

Learning a Distance Metric from a Network

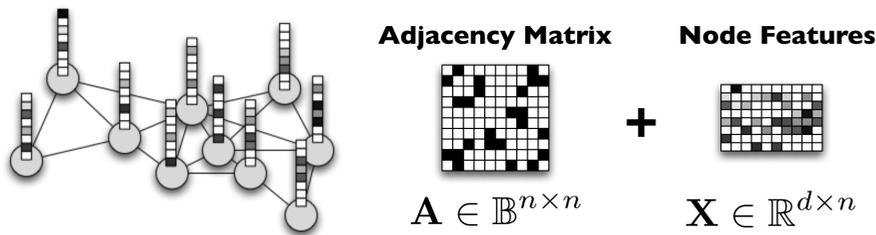
Blake Shaw
Foursquare

Bert Huang
University of Maryland

Tony Jebara
Columbia University

Introduction

Real-world network data often consists of **both** node features and connectivity. Can we learn a metric that relates features to links?



Examples:

Wikipedia: node features are word-counts and links are hyperlinks

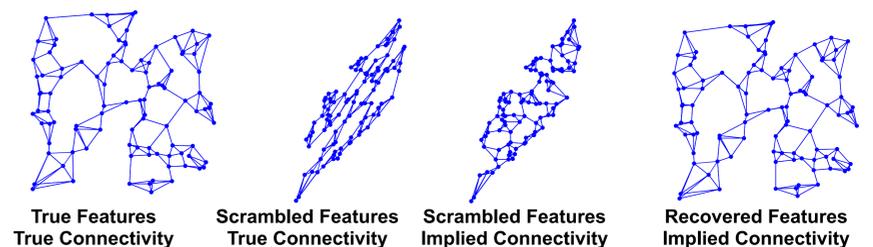
Facebook: node features are profiles and links are friendships

Foursquare: node features are places people visit and links are friendships

- While homophily is expected in natural networks, nodes do not simply connect based on similarity of their features alone
- Modeling independent links is insufficient, so one must account for the inherent topology of the network
- We propose learning a distance metric from large social networks that captures relationships between node features and the structure of the network: Structure Preserving Metric Learning (SPML)

Background

- Current metric learning algorithms are used for supervised tasks like classification [Chechnik et al. '10, Weinberger et al. '10]
- These methods push away “class impostors”
- SPML pushes away “graph impostors”
- Following intuition from Structure Preserving Embedding [Shaw and Jebara '09], SPML finds a metric that is structure preserving
- Given adjacency matrix \mathbf{A} and node features \mathbf{X} , a distance metric parametrized by $\mathbf{M} \in \mathbb{R}^{d \times d}$ is **structure preserving** with respect to a connectivity algorithm \mathcal{G} if $\mathcal{G}(\mathbf{X}, \mathbf{M}) = \mathbf{A}$



Structure Preserving Metric Learning

- Regularizing \mathbf{M} by penalizing its Frobenius norm, SPML objective is a semidefinite program (SDP), which is too expensive for large networks
- Instead we rewrite the objective function over a large set of triplet constraints and optimize via stochastic gradient descent

$$f(\mathbf{M}) = \frac{\lambda}{2} \|\mathbf{M}\|_F^2 + \frac{1}{|S|} \sum_{(i,j,k) \in S} \max(D_{\mathbf{M}}(\mathbf{x}_i, \mathbf{x}_j) - D_{\mathbf{M}}(\mathbf{x}_i, \mathbf{x}_k) + 1, 0),$$

where $S = \{(i, j, k) \mid A_{ij} = 1, A_{ik} = 0\}$

- Using the distance transformation

$$D_{\mathbf{M}}(\mathbf{x}_i, \mathbf{x}_j) = \mathbf{x}_i^T \mathbf{M} \mathbf{x}_i + \mathbf{x}_j^T \mathbf{M} \mathbf{x}_j - \mathbf{x}_i^T \mathbf{M} \mathbf{x}_j - \mathbf{x}_j^T \mathbf{M} \mathbf{x}_i,$$

constraints can be written using a sparse matrix $\mathbf{C}^{(i,j,k)}$, where

$$C_{jj}^{(i,j,k)}, C_{ik}^{(i,j,k)}, C_{ki}^{(i,j,k)} = 1, \text{ and } C_{ij}^{(i,j,k)}, C_{ji}^{(i,j,k)}, C_{kk}^{(i,j,k)} = -1$$

- The subgradient of f at \mathbf{M} is then

$$\nabla f = \lambda \mathbf{M} + \frac{1}{|S|} \sum_{(i,j,k) \in S_+} \mathbf{X} \mathbf{C}^{(i,j,k)} \mathbf{X}^T,$$

where $S_+ = \{(i, j, k) \mid D_{\mathbf{M}}(\mathbf{x}_i, \mathbf{x}_j) - D_{\mathbf{M}}(\mathbf{x}_i, \mathbf{x}_k) + 1 > 0\}$

- We perform stochastic subgradient descent by randomly sampling a mini-batch of triplets at each iteration
- Theorem: This method does not scale with the size of the network, only the desired approximation error!!!

Algorithm 3 Structure preserving metric learning with nearest neighbor constraints and optimization with projected stochastic subgradient descent

Input: $\mathbf{A} \in \mathbb{B}^{n \times n}$, $\mathbf{X} \in \mathbb{R}^{d \times n}$, and parameters λ, T, B

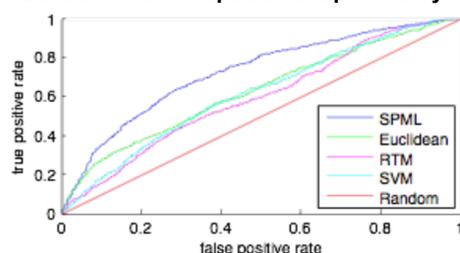
- $\mathbf{M}_1 \leftarrow \mathbf{I}_d$
- for** t from 1 to $T - 1$ **do**
- $\eta_t \leftarrow \frac{1}{\lambda t}$
- $\mathbf{C} \leftarrow \mathbf{0}_{n,n}$
- for** b from 1 to B **do**
- $(i, j, k) \leftarrow$ Sample random triplet from $S = \{(i, j, k) \mid A_{ij} = 1, A_{ik} = 0\}$
- if** $D_{\mathbf{M}_t}(\mathbf{x}_i, \mathbf{x}_j) - D_{\mathbf{M}_t}(\mathbf{x}_i, \mathbf{x}_k) + 1 > 0$ **then**
- $\mathbf{C}_{jj} \leftarrow \mathbf{C}_{jj} + 1, \mathbf{C}_{ik} \leftarrow \mathbf{C}_{ik} + 1, \mathbf{C}_{ki} \leftarrow \mathbf{C}_{ki} + 1$
- $\mathbf{C}_{ij} \leftarrow \mathbf{C}_{ij} - 1, \mathbf{C}_{ji} \leftarrow \mathbf{C}_{ji} - 1, \mathbf{C}_{kk} \leftarrow \mathbf{C}_{kk} - 1$
- end if**
- end for**
- $\nabla_t \leftarrow \mathbf{X} \mathbf{C} \mathbf{X}^T + \lambda \mathbf{M}_t$
- $\mathbf{M}_{t+1} \leftarrow \mathbf{M}_t - \eta_t \nabla_t$
- Optional: $\mathbf{M}_{t+1} \leftarrow [\mathbf{M}_{t+1}]^+$ {Project onto the PSD cone}
- end for**
- return** \mathbf{M}_T

Experiments

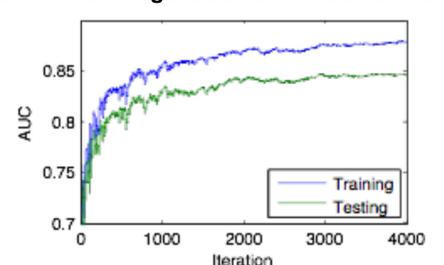
- Learn a metric using 80% of nodes as training and evaluate prediction of links for the held-out 20%, scoring AUC of ranking
- Data sources: Wikipedia data, Facebook data, Foursquare data
- Compare with existing methods: Euclidean distance, Relational Topic Models (RTM), and Support Vector Machines (SVM)

	n	m	d	Euclidean	RTM	SVM	SPML	
Wikipedia	Graph Theory	223	917	6695	0.624	0.591	0.610	0.722
	Philosophy Concepts	303	921	6695	0.705	0.571	0.708	0.707
	Search Engines	269	332	6695	0.662	0.487	0.611	0.742
	Philosophy Crawl	100,000	4,489,166	7702	0.547	-	-	0.601
FB	Harvard	1937	48,980	193	0.764	0.562	0.839	0.854
	MIT	2128	95,322	173	0.702	0.494	0.784	0.801
	Stanford	3014	147,516	270	0.718	0.532	0.784	0.808
	Columbia	3050	118,838	251	0.717	0.519	0.796	0.818
4SQ	Foursquare	83	4322	24082	0.760	0.501	0.710	0.829

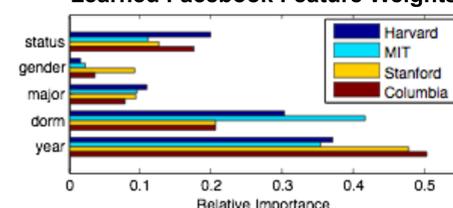
ROC Curve for Wikipedia "Graph Theory Topics"



Convergence of Stochastic SPML



Learned Facebook Feature Weights



Extensions

- To alleviate the cost of optimizing over the full \mathbf{M} matrix in high dimensional problems, we can limit the optimization to allow nonzero entries only along the diagonal of \mathbf{M}
- Alternatively, we can optimize a fixed-sized, low-rank factorization of \mathbf{M} by rewriting $\mathbf{M} = \mathbf{L} \mathbf{L}^T$ and optimizing \mathbf{L}

- We can simultaneously learn feature-dependent *degree preference functions* which adds dependency between the node feature and its structural degree (see NIPS '11 workshop talk)

Learning a Degree-Augmented Distance Metric from a Network

Bert Huang, Blake Shaw, Tony Jebara

at NIPS Workshop -- Beyond Mahalanobis: Supervised Large-Scale Learning of Similarity