


Network Science of Teams: Characterization, Prediction, and Optimization

Liangyue Li, Hanghang Tong

Arizona State University

Slides can be downloaded from:
<http://www.public.asu.edu/~liangyue/team-tutorial.html>

Shift from Individuals to Teams



Teams increasingly dominate solo authors in the production of knowledge

Teams Are Everywhere

1. Film Crew



2. Sports Team



3. Sales Team



4. Research Team



5. Military Team



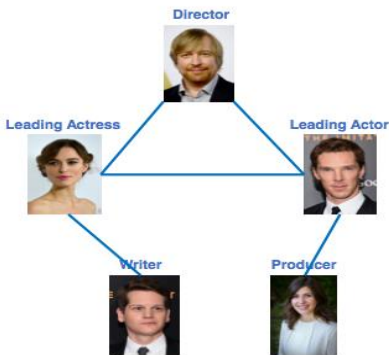
6. Development Team




Wuchty, Stefan, Ben Jones, and Brian Uzzi. "The Increasing Dominance of Teams in the Production of Knowledge," Science, May 2007, 316:1036-1039.

Networks Are Everywhere in Teams

1. Film Crew



2. Sports Team



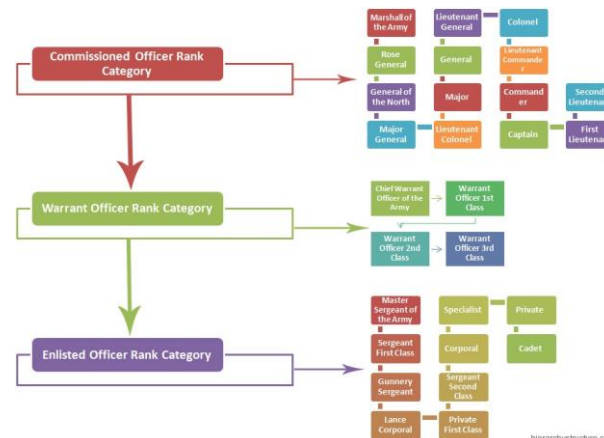
3. Sales Team



4. Research Team



5. Military Team

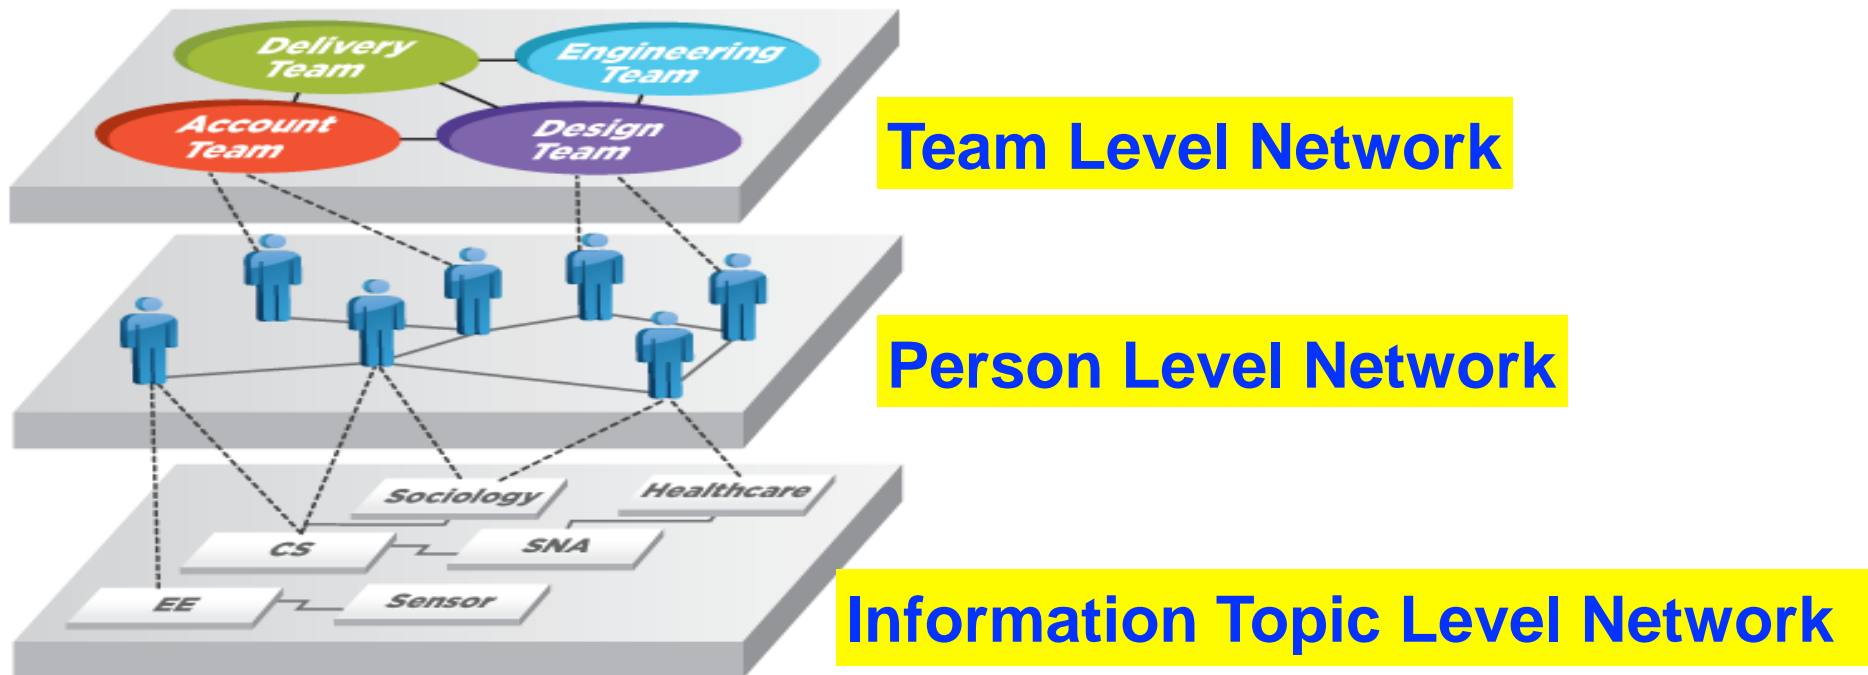


6. Development Team



Network Science of Teams

People collaborate as a team to collectively perform some complex tasks




Research Questions

- Q1: What do high-performing teams share in common? [Uzzi+Science13]
- Q2: How to foresee the success at an early stage? [Wang+Science13]
- Q3: What's the optimal design for a team in the context of networks? [Lappas+KDD09, Rangapuram+WWW13]

- S. Wuchty, B. Jones, and B. Uzzi. The Increasing Dominance of Teams in the Production of Knowledge, Science, 2007
- D. Wang, C. Song, and A.-L. Barabasi. Quantifying long-term scientific impact. Science, 342(6154): 127-132, 2013.
- T. Lappas, K. Liu, and E. Terzi. Finding a team of experts in social networks. In KDD, pages 467–476, 2009.
- S. S. Rangapuram, T. Buhler, and M. Hein. Towards realistic team formation in social networks based on densest subgraphs. WWW


Roadmap

- Motivations and Background
-  Part I: Team Performance Characterization
- Part II: Team Performance Prediction
- Part III: Team Performance Optimization
- Part IV: Open Challenges
- Demo

Part I: Team Performance Characterization

- Collective Intelligence
- Virtual Teams in online games
- Network in Sports Teams
- Network in Github Teams

Individual Intelligence



- Spearman's g
 - Individuals take a diverse set of cognitive tasks
 - The first factor extracted in a factor analysis of these scores accounts for 30% to 50% of the variance

Collective Intelligence

- Definition: general ability of the group to perform a wide variety of tasks
- Question: Is there a single factor for groups?

Study 1

- 40 groups spend five hours together in the laboratory
- Work together on a diverse set of tasks, plus a more complex criterion task
- Also measured individual intelligence

Example Tasks


Task	Description	Scoring
Generate	Brainstorming. Come up with as many uses for a brick as possible.	Scored on quantity and quality of ideas.
Choose	Intellective. Members answer a set of Raven's Matrices questions as a group.	Scored on correctness.
Negotiate	Devise a shopping trip using a shared car so that all members can get as many of their items at the best places possible.	Cumulative score of all group members.
Execute	Typing task. Members must collectively type difficult text into a shared online document.	Scored on number of words typed minus errors and skipped words in limited time period.

Study 1

- Average inter-item correlation = .28
- First principal component accounts for 43% of variance
- Factor loadings on the first factor are used to calculate *c* score – strongly predicts the performance on the criterion task
- Avg and max individual intelligence not predictive of criterion task performance

Study 2

- 152 groups ranging from 2-5 members
- Replicate findings using broader tasks



But what can predict c

- Average social perceptiveness

Playful

Comforting

Irritated

Bored



“Reading the Mind in the Eyes” Baron-Cohen et al., 2001

Woolley, Anita Williams, et al. "Evidence for a collective intelligence factor in the performance of human groups." science 330.6004 (2010): 686-688

But what can predict c

- The proportion of females positively correlate with c
 - Might be mediated by social perceptiveness
- The variance in the number of speaking turns negatively correlate with c

Virtual Teams

- Does collective intelligence exist in virtual teams where face-to-face interaction is not available?
 - Multiplayer Online Battle Arena (MOBA) teams

League of Legends



- A match is between two five-person teams
- Matchmaking algorithms vs. self-organize
- A team's goal is to destroy the opponent team's base

Study Hypotheses

- **H1:** CI will predict team performance in League of Legends
- **H2:** social perceptiveness and proportion of woman will be positively associated with CI in League of Legends teams
- **H3:** CI will not be associated with equality of contribution to conversation or decision making in LOL teams.

Method

- Data for CI, game performance, team characteristics
 - All team members individually completed a questionnaire (demographic, psychological variables, cognition, affect)
 - Test of Collective Intelligence
 - In game data (performance metrics, play history, statistics)

Sample

- Research advertisement on official community board
- 248 teams completed all components
- 97% male, avg age is 22


Results

- CI factor analysis
 - Factor analysis of scores on all tasks in TCI yielded one factor accounting for 28.28% of the variance

H1:CI and Game Performance

	MMR at Time of Study		MMR after 6 Months	
	Step 1	Step 2	Step 1	Step 2
Individual Play Time	.30***	.32***	.27***	.28***
Team Play Time	-.22***	-.22***	-.21**	-.22***
Collective Intelligence		.14*		.15*
R^2	.14	.16	.11	.14
R^2 Change		.02*		.02*

H1:CI and Game Performance




H2: Women, Social Perceptiveness and CI

- CI is positively correlated with the number of woman in the team ($r=0.18$, $p=0.005$)
- CI is positively correlated with social perceptiveness ($r=0.14$, $p=0.03$)

H3: Communication Processes and CI

- Standard deviation of chat lines and chat word count, is not significantly correlated with CI
- CI negatively correlates with
 - perceived equality in decision making,
 - frequency of game-specific communication
 - strategy-related process
 - team learning behavior

Network in Sports Teams



Flow Network:

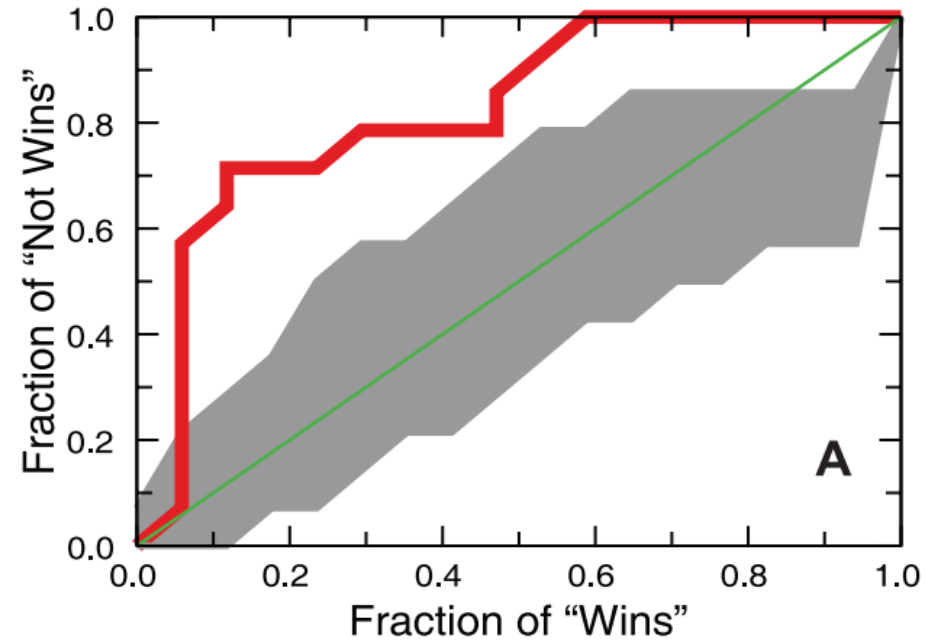
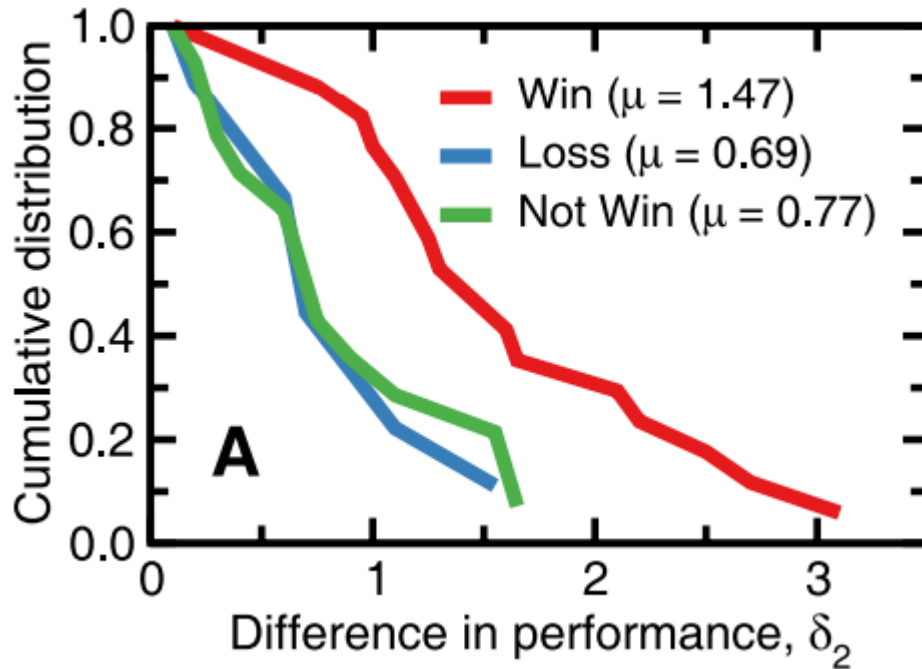
Node: players

Arc weights: passing success rate btw two players

Team Performance

- Match performance of player: normalized value of log of the player's betweenness centrality
- Team performance: avg performance of the top k players
- Difference between two teams indicate winning probability

Results




Network Structure in Github

- Network Construction
 - Project-project network: two projects are connected if they share at least one developer
 - Developer-developer network: two developers are connected if they work together in at least one project

Github Data

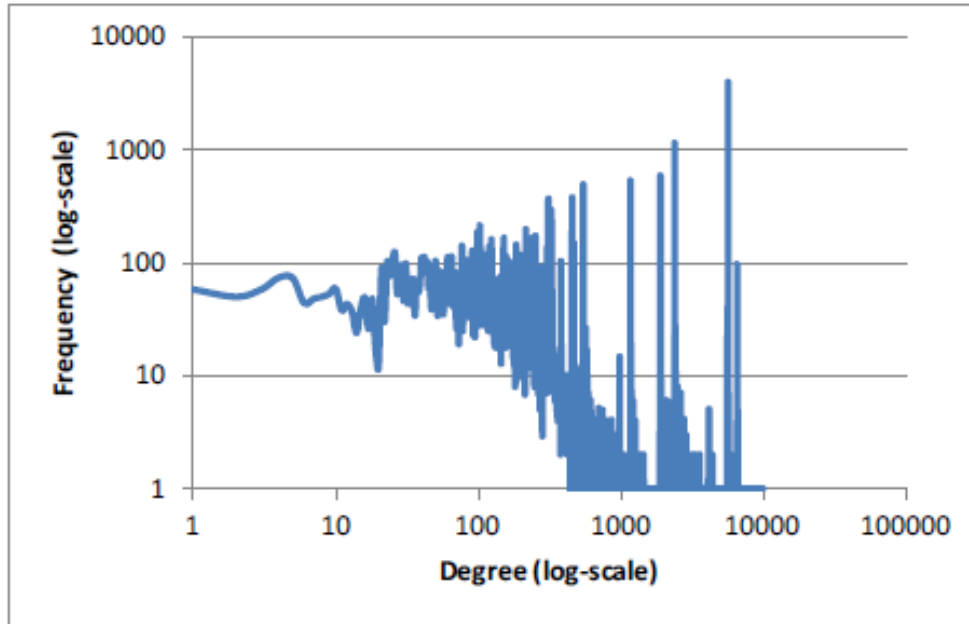
- 100,000 projects retrieved from GitHub API
- 1,161,522 edges in the project-project network
- 23,678,455 edges in the developer-developer network

Project-project network



The diameter of the largest connected component: 9
Avg shortest path: 3.7

Developer-developer network



The diameter of the largest connected component: 5

Avg shortest path: 2.47

-> compare with avg shortest path of Facebook: 4.7

Social coding enables substantially more collaborations among developers

Influential Projects

Project url	PageRank
https://github.com/mxcl/homebrew	0.0009862
https://github.com/rails/rails	0.0006378
https://github.com/lifo/docrails	0.0006370
https://github.com/joyent/node	0.0002161
https://github.com/rubinius/rubinius	0.0001678

Table I
TOP 5 MOST INFLUENTIAL PROJECTS

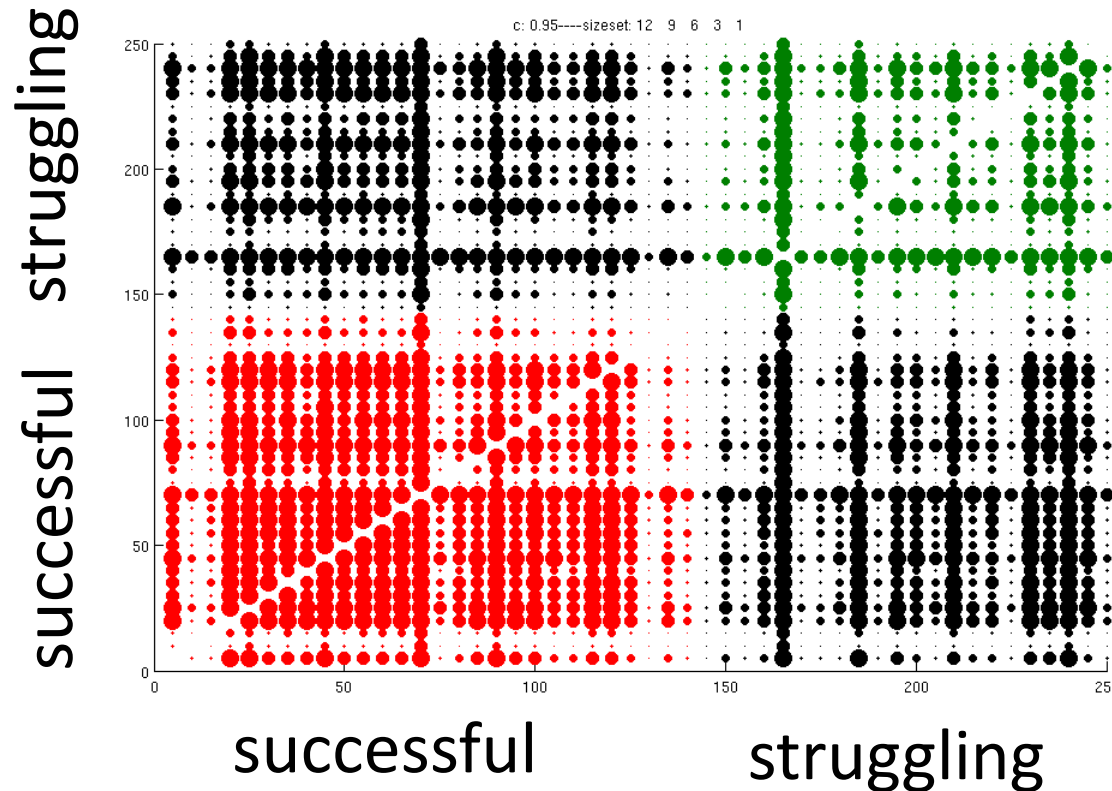
Influential Developers

Developer		PageRank
Joshua Peek	josh[AT]joshpeek.com	0.00009536
Aman Gupta	aman[AT]tmm1.net	0.00008860
Steve Richert	steve.richert[AT]gmail.com	0.00008850
Michael Klishin	michaelklishin[AT]me.com	0.00008170
Josh Kalderimis	josh.kalderimis[AT]gmail.com	0.00008163

Table II
TOP 5 MOST INFLUENTIAL DEVELOPERS

The Effect of Team Network Connectivity


Pair-wised team similarity



"Happy families resemble each other; each unhappy family is unhappy in its own way."

- Leo Tolstoy, Russian writer

Roadmap

- Motivations and Background
- Part I: Team Performance Characterization
-  Part II: Team Performance Prediction
- Part III: Team Performance Optimization
- Open Challenges
- Demo

Part II: Team Performance Prediction

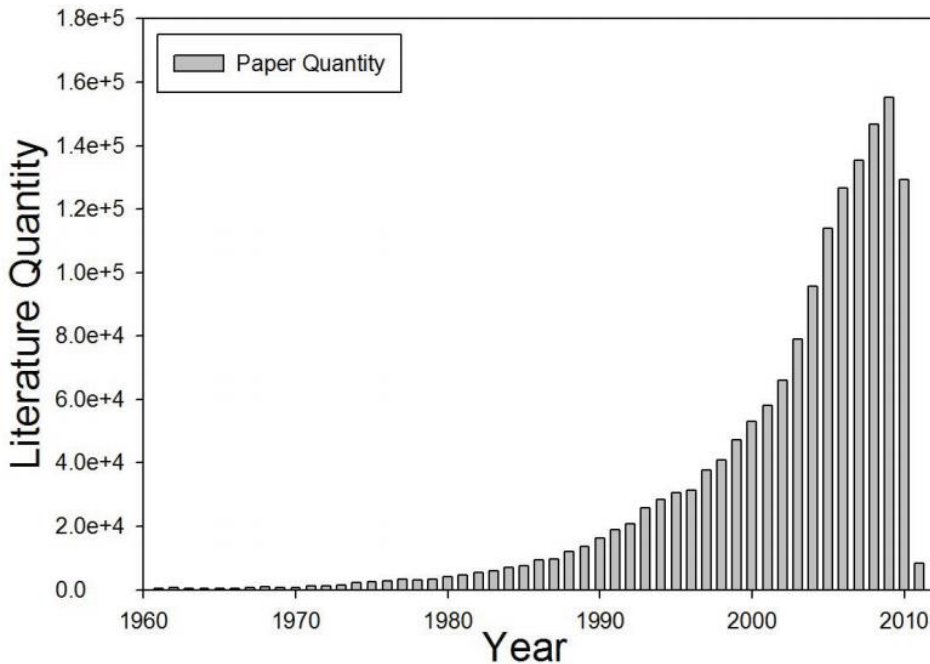
- Citation Count Prediction
- Mechanistic Model for Scientific Impact
- Long-term Performance
- Performance Trajectory
- Joint Modeling of Parts and Whole

Scientific Teams

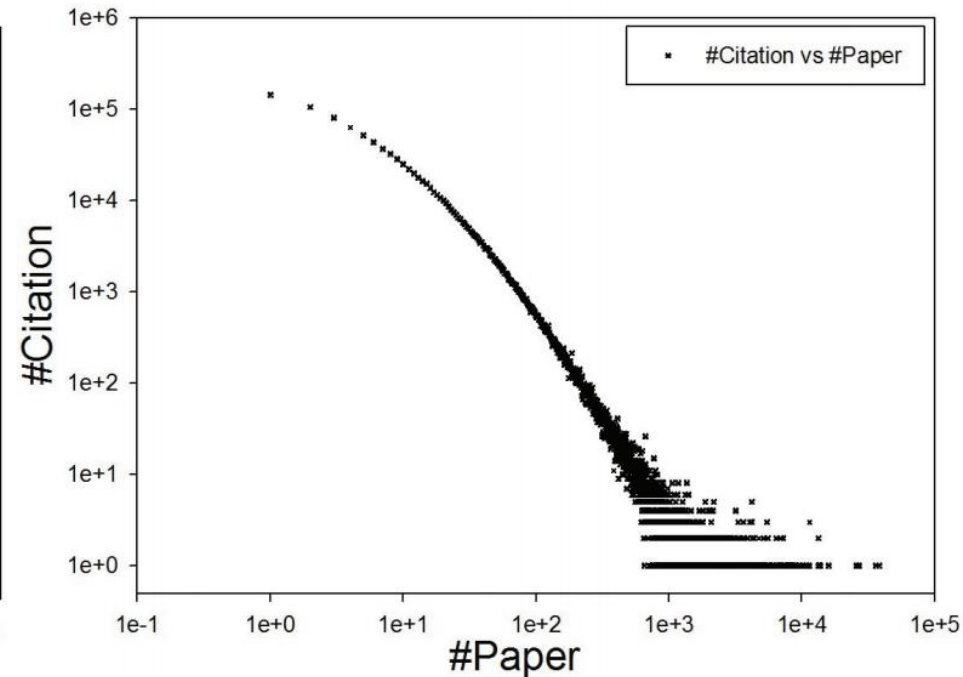


- Science of science
 - Prediction of future impact of scientific works
- Implications
 - Research grants evaluation
 - Scholarly awards dispensing

Scientific Impact



(a). The growing volume of literatures.



(b). Distribution of literature citation.

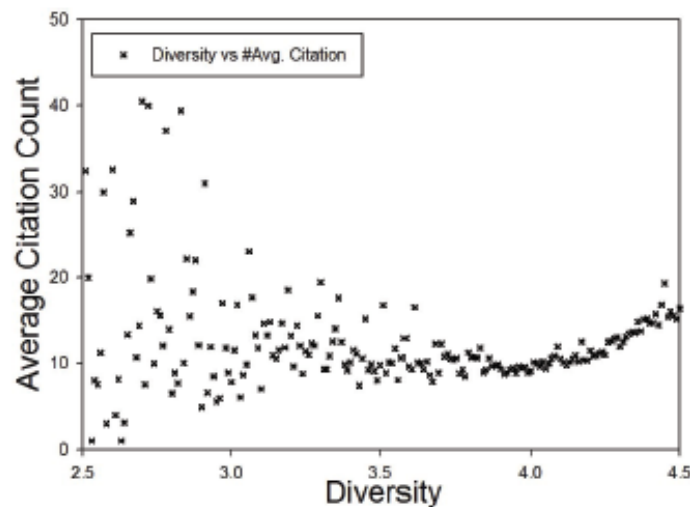
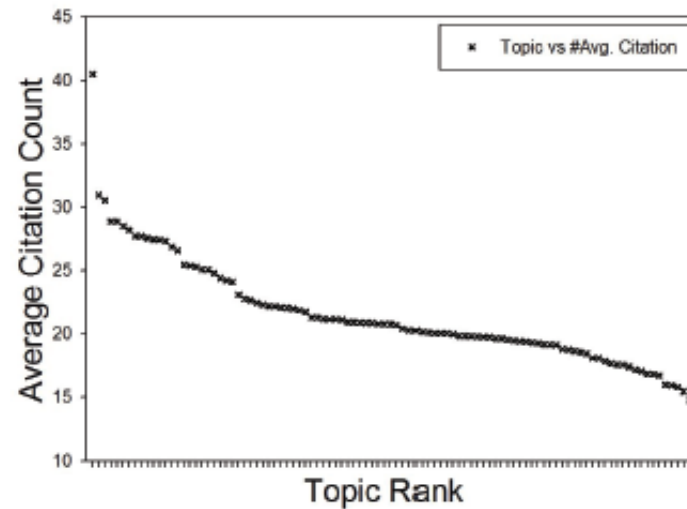
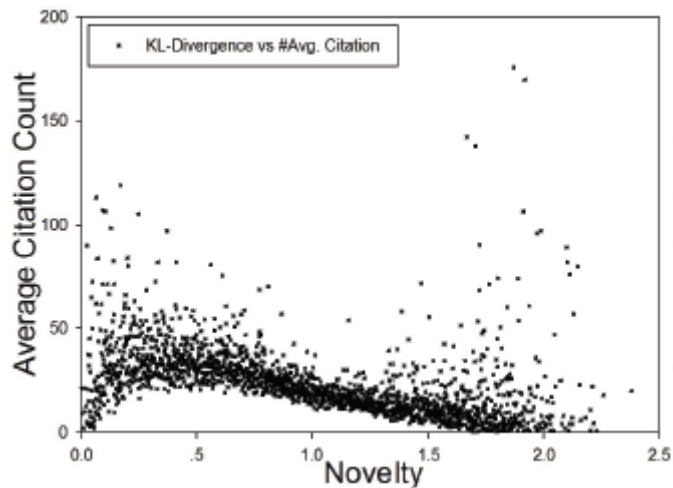
Factors driving scientific impact

- Content
- Author
 - Collaboration social network
- Venue
- Temporal

Content Features

- Novelty: difference between a particular paper and the other publications
- Topic Rank: popular topics accumulate more citation counts than unpopular ones
- Diversity: the breadth of an article from its topic distributions

Content Features



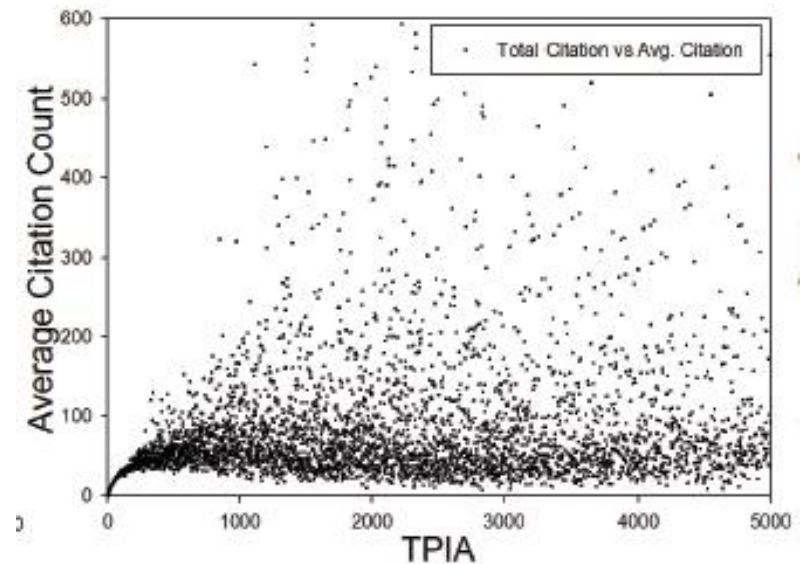
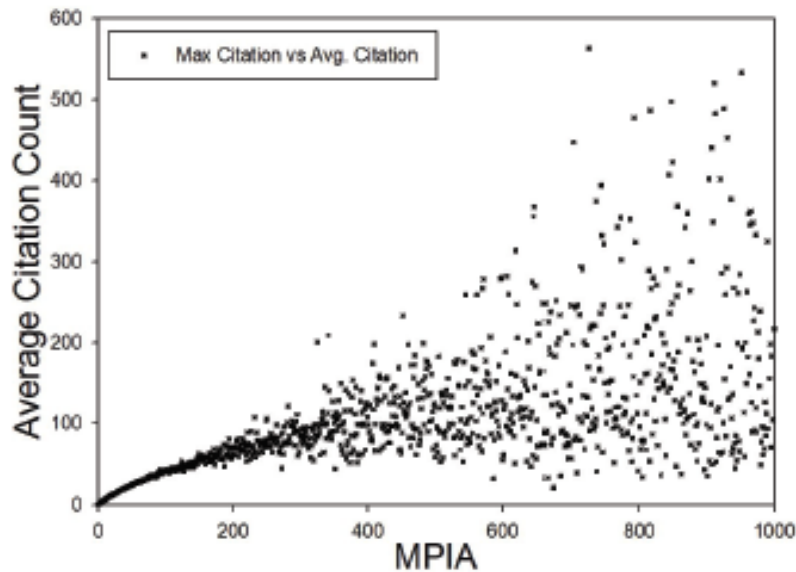
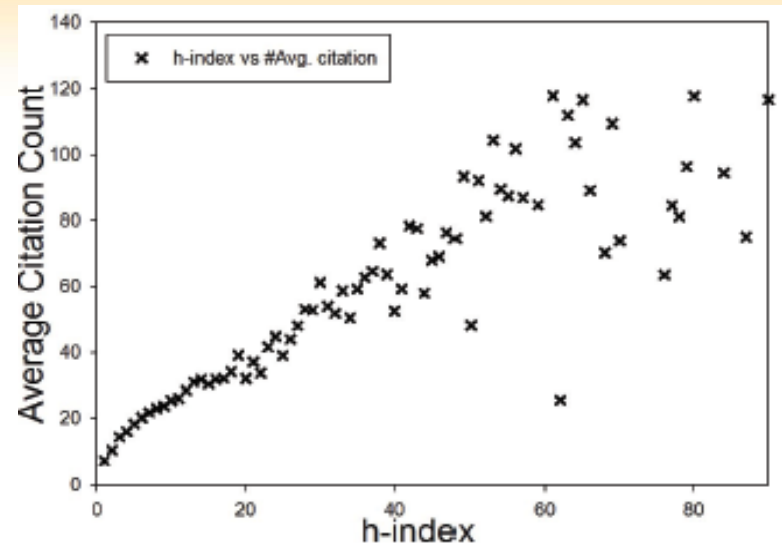

Author Features

- Author Rank: “fame” of an author ensures the amount of citations
- H-index
- Past influence of authors
 - Maximum past influence
 - Total past influence
- Productivity: the number of published papers

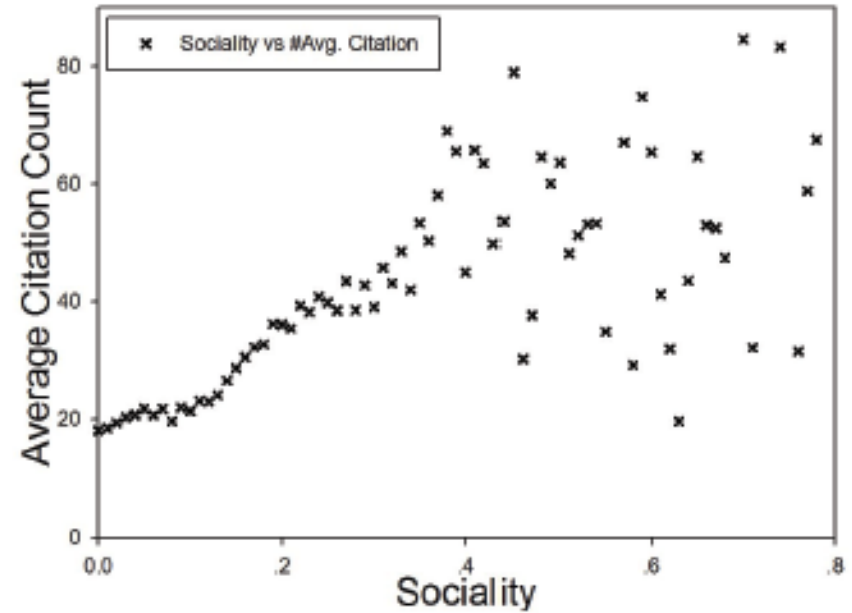
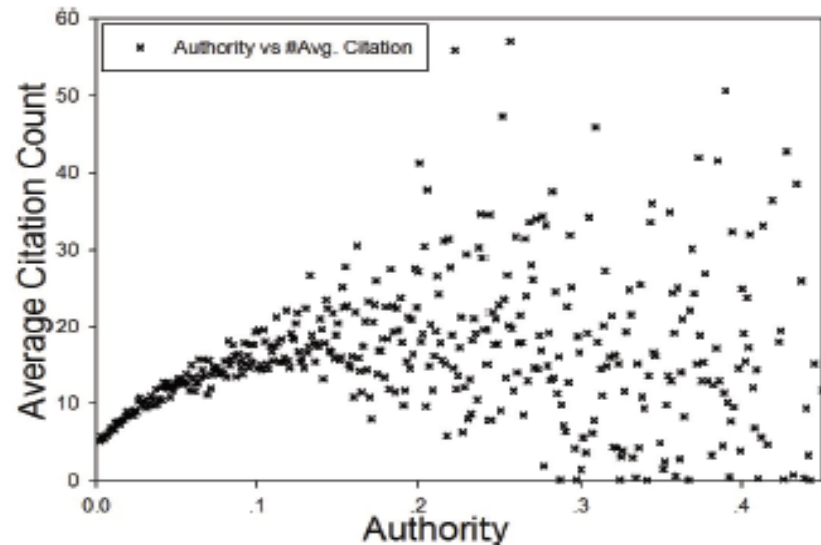
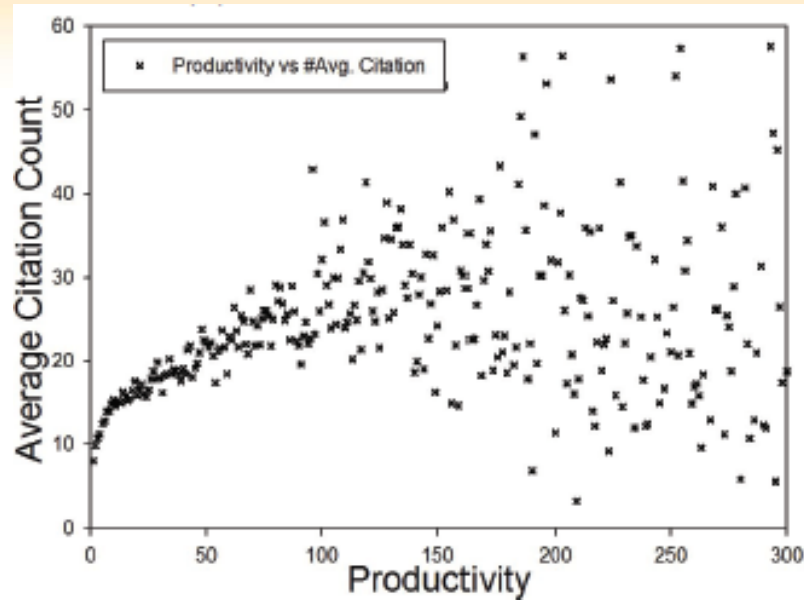
Author Features – con't

- Sociality: PageRank-like measure in co-author network
- Authority: PageRank-like measure in paper citation network and transmit paper authority to all its authors
- Versatility: topic breadth of an author's research

Author Features



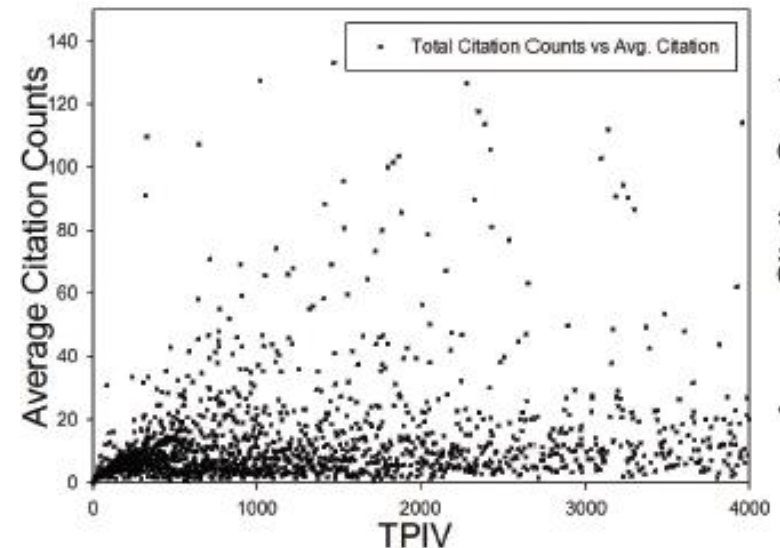
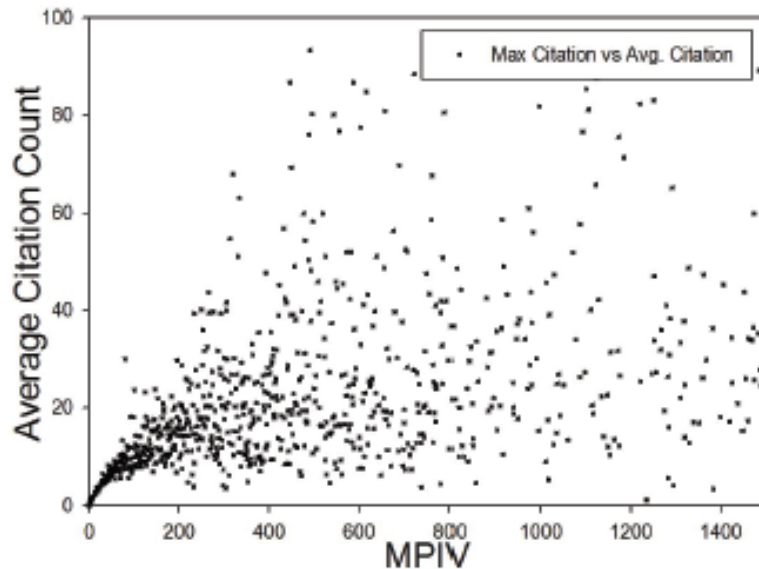
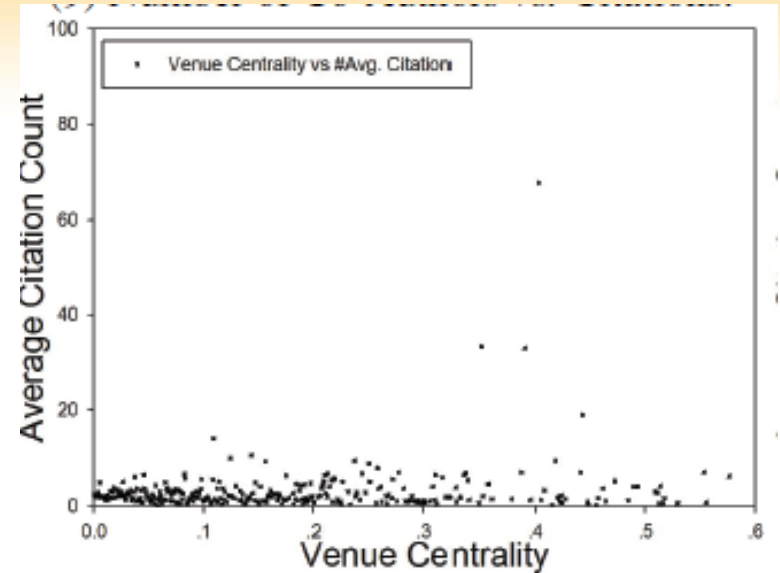

Author Features



Venue Features


- Venue Rank: prestigious venues attract more focus
- Venue Centrality: PageRank-like measure in the venue citation network
- Past Influence of venues:
 - Maximum past influence
 - Total past influence

Venue Features



Temporal Feature

- Recency: the number of years since the article was published



Data Description

- AMiner (<https://aminer.org/citation>)
 - 1,558,499 papers in CS
 - 916,946 researchers (from 1960-2010)
 - Co-author network (3,063,257 edges)
 - Citation network (20,083,947 edges)

Set-up

- Test set: 10,000 papers from year 2009
- For training, only use features available up to year 2008
- Evaluation Metric
 - Coefficient of determination R^2
 - $$R^2 = \frac{\sum(\hat{y} - \bar{y})^2}{\sum(y - \bar{y})^2}$$

Predictive Models

- kNN
- Linear Regression
- Support Vector Regression
- Classification and Regression Tree (CART)
- Gaussian Process Regression (GPR)

Performance Comparisons

	1-Year FIP ($\Delta t=1$)			5-Year FIP ($\Delta t=5$)			10-Year FIP ($\Delta t=10$)		
Methods	FData	RData	Combined	FData	RData	Combined	FData	RData	Combined
kNN	0.515	0.311	0.593	0.681	0.268	0.734	0.649	0.161	0.767
LR	0.625	0.479	0.692	0.798	0.134	0.811	0.885	0.123	0.912
SVR	0.590	0.268	0.644	0.723	0.162	0.771	0.813	0.111	0.861
CART	0.679	0.441	0.713	0.797	0.203	0.834	0.852	0.128	0.905
GPR	0.601	0.349	0.668	0.823	0.153	0.869	0.894	0.130	0.927

Accuracy increase as ∇t increases

Non-linear regression achieves better performance


GPR performs the best

Feature Analysis



	FData		RData	
Feature	+Add	−Drop	+Add	−Drop
Novelty	0.059	0.754	0.066	0.751
T.Rank	0.079	0.783	0.135	0.678
Diversity	0.157	0.661		
A.Rank	0.593	0.406	0.227	0.626
H-Index	0.244	0.611	0.186	0.663
Productivity	0.198	0.652	0.187	0.684
MPIA	0.585	0.419	0.363	0.596
TPIA	0.048	0.805	0.037	0.811
NOCA	0.056	0.794	0.158	0.643
Sociality	0.249	0.597	0.181	0.632
Authority	0.155	0.668	0.178	0.615
Versatility	0.160	0.649	0.139	0.665
Recency	0.101	0.738		
V.Rank	0.337	0.603	0.225	0.648
V.Centrality	0.049	0.793	0.067	0.776
MPIV	0.329	0.616	0.196	0.667
TPIV	0.023	0.815	0.021	0.823


Lack of predictability in citation patterns



Citation history of 463,348 papers extracted from the *Physical Review* corpus

Preferential attachment


- Highly cited papers are more likely to be cited again



Temporal Citation Trend

- Long-term decay follows a log-normal survival probability

$$P_i(t) = \frac{1}{\sqrt{2\pi}\sigma_i t} \exp \left[-\frac{(\ln t - \mu_i)^2}{2\sigma_i^2} \right]$$



Fitness η of a paper

- The paper's importance relative to its peers

Mechanistic Model

- The probability that paper i is cited at time t after publication is

$$\Pi_i(t) \sim \eta_i c_i^t P_i(t)$$

- Solving for the cumulative number of citations acquired by paper i at time t

$$c_i^t = m \left[e^{\frac{\beta \eta_i}{A} \Phi\left(\frac{\ln t - \mu_i}{\sigma_i}\right)} - 1 \right] \equiv m \left[e^{\lambda_i \Phi\left(\frac{\ln t - \mu_i}{\sigma_i}\right)} - 1 \right] \quad (3)$$





$$\begin{aligned} \tilde{c} &= \Phi(\tilde{t}) \\ \tilde{c} &= \frac{\ln\left(1 + \frac{c_i^t}{m}\right)}{\lambda_i} \\ \tilde{t} &= (\ln t - \mu_i)/\sigma_i \end{aligned}$$



where

$$\Phi(x) \equiv (2\pi)^{-1/2} \int_{-\infty}^x e^{-y^2/2} dy \quad (4)$$

Model's validity



Predicting future impact




Measure the fraction of papers that fall within the envelope for all PR papers published in 1960

With $T_{\text{train}}=5$, 6.5% left the envelope 30 years later

Performance Prediction: Setup

- **Given:** Initial Performance of a team
- **Predict:**
 - (1) Long-Term Performance [KDD15]
 - (2) Performance Trajectory [SDM16]




- L. Li, and H. Tong: The Child is Father of the Man: Foresee the Success at the Early Stage. KDD 2015: 655-664
- L. Li, H. Tong, J. Tang and W. Fan: "iPath: Forecasting the Pathway to Impact". SDM 2016

Performance Prediction: Challenges


- C1: Scholarly feature design
- C2: Non-linearity
- C3: Domain heterogeneity
- C4: Dynamics

C1: Scholarly Feature Design




Obs.: Adding content features brings little improvement

C2: Non-linearity



Obs.: Non-linear methods > linear ones


C3: Domain heterogeneity



Obs.: Impact of scientific work from different domains behaves differently


C4: Dynamics

arXiv monthly submission rates



Q: How to quickly update the predictive model?

iBall — Formulations



■ Optimization Formulation

Within-Domain Model

$$\min_{\mathbf{w}^{(i)}, i=1, \dots, n_d} \sum_{i=1}^{n_d} \mathcal{L}[f(\mathbf{X}^{(i)}, \mathbf{w}^{(i)}), \mathbf{Y}^{(i)}] + \lambda \sum_{i=1}^{n_d} \Omega(\mathbf{w}^{(i)})$$
$$+ \theta \sum_{i=1}^{n_d} \sum_{j=1}^{n_d} \underbrace{\mathbf{A}_{ij} g(\mathbf{w}^{(i)}, \mathbf{w}^{(j)})}_{\text{Cross-Domain Consistency}}$$


Cross-Domain Consistency

■ Remarks

- **Within-Domain Model**: regression/classification, linear/non-linear
- **Cross-Domain Consistency**: similar domains have similar models

Question: how to instantiate such consistency?

iBall — linear formulation




Details:

$$\min_{\mathbf{w}^{(i)}, i=1, \dots, n_d} \sum_{i=1}^{n_d} \|\mathbf{X}^{(i)} \mathbf{w}^{(i)} - \mathbf{Y}^{(i)}\|_2^2 + \lambda \sum_{i=1}^{n_d} \|\mathbf{w}^{(i)}\|_2^2 + \theta \sum_{i=1}^{n_d} \sum_{j=1}^{n_d} \underbrace{\mathbf{A}_{ij}}_{\text{similarity}} \|\mathbf{w}^{(i)} - \mathbf{w}^{(j)}\|_2^2$$

Intuitions: similar domains (large \mathbf{A}_{ij})

→ same feature has similar effect (small $\|\mathbf{w}^{(i)} - \mathbf{w}^{(j)}\|_2^2$)

iBall — non-linear formulation



Details:

$$\min_{\mathbf{w}^{(i)}, i=1, \dots, n_d} \sum_{i=1}^{n_d} \|\mathbf{K}^{(i)} \mathbf{w}^{(i)} - \mathbf{Y}^{(i)}\|_2^2 + \lambda \sum_{i=1}^{n_d} \mathbf{w}^{(i)'} \mathbf{K}^{(i)} \mathbf{w}^{(i)} + \theta \sum_{i=1}^{n_d} \sum_{j=1}^{n_d} \mathbf{A}_{ij} \|\underbrace{\mathbf{K}^{(i)} \mathbf{w}^{(i)}}_{\text{Predicted output (domain } i \rightarrow \text{domain } i)}} - \underbrace{\mathbf{K}^{(ij)} \mathbf{w}^{(j)}}_{\text{Predicted output (domain } j \rightarrow \text{domain } i)}\|_2^2$$

Predicted output
(domain $i \rightarrow$ domain i)

Predicted output
(domain $j \rightarrow$ domain i)

Intuitions: similar domains (large \mathbf{A}_{ij})

→ similar predicted outputs (small $\|\mathbf{K}^{(i)} \mathbf{w}^{(i)} - \mathbf{K}^{(ij)} \mathbf{w}^{(j)}\|_2^2$)

iBall — Closed-form Solutions

- Closed-form Solution

$$\mathbf{w} = \mathbf{S}^{-1} \mathbf{Y}$$

- iBall — linear:

$$\mathbf{w} = [\mathbf{w}^{(1)}; \dots; \mathbf{w}^{(k)}] \quad \mathbf{Y} = [\mathbf{X}^{(1)'} \mathbf{Y}^{(1)}; \dots; \mathbf{X}^{(k)'} \mathbf{Y}^{(k)}]$$

i-th block column

j-th block column

$$\mathbf{S} = \begin{bmatrix} \dots & \dots & \dots \\ \dots & \mathbf{X}^{(i)'} \mathbf{X}^{(i)} + (\theta \sum_{j=1}^k \mathbf{A}_{ij} + \lambda) \mathbf{I} & -\theta \mathbf{A}_{ij} \mathbf{I} \\ \dots & \dots & \dots \end{bmatrix}$$

i-th block row

Time Complexity: $O((dk)^3)$

d : # of features; k : # of domains
(dk : in the order of 10 or 100)



iBall — Closed-form Solutions

- Closed-form Solution

$$\mathbf{w} = \mathbf{S}^{-1} \mathbf{Y}$$

- iBall — non-linear:

$$\mathbf{w} = [\mathbf{w}^{(1)}; \dots; \mathbf{w}^{(k)}] \quad \mathbf{Y} = [\mathbf{Y}^{(1)}; \dots; \mathbf{Y}^{(k)}]$$

i-th block column

j-th block column

$$\mathbf{S} = \begin{bmatrix} \dots & \dots & \dots \\ \dots & (1 + \theta \sum_{j=1}^k \mathbf{A}_{ij}) \mathbf{K}^{(i)} + \lambda \mathbf{I} & -\theta \mathbf{A}_{ij} \mathbf{K}^{(ij)} \\ \dots & \dots & \dots \end{bmatrix}$$

i-th block row

Time Complexity: $O(n^3)$

n : total # of training examples
(in the order of millions)



iBall — Scale-up with Dynamic Update

- **Key idea #1:** Approx **S** by low-rank approx
- **Details:**

$$\begin{aligned} \mathbf{S}_{t+1} \approx \mathbf{U}_{t+1} \mathbf{\Lambda}_{t+1} \mathbf{U}_{t+1}' &\xrightarrow{\mathbf{w}_{t+1}} \mathbf{S}_{t+1}^{-1} \mathbf{Y}_{t+1} \\ &= \mathbf{U}_{t+1} \mathbf{\Lambda}_{t+1}^{-1} \mathbf{U}_{t+1}' \mathbf{Y}_{t+1} \end{aligned}$$

(Overall: $O(n^2 r)$) (Overall: $O(nr)$)

- **Complexity:** $O(n^3) \rightarrow O(n^2 r + nr)$
- **Benefit:** avoid matrix inverse


Question: how to avoid re-computing low-rank approx at each time step?

iBall — Scale-up with Dynamic Update

- **Key idea #2:** Incrementally update the low rank structure of \mathbf{S}


- **Details:**

white: zeros
blue: old at t
pink: new at $t+1$




\mathbf{S}_{t+1}

=



$\tilde{\mathbf{S}}_t$

+




$\Delta \mathbf{S}$

(low rank, sparse)

- **Complexity:** $O(n^2 r) \rightarrow O((n + m)(r^2 + r'^2)), r \ll n$
- **Benefit:** avoid re-computing low-rank approx



Paper Citation Prediction Performance



Datasets: AMiner (2,243,976 papers, 1,274,360 authors, 8,882 venues)




Error Analysis



Obs.: bright region at $x = y$


- L. Li, and H. Tong: The Child is Father of the Man: Foresee the Success at the Early Stage. KDD 2015

Running Time Comparison




Obs.: iBall-fast outperforms other non-linear methods

Quality vs. Speed



Obs.: iBall-fast: good trade-off between quality and speed

iBall: Summary



- **Goal:** predict long-term impact of scholarly entities
- **Solutions:** joint predictive model (**iBall**)

<i>Challenges</i>	C1 <i>feature design</i>	C2 <i>non-linearity</i>	C3 <i>domain-heterogeneity</i>	C4 <i>dynamics</i>
<i>Tactics</i>	<i>first 3 years' citation</i>	<i>kernel trick</i>	<i>domain consistency</i>	<i>low-rank approximation</i>

- **Results:**
 - iBall joint models better than separate versions
 - iBall-fast updates efficiently and accurately

Foresee the Pathway to Impact



Geoffrey Hinton

Learning representations by back-propagating errors

Authors David E Rumelhart, Geoffrey E Hinton, Ronald J Williams


Publication date 1986

Journal Nature

Volume 323

Pages 533-536

Predict



Implications of forecasting the pathway to impact

- Tracking research frontier
- Invoking early intervention

Question: how to foresee the impact pathway at the early stage?

Modeling Scientific Impact

- **Effective scholarly feature design**
[Yan+CIKM11]
- **Mechanistic model for the citation dynamics of individual papers**
[Wang+Science13]
- ***iBall*- Joint Predictive Model for long-term impact prediction** [Li+KDD15]

All for Point Prediction




Challenges

- C1: Output Space -- Correlation

- Possible solution: multi-label/task learning

- Challenge: correlation unknown



- C2: Parameter Space -- Smoothness


- Possible solution: linear dynamic system

- Challenge: transition process unknown

Design Objectives

- D1: Prediction Consistency (for C1)


- **Exploit** the correlation in output space
- **Infer** the impact relation structure



- D2: Parameter Smoothness (for C2)

- **Apply** linear transition to adjacent parameters
- **Learn** the linear transition process

iPath -- Formulations




■ Optimization Formulations

$$\begin{aligned}
 \min_{\mathbf{W}, \mathbf{B}, \mathbf{A}} & \left[\mathcal{L}[f(\mathbf{X}, \mathbf{W}), \mathbf{Y}] \right] + \left[\alpha \sum_{i=1}^l \sum_{j=1}^l \mathbf{A}_{ij} g(\mathbf{w}_i, \mathbf{w}_j) \right] \\
 & + \left[\beta \sum_{t=2}^l \|\mathbf{w}_t - \mathbf{B} \mathbf{w}_{t-1}\|_2^2 \right] \text{Parameter Smoothness} \\
 & + \underbrace{\gamma \|\mathbf{B} - \mathbf{I}\|_F^2 + \delta \sum_{i=1}^l \Omega(\mathbf{w}_i) + \epsilon \|\mathbf{A} - \mathbf{A}_0\|_F^2}_{\text{Regularizations}}
 \end{aligned}$$

■ Remarks

- **Prediction Consistency**: similar impacts have similar models
- **Parameter Smoothness**: model parameters at adjacent time steps have linear transformation

iPath – linear formulation



■ Details:

$$\min_{\mathbf{W}, \mathbf{B}, \mathbf{A}} \|\mathbf{X}\mathbf{W} - \mathbf{Y}\|_F^2 + \alpha \sum_{i=1}^l \sum_{j=1}^l \mathbf{A}_{ij} \|\mathbf{X}\mathbf{w}_i - \mathbf{X}\mathbf{w}_j\|_2^2$$

$$+ \beta \sum_{t=2}^l \|\mathbf{w}_t - \mathbf{B}\mathbf{w}_{t-1}\|_2^2 + \gamma \|\mathbf{B} - \mathbf{I}\|_F^2$$


$$+ \delta \sum_{i=1}^l \|\mathbf{w}_i\|_2^2 + \epsilon \|\mathbf{A} - \mathbf{A}_0\|_F^2$$

■ Intuition:

Similar impacts (large \mathbf{A}_{ij})

➔ Similar Predictions (small $\|\mathbf{X}\mathbf{w}_i - \mathbf{X}\mathbf{w}_j\|_2^2$)

iPath – non-linear formulation



■ Details:

$$\min_{\mathbf{W}, \mathbf{B}, \mathbf{A}} \|\mathbf{KW} - \mathbf{Y}\|_F^2 + \alpha \sum_{i=1}^l \sum_{j=1}^l \boxed{\mathbf{A}_{ij} \|\mathbf{Kw}_i - \mathbf{Kw}_j\|_2^2} \\ + \beta \sum_{t=2}^l \|\mathbf{w}_t - \mathbf{Bw}_{t-1}\|_2^2 + \gamma \|\mathbf{B} - \mathbf{I}\|_F^2 \\ + \delta \sum_{i=1}^l \mathbf{w}_i' \mathbf{Kw}_i + \epsilon \|\mathbf{A} - \mathbf{A}_0\|_F^2$$

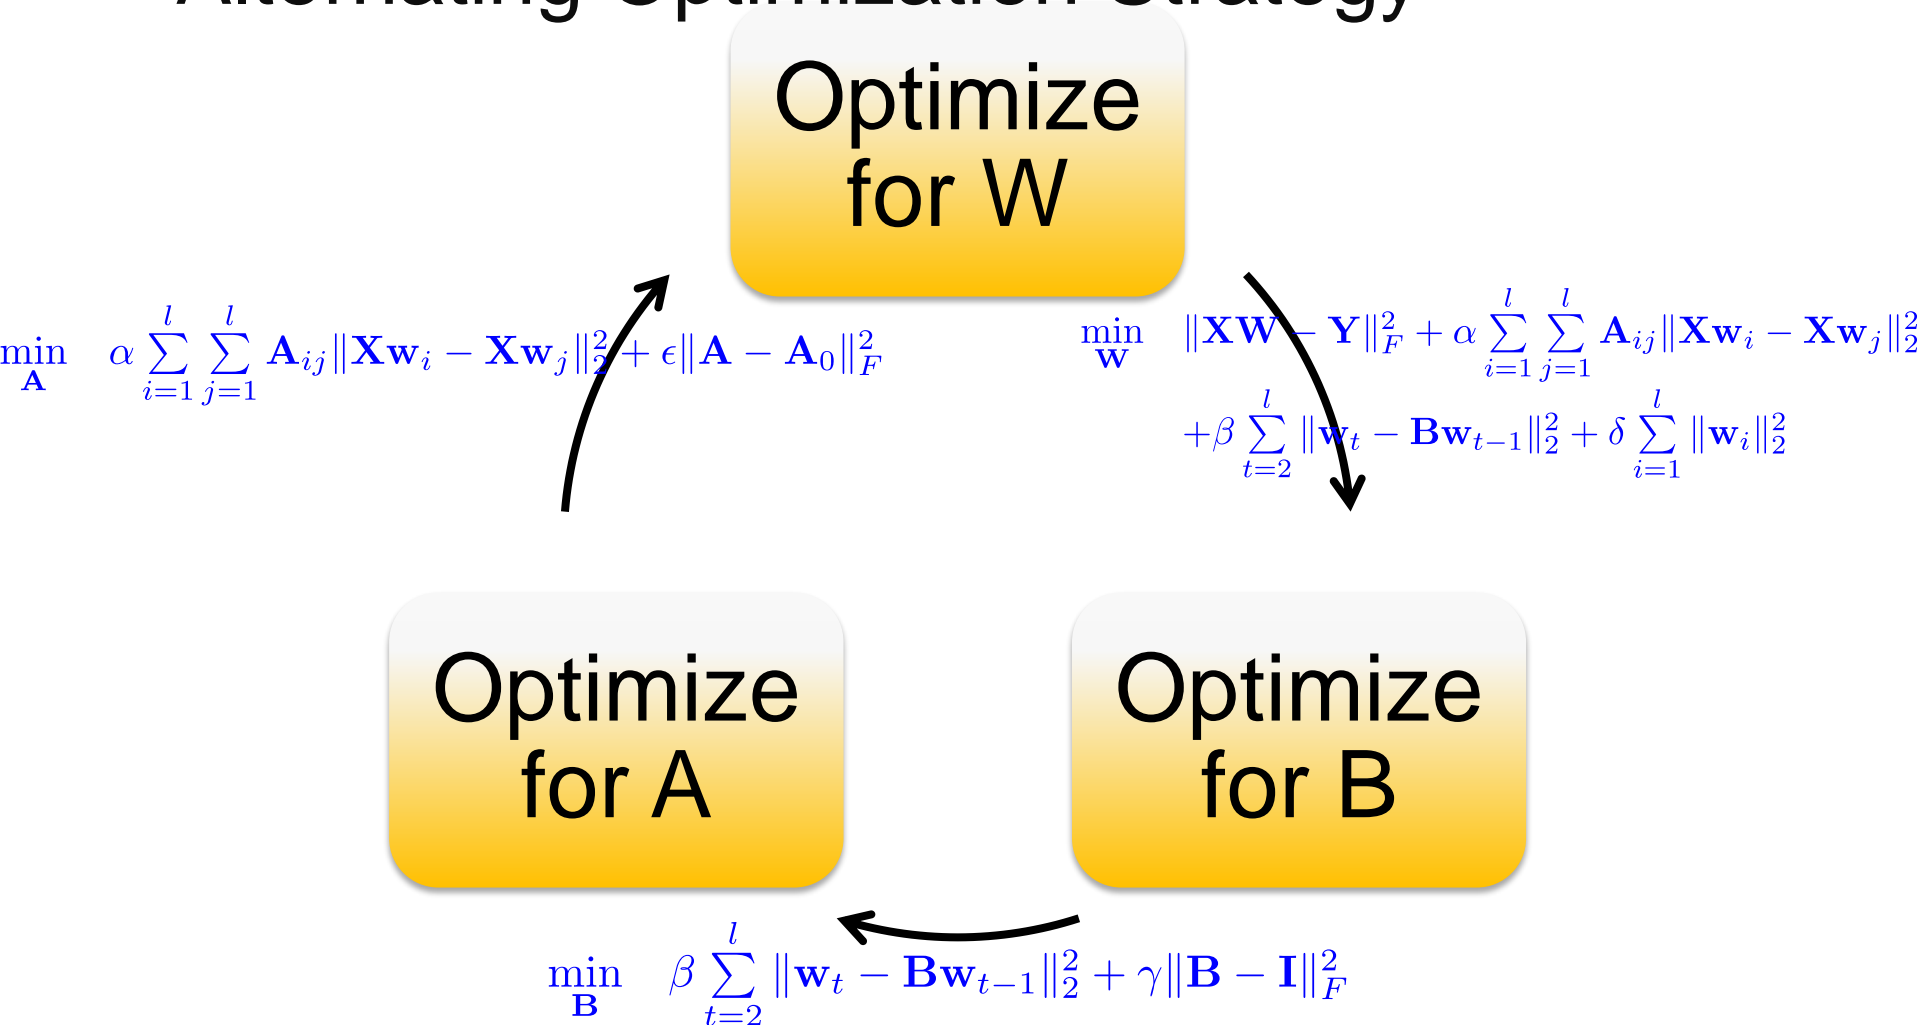
■ Intuition:

Similar Impacts (large \mathbf{A}_{ij})

➡ Similar Predictions (small $\|\mathbf{Kw}_i - \mathbf{Kw}_j\|_2^2$)

iPath – Optimization Solutions


- Alternating Optimization Strategy



Experiment Setup


- **Datasets:** *AMiner* (2,243,976 papers, 1,274,360 authors, 8,882 venues)
- **Task:** Observing the first 5 years' citations, predict yearly citations from year 6 – 15
- **Evaluation Metric:** Root Mean Squared Error (RMSE)

Paper Impact Pathway Prediction




Obs: iPath-ker performs the best among all the competitors

Author Impact Pathway Prediction




Obs: iPath-ker performs the best among all the competitors

Sensitivity Analysis



RMSE vs. α



RMSE vs. β

Obs: iPath is stable in a large range of parameter spaces

Performance Gain Analysis

RMSE	Paper Impact	Author Impact
	9.602	11.608
① + ②	9.507	11.548
① + ② + ③	9.335	11.489
① + ② + ③ + ④	9.171	11.391

$$\begin{aligned}
 \min_{\mathbf{W}, \mathbf{B}, \mathbf{A}} & \left[\textcircled{1} \text{Basic form} \right] + \left[\alpha \sum_{i=1}^l \sum_{j=1}^l \textcircled{2} \text{relation} \mathbf{A}_{ij} \|\mathbf{K} \mathbf{w}_i - \mathbf{K} \mathbf{w}_j\|_2^2 \right] \\
 & + \left[\beta \sum_{t=2}^l \textcircled{3} \text{transition} \|\mathbf{w}_t - \mathbf{B} \mathbf{w}_{t-1}\|_2^2 \right] + \gamma \|\mathbf{B} - \mathbf{I}\|_F^2 \\
 & + \delta \sum_{i=1}^l \mathbf{w}_i' \mathbf{K} \mathbf{w}_i + \left[\epsilon \|\mathbf{A} - \mathbf{A}_0\|_F^2 \right] \textcircled{4} \text{inferring}
 \end{aligned}$$


Obs: relation, transition and inferring are all beneficial in improving the prediction

Robustness to Noise in A_0

$$\min_{\mathbf{W}, \mathbf{B}, \mathbf{A}} \|\mathbf{X}\mathbf{W} - \mathbf{Y}\|_F^2 + \alpha \sum_{i=1}^l \sum_{j=1}^l \mathbf{A}_{ij} \|\mathbf{X}\mathbf{w}_i - \mathbf{X}\mathbf{w}_j\|_2^2$$


$$+ \beta \sum_{t=2}^l \|\mathbf{w}_t - \mathbf{B}\mathbf{w}_{t-1}\|_2^2 + \gamma \|\mathbf{B} - \mathbf{I}\|_F^2$$

$$+ \delta \sum_{i=1}^l \|\mathbf{w}_i\|_2^2 + \epsilon \|\mathbf{A} - \mathbf{A}_0\|_F^2$$





Obs: iPath degenerates gradually with the noise level

iPath: Summary




- **Goals:** predict the pathway to impact
- **Solutions:** *i*Path prediction model
 - Design objectives:
 - Prediction Consistency →
 - Parameter Smoothness →
 - Results:
 - Lower error than competitors
 - Robust to noise in impact relations



From the Ancient Philosophy

"The whole is greater
than the sum of its parts."
-Aristotle



The whole is greater than the sum of its parts. -- Aristotle

- **Whole:** a collection of parts
- **Parts:** individual elements
- **Aristotle's hypothesis:**
 - whole > sum of parts


Part-Whole in Team Science



Research Team



Sports Team



Film Crew



Sales Team

Whole – Team
Parts – Team members

Part-Whole Beyond Teams



Autonomous System

Whole: system

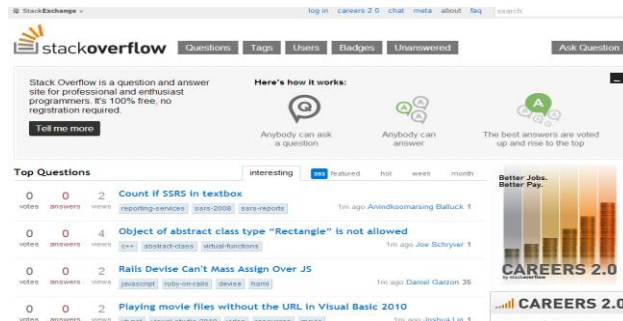
Parts: individual drones



Stock Market

Whole: DJIA

Parts: individual stock




Community Question Answering

Whole: question

Parts: individual answers

Parts of a passenger jet airplane



© 2010 Encyclopedia Britannica, Inc.

System Reliability

Whole: system

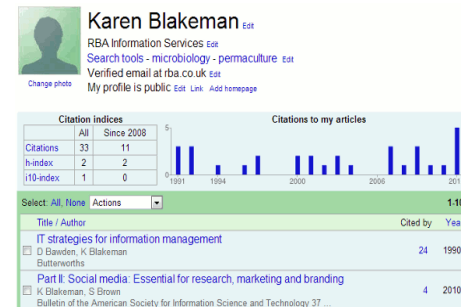
Parts: individual component

Outcome of Part-Whole



Whole: Team
Part: Members

Whole outcome: Team's performance
Part outcome: each member's performance



Whole: Researcher
Part: Publications

Whole outcome: h-index
Part outcome: #citations of publications

Question: how can we predict the outcome of whole/parts?




Predict the Part-Whole Outcomes

- Existing Algorithmic Work
 - Separate models for parts and whole
 - Joint **linear** models
- Aristotle's hypothesis: $\text{whole} > \text{sum}(\text{parts})$
- Question: how to jointly predict part/whole
 - by leveraging the part-whole relationship *beyond* the linear models?

Challenges -- Modeling

- **Nonlinear** Part-whole Relationship


- **Max**: impact of a question is strongly correlated with that of the *best* answer



- **Min**: classic Wooden Bucket Theory
- **Sparsity**: team performance often dominated by a *few top-performing* team members

Challenges – Modeling (con't)

- Part-part Interdependency
 - Parts are connected via underlying network
 - Impact of such interdependency on outcomes
- Hypothesis-1: similar parts -> similar contribution to whole
- Hypothesis-2: similar parts -> similar parts outcome



Question: how can we utilize

1. nonlinear part-whole relationship
2. part-part interdependency

Challenges -- Algorithm

Non-linearity

+

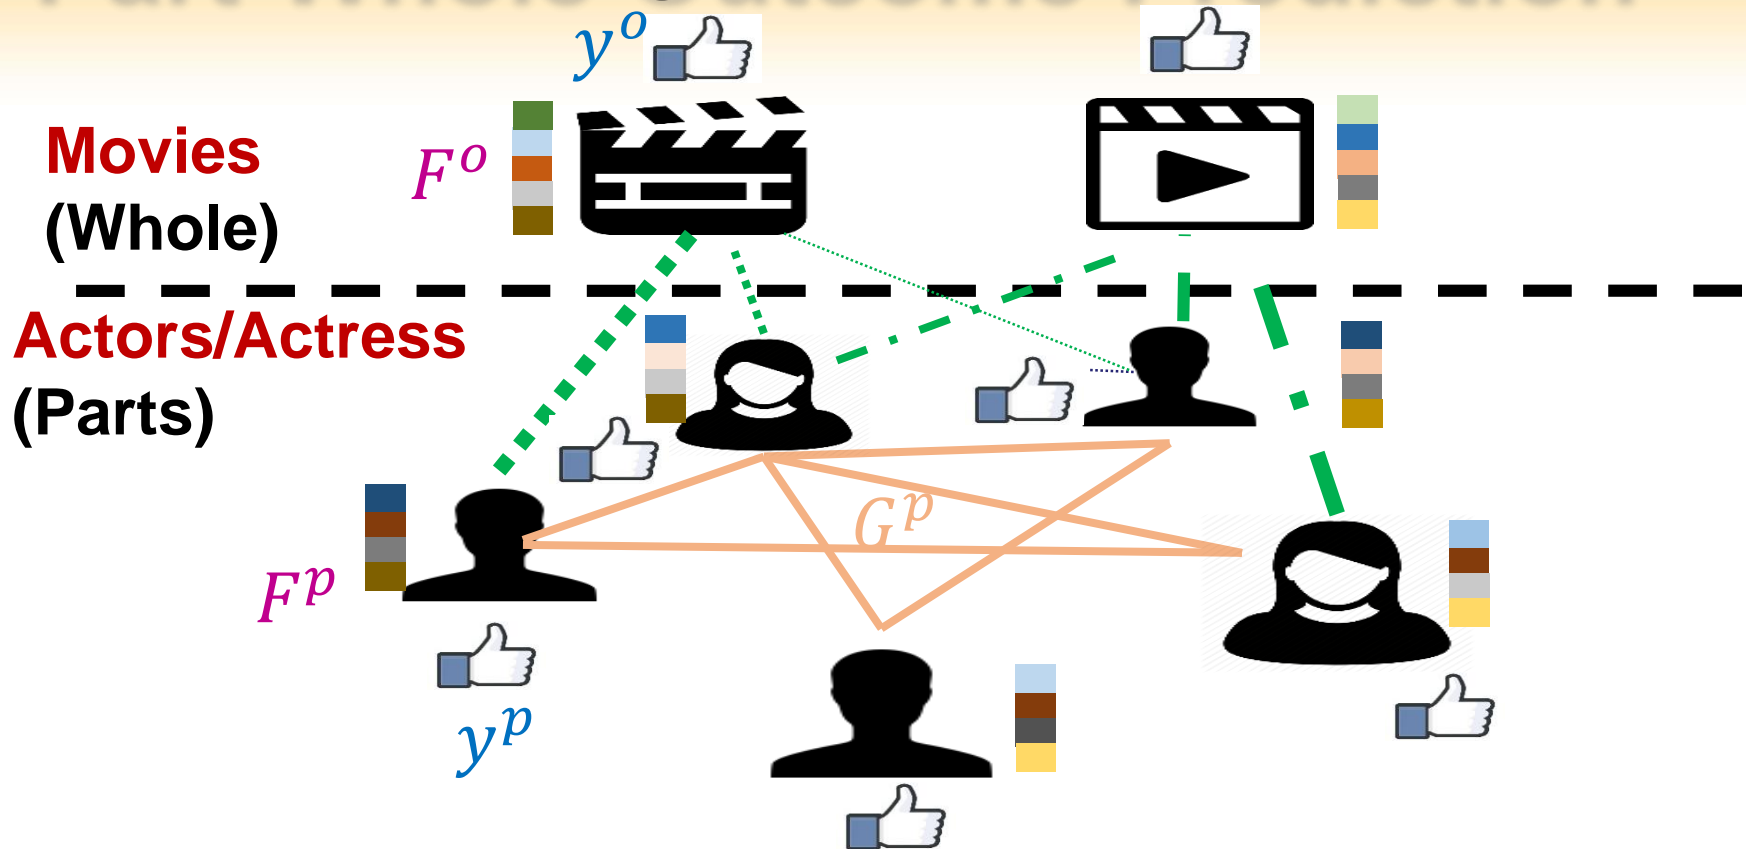
Interdependency



high complexity

Question: how to scale up the computation?

Part-Whole Outcome Prediction



- Given:**
1. feature matrix for whole/part F^o / F^p
 2. outcome vector for whole/part y^o / y^p
 3. whole to part mapping ϕ
 4. parts' network G^p (optional)

Predict: outcome of new whole/parts

A Generic Joint Prediction Framework -- PAROLE

■ Formulation

$$\min J = J_o + J_p + J_{po} + J_{pp} + J_r$$


$$= \frac{1}{n_o} \sum_{i=1}^{n_o} L[f(F^o(i, :), w^o), y^o(i)]$$

$$+ \frac{1}{n_p} \sum_{i=1}^{n_p} L[f(F^p(i, :), w^p), y^p(i)]$$

$$+ \frac{\alpha}{n_o} \sum_{i=1}^{n_o} h(f(F^o(i, :), w^o), \text{Agg}(\phi(o_i)))$$

$$+ \frac{\beta}{n_p} \sum_{i=1}^{n_p} \sum_{j=1}^{n_p} G_{ij}^p g(f(F^p(i, :), w^p), f(F^p(j, :), w^p))$$

$$+ \gamma(\Omega(w^o) + \Omega(w^p))$$



J_o : Predictive Model for Whole


J_p : Predictive Model for Part

J_{po} : Part-whole Relationship

J_{pp} : Part-part Interdependency

J_r : parameter regularizer

Modeling Part-Whole Relationship




- **Overview:** for each whole entity o_i , define

$$\underline{e_i} = \mathbf{F}^o(i, :) \mathbf{w}^o - Agg(o_i)$$

- e_i : Measure the difference between
 - predicted whole outcome using whole feature
 - predicted whole outcome using aggregated parts outcome
- **Key idea:** model part-whole relations by
 - Different loss functions on e_i
 - Different aggregation functions $Agg(\cdot)$

Overview



- **Intuition:** whole \leftarrow (weighted) sum of parts

- **Details:**

$$e_i = F^o(i, :)w^o - Agg(o_i)$$

$$Agg(o_i) = \sum_{j \in \phi(o_i)} a_j^i F^p(j, :)w^p$$

- a_j^i : weight of part j 's contribution to the whole o_i 's outcome

- **Remark:**


- Characterize part-whole relationships
 - Use different loss functions on e_i
 - Use different norms on a_i

Linear Part-Whole Relation



- **Intuition:** Whole \leftarrow linear combination of parts
 - some parts play more important roles than the others in contributing to the whole outcome
- **Details:** $J_{po} = \frac{\alpha}{2n_o} \sum_{i=1}^{n_o} e_i^2$
- **Remark:**
 - $a_j^i = 1$: the whole is the sum of its parts
 - $a_j^i = \frac{1}{|o_i|}$: average coupling

Sparse Part-Whole Relation



- **Intuition:** Whole \leftarrow a few parts
 - some parts have little or no effect on the whole outcome
- **Details:** $J_{po} = \frac{\alpha}{n_o} \sum_{i=1}^{n_o} \left(\frac{1}{2} e_i^2 + \lambda |\mathbf{a}_i|_1 \right)$
- **Remark:**
 - The l_1 norm can shrink some part contributions a_j^i to exactly zero
 - **Nonlinear** part-whole relation


Ordered Sparse Part-Whole Relation

- **Intuition:** Whole \leftarrow a few top parts
 - team performance is determined by not only a few key members, but also the structural hierarchy between them
- **Details:** $J_{po} = \frac{\alpha}{n_o} \sum_{i=1}^{n_o} (\frac{1}{2} e_i^2 + \lambda \Omega_w(\mathbf{a}_i))$
 - $\Omega_w(x) = \sum_{i=1}^n |x|_{[i]} w_i = \mathbf{w}^T |\mathbf{x}|_{\downarrow}$: **ordered weighted l_1 norm**
 - $w \in \mathcal{K}_{m+}$: vector of non-increasing non-negative weights

Robust Part-Whole Relation

- **Intuition:** Whole \leftarrow parts that are not outliers
 - squared loss is sensitive to outliers
- **Solution:** robust regression model
- **Details:** $J_{po} = \frac{\alpha}{n_o} \sum_{i=1}^{n_o} \rho(e_i)$
 - $\rho(\cdot)$ is robust estimator

Method \ Case	$ e \leq t$	$ e > t$
Huber $\rho_H(e)$	$\frac{1}{2}e^2$	$t e - \frac{1}{2}t^2$
Bisquare $\rho_B(e)$	$\frac{t^2}{6} \left\{ 1 - \left[1 - \left(\frac{e}{t} \right)^2 \right]^3 \right\}$	$\frac{t^2}{6}$



Maximum Part-Whole Relation



- **Intuition:** $\text{Whole} \leftarrow \max(\text{parts})$
 - team performance dominated by the best team member/leader
- **Details:**
 - $\text{Agg}(o_i) = \max(\text{parts}' \text{ outcome})$ [not differentiable]
 - Soft max function: $\max(x_1, x_2, \dots, x_n) \approx \ln(\exp(x_1) + \exp(x_2) + \dots + \exp(x_n))$
 - Aggregation: $\text{Agg}(o_i) = \ln(\sum_{j \in \phi(o_i)} \exp(F^p(j, :)w^p))$

$$J_{po} = \frac{\alpha}{2n_o} \sum_{i=1}^{n_o} e_i^2$$

Summarize Part-Whole Relations


Name	$Agg(o_i)$ Aggregation of parts	J_{po} Sub-objective	Remark
Maximum	$\ln(\sum \exp(F^p(j, :)w^p))$	$\frac{\alpha}{2n_o} \sum e_i^2$	Nonlinear Whole \leftarrow max(parts)
Linear	$\sum a_j^i F^p(j, :)w^p$	$\frac{\alpha}{2n_o} \sum e_i^2$	Linear Whole \leftarrow linear combination of parts
Sparse	$\sum a_j^i F^p(j, :)w^p$	$\frac{\alpha}{n_o} \sum (\frac{1}{2} e_i^2 + \lambda a_i _1)$	Nonlinear Whole \leftarrow a few parts
Ordered Sparse	$\sum a_j^i F^p(j, :)w^p$	$\frac{\alpha}{n_o} \sum (\frac{1}{2} e_i^2 + \lambda \Omega_w(a_i))$	Nonlinear Whole \leftarrow a few top parts
Robust	$\sum a_j^i F^p(j, :)w^p$	$\frac{\alpha}{n_o} \sum \rho(e_i)$	Nonlinear Whole \leftarrow parts that are not outliers

Modeling Part-Part Interdependency

- **Effect on the whole outcome**
 - **Intuition:** closely connected parts might have similar contribution to the whole outcome

- **Details:**

$$\mathcal{J}_{po} = \frac{\alpha}{n_o} \sum_{i=1}^{n_o} \left[\frac{1}{2} e_i^2 + \lambda |\mathbf{a}_i|_1 + \frac{1}{2} \sum_{k,l \in \phi(o_i)} G_{kl}^p (a_k^i - a_l^i)^2 \right]$$



- Similar parts (**large G_{kl}^p**)
→ similar contributions ($a_k^i \approx a_l^i$)


Modeling Part-Part Interdependency

- **Effect on the parts outcome**

- **Intuition:** closely connected parts might share similar outcomes themselves

- **Details:**


$$\mathcal{J}_{pp} = \frac{\beta}{2n_p} \sum_{i=1}^{n_p} \sum_{j=1}^{n_p} G_{ij}^p (\mathbf{F}^p(i, :) \mathbf{w}^p - \mathbf{F}^p(j, :) \mathbf{w}^p)^2$$



- Similar parts (**large G_{ij}^p**)

→ similar predicted outcomes ($F^p(i, :)w^p \approx F^p(j, :)w^p$)

Optimization Solution



■ Formulation:

$$J = J_o(w^o) + J_p(w^p) + J_{po}(w^o, w^p, a_j^i) + J_{pp}(w^p) + J_r(w^o, w^p)$$

■ Observation:

- not jointly convex w.r.t. w^o, w^p, a_i^j
- Convex w.r.t. to one block while fixing others

■ Solution: block coordinate descent

Block Coordinate Descent

- Three coordinate blocks: w^o, w^p, a_j^i
- Update one block while fixing others
- Update each block
 - (proximal) gradient descent

	$\frac{\partial J_{po}}{\partial w^o}$	$\frac{\partial J_{po}}{\partial w^p}$	$\frac{\partial J_{po}}{\partial a_i}$ or proximal gradient update
Maximum Agg	$\frac{\alpha}{n_o} \sum_{i=1}^{n_o} e_i (F^o(i, :))'$	$\frac{\alpha}{n_o} \sum_{i=1}^{n_o} e_i \frac{\sum_{j \in \phi(o_i)} (F^p(j, :))' \tilde{y}_i^p}{\sum_{j \in \phi(o_i)} \tilde{y}_i^p}$	N/A
Linear Agg	$\frac{\alpha}{n_o} (F^o)' (F^o w^o - M F^p w^p)$	$-\frac{\alpha}{n_o} (F^p)' M' (F^o w^o - M F^p w^p)$	$e_i (-F^p(\phi(o_i), :) w^p) + L_i^p a_i$
Sparse Agg	$\frac{\alpha}{n_o} (F^o)' (F^o w^o - M F^p w^p)$	$-\frac{\alpha}{n_o} (F^p)' M' (F^o w^o - M F^p w^p)$	$z = a_i - \tau [e_i (-F^p(\phi(o_i), :) w^p) + L_i^p a_i]$ $a_i \leftarrow \text{prox}_{\lambda \tau l_1}(z)$
Order Sparse Agg	$\frac{\alpha}{n_o} (F^o)' (F^o w^o - M F^p w^p)$	$-\frac{\alpha}{n_o} (F^p)' M' (F^o w^o - M F^p w^p)$	$z = a_i - \tau [e_i (-F^p(\phi(o_i), :) w^p) + L_i^p a_i]$ $a_i \leftarrow \text{prox}_{\lambda \tau \Omega_w}(z)$
Robust Agg	$\frac{\alpha}{n_o} \sum_{i=1}^{n_o} \frac{\partial \rho(e_i)}{\partial e_i} F^o(i, :)'$	$\frac{\alpha}{n_o} \sum_{i=1}^{n_o} \frac{\partial \rho(e_i)}{\partial e_i} (-\sum_{j \in \phi(o_i)} a_j F^p(j, :))'$	$\frac{\alpha}{n_o} \left[\frac{\partial \rho(e_i)}{\partial e_i} (-F^p(\phi(o_i), :) w^p) + L_i^p a_i \right]$


Optimization Properties

■ Convergence and Optimality

- Under mild conditions, the optimization alg converges to a coordinate-wise minimum point

■ Complexity

- The alg scales linearly w.r.t. the size of part-whole graph in both time and space



Complexity: $O(n_o d_o + n_p d_p + m_{po} + m_{pp})$
 n_o : #whole entities
 n_p : #part entities
 m_{po} : #links from whole to parts
 m_{pp} : #links in part-part network
 d_o, d_p : feature dimension of whole, parts

details

THEOREM 4.1. As long as $-\gamma$ is not an eigenvalue of $\frac{\partial^2}{\partial \mathbf{w}^2} f^*$ or $\frac{\partial^2}{\partial \mathbf{w}^2} f^* + \beta \mathbf{I}$, $\mathcal{L}(\mathbf{w}^t) \leq \frac{\beta}{2} \mathbf{F}(\mathbf{w}^t) \mathbf{M} \mathbf{F}^T$, Algorithm 1 converges to a coordinate-wise minimum point.

PROOF. It is not hard to see that our objective function \mathcal{L} satisfies the structure of f , with $\mathcal{J}_{\mathcal{L}}$ corresponding to $f(\mathbf{w}^t, \mathbf{w}^t)$, $\mathcal{J}_{\mathcal{L}}^T(\mathbf{w}^t, \mathbf{w}^t) = -\gamma \mathbf{I}$, $\mathcal{J}_{\mathcal{L}} = \frac{\beta}{2} \mathbf{F}(\mathbf{w}^t) \mathbf{M} \mathbf{F}^T$ and the rest of the terms forming $f(\mathbf{w}^t)$.

Observing that \mathcal{L} is a continuous function on its domain for all the part-whole relationships introduced in Sec. 3.2, Assumption (B1) is satisfied. Next we show Assumption (B2) also holds using linear aggregation as an example, which can be adapted to other relationships. Let us first fix the blocks \mathbf{w}^p and various \mathbf{a}_i^j . We are left with a function of \mathbf{w}^o as $f(\mathbf{w}^o) = \frac{\beta}{2} \mathbf{z}_o^T \mathbf{z}_o + \mathbf{F}^T(\mathbf{I} - \gamma \mathbf{w}^o) \mathbf{F} \mathbf{w}^o + \frac{\beta}{2} \|\mathbf{w}^o\|^2 + \text{const}$, which is convex and thus quasiconvex. Recall that a function is called *hemivariante* if it is not constant on any line segments. We proceed using proof by contradiction and assume there exist \mathbf{w}_1^o and \mathbf{w}_2^o such that $\forall t \in [0, 1]$, the following holds:

$$g(t) \equiv f(t\mathbf{w}_1^o + (1-t)\mathbf{w}_2^o) = \text{a constant}$$

Take the derivative of $g(t)$ w.r.t. t , we have

$$\frac{dg(t)}{dt} = \left[\left(\frac{\beta}{2} + \frac{1}{n_o} \right) \mathbf{F}^T \mathbf{F} + \gamma \mathbf{F}(\mathbf{w}_1^o + (1-t)\mathbf{w}_2^o) + \frac{1}{n_o} (\mathbf{F}^T)^T \gamma^T \right] \cdot \frac{d}{dt} (t\mathbf{w}_1^o + (1-t)\mathbf{w}_2^o) = 0$$



This holds for $\forall t \in [0, 1]$ and since \mathbf{w}_1^o and \mathbf{w}_2^o are distinct, we have

$$\left(\frac{\beta}{2} + \frac{1}{n_o} \right) \mathbf{F}^T \mathbf{F} + \gamma \mathbf{F}(\mathbf{w}_1^o + (1-t)\mathbf{w}_2^o) = \frac{1}{n_o} (\mathbf{F}^T)^T \gamma^T + \frac{\beta}{2} \mathbf{F}^T \mathbf{M} \mathbf{F}^T \mathbf{w}^o$$

When the eigenvalues of $\frac{\partial^2}{\partial \mathbf{w}^2} f^*$ do not take value of $-\gamma$, the left matrix $(\frac{\beta}{2} + \frac{1}{n_o}) \mathbf{F}^T \mathbf{F} + \gamma \mathbf{F}(\mathbf{w}_1^o + (1-t)\mathbf{w}_2^o)$ is of full rank. As a result, $\mathbf{w}_1^o + (1-t)\mathbf{w}_2^o$ can only take an unique value, making $\mathbf{w}_1^o = \mathbf{w}_2^o$, a contradiction. So $f(\mathbf{w}^o)$ is hemivariante.


Next, let us fix \mathbf{w}^o and various \mathbf{a}_i^j and denote the function of \mathbf{w}^p as $f(\mathbf{w}^p) = \frac{\beta}{2} \mathbf{z}_p^T \mathbf{z}_p + \mathbf{F}^T(\mathbf{I} - \gamma \mathbf{w}^o) \mathbf{F} \mathbf{w}^p + \frac{\beta}{2} \sum_{i=1}^{n_p} \sum_{j=1}^{n_o} \mathcal{C}_{ij}^T(\mathbf{F}^T(\mathbf{I} - \gamma \mathbf{w}^o) \mathbf{F} \mathbf{w}^p + \frac{\beta}{2} \mathbf{z}_p^T \mathbf{z}_p) + \frac{\beta}{2} \|\mathbf{w}^p\|^2 + \text{const}$. We still use proof by contradiction and assume

Datasets

Data	Whole	Part	#Whole	#Part
Math	Question (#votes)	Answer (#votes)	16,638	32,876
SO	Question (#votes)	Answer (#votes)	1,966,272	4,282,570
DBLP	Author (h-index)	Paper (#citation)	234,681	129,756
Movie	Movie (# )	Actors/Actress (# )	5,043	37,365

- **Setup:** sort whole in chronological order, gather first x percent and corresponding parts as training, test on last 10%
- **Metric:** root mean squared error (RMSE)

Outcome Prediction Performance




Observations

1. Joint prediction models > separate models
2. Non-linear part-whole relationships (max, Huber, Bisquare, Lasso, OWL) > linear relationships (Sum, Linear)
3. Lasso and OWL are the best methods in most cases

Effect of part-part interdependency


Movie



- PAROLE-Basic – no network information
- Part-part interdependency on whole outcome and parts outcome both boost the performance

Convergence Analysis


SO



- PAROLE converges fast (25-30 iterations)

Parameter Sensitivity


Movie



- α controls importance of part-whole relation
- β controls importance of part-part interdependency
- Stable in a relatively large parameter space

Scalability of PAROLE



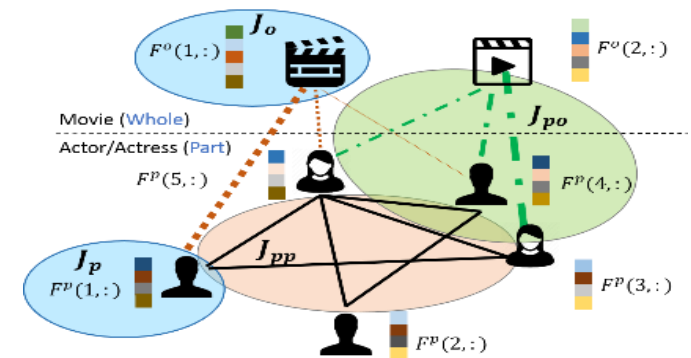
SO



- PAROLE scales linearly w.r.t. part-whole graph size

Conclusions -- PAROLE

- **Goals:** leverage potentially non-linear part-whole relationships for outcome prediction
- **Solutions:** PAROLE
 - **Modeling**
 - Part-whole relationship
 - Part-part interdependency
 - **Optimization**
 - Block coordinate descent
 - Converges to a coordinate-wise minimum point
 - Scales linearly w.r.t. the part-whole graph size



Roadmap

- Motivations and Background
- Part I: Team Performance Characterization
- Part II: Team Performance Prediction
- ➔ Part III: Team Performance Optimization
- Open Challenges
- Demo

Part III: Team Performance Optimization

- Team Formation and its variants
- Team Member Replacement
- Team Enhancement

Simple Team formation Problem

- Input:
 - A **task** T , consisting of a set of skills
 - A set of candidate experts each having a **subset of skills**

$T = \{\text{algorithms}, \text{java}, \text{graphics}, \text{python}\}$

Alice {algorithms}	Bob {python}	Cynthia {graphics, java}	David {graphics}	Eleanor {graphics, java, python}
------------------------------	------------------------	------------------------------------	----------------------------	--

- **Problem:** Given a **task** and a **set of experts**, find the smallest subset (**team**) of experts that together have all the required skills for the task

Set Cover

- The Set Cover problem:
 - We have a universe of elements $U = \{x_1, \dots, x_N\}$
 - We have a collection of subsets of U , $\mathcal{S} = \{S_1, \dots, S_n\}$, such that $\bigcup_i S_i = U$
 - We want to find the smallest sub-collection $\mathcal{C} \subseteq \mathcal{S}$ of \mathcal{S} , such that $\bigcup_{S_i \in \mathcal{C}} S_i = U$
 - The sets in \mathcal{C} cover the elements of U

Coverage

- The Simple Team Formation Problem is a just an instance of the **Set Cover** problem
 - **Universe** U of elements = Set of all **skills**
 - Collection S of **subsets** = The set of **experts** and the subset of skills they possess.

$T = \{\text{algorithms}, \text{java}, \text{graphics}, \text{python}\}$

Alice

{algorithms}

Bob

{python}

Cynthia

{graphics, java}

David

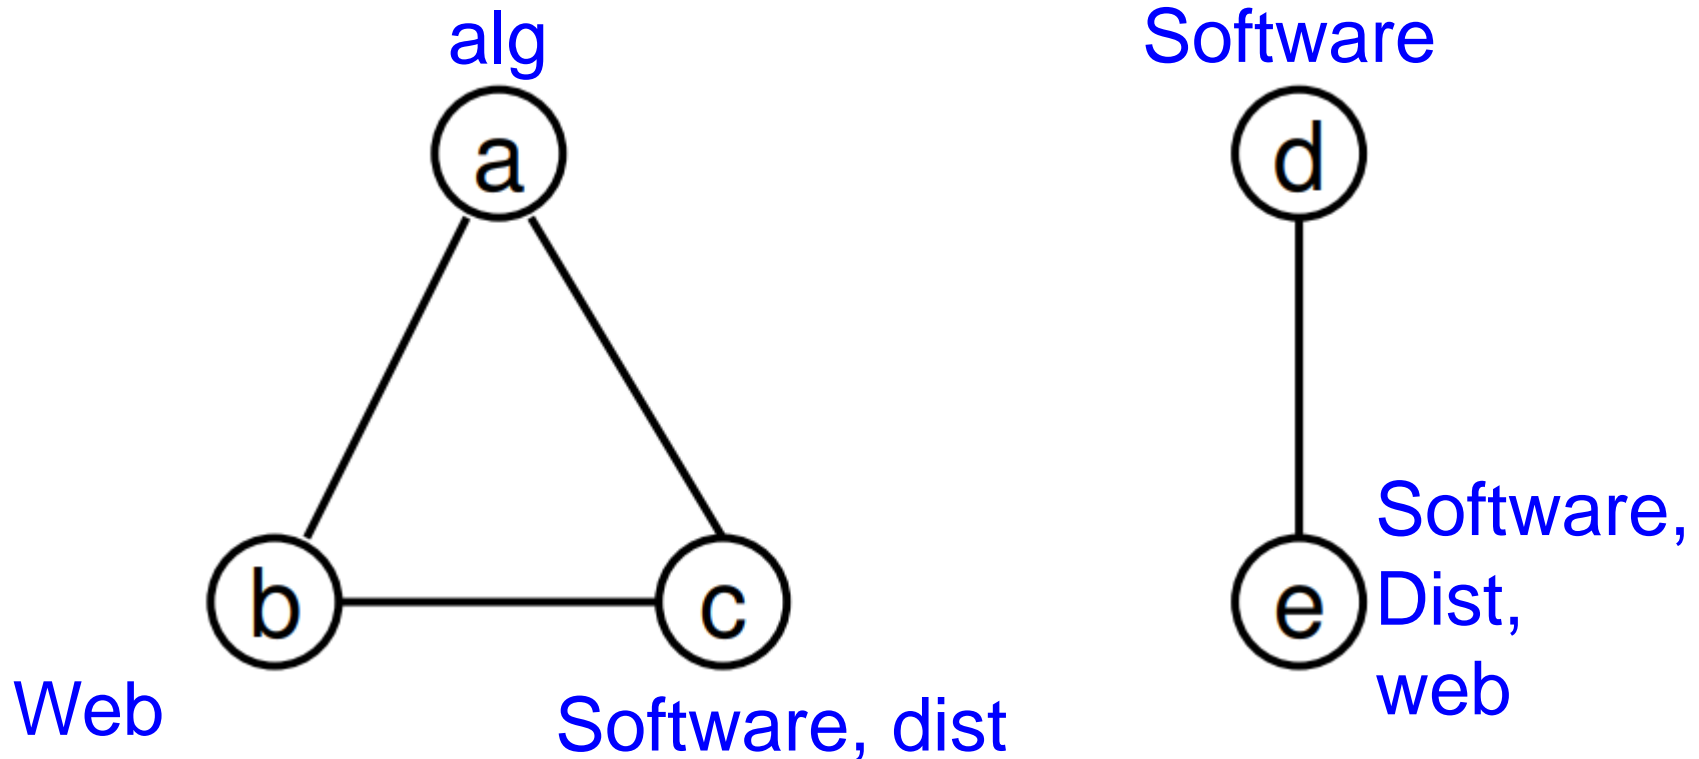
{graphics}

Eleanor

{graphics, java, python}

Team Formation with Networks

- $T = \{\text{algorithms, software engineering, distributed systems, web programming}\}$



Problem Definition

■ Given:

- Task requiring a set of skills
- Set of individuals
- Skills possessed by each individual
- Graph of communication cost between individuals

■ Find

- A subset of individuals containing all required skills with minimized communication cost

Communication cost


- Diameter ($CC-R$)
 - Diameter of the subgraph of the selected individuals
 - NP-complete (reduce to Multiple-Choice Cover)
- Minimum Spanning Tree ($CC-MST$)
 - Cost of the MST on the subgraph of the selected individuals
 - NP-complete (reduce to Group Steiner Tree)

Algorithm for Diameter-TF

- For every skill a required by the task T , compute $S(a)$: the individuals with a
- Pick the skill a_{rarest} with lowest cardinality
- Among all candidates from the set $S(a_{rarest})$, pick the one that leads to the smallest diameter

The RarestFirst algorithm

$T = \{\text{algorithms}, \text{java}, \text{graphics}, \text{python}\}$



Skills:

algorithms

graphics

java

python


$\alpha_{\text{rare}} = \text{algorithms}$

$S_{\text{rare}} = \{\text{Bob}, \text{Eleanor}\}$

Diameter = 2

The **RarestFirst** algorithm

$T = \{\text{algorithms}, \text{java}, \text{graphics}, \text{python}\}$



Skills:

algorithms

graphics

java

python

$\alpha_{\text{rare}} = \text{algorithms}$

$S_{\text{rare}} = \{\text{Bob}, \text{Eleanor}\}$

Diameter = 1

Algorithm for MST-TF

- CoverSteiner

- $X_0 \leftarrow \text{GreedyCover}$

- Add individuals with most uncovered skills

- $X' \leftarrow \text{SteinerTree}(G, X_0)$

- 1: $\mathcal{X}' \leftarrow v$, where v is a random node from \mathcal{X}_0 .
- 2: **while** $(\mathcal{X}_0 \setminus \mathcal{X}') \neq \emptyset$ **do**
- 3: $v^* \leftarrow \arg \min_{u \in \mathcal{X}_0 \setminus \mathcal{X}'} d(u, \mathcal{X}')$
- 4: **if** $\text{Path}(v^*, \mathcal{X}') \neq \emptyset$ **then**
- 5: $\mathcal{X}' \leftarrow \mathcal{X}' \cup \{\text{Path}(v^*, \mathcal{X}')\}$
- 6: **else**
- 7: Return Failure

Another algorithm for MST-TF

- 1: $H \leftarrow \text{EnhanceGraph}(G, T)$
- 2: $\mathcal{X}_H \leftarrow \text{SteinerTree}(H, \{Y_1, \dots, Y_k\})$
- 3: $\mathcal{X}' \leftarrow \mathcal{X}_H \setminus \{Y_1, \dots, Y_k\}$

EnhanceGraph:

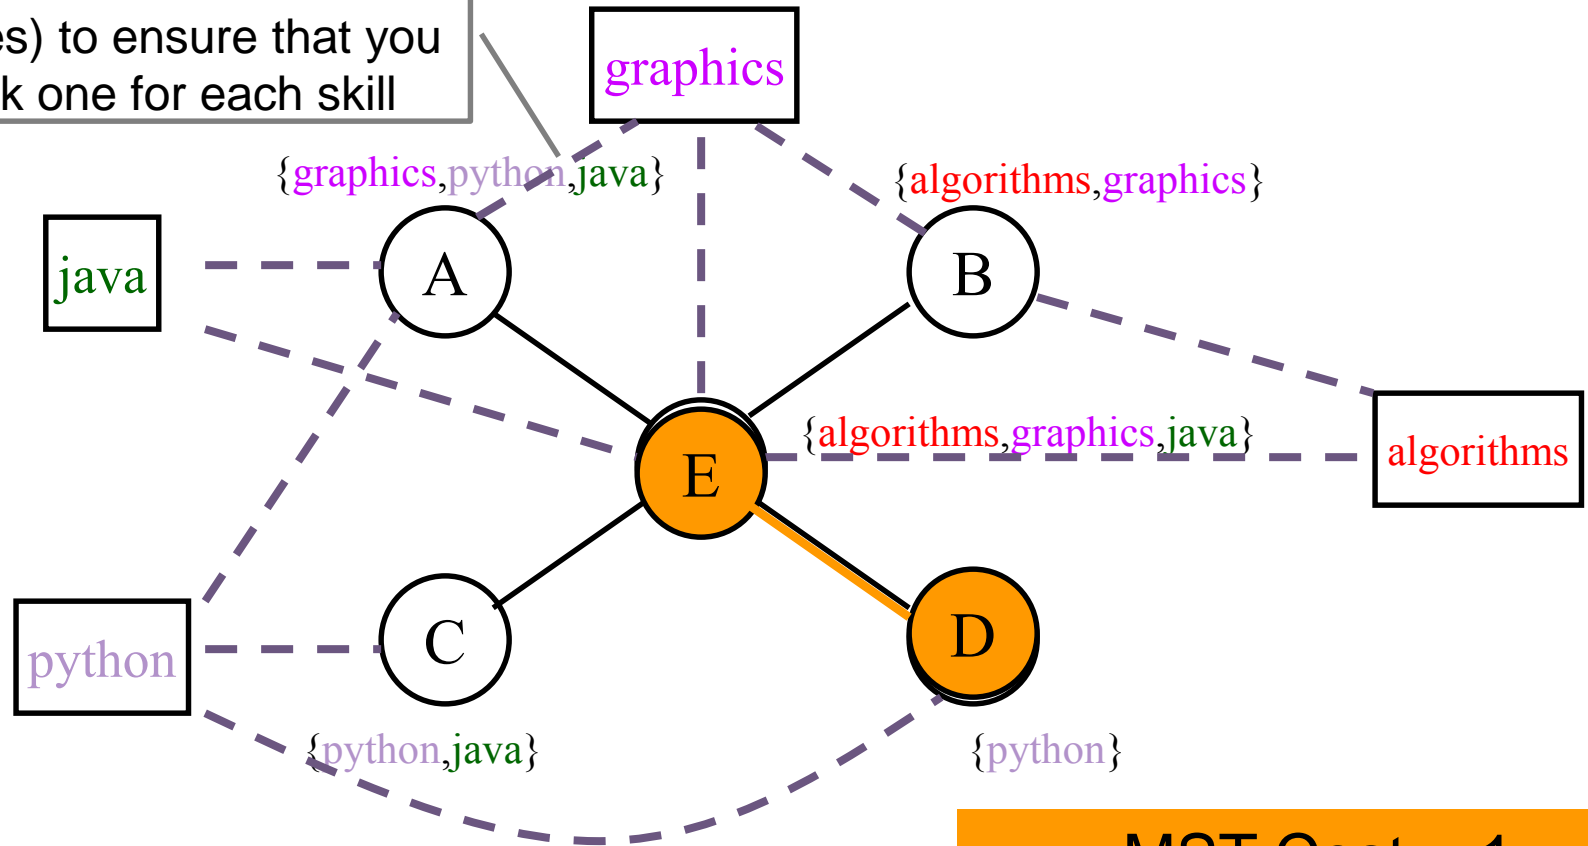
For every skill a_j in T

1. Create an additional node Y_j
2. Connect Y_j to all individuals with a_j with large weight

The EnhancedSteiner algorithm

Put a large weight on the new edges (more than the sum of all edges) to ensure that you only pick one for each skill

$T = \{\text{algorithms}, \text{java}, \text{graphics}, \text{python}\}$




MST Cost = 1


Experimental Evaluation

- DBLP: papers in database, data mining, AI, theory
- Skills derived from common terms in paper titles
- Communication weights determined by co-authorship
- 5509 individuals, 1792 skills
- Tasks generated with 2 to 20 skills
- Average over 100 combinations

Communication Cost




(a) CC-R cost




(b) CC-MST cost

Cardinality of the team



(a) Cardinality of the team.

Connectivity of the team



(b) Number of disconnected teams.

Case study on 10 papers

Rank	Paper title
1	The anatomy of a large-scale hypertextual Web search engine
2	Fast algorithms for mining association rules
3	Mining association rules between sets of items in large databases
4	Text categorization with support vector machines: Learning with many relevant features
5	Conditional random fields: Probabilistic models for segmenting and labeling sequence data
6	Mining frequent patterns without candidate generation
7	A survey of approaches to automatic schema matching
8	Automatic subspace clustering of high dimensional data for data mining applications
9	Models and issues in data stream systems
10	NiagaraCQ: A Scalable Continuous Query System for Internet Databases

Case Study Results

Rank	Actual authors	RarestFirst result	EnhancedSteiner result
1	S. Brin, L. Page	Paolo Ferragina, Patrick Valduriez, H. V. Jagadish, Alon Y. Levy, Daniela Florescu Divesh Srivastava, S. Muthukrishnan	P. Ferragina, J. Han, H. V. Jagadish, Kevin Chen-Chuan Chang, A. Gulli, S. Muthukrishnan, Laks V. S. Lakshmanan
2	R. Agrawal, R. Srikant	R. Agrawal	Philip S. Yu
3	R. Agrawal, T. Imielinski, A. N. Swami	Philip S. Yu	Wei Wang, Philip S. Yu
4	T. Joachims	Wei-Ying Ma, Gui-Rong Xue, H. Liu, J. Han, H. Lu, Z. Chen, Q. Yang, H. Cheng	J. Han, H. Lu, Wei-Ying Ma, Z. Chen, H. Liu, Gui-Rong Xue, Q. Yang
5	J. Lafferty, F. Pereira, A. McCallum	A. McCallum	A. McCallum
6	J. Han, J. Pei, Y. Yin	F. Bonchi	A. Gionis, H. Mannila, R. Motwani
7	E. Rahm, P. A. Bernstein	C. Bettini, R. Agrawal, Kevin Chen-Chuan Chang, T. Imielinski, H. Garcia-Molina, D. Barbara, S. Jajodia	C. Bettini, P. A. Bernstein, H. Garcia-Molina, S. Jajodia, D. Maier, D. Barbara
8	R. Agrawal, J. Gehrke, D. Gunopulos, P. Raghavan	D. Gunopulos, R. Agrawal	R. Agrawal, D. Gunopulos
9	B. Babcock, S. Babu, M. Datar, R. Motwani, J. Widom	M. T. Oszu	H. V. Jagadish, D. Srivastava
10	J. Chen, D. J. DeWitt, F. Tian, Y. Wang	Donald Kossmann, David J. DeWitt, Michael J. Franklin, Michael J. Carey	M. J. Carey, M. J. Franklin, D. Kossmann, D. J. DeWitt

Steaming Tasks

- Steam of tasks arriving online
- Create teams on-the-fly for each task
 - Teams should be fit for the tasks
 - Allocation should be fair to people

Balanced Task Covering

$$\min L \quad (1)$$

$$\sum_{j=1}^n \mathbf{p}_{\ell}^j X_{ji} \geq \mathbf{J}_{\ell}^i, \quad \forall i = 1, \dots, k, \ell = 1, \dots, m \quad (2) \quad \text{All tasks are executed}$$

$$\sum_{i=1}^k X_{ji} \leq L, \quad \forall j = 1, \dots, n \quad (3) \quad \text{Load balancing}$$

$$L \geq 0, X_{ji} \in \{0, 1\} \quad (4)$$

Online TF in Social Networks

- Forming teams that can accomplish the specified tasks while optimizing:
 - **Load:** number of tasks one expert participates
 - **Coordination Cost:**
 - Steiner tree
 - Diameter

Balanced Social Task

$$\min \max_{i \in \mathcal{P}} L(\mathbf{p}^i)$$

Load balancing

$$\mathbf{cov}(\mathbf{J}^j, \mathbf{q}^j) = 1$$

$$\forall j \in \mathcal{J}$$

All tasks are executed

$$c(Q^j) \leq B$$

$$\forall j \in \mathcal{J}$$

Bound on the
communication cost

Realistic Team Formation

- Realistic Requirements
 - Inclusion of a designated team leader and/or a group of experts
 - Skill requirement
 - Team size, or team cost
 - Locality of the team, e.g., in a geographical sense

Measure of collaborative compatibility

- Generalized form of subgraph density

- $density(C) = \frac{assoc(C)}{vol_g(C)} = \frac{\sum_{i,j \in C} w_{ij}}{\sum_{i \in C} g_i}$

- Strict monotonicity

- Robustness

Problem Formulation

$$\max_{C \subseteq V} \frac{\text{assoc}(C)}{\text{vol}_g(C)}$$

subject to : $S \subseteq C$

Required inclusion

Skill Requirement $\kappa_j \leq \text{vol}_{M_j}(C) \leq \iota_j, \quad \forall j \in \{1, \dots, p\}$

$$|C| \leq b$$

Team size

$$\text{vol}_c(C) \leq B$$

Budget constraint

$$\text{dist}(u, v) \leq d_0, \quad \forall u, v \in C,$$

Team locality

Churn of A Team Member

- **Case 1:** Employee resigns in a sales team
- **Case 2:** Task force down in a SWAT team
- **Case 3:** Rotation tactic between benches in NBA team

Q: How to find the best alternative when a team member leaves?

- L. Li, H. Tong, N. Cao, K. Ehrlich, Y.-R. Lin and N. Buchler: Replacing the Irreplaceable: Fast Algorithms for Team Member Recommendation, WWW 2015
- N. Cao, Y.-R. Lin, L. Li, H. Tong: g-Miner: Interactive Visual Group Mining on Multivariate Graphs, ACM CHI 2015
- System prototype & video demo: <http://team-net-work.org>

Team Member Replacement


Problem Definition:

Given: (1) A labelled social network $G := \{A, L\}$
(2) A team $G(\mathcal{T})$
(3) A team member $p \in \mathcal{T}$

Adj. Matrix

Skill Indicator

Recommend: A “best” alternative $q \notin \mathcal{T}$ to replace the person p 's role in the team $G(\mathcal{T})$



Q: who is a good candidate to replace the person to leave

DM

VIS

DB


NLP

AI

SYSTEM

MULTIMEDIA

Social Science Literature





- Team members prefer to work with people they have worked before [Hinds+OBHDP00]
- Distributed teams perform better when members know each other [Cummings+CSCW08]
- Specific communication patterns amongst team members are critical for performance [Cataldo+CHI12]

Conjecture: The similarity should be measured in the context of the team itself

Design Objectives


Objective 1: A good candidate should have a similar skill set




New team would have a similar skill set as the old team to continue to complete the task

Design Objectives

Objective 2: A good candidate should have a similar network structure




Skill Set: DM VIS DB NLP AI SYSTEM MULTIMEDIA



New team would have a similar network structure as the old team to collaborate effectively


Design Objectives

The skill and structure match should be fulfilled simultaneously!




New team would have similar skill and communication configuration for each sub-task

Random Walk based Graph Kernel



Graph 1




Graph 2


Details:

1. Compare similarity of every pair of nodes from each graph
— Eg: (1,2) vs (a, j) \rightarrow less similar
 (1,5) vs (a,e) \rightarrow more similar
2. Node pair similarity is measured by random walks
3. Two graphs are similar if they share many similar node pairs

Random Walk based Graph Kernel



Team 1




Team 2

Remarks:

- Incorporates both **attributes** and **structures** similarity
- Ideal fit for our two design objectives **simultaneously**

Kronecker Product Graph w/o Attribute

Graph Illustration



Matrix Description

$$A_1 = \begin{bmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix}$$

$$A_2 = \begin{bmatrix} 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \end{bmatrix}$$

$$A_1 \otimes A_2 =$$

Kronecker product

$$\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

One Random Walk on A_1

+

One Random Walk on A_2

= One Random Walk on $A_1 \otimes A_2 = A_x$

RW Graph Kernel — Formulation

Taking expectations instead of summing

$$\begin{aligned}\text{Ker}(G_1, G_2) &= \sum_k c^k q'_x (L_x A_x)^k L_x p_x \\ &= q'_x (I - c L_x A_x)^{-1} L_x p_x\end{aligned}$$


Attribute Indicator

- **Computational cost** ($A_x: t^2 \times t^2$)
 - *Exact methods*: [Vishwanathan+JMLR2010]
 - $O(t^6)$ - Direct computation
 - $O(t^3)$ - Sylvester equation
 - *Approx methods*: $O(t^2 r^4 + m r + r^6)$ [Kang+SDM12]

TEAMREP-BASIC

Find a new member q not in the current team that satisfies:

$$q = \arg \max_{j, j \notin \mathcal{T}} \text{Ker}(G(\mathcal{T}), G(\mathcal{T}_{p \rightarrow j}))$$

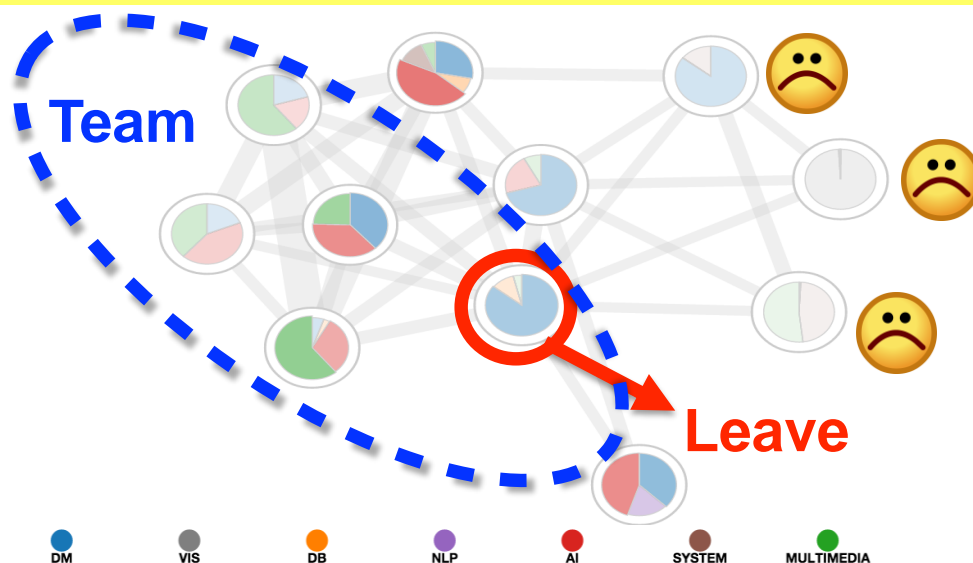


One graph kernel computation for every possible candidate

- **Challenge:** need to compute many graph kernel
overall complexity: $O(nt^3)$
- **Questions:**
 - **Q1:** how to reduce the number of graph kernels
 - **Q2:** how to speed up the computation for each graph kernel


Scale-up: Candidate Filtering

Pruning Strategy: Filter out all the candidates w/o any connections to any of the rest team members.



- **Theorem:** The pruning is safe: won't miss any potentially good replacement
- **Benefit:** The number of graph kernel computations is reduced to $O(\text{size of the neighborhood of } T)$ $O(\sum_{i \in \mathcal{T}/p} d_i)$


Speedup — Observation



Observation:

Many redundancies — *the nodes and edges within the rest team members remain the same*

Speedup — Approx Approach



The common part is the adjacency matrix of the rest team members

Speedup — Approx Approach

$$\text{Ker} \left(\begin{array}{|c|c|} \hline \text{red} & \text{cyan} \\ \hline \text{cyan} & \text{white} \\ \hline \end{array}, \begin{array}{|c|c|} \hline \text{red} & \text{orange} \\ \hline \text{orange} & \text{white} \\ \hline \end{array} \right)$$

$$\approx y'(1 - cL_{\times}(X_1Y_1) \otimes (X_2Y_2))^{-1}L_{\times}x$$

$$= y'L_{\times}x + cy'L_{\times}(X_1 \otimes X_2)\boxed{M}(Y_1 \otimes Y_2)L_{\times}x$$

$$M = (I - c(\sum_{j=1}^l Y_1 L_1^{(j)} X_1 \otimes Y_2 L_2^{(j)} X_2))^{-1}$$

$$M \text{ is of size } (r+2)^2 \times (r+2)^2$$

$$\text{Time Complexity: } O((\sum_{i \in \mathcal{T}/p} d_i)(lt^2r + r^6))$$

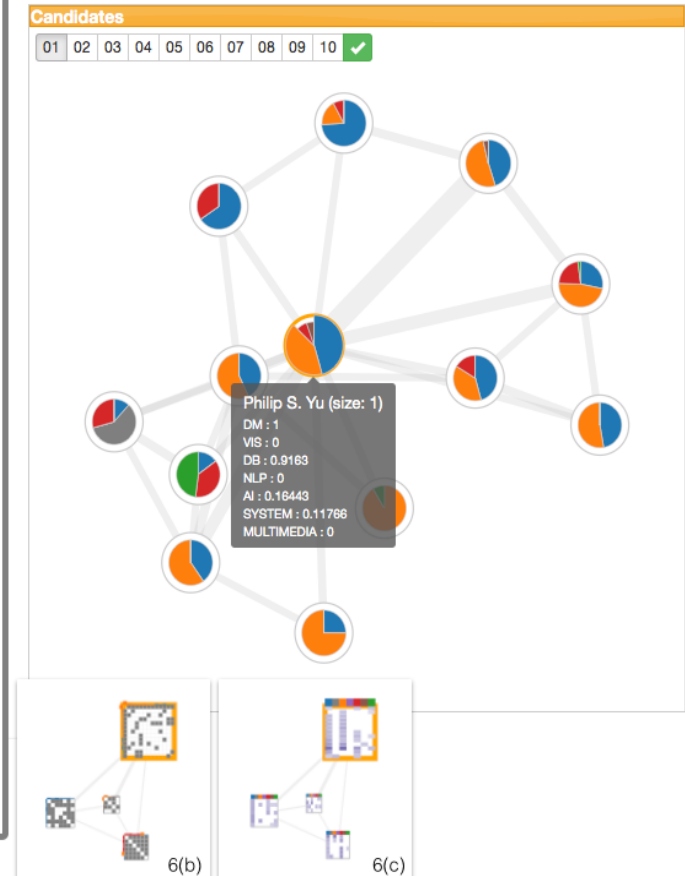
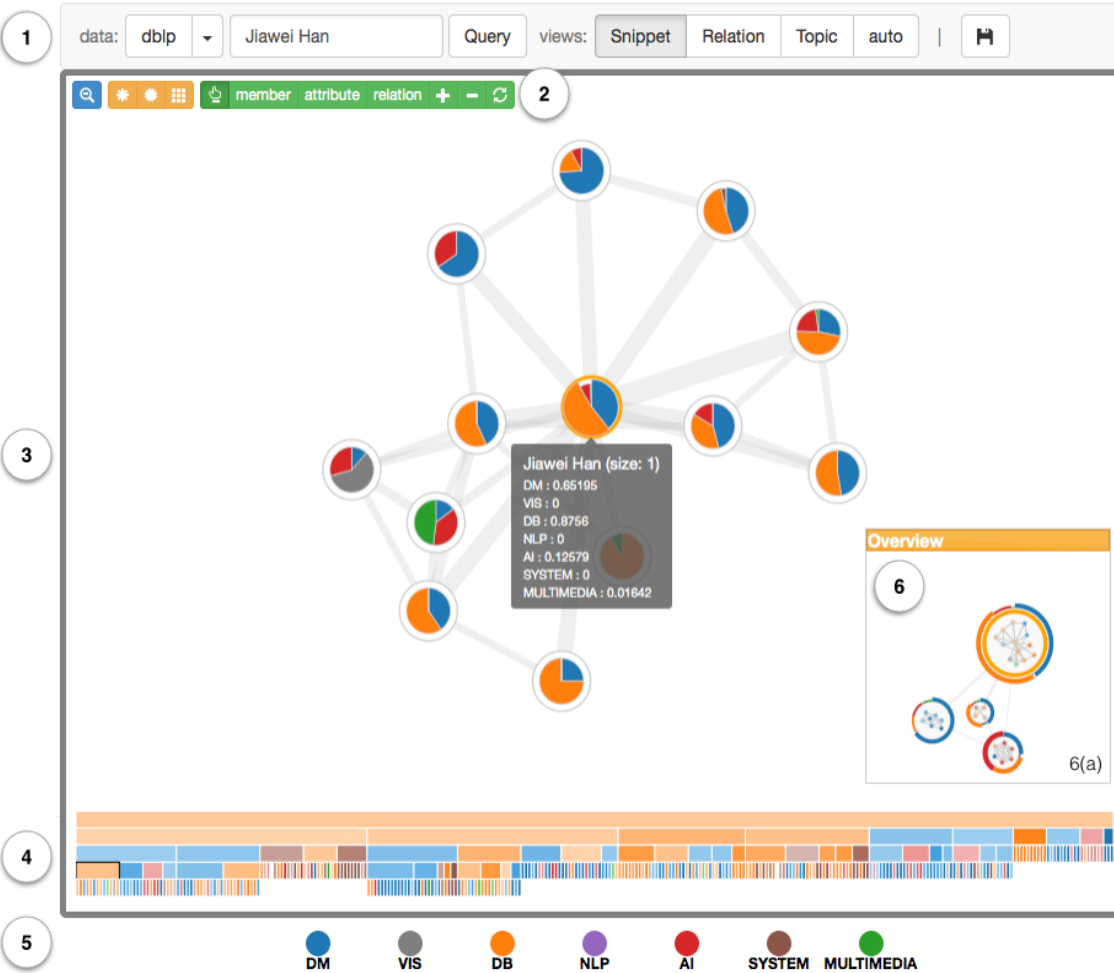
$$[\sum_{i \in \mathcal{T}/p} d_i \ll n, r \ll t]$$

$$\text{Original Complexity: } O(nt^3)$$

Prototype Systems

Questions


- **Q1:** How effective is skill + structure?
- **Q2:** How fast is pruning?
- **Q3:** How fast is proposed solution?
- **Q4:** How is the scalability?



prototype: <http://team-net-work.org>


- Nan Cao, Yu-Ru Lin, Liangyue Li, Hanghang Tong. "g-Miner: Interactive Visual Group Mining on Multivariate Graphs", ACM CHI 2015.

User Studies



Our method achieves the best average recall, precision and R@1

Application in Author Alias Prediction




Our method achieves the highest accuracy

Author Alias: *Alexander J. Smola* vs. *Alex J. Smola*

Speed-up by Pruning

Questions

- Q1: How effective is skill + structure?
- Q2: How fast is pruning?
- Q3: How fast is proposed solution?
- Q4: How is the scalability?




Pruning has dramatic speed improvement


Further Speed-up

Questions

- Q1: How effective is skill + structure?
- Q2: How fast is pruning?
- Q3: How fast is proposed solution?
- Q4: How is the scalability?



Exact methods




Approximate methods

Exploiting redundancy leads to additional speed-up!


Scalability

Questions

- Q1: How effective is skill + structure?
- Q2: How fast is pruning?
- Q3: How fast is proposed solution?
- Q4: How is the scalability?



TEAMREP-FAST-EXACT

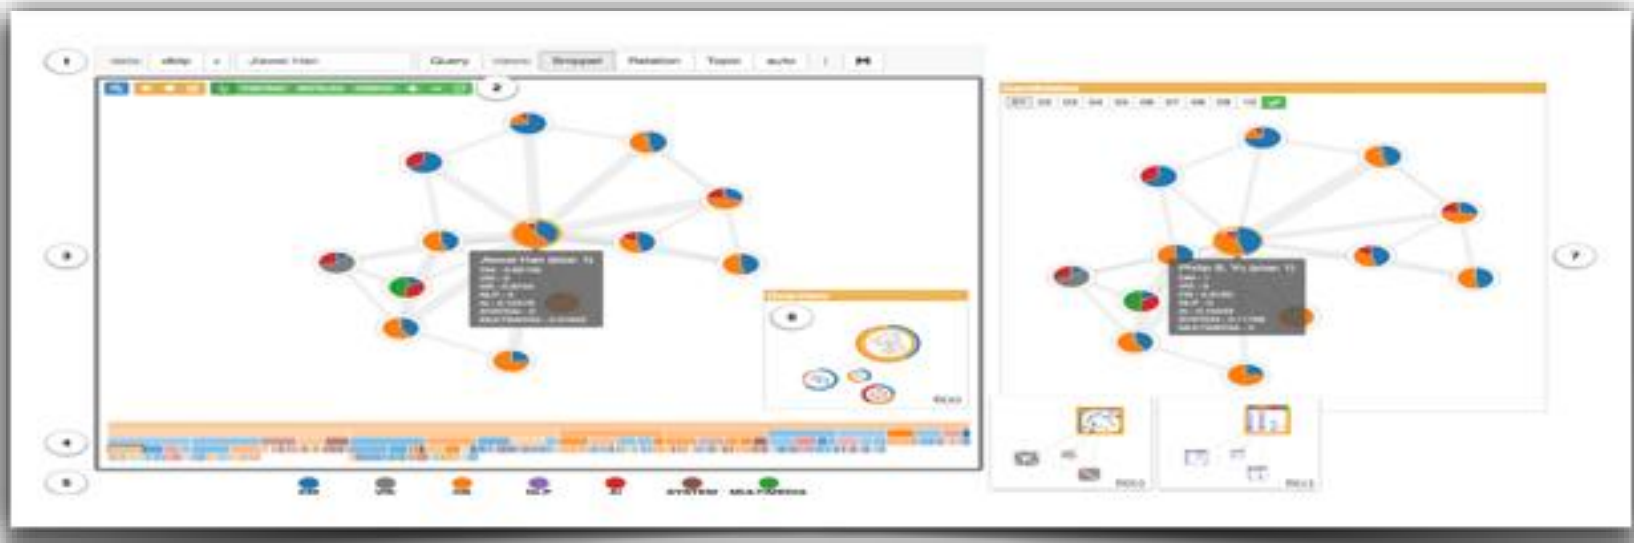


TEAMREP-FAST-APPROX

Our fast solutions scale sub-linearly

Team Member Replacement - Summary

- **Problem Def:** Team Member Replacement
- **Design Objectives:** skill + structural matching
- **Solutions:** graph kernel and fast algorithms
- **Prototype Systems:** <http://team-net-work.org/>



- L. Li, H. Tong, N. Cao, K. Ehrlich, Y.-R. Lin and N. Buchler: Replacing the Irreplaceable: Fast Algorithms for Team Member Recommendation, WWW 2015
- N. Cao, Y.-R. Lin, L. Li, H. Tong: g-Miner: Interactive Visual Group Mining on Multivariate Graphs, ACM CHI 2015

Beyond Team Member Replacement

■ Team Shrinkage

- If we need to reduce the size of an existing team (e.g., for the purpose of cost reduction), who shall leave the team?

■ Team Expansion


- If the team leader perceives the need to enhance certain expertise of the entire team, who shall we bring into the team?

■ Team Conflict Resolution

- If the team leader sees a conflict between certain team members, how shall we resolve it?

Key Idea: Solve all these team enhancement scenarios by team member replacement !

Team Expansion




Team Expansion – Case Study

- Expand the organizing committee of KDD 2013 by hiring some
 - strong expertise in AI
 - collaborated with as many existing committee members as possible
- Top five candidate:
 - *Qiang Yang, Zoubin Ghahramani, Eric Horvits, Thomas Dietteirich, Raymond J. Mooney*

Team Shrinkage


- Select teams with over 10 members and manually inject a “noisy” individual
 - Connect to all team members w/ random weights
 - Random skill vector
 - “best” candidate to leave the team

Team Shrinkage -- Results




Team Conflict Resolution

- E.g., Bob has a conflict with Alice
 - Replace either
 - Remove either



Roadmap

- Motivations and Background
- Part I: Team Performance Characterization
- Part II: Team Performance Prediction
- Part III: Team Performance Optimization
-  Part IV: Open Challenges
- Demo

Open Challenges

- Prediction Explanation
- Optimization Explanation
- Multiple Teams Optimization

Prediction Explanation

- **Observations:**

- Predictive models are mostly black-box or too complicated to understand reasons behind
- Interpretable models cannot achieve satisfactory prediction accuracy

- **Goal:**

- Provide explanations to performance prediction
- Assess trust of models


What an explanation looks like

From: Keith Richards
Subject: Christianity is the answer
NTTP-Posting-Host: x.x.com

I think Christianity is the one true religion.
If you'd like to know more, send me a note

atheism

christian



Appear in 21% of training examples, almost always in atheism

Appears in 11% of training examples, **always** in atheism


Why did this happen?
How do I fix it?

→ Will not generalize
→ Don't trust this model!

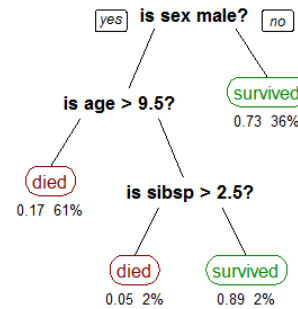
Three must-haves for a good explanation

Interpretable

- Humans can easily interpret reasoning



Definitely
not interpretable



Potentially
interpretable


Three must-haves for a good explanation

Interpretable

- Humans can easily interpret reasoning

Faithful

- Describes how this model actually behaves



Three must-haves for a good explanation

Interpretable

- Humans can easily interpret reasoning


Faithful

- Describes how this model actually behaves

Model agnostic

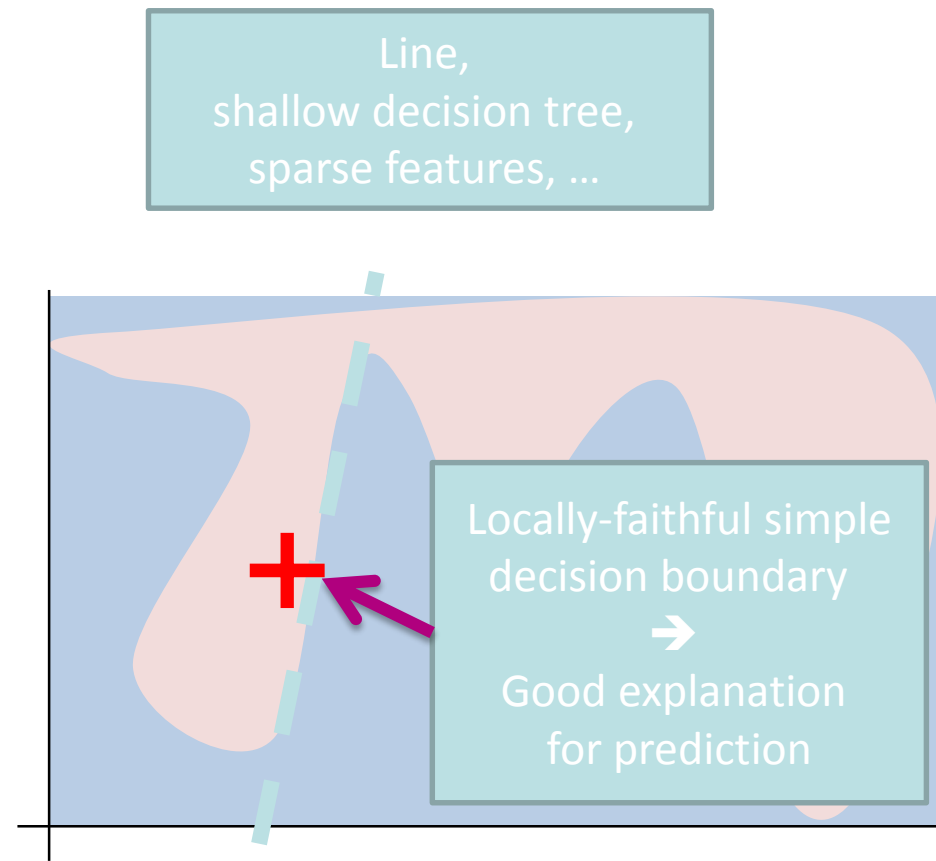
- Can be used for *any* ML model

Can explain
this mess 😊




LIME – Key Ideas

1. Pick a model class interpretable by humans
 - Not globally faithful... ☹️
2. Locally approximate global (blackbox) model
 - Simple model globally bad, but locally good





Using LIME to explain a complex model's prediction for input x_i

1. Sample points around x_i
2. Use complex model to predict labels for each sample
3. Weigh samples according to distance to x_i
4. Learn new simple model on weighted samples
5. Use simple model to explain



Interpretable representations

x (embeddings)



Model

This is what we perturb, and this is what we use in the interpretable approximation


x' (words)

This is a horrible
movie.



Interpretable representation: images

x (3 color channels /
pixel)



Model

x' (contiguous superpixels)



Human

Explaining Google's Inception NN



$$P(\text{🎸}) = 0.32$$



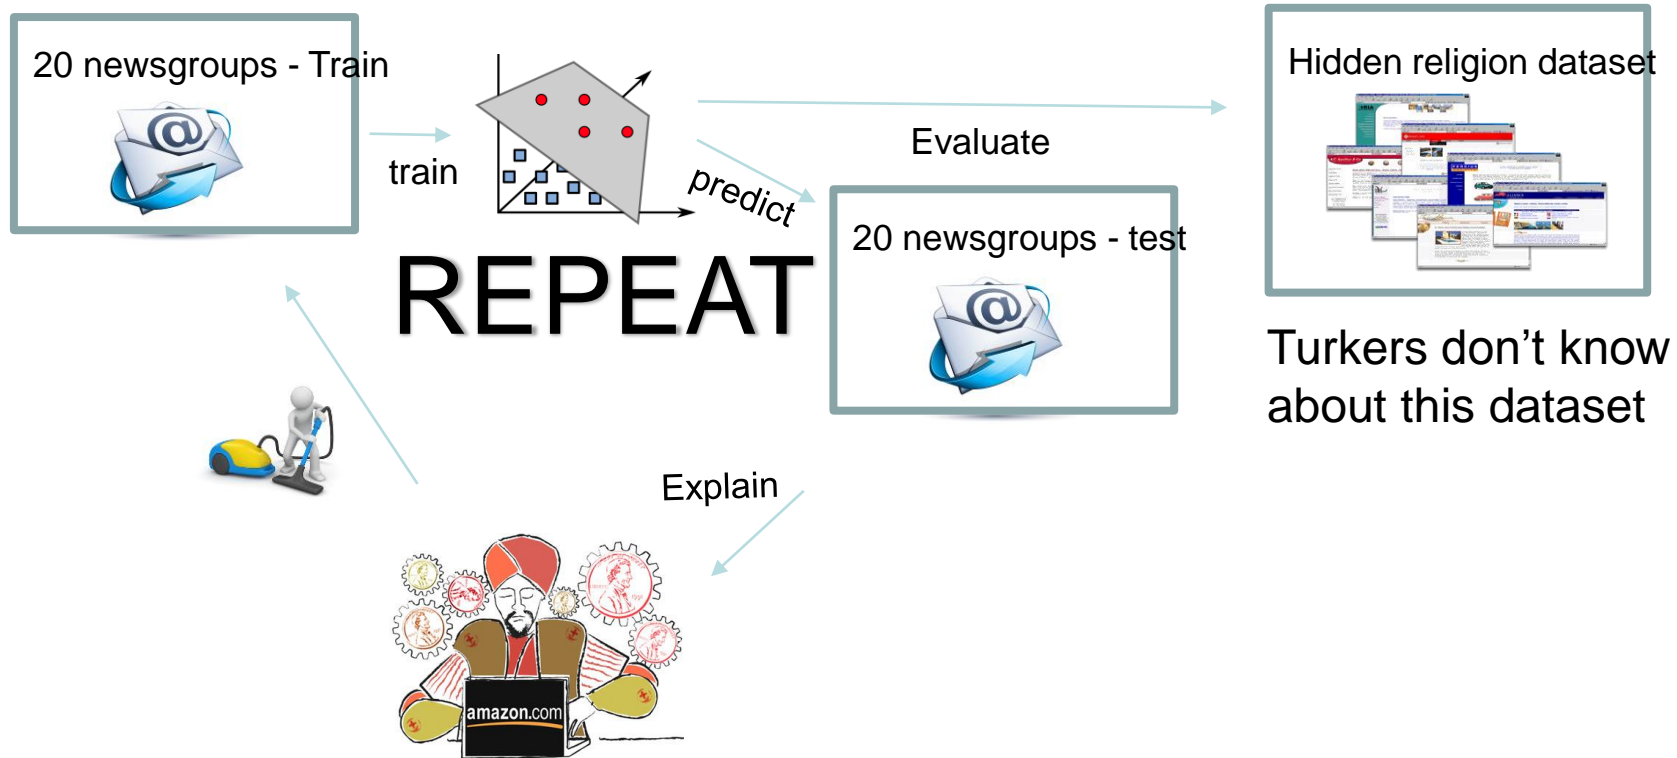
$$P(\text{🎸}) = 0.24$$



$$P(\text{🐶}) = 0.21$$



Fixing bad classifiers



Fixing bad classifiers

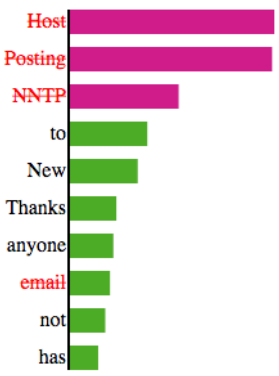
- Turkers click on 'useless' words for the task in each round

Example #5 of 10

True Class: ● Atheism

[Instructions](#) [Previous](#) [Next](#)

Words that the algorithm considers important.



Word	Importance (Bar Length)	Topic (Color)
Host	High	Christianity (Green)
Posting	High	Christianity (Green)
NNTP	Medium	Christianity (Green)
to	Low	Atheism (Magenta)
New	Low	Atheism (Magenta)
Thanks	Low	Atheism (Magenta)
anyone	Low	Atheism (Magenta)
email	Low	Atheism (Magenta)
not	Low	Atheism (Magenta)
has	Low	Atheism (Magenta)

Bar length indicates importance, and color indicates to which topic: Christianity (green) or Atheism (Pink).

Please click on the words (right next to the bars) that you think the algorithm is using incorrectly, because they are not important to distinguish between Atheism and Christianity. They should be red and crossed off after you click them.

Document

From: johnchad@triton.unm.edu (jchadwic)
Subject: Another request for Darwin Fish
Organization: University of New Mexico, Albuquerque
Lines: 11
NNTP-Posting-Host: triton.unm.edu


Hello Gang,

There have been some notes recently asking where to obtain the DARWIN fish.
This is the same question I have and I have not seen an answer on the net. If anyone has a contact please post on the net or email me.

Thanks,

john chadwick
johnchad@triton.unm.edu
or

Fixing bad classifiers



Explain through examples

- Consider the following learning task

Training points: z_1, \dots, z_n

Loss: $L(z_i, \theta)$

Params: $\hat{\theta} \triangleq \arg \min_{\theta \in \Theta} \frac{1}{n} \sum_{i=1}^n L(z_i, \theta)$

- Upweighting a training example

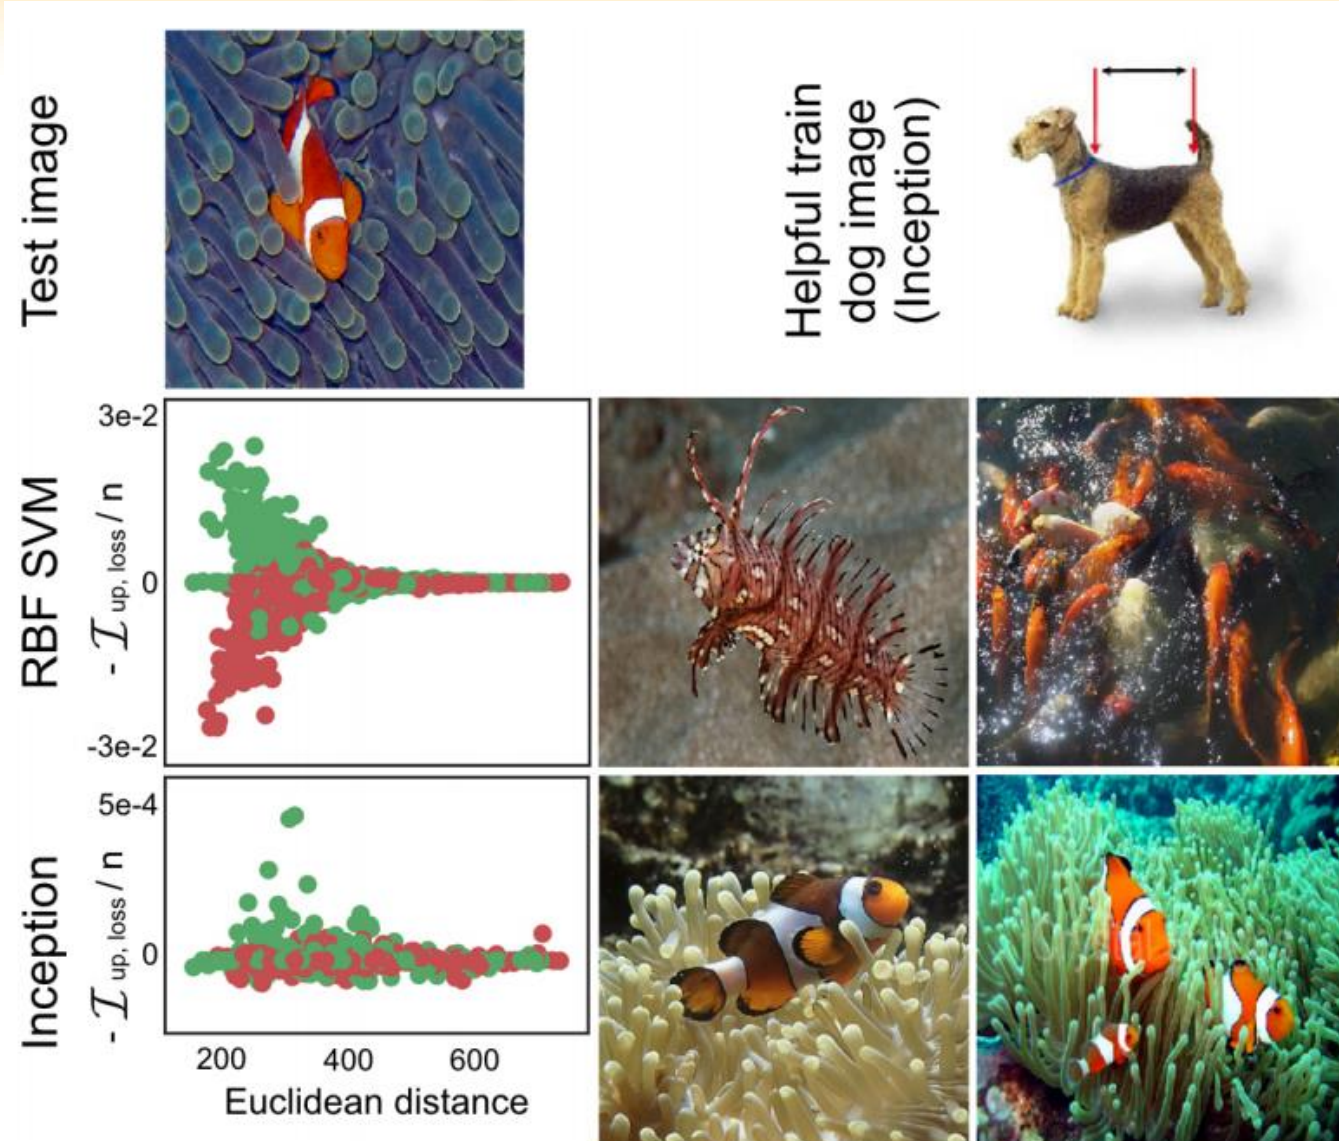
$$\mathcal{I}_{\text{up,params}}(z) \triangleq \left. \frac{d\hat{\theta}_{\epsilon,z}}{d\epsilon} \right|_{\epsilon=0} = -H_{\hat{\theta}}^{-1} \nabla_{\theta} L(z, \hat{\theta})$$

Explain through examples



- Effect on the loss of a test example

$$\begin{aligned}\mathcal{I}_{\text{up,loss}}(z, z_{\text{test}}) &\triangleq \left. \frac{dL(z_{\text{test}}, \hat{\theta}_{\epsilon, z})}{d\epsilon} \right|_{\epsilon=0} \\ &= -\nabla_{\theta} L(z_{\text{test}}, \hat{\theta})^{\top} H_{\hat{\theta}}^{-1} \nabla_{\theta} L(z, \hat{\theta})\end{aligned}$$

Understanding Model Behavior



Fixing Mislabelled Examples



Optimization Explanation

- **Observation:** existing work focus on recommending candidates for optimization
- **Goal:** provide explanations for team optimization algorithms
 - Convince the manager to make appropriate decisions
 - Example explanations for replacement
 - The candidate also participates in the key subtasks that the person leaving is involved in

Multiple Teams Optimization

- How to optimally shrink one team while expanding another team?
- How to recruit a new player from several other teams?
- Enhance all teams within an organization and/or form new teams by collectively imposing a series of team enhancement operations

Data

- *AMiner*: <https://aminer.org/data>
- *Semantic Scholar*:
<http://labs.semanticscholar.org/corpus/>
- *MovieLens*:
<https://grouplens.org/datasets/movielens/>
- *NBA*: <https://www.basketball-reference.com>
- *Github*: <https://www.githubarchive.org/>

Resources

- Project Website: <http://team-net-work.org/>
for papers, code, slides
- Prototype System: <http://team-net-work.org/system.html>

References

- Wuchty, Stefan, Ben Jones, and Brian Uzzi. "The Increasing Dominance of Teams in the Production of Knowledge," *Science*, May 2007, 316:1036-1039.
- Woolley, Anita Williams, et al. "Evidence for a collective intelligence factor in the performance of human groups." *science* 330.6004 (2010): 686-688
- Rui Yan, Jie Tang, Xiaobing Liu, Dongdong Shan, and Xiaoming Li. 2011. Citation count prediction: learning to estimate future citations for literature. *CIKM*, 2011.
- D. Wang, C. Song, and A.-L. Barabasi. Quantifying long-term scientific impact. *Science*, 342(6154): 127-132, 2013.
- T. Lappas, K. Liu, and E. Terzi. Finding a team of experts in social networks. In *KDD*, pages 467–476, 2009.
- S. S. Rangapuram, T. Buhler, and M. Hein. Towards realistic team formation in social networks based on densest subgraphs. *WWW* 2013.
- Kim, Young Ji, et al. "What Makes a Strong Team?: Using Collective Intelligence to Predict Team Performance in League of Legends." *CSCW*. 2017.

References

- Duch, Jordi, Joshua S. Waitzman, and Luís A. Nunes Amaral. "Quantifying the performance of individual players in a team activity." PloS one 5.6 (2010): e10937.
- F. Thung, T. F. Bissyande, D. Lo, and L. Jiang, Network Structure of Social Coding in GitHub. CSMR 2013
- Yuxiao Dong, Reid A. Johnson, and Nitesh V. Chawla. 2015. Will This Paper Increase Your h-index?: Scientific Impact Prediction. WSDM, 2015.
- L. Li, and H. Tong: The Child is Father of the Man: Foresee the Success at the Early Stage. KDD 2015: 655-664
- L. Li, H. Tong, J. Tang and W. Fan: "iPath: Forecasting the Pathway to Impact". SDM 2016
- Liangyue Li, Hanghang Tong, Yong Wang, Conglei Shi, Nan Cao and Norbou Buchler. Is the Whole Greater Than the Sum of Its Parts? KDD, 2017.
- Aris Anagnostopoulos, Luca Becchetti, Carlos Castillo, Aristides Gionis, and Stefano Leonardi. 2010. Power in unity: forming teams in large-scale community systems. CIKM, 2010.

References

- Duch, Jordi, Joshua S. Waitzman, and Luís A. Nunes Amaral. "Quantifying the performance of individual players in a team activity." PloS one 5.6 (2010): e10937.
- Aris Anagnostopoulos, Luca Becchetti, Carlos Castillo, Aristides Gionis, and Stefano Leonardi. Online Team Formation in Social Networks. WWW, 2012.
- Syama Sundar Rangapuram , Thomas Bühler , Matthias Hein, Towards realistic team formation in social networks based on densest subgraphs. WWW, 2013.
- N. Cao, Y.-R. Lin, L. Li, H. Tong: g-Miner: Interactive Visual Group Mining on Multivariate Graphs, ACM CHI 2015
- L. Li, H. Tong, N. Cao, K. Ehrlich, Y.-R. Lin and N. Buchler: Enhancing Team Composition in Professional Networks: Problem Definitions and Fast Solutions, TKDE, 2016
- Marco Tulio Ribeiro, Sameer Singh, Carlos Guestrin. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. KDD, 2016.
- Koh PW and Liang Percy Liang. Understanding black-box predictions via influence functions. ICML, 2017.