000), 199–222. <https://doi.org/10.1137/S1052623498344720>
- [4] Zhaojun Bai, James Demmel, Jack Dongarra, Axel Ruhe, and Henk van der Vorst. 2000. *Templates for the solution of algebraic eigenvalue problems*. SIAM. <https://doi.org/10.1137/1.9780898719581>
- [5] Nathan Bell, Luke Olson, and Jacob Schroder. 2011. PyAMG: Algebraic Multigrid Solvers in Python v2.0. <http://pyamg.github.io/>
- [6] Yoshua Bengio, Olivier Delalleau, and Nicolas Le Roux. 2006. Label propagation and quadratic criterion. In *Semi-supervised learning*, Olivier Chapelle, Bernhard Schölkopf, and Alexander Zien (Eds.). MIT Press, 193–216. <https://doi.org/10.7551/mitpress/9780262033589.003.0011>
- [7] Smriti Bhagat, Graham Cormode, and S. Muthukrishnan. 2011. Node Classification in Social Networks. In *Social Network Data Analytics*, Charu C. Aggarwal (Ed.). Springer, 115–148. http://doi.org/10.1007/978-1-4419-8462-3_5
- [8] Smriti Bhagat, Graham Cormode, and Irina Rozenbaum. 2009. Applying Link-Based Classification to Label Blogs. In *Advances in Web Mining and Web Usage Analysis (SNAKDD 2007) (LNCS)*, Vol. 5439. Springer, 97–117. https://doi.org/10.1007/978-3-642-00528-2_6
- [9] Sourav S. Bhowmick and Boon Siew Seah. 2016. Clustering and Summarizing Protein-Protein Interaction Networks: A Survey. *TKDE* 28, 3 (2016), 638–658. <http://dx.doi.org/10.1109/TKDE.2015.2492559>
- [10] Christopher M Bishop. 2006. *Pattern recognition and machine learning*. Springer, New York. <https://www.springer.com/gp/book/9780387310732>
- [11] Soumen Chakrabarti. 2007. Dynamic personalized pagerank in entity-relation graphs. In *WWW*. 571–580. <https://doi.org/10.1145/1242572.1242650>
- [12] Xiaowei Chen, Yongkun Li, Pinghui Wang, and John C. S. Lui. 2016. A General Framework for Estimating Graphlet Statistics via Random Walk. *PVLDB* 10, 3 (2016), 253–264. <https://doi.org/10.14778/3021924.3021940>
- [13] David L. Donoho, Arian Maleki, and Andrea Montanari. 2009. Message-passing algorithms for compressed sensing. *PNAS* 106, 45 (2009), 18914–18919. <http://doi.org/10.1073/pnas.0909892106>
- [14] Gal Elidan, Ian McGraw, and Daphne Koller. 2006. Residual Belief Propagation: Informed Scheduling for Asynchronous Message Passing. In *UAI*. 165–173. <https://doi.org/doi/10.5555/3020419.3020440>
- [15] Dhivya Eswaran, Stephan Günnemann, Christos Faloutsos, Disha Makhija, and Mohit Kumar. 2017. ZooBP: Belief Propagation for Heterogeneous Networks. *PVLDB* 10, 5 (2017), 625–636. <https://doi.org/10.14778/3055540.3055554>
- [16] Wolfgang Gatterbauer. 2014. Semi-Supervised Learning with Heterophily. *CoRR* abs/1412.3100 (2014). <https://arxiv.org/abs/1412.3100>
- [17] Wolfgang Gatterbauer. 2017. The Linearization of Belief Propagation on Pairwise Markov Random Fields. In *AAAI*. 3747–3753. <https://arxiv.org/pdf/1502.04956>
- [18] Wolfgang Gatterbauer, Stephan Günnemann, Danai Koutra, and Christos Faloutsos. 2015. Linearized and Single-Pass Belief Propagation. *PVLDB* 8, 5 (2015), 581–592. <https://doi.org/doi/10.14778/2735479.2735490>
- [19] Johannes Gehrke, Paul Ginsparg, and Jon Kleinberg. 2003. Overview of the 2003 KDD Cup. *SIGKDD Explor. Newsl.* 5, 2 (Dec. 2003), 149–151. <https://doi.org/10.1145/980972.980992>
- [20] Andrew B. Goldberg, Xiaojin Zhu, and Stephen J. Wright. 2007. Dissimilarity in Graph-Based Semi-Supervised Classification. In *AISTATS*. 155–162. <http://www.jmlr.org/proceedings/papers/v2/goldberg07a.html>
- [21] William L. Hamilton, Rex Ying, and Jure Leskovec. 2017. Inductive Representation Learning on Large Graphs. *USA*, 1025–1035. <https://dl.acm.org/doi/10.5555/3294771.3294869>
- [22] Ki-ichiro Hashimoto. 1989. Zeta functions of finite graphs and representations of p-adic groups. In *Automorphic forms and geometry of arithmetic varieties*. Elsevier, 211–280. <https://doi.org/10.1016/B978-0-12-330580-0.50015-X>
- [23] T.H. Haveliwala. 2003. Topic-sensitive pagerank: A context-sensitive ranking algorithm for web search. *TKDE* 15, 4 (2003), 784–796. <https://dl.acm.org/10.1109/TKDE.2003.1208999>
- [24] Taher Haveliwala, Sepandar Kamvar, and Glen Jeh. 2003. *An Analytical Comparison of Approaches to Personalizing PageRank*. Technical Report 2003-35. Stanford InfoLab, Stanford. <http://ilpubs.stanford.edu:8090/596/>
- [25] Raphael Hoffmann, Congle Zhang, Xiao Ling, Luke S. Zettlemoyer, and Daniel S. Weld. 2011. Knowledge-Based Weak Supervision for Information Extraction of Overlapping Relations. In *ACL*. 541–550. <https://www.aclweb.org/anthology/P11-1055/>
- [26] Eric Jones, Travis Oliphant, Pearu Peterson, et al. 2001–. SciPy: Open source scientific tools for Python. <http://www.scipy.org/>.
- [27] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *ICLR (Poster)*. <https://openreview.net/forum?id=SJU4ayYgl>
- [28] Daphne Koller and Nir Friedman. 2009. *Probabilistic Graphical Models: Principles and Techniques - Adaptive Computation and Machine Learning*. The MIT Press. <https://doi.org/doi/book/10.5555/1795555>
- [29] Danai Koutra, Tai-You Ke, U. Kang, Duen Horng Chau, Hsing-Kuo Kenneth Pao, and Christos Faloutsos. 2011. Unifying Guilt-by-Association Approaches: Theorems and Fast Algorithms. In *ECML/PKDD (2)*. 245–260. http://dx.doi.org/10.1007/978-3-642-23783-6_16
- [30] Krishna Kumar P., Paul Langton, and Wolfgang Gatterbauer. 2020. Factorized graph representations for semi-supervised learning from sparse data. *CoRR* abs/2003.02829 (2020). <https://arxiv.org/abs/2003.02829>
- [31] Florent Krzakala, Cristopher Moore, Elchanan Mossel, Joe Neeman, Allan Sly, Lenka Zdeborová, and Pan Zhang. 2013. Spectral redemption in clustering sparse networks. *PNAS* 110, 52 (2013), 20935–20940. <http://dx.doi.org/10.1073/pnas.1312486110>
- [32] Chul-Ho Lee, Xin Xu, and Do Young Eun. 2012. Beyond Random Walk and Metropolis-hastings Samplers: Why You Should Not Backtrack for Unbiased Graph Sampling. In *SIGMETRICS*. 319–330. <http://dx.doi.org/10.1145/2254756.2254795>
- [33] Jiongqian Liang, Deepak Ajwani, Patrick K. Nicholson, Alessandra Sala, and Srinivasan Parthasarathy. 2016. What Links Alice and Bob?: Matching and Ranking Semantic Patterns in Heterogeneous Networks. In *WWW*. 879–889. <http://doi.acm.org/10.1145/2872427.2883007>
- [34] Frank Lin and William W. Cohen. 2010. Semi-Supervised Classification of Network Data Using Very Few Labels. In *ASONAM*. 192–199. <http://dx.doi.org/10.1109/ASONAM.2010.19>
- [35] Qing Lu and Lise Getoor. 2003. Link-based Classification. In *ICML*. 496–503. <http://www.aaai.org/Library/ICML/2003/icml03-066.php>
- [36] Travis Martin, Xiao Zhang, and M. E. J. Newman. 2014. Localization and centrality in networks. *Phys. Rev. E* 90 (Nov 2014), 052808. Issue 5. <http://doi.org/10.1103/PhysRevE.90.052808>

- [37] Julian J. McAuley and Jure Leskovec. 2012. Image Labeling on a Network: Using Social-Network Metadata for Image Classification. In *ECCV*. 828–841. https://doi.org/10.1007/978-3-642-33765-9_59
- [38] Mary McGlohon, Stephen Bay, Markus G. Anderle, David M. Steier, and Christos Faloutsos. 2009. SNARE: a link analytic system for graph labeling and risk detection. In *KDD*. 1265–1274. <http://doi.acm.org/10.1145/1557019.1557155>
- [39] Mike Mintz, Steven Bills, Rion Snow, and Daniel Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *ACL*. 1003–1011. <https://www.aclweb.org/anthology/P09-1113/>
- [40] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. 2012. *Foundations of machine learning*. MIT Press, Cambridge, MA. <https://cs.nyu.edu/~mohri/mlbook/>
- [41] Joris M. Mooij and Hilbert J. Kappen. 2007. Sufficient Conditions for Convergence of the Sum-Product Algorithm. *IEEE Transactions on Information Theory* 53, 12 (2007), 4422–4437. <http://doi.org/10.1109/TIT.2007.909166>
- [42] Christopher Moore, Xiaoran Yan, Yaojia Zhu, Jean-Baptiste Rouquier, and Terran Lane. 2011. Active learning for node classification in assortative and disassortative networks. In *KDD*. 841–849. <http://doi.acm.org/10.1145/2020408.2020552>
- [43] John Moore and Jennifer Neville. 2017. Deep Collective Inference. In *AAAI*. 2364–2372. <http://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14650>
- [44] Kevin P. Murphy. 2012. *Machine learning: a probabilistic perspective*. MIT Press, Cambridge, MA. <https://mitpress.mit.edu/books/machine-learning-1>
- [45] DA Olteanu and Maximilian Schleich. 2016. Factorized databases. *SIGMOD Record* 45, 2 (2016). <https://doi.org/10.1145/3003665.3003667>
- [46] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1999. *The PageRank Citation Ranking: Bringing Order to the Web*. TR 1999-66. Stanford InfoLab. <http://ilpubs.stanford.edu:8090/422/>
- [47] Jia-Yu Pan, Hyung-Jeong Yang, Christos Faloutsos, and Pinar Duygulu. 2004. Automatic Multimedia Cross-modal Correlation Discovery. In *KDD*. 653–658. <https://doi.org/10.1145/1014052.1014135>
- [48] Shashank Pandit, Duen Horng Chau, Samuel Wang, and Christos Faloutsos. 2007. Netprobe: a fast and scalable system for fraud detection in online auction networks. In *WWW*. 201–210. <http://doi.acm.org/10.1145/1242572.1242600>
- [49] Pedregosa et. al. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830. <http://scikit-learn.org>.
- [50] Leto Peel. 2017. Graph-based semi-supervised learning for relational networks. In *ICDM*. SIAM, 435–443. <https://doi.org/10.1137/1.9781611974973.49>
- [51] Alexander Ratner, Stephen H. Bach, Henry R. Ehrenberg, Jason Alan Fries, Sen Wu, and Christopher Ré. 2017. Snorkel: Rapid Training Data Creation with Weak Supervision. *PVLDB* 11, 3 (2017), 269–282. <https://doi.org/10.14778/3157794.3157797>
- [52] Prithviraj Sen and Lise Getoor. 2007. *Link-based classification*. Technical Report. University of Maryland Technical Report CS-TR-4858. <https://drum.lib.umd.edu/bitstream/handle/1903/4298/report.pdf>
- [53] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Gallagher, and Tina Eliassi-Rad. 2008. Collective Classification in Network Data. *AI Magazine* 29, 3 (2008), 93–106. <https://doi.org/10.1609/aimag.v29i3.2157>
- [54] Shilad Sen, Jesse Vig, and John Riedl. 2009. Tagommenders: connecting users to items through tags. In *WWW*. 671–680. <https://doi.org/10.1145/1526709.1526800>
- [55] Laura M Smith, Linhong Zhu, Kristina Lerman, and Zornitsa Kozareva. 2013. The role of social media in the discussion of controversial topics. In *SocialCom*. IEEE, 236–243. <https://doi.org/10.1109/SocialCom.2013.41>
- [56] Amarnag Subramanya and Partha Pratim Talukdar. 2014. Graph-Based Semi-Supervised Learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 8, 4 (2014), 1–125. <http://doi.org/10.2200/S00590ED1V01Y201408AIM029>
- [57] Lubos Takac and Michal Zabovsky. 2012. Data analysis in public social networks. <https://snap.stanford.edu/data/soc-pokec.pdf>
- [58] Leo Torres, Pablo Suárez-Serrato, and Tina Eliassi-Rad. 2019. Non-backtracking cycles: length spectrum theory and graph mining applications. *Applied Network Science* 4, 1 (2019), 41:1–41:35. <https://doi.org/10.1007/s41109-019-0147-y>
- [59] S. van der Walt, S. C. Colbert, and G. Varoquaux. 2011. The NumPy Array: A Structure for Efficient Numerical Computation. *Computing in Science Engineering* 13, 2 (March 2011), 22–30. <http://docs.scipy.org/>.
- [60] Fei Wang and Changshui Zhang. 2008. Label Propagation through Linear Neighborhoods. *TKDE* 20, 1 (2008), 55–67. <https://doi.org/10.1109/TKDE.2007.190672>
- [61] Yair Weiss. 2000. Correctness of Local Probability Propagation in Graphical Models with Loops. *Neural Computation* 12, 1 (2000), 1–41. <http://doi.org/10.1162/089976600300015880>
- [62] I. H Witten, Eibe Frank, and Mark A Hall. 2011. *Data mining: practical machine learning tools and techniques* (3rd ed ed.). Morgan Kaufmann, Burlington, MA. <http://www.sciencedirect.com/science/book/9780123748560>
- [63] Mingrui Wu and Bernhard Schölkopf. 2007. Transductive Classification via Local Learning Regularization. In *AISTATS*. 628–635. <http://proceedings.mlr.press/v2/wu07a.html>
- [64] Dengyong Zhou, Olivier Bousquet, Thomas Navin Lal, Jason Weston, and Bernhard Schölkopf. 2003. Learning with Local and Global Consistency. In *NIPS*. 321–328. <https://dl.acm.org/doi/10.5555/2981345.2981386>
- [65] Xiaojin Zhu. 2005. *Semi-Supervised Learning Literature Survey*. Technical Report 1530. Computer Sciences, University of Wisconsin-Madison. <https://minds.wisconsin.edu/handle/1793/60444>
- [66] Xiaojin Zhu, Zoubin Ghahramani, and John D. Lafferty. 2003. Semi-Supervised Learning Using Gaussian Fields and Harmonic Functions. In *ICML*. 912–919. <https://dl.acm.org/doi/10.5555/3041838.3041953>
- [67] Xiaojin Zhu, Jaz Kandola, John Lafferty, and Zoubin Ghahramani. 2006. Graph Kernels by Spectral Transforms. In *Semi-supervised learning*. MIT Press, 277–291. <https://doi.org/10.7551/mitpress/9780262033589.001.0001>