

A Single New Algorithm for Many New Applications

Or Learning to Make Predictions in Networks

And Do Social Networks Really Exist?

Charles Elkan
University of California, San Diego

NetSci satellite, June 4, 2013

Research with Aditya Menon


- Predicting labels for dyadic data,
Menon, Elkan, ECML 2010. Best paper selection.
- A log-linear model with latent features for dyadic prediction,
Menon, Elkan, ICDM 2010.
- Link prediction via matrix factorization,
Menon, Elkan, ECML 2011.
- Response prediction using collaborative filtering with hierarchies
and side-information,
Menon, Chitrapura, Garg, Agarwal, Kota, KDD 2011.

Outline

- 1 Part I: Ten related prediction tasks
- 2 The LFL method
- 3 Link prediction in networks
- 4 Experiments
- 5 Part II: Do social networks really exist?

(Ten tasks) 1: Link prediction

- Given known friendship edges, predict unknown edges.



- Application: Estimate real-world connections.
- Method: Count shared neighbors, shortest paths, etc.

2: Collaborative filtering

- Given ratings of some movies by users, predict other ratings.



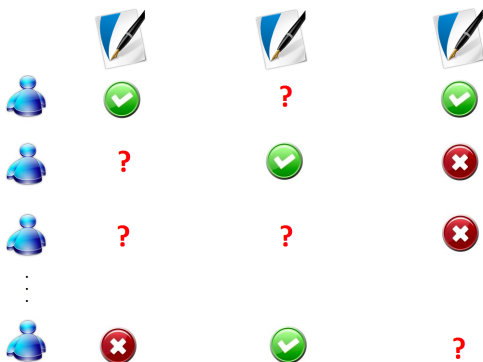
...



- Application: Netflix. Method: Factorize matrix of ratings.

3: Suggesting citations

- Each author has cited (or disliked!) certain papers. Which other papers should s/he read?




- Application: *Collaborative Topic Modeling for Recommending Scientific Articles*, Chong Wang and David Blei, KDD 2011.
- Method: A specialized graphical model.

4: Protein-protein interaction

- Experiments indicate which proteins form complexes together.

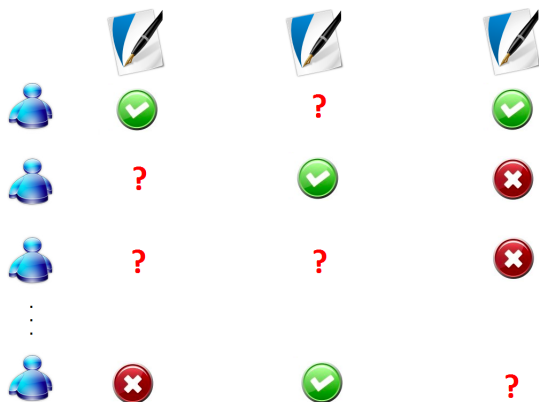
Interactome of *C. elegans* from proteinfunction.net



- Application: Augment experimental data, fix mistakes.

5: Gene-protein regulation




















- Experiments indicate which proteins switch on/off which genes.



- Application: Designing bacteria to convert waste into fuel?
- Popular method: Support vector machines (SVMs).

6: Item response theory

- Given answers by students to exam questions, predict performance on other questions.

			
			
			
			
⋮			
			

- Applications: Adaptive testing, adaptive education.
- Popular method: Latent trait models (since the 1950s).

7: Compatibility prediction for couples

- Given answers to questionnaires, predict successful dates.



- Application: eHarmony matchmaking service.
- Popular method: Learn a Mahalanobis distance metric.

8: Detecting confidentiality violations


- Thousands of employees access thousands of private records.



- Which accesses are legitimate, and which are snooping?
- Applications: Email providers, medical clinics.

9: Analyzing legal decision-making




















- In the U.S., each appeals case is assigned three judges randomly.



- How would other judges have voted? What is the probability of a different verdict?

10: Predicting behavior of shoppers


- Customer actions include { *view product, add to cart, finish purchase, write review, request refund, etc.* }

	 	 	 
			
			
⋮			
			

- New method: LFL (latent factor log linear) model.

Dyadic prediction in general

- Given labels for some pairs of entities (some **dyads**), predict labels for other dyads.



- The graph may be bipartite or not.
- Edges have any discrete set of labels. Existence is a label.
- Popular method: Depends on research community!

Latent feature models

- For simplicity, talk about **users**, **movies**.
- Associate **latent vector** values with each user and movie.
- Each **rating** is the dot-product of two latent vectors.
- **Learn** best predictive vector for each user; for each movie.



- Latent features function like explicit features.

Outline

- 1 Part I: Ten related prediction tasks
- 2 The LFL method**
- 3 Link prediction in networks
- 4 Experiments
- 5 Part II: Do social networks really exist?

What we want for dyadic prediction

- What if labels (ratings) are not numerical?
 - ▶ Link types may be { *friend, colleague, family* } etc.
- Predictions are pointless unless used to make decisions.
 - ▶ Need **probabilities** of labels e.g. $p(5 \text{ stars} | \text{user, movie})$
 - ▶ Probabilities let us make **optimal** decisions.
- What if a user has no ratings, but has **side-information**?
 - ▶ Use data from both **latent** and **explicit** features.

What's new

- Using **both** explicit and latent features.
- Allowing **any** set of labels.
- Solving a **predictive**, not descriptive, problem.
- Providing **well-calibrated** probabilities.
- Inferring accurate models from **unbalanced** data.
- Scaling to **100 million** edges.
- Unifying disparate problems in a **single framework**.

The log-linear framework

- A **log-linear** model for inputs $x \in \mathcal{X}$ and labels $y \in \mathcal{Y}$ assumes

$$p(y|x; w) = \exp \left(\sum_{i=1}^n w_i f_i(x, y) \right) / Z$$

- Predefined **feature functions** $f_i : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$.
- Trained weight vector w .
- Useful general foundation for predictive models:
 - ▶ Models **probabilities** of labels y given an example x
 - ▶ Purely **predictive**: no attempt to model x
 - ▶ Combines all feature types f_i correctly.

The LFL method

- Log-linear model with latent **and** explicit features:

$$p(y|(r, c); w) = \exp \left(\sum_{k=1}^K \alpha_{rk}^y \beta_{ck}^y + (v^y)^T s_{rc} + u_r^T V^y m_c \right) / Z$$

α_r^y and β_c^y are **latent feature** vectors for each y, r, c

- ▶ K is number of latent features

- Practical issues:


- ▶ Fix a **base label** y for identifiability.
- ▶ **Baseline** terms for each user and movie are important.
- ▶ Use L_2 regularization.
- ▶ Train with stochastic gradient descent (SGD).

Outline

- 1 Part I: Ten related prediction tasks
- 2 The LFL method
- 3 Link prediction in networks**
- 4 Experiments
- 5 Part II: Do social networks really exist?

Link prediction

- Given a **partially observed** graph, predict whether or not edges exist for dyads with **unknown** status.



Traditional methods for predicting links

- Classically, use **non-learning** functions $a(x, y)$ of dyads (x, y) :
 - ▶ Katz measure:

$$a(x, y) = \sum_{l=1}^{\infty} \beta^l \# \text{paths}(x, y, \text{length } l)$$

- ▶ Preferential attachment: $a(x, y) = \text{degree}(x) \cdot \text{degree}(y)$
- ▶ Adamic-Adar:

$$a(x, y) = \sum_{z \in \Gamma(x) \cap \Gamma(y)} \frac{1}{\log \text{degree}(z)}$$

where $\Gamma(x)$ is the set of neighbors of node x .

Latent feature approach

- The **identity** of each node influences its linking behavior.
- Identity determines values for **latent** features.
- Learning **discovers** these values.

- Nodes can also have **side-information** = explicit features.
 - ▶ For author-author linking: words written in papers, etc.
- Edges can also have side-information.
 - ▶ For country-country conflict: geographic distance, trade volume, etc.

LFL method

- LFL model for binary link prediction has parameters
 - ▶ **latent vectors** $\alpha_i \in \mathbb{R}^k$ for each node i
 - ▶ **multiplier matrix** $\Lambda \in \mathbb{R}^{k \times k}$
 - ▶ **weight matrix** $W \in \mathbb{R}^{d \times d}$ for combining node features
 - ▶ **weight vector** $v \in \mathbb{R}^{d'}$ for edge features.
- Values of parameters are learned from data.
- Including node and edge **side-information** x_i and z_{ij}

$$p(\text{edge}|i, j) = \sigma(\alpha_i^T \Lambda \alpha_j + x_i^T W x_j + v^T z_{ij})$$

where $\sigma(x) = 1/(1 + \exp(-x))$.

Challenge: Capture structures beyond clusters

- Networks contain modules that are **not** communities!
- Example: A planet with moons, also called **hub and spokes**.
 - ▶ Let the planet be node p and let a moon be node m .
 - ▶ For latent attribute i , let $\Lambda_{ii} = -1$.
 - ▶ Set $\alpha_{pi} = +1$ for the planet, $\alpha_{mi} = -1$ for each moon.
 - ▶ Then $p(\text{edge}|p, m) \gg 0$ and $p(\text{edge}|m_1, m_2) \approx 0$.
- LFL also handles **overlapping** modules, communities, etc.
- And **multiplex** networks with multiple edge types.

Challenge: Class imbalance

- Vast majority of dyads do not link with each other.
- Models trained to maximize accuracy are suboptimal.
 - ▶ **Sampling** is popular, but loses information.
 - ▶ **Weighting** is merely heuristic.
- AUC (area under ROC curve) is standard performance measure.
- For a random pair of positive and negative examples, AUC is the probability that the positive one has higher score.
 - ▶ Not influenced by relative size of positive and negative classes.

Optimizing AUC

- AUC counts concordant pairs: $\sum_{p \in +, q \in -} \mathbf{1}[f_p - f_q > 0]$.
- Train **latent features** to maximize approximation to AUC:

$$\min_{\alpha, \Lambda, W, v} \sum_{(i, j, k) \in D} \ell(p(\text{edge}|i, j) - p(\text{edge}|i, k), 1) + \Omega(\alpha, \Lambda, W, v)$$

with $D = \{(i, j, k) : \text{edge}_{ij} = 1, \text{edge}_{ik} = 0\}$.


- Using stochastic gradient descent, a fraction of one epoch is enough for convergence.

Outline

- 1 Part I: Ten related prediction tasks
- 2 The LFL method
- 3 Link prediction in networks
- 4 Experiments**
- 5 Part II: Do social networks really exist?

Link prediction with multiple classes

- The *Alywarra* dataset has multiplex kinship relations { *brother, sister, father, ...* } between 104 people.




- LFL outperforms Bayesian models, even infinite ones.
 - ▶ MMSB, IRM assume interactions set by cluster membership.
 - ▶ IBP has binary latent features.

Six diverse link prediction datasets

dataset	nodes	$ T^+ $	$ T^- $	positive:negative	mean degree
Prot-Prot	2617	23710	6,824,979	1 : 300	9.1
Metabolic	668	5564	440,660	1 : 80	8.3
NIPS	2865	9466	8,198,759	1 : 866	3.3
Condmat	14230	2392	429,232	1 : 179	0.17
Conflict	130	320	16580	1 : 52	2.5
PowerGrid	4941	13188	24,400,293	1 : 2000	2.7

- Protein-protein interactions with 76 features per protein [Noble].
- Metabolic pathways of *S. cerevisiae* from the KEGG/PATHWAY database. Three explicit feature vectors per protein: 157*D* phylogenetic information, 145*D* gene expression information, 23*D* gene location information.
- NIPS co-authorship: Each node has a 14035*D* bag-of-words feature vector per node: words used by author in publications. LSI reduces to 100*D*.
- Co-author network of condensed-matter physicists [Newman].
- International military disputes between countries [MID 3.0]. Three features per country: population, GDP and polity. Six features per dyad, e.g. geographical distance.
- U.S. electric power grid network [Watts and Strogatz].


Latent features versus non-learning methods



- LFL (green) always has highest accuracy, because each non-learning method assumes a **single specific** structure type.
- LFL **learns** to model **multiple** types, both community and not.

Learning curves

- Unsupervised scores depend on individual edges. Latent features are holistic, hence predictive with fewer known edges.
- For the military conflicts dataset:



Outline

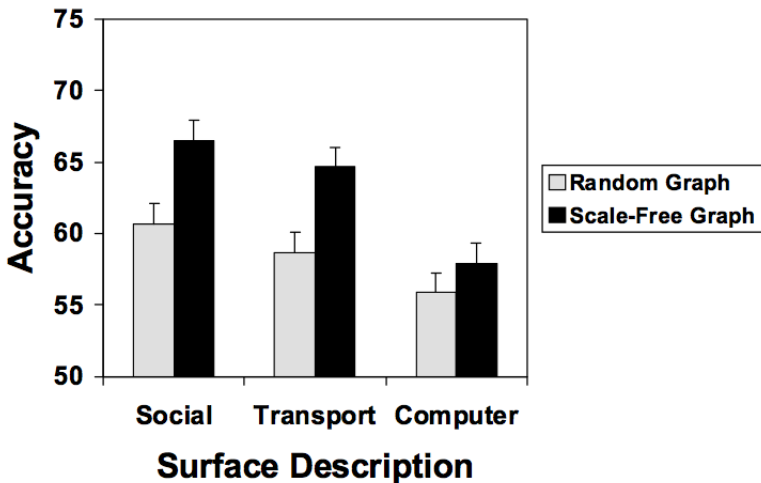
- 1 Part I: Ten related prediction tasks
- 2 The LFL method
- 3 Link prediction in networks
- 4 Experiments
- 5 Part II: Do social networks really exist?

Question: Are networks special?

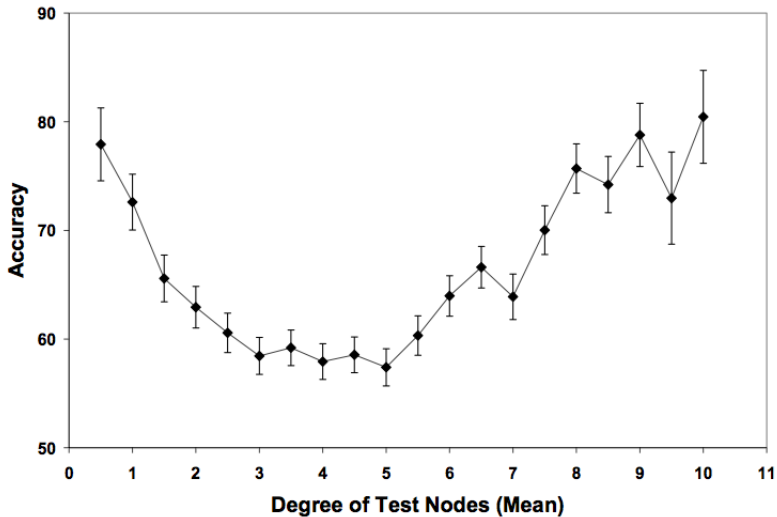
- Claim: Social networks are **not** cognitively special.
- Experimentally:
 - ▶ Humans are **bad** at learning network structures.
 - ▶ We learn social networks **no better** than non-social networks.
 - ▶ We do **not** need to know network structures explicitly.

What do humans learn?

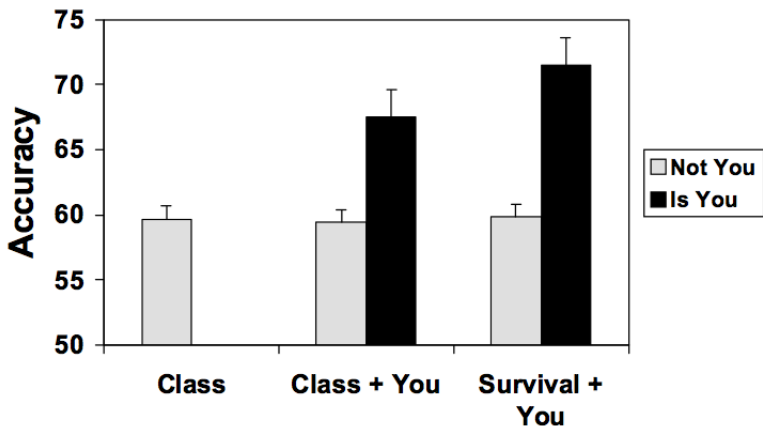
- Source: *Acquisition of Network Graph Structure* by Jason Jones, Ph.D. thesis, Dept of Psychology, UCSD, November 2011.
- My interpretation, not necessarily the author's.




- Variation is due to different memorability of node names.
- Humans learn **social** networks **no better** than other networks.



- Humans are accurate **only** on nodes with low or high degree.



- Humans learn edges involving **themselves** better than edges involving two other people.



- Humans do **not** memorize edges. Accuracy **plateaus** at low levels.

Summary of human learning

- A person learns an edge in a network well **only if**
 - ▶ the edge involves him/herself, **or**
 - ▶ one node of the edge has low **or** high degree.
- Conclusion: Humans do **not** naturally learn network structures.
- Hypothesis: Humans learn **unary** characteristics of other people:
 - ▶ whether another person is a loner or gregarious,
 - ▶ whether a person is a friend or rival **of oneself**,
 - ▶ etc.
- These unary characteristics are latent or explicit feature values.
- Matrix factorization methods also learn feature values for nodes.

Six degrees of separation

- Conclusion: Humans do **not** naturally learn specific edges in a social network.
- Observation: Humans do not need to know these edges.
- Consider Milgram's famous experiment sending letters:

References

- [Predicting labels for dyadic data](#),
Menon, Elkan, ECML 2010. Best paper selection.
- [A log-linear model with latent features for dyadic prediction](#),
Menon, Elkan, ICDM 2010.
- [Link prediction via matrix factorization](#),
Menon, Elkan, ECML 2011.
- [Response prediction using collaborative filtering with hierarchies and side-information](#),
Menon, Chitrapura, Garg, Agarwal, Kota, KDD 2011.