(A Few) Key Lessons Learned Building Recommender Systems For Large-scale Social Networks

Christian Posse

Principal Data Scientist and Product Manager at LinkedIn

San Francisco Bay Area | Information Technology and Services

World’s Largest Professional Network

175M+

62% non U.S.

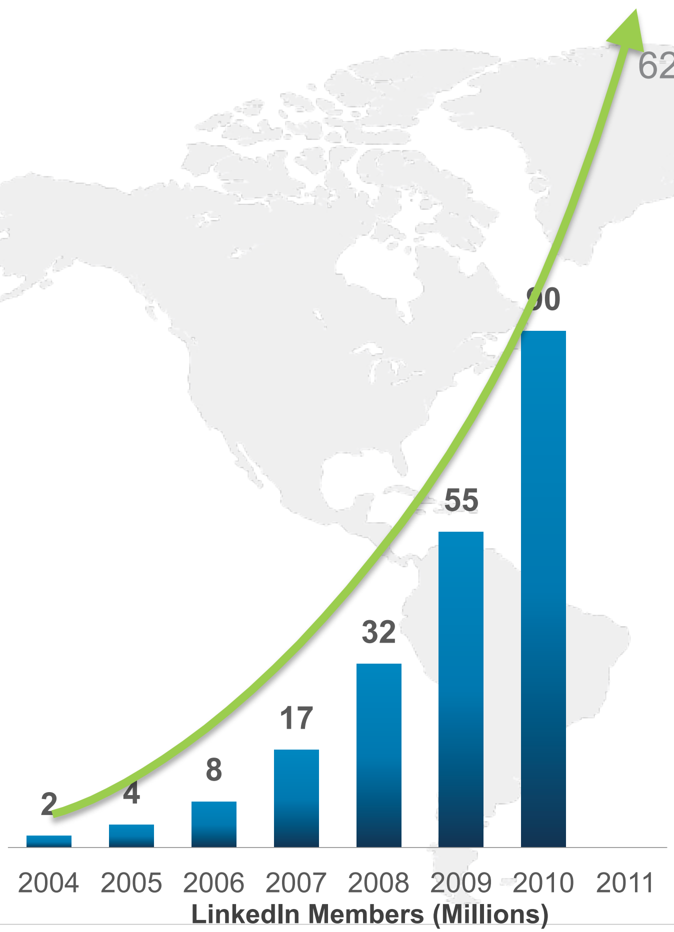

2/sec

25th

Most visit website worldwide (Comscore 6-12)

>2M

Company pages

85%

Fortune 500 Companies use LinkedIn to hire

LinkedIn Homepage

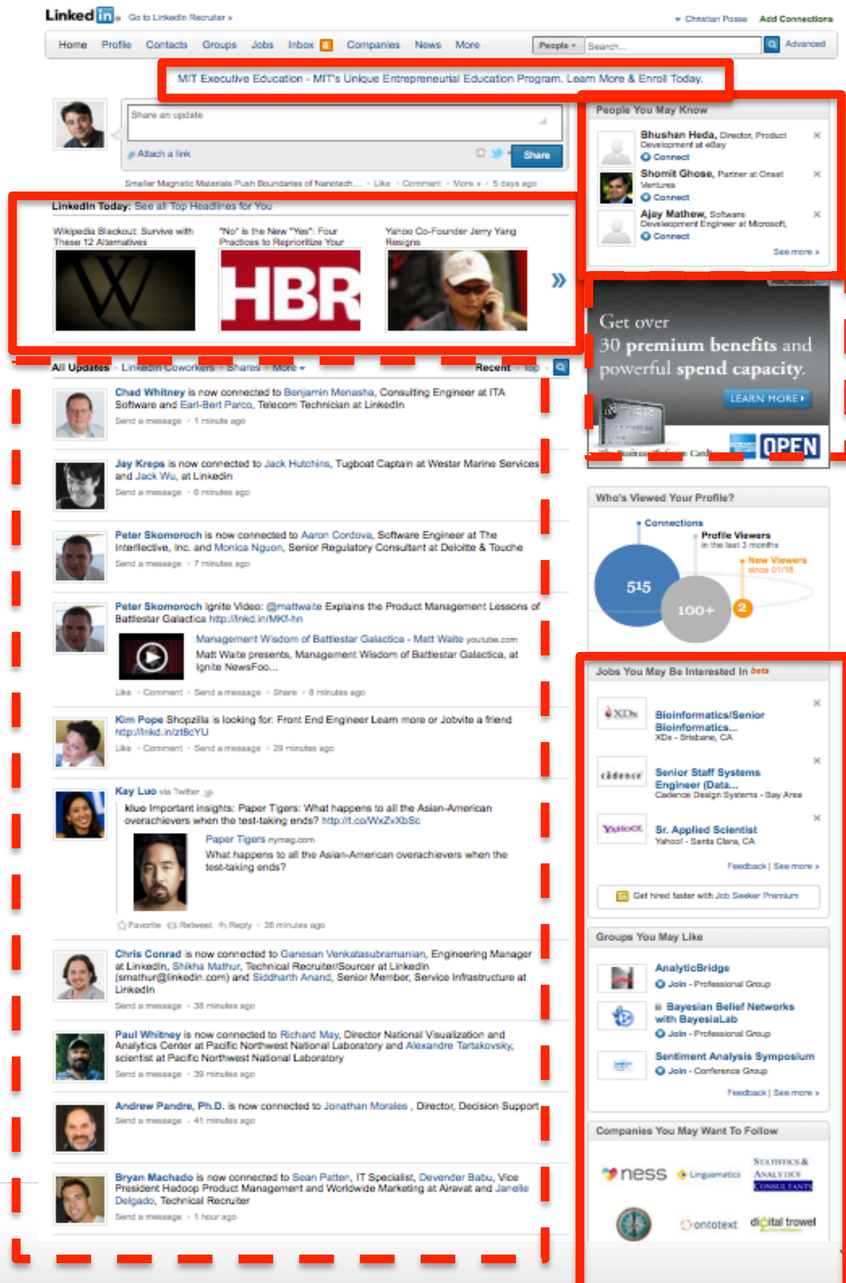

Powered by Recommendations!

The Recommendations Opportunity

Jobs You May Be Interested In

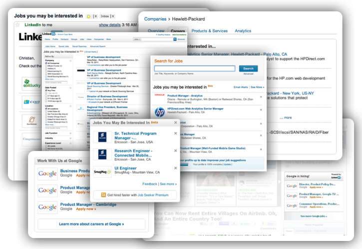

Talent Match

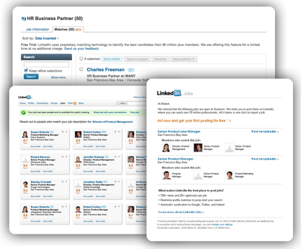

CAP

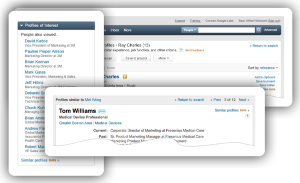

Similar Profiles

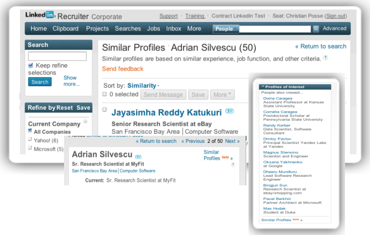

Related search

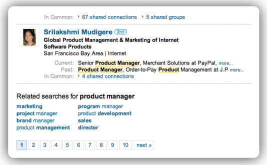

Profile browse maps

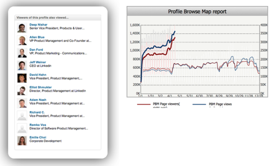

Connections

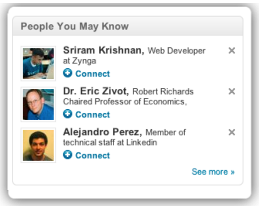

Companies

Recommendations, similar companies search, peer companies, and company browse maps, company products and services browse maps

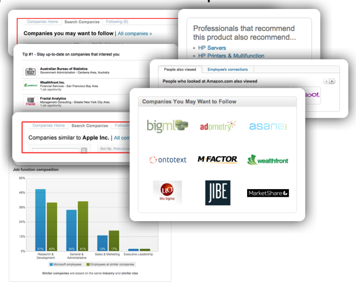

Network updates

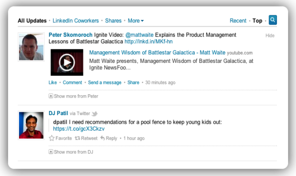

Jobs browse maps

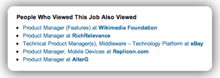

Ad matching engine

pCTR = f(member, creative, advertiser, context, inventory, OCTR)

Events You May Be Interested In

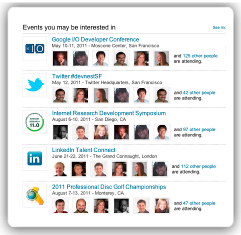

Groups

Recommendations, similar groups search

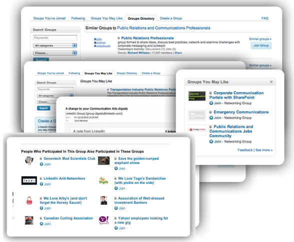

Similar jobs

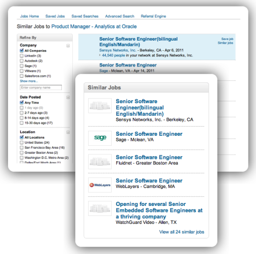

Referral Engine

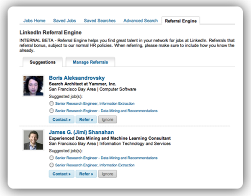

News

What are they worth? Think 50%

- > 50% of connections are from recommendations (PYMK)
- > 50% of job applications are from recommendations (JYMBII)
- > 50% of group joins are from recommendations (GYML)

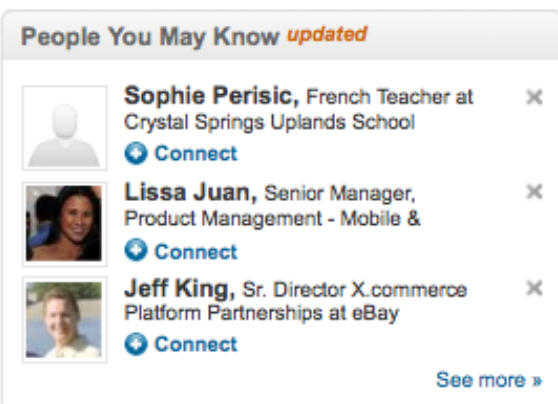

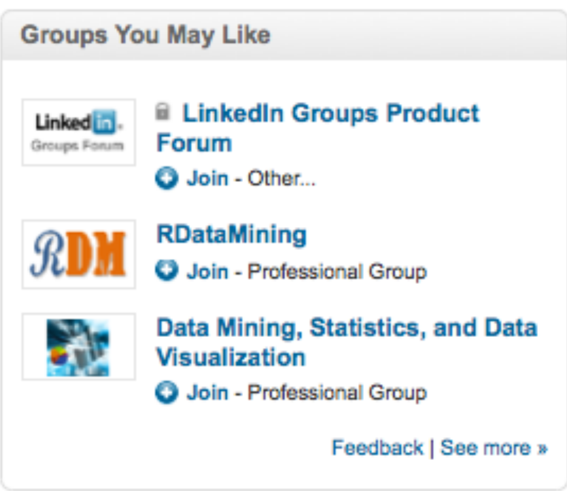

What is a Recommender System?

A Recommender selects a product that if acquired by the “buyer” maximizes value of both “buyer” and “seller” at a given point in time

Lesson 1 : Recommendations must make strategic sense…

- Conflicts of interest between ‘buyers’
  - E.g. job posters vs. job seekers

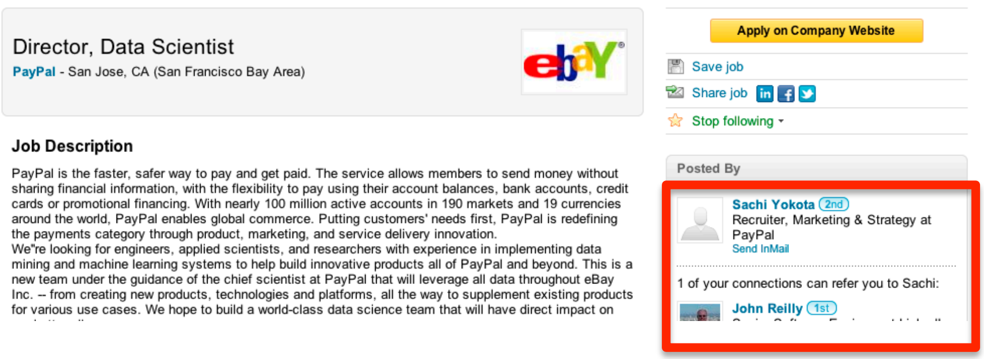

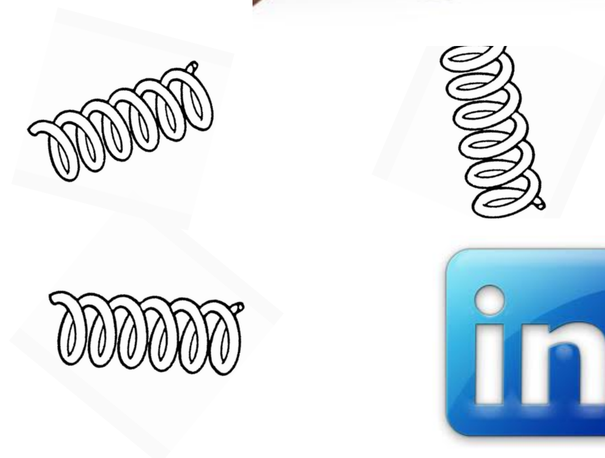

- What is the (long-term) value of an action?
  - How do we compare an invitation from a job application from a group join or the reading of a news?

Ingredients of a Recommender System

A Recommender processes information and transforms it into actionable knowledge

Data (Feature Engineering)

Algorithms

Business Logic and Analytics

User Experience

Lesson 2: User Experience Matters Most

1. Understand **user intent**
   → Define, model, leverage
2. Be in the **user flow**
3. Optimize **location** on the page
4. Set **right** expectations (“You May…”)
5. **Explain** recommendations
6. **Interact** with the user
7. Leverage **Social Referral**

Contextualization

Messaging

C+M

User Intent, User Flow, Location, Message

Job recommendation use cases

User experience	Optimization	Impact on job application rate
User intent / location	Homepage personalization	**2.5X**
User flow / user intent	Before vs. after having applied to a job	**7X**
User flow	LinkedIn homepage vs. Jobs homepage	**10X**
Location	Center rail vs. right rail	**5X**
Message	Followers vs. leaders	**2X**

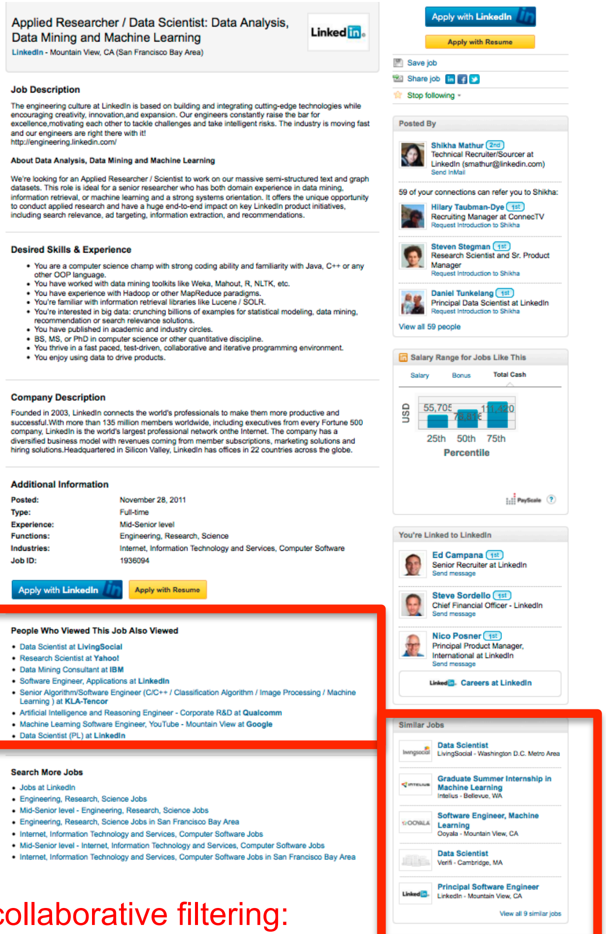

Item based collaborative filtering:
→ Follower audience

Content based:
→ Leader audience

User Intent Modeling

1. Intent labeling
    - Job seeker, recruiter, outbound professional, content consumer, content creator, networker, profile curator

2. Build normalized propensity model for each intent

Job Seeker Intent

- Features
  - Behavior: job searches, views & applications, job related email replies, has a job seeker sub, …
  - Social graphs: # of colleagues in network who recently left, …
  - Content: title, industry, anniversary date, …
- Propensity Score
  - Parametric accelerated failure time survival model
    - $log\ T_i = \Sigma_k \beta_k x_{ik} + \sigma \varepsilon_i$
  - Score: P(switch job next month)
- Evaluation
  - Gold standard
  - Directional validation

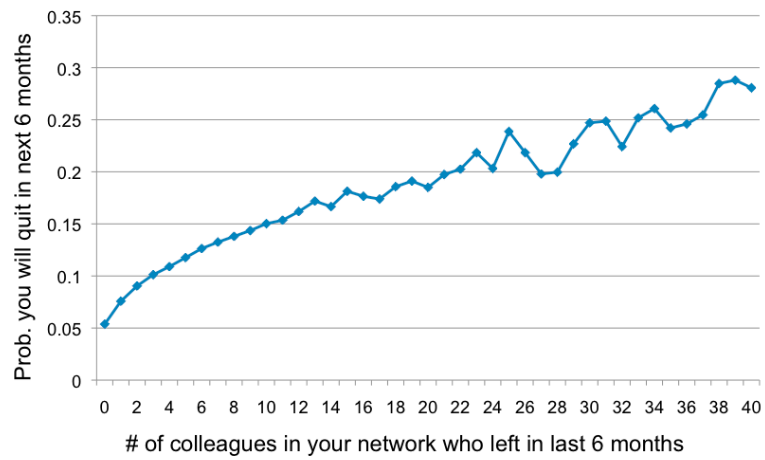

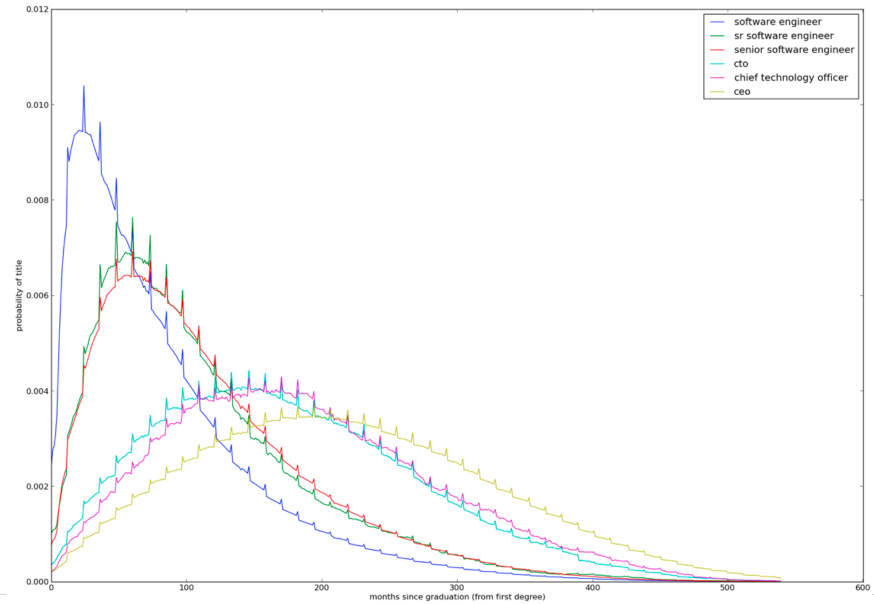

Lesson 3: Recommendations often cater to multiple competing objectives

Suggest relevant groups … that one is more likely to participate in

TalentMatch
(Top 24 matches of a posted job for sale)

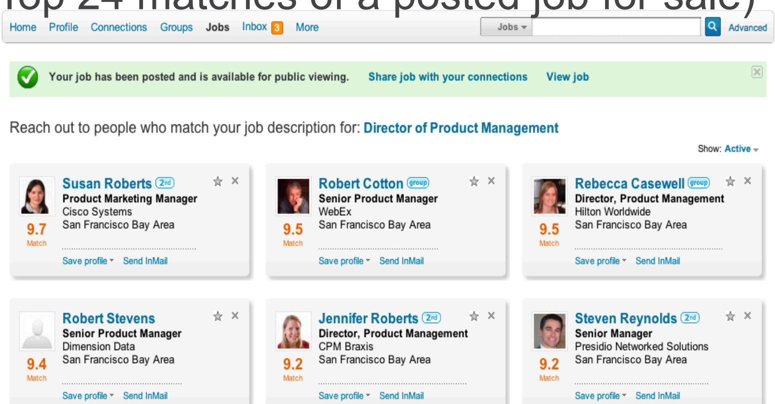

Suggest skilled candidates … who will likely respond to hiring managers inquiries

Semantic + engagement objectives

Multiple Objective Optimization

Constraint optimization problem

1. Rank top K’ > K results wrt to primary objective (e.g. relevance)
2. Perturb the ranking with a parametric function which leads to the inclusion of secondary objective(s) (e.g. engagement)
   - Measure the perturbation using a distance function wrt primary objective
   - Create a framework to quantify the tradeoff between the objectives

TalentMatch Use Case I

- Proxy for likelihood to answer hiring managers inquiries:
  - Job seekers are 16X more likely to answer than non-seekers

→ Increase % of job seekers in in TalentMatch results

- Control for TalentMatch booking rate and sent emails rate

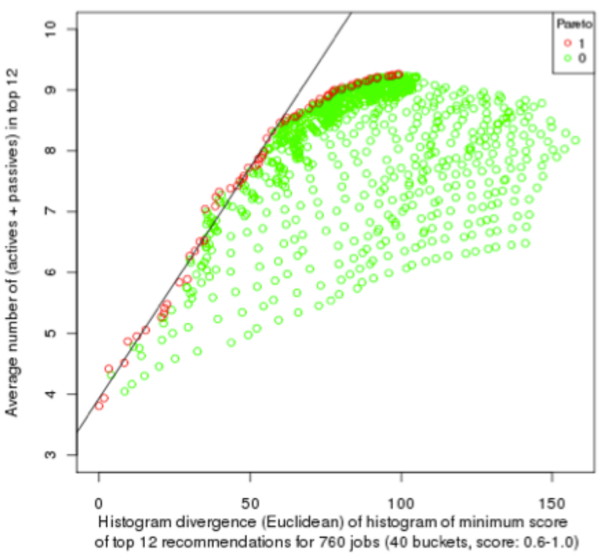

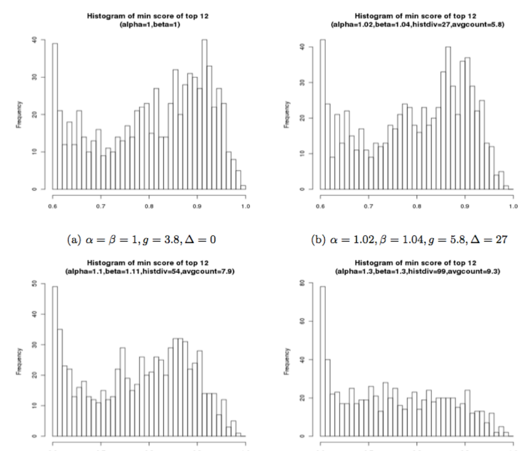

Distribution of min TalentMatch scores over 1 month of jobs posted on LinkedIn

TalentMatch Use Case II

Table 2: InMail response rate ratio of treatment group over control group. Data collected over a period of four weeks around May/27/2012-June/17/2012.

Treatment	Treatment/Control	95% CI
$\alpha = \beta = 1.15$	1.42	0.95 - 2.01
$\alpha = \beta = 1.07$	1.22	0.94 - 1.59

Table 3: Booking rate ratio of treatment group over control group. Data collected over a period of four weeks around May/27/2012-June/17/2012.

Treatment	Treatment/Control	95% CI
$\alpha = \beta = 1.15$	0.96	0.81 - 1.15
$\alpha = \beta = 1.07$	1.00	0.85 - 1.20

Table 4: Inmails/booking ratio of treatment group over control group. For perspective, the control group has approximately 2 InMails/booking. Data collected over a period of four weeks around May/27/2012-June/17/2012.

Treatment	Treatment/Control	95% CI
$\alpha = \beta = 1.15$	1.25	1.07 - 1.51
$\alpha = \beta = 1.07$	1.31	1.11 - 1.57

Data Sources for Feature Engineering

Social Graphs

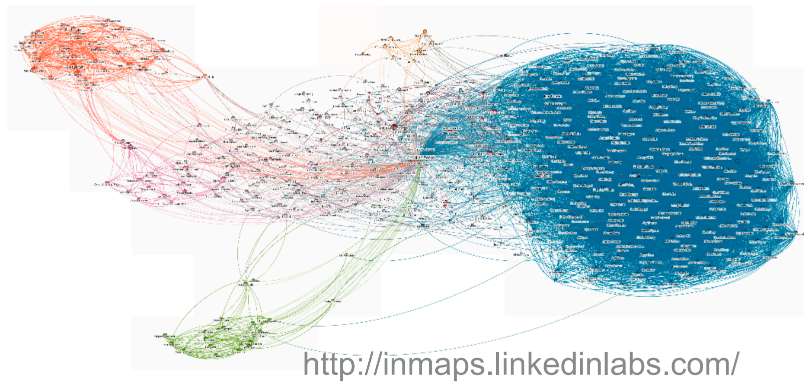

http://inmaps.linkedinlabs.com/

Content

. . .

Behavior

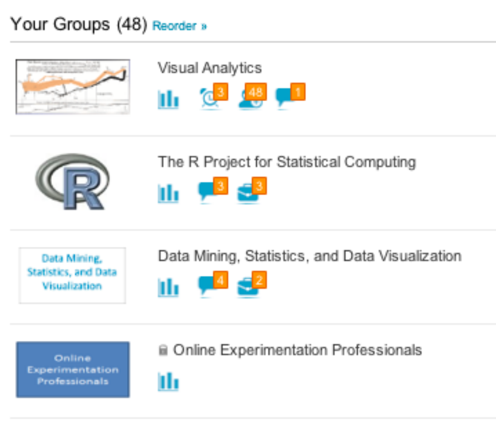

Lesson 4: The Unreasonable Effectiveness of Big Data

- Data jujitsu: slice and dice (smartly), then count
- Rich features engineered from profiles jujitsu
  - Job tenure distributions
  - Job transition probabilities
  - Related titles, industries, companies
  - Profile & job seniority
    - Impact on job recommendations:
      - 25% lift in views, viewers, applications, applicants
      - 90% drop in negative feedback

Lesson 4: The Unreasonable Effectiveness of Big Data (cont.)

- Region stickiness
- Related regions
  - Impact on job recommendations: 20% lift in views, viewers, applications, applicants
- Individuals propensity to migrate

Most vs Least sticky regions*

*what percentage of people stay in the same region when switching companies

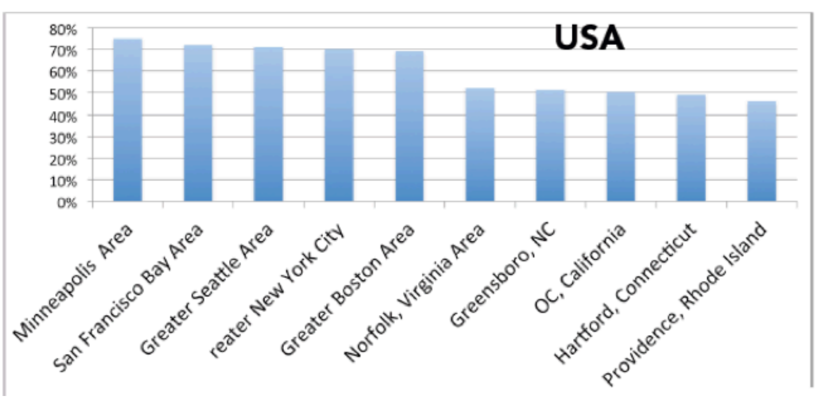

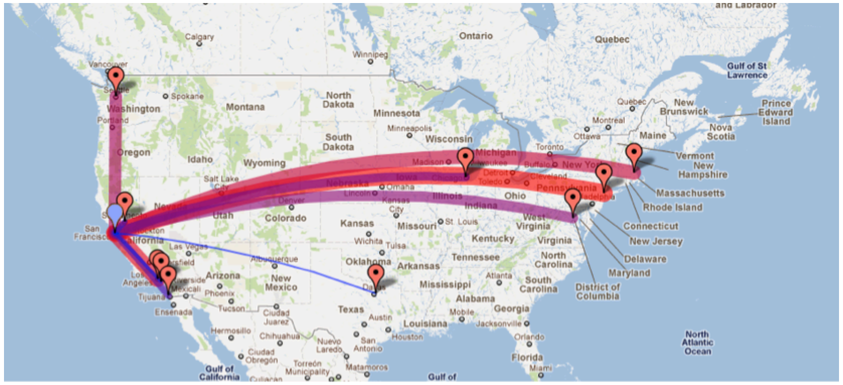

Lesson 5: Data Labeling Can Be Daunting

- Historical data
  - Similar objective
  - Unrelated processes, e.g., same session search selection
    - reduce presentation bias, position bias
    - What about intent bias?
- Random suggestions
  - Great with ads, company follows
  - Not for products with high cost of bad recommendations
    - jobs, alumni groups, …
  - Not for similar recommendations
- Crowdsourced manual labeling
  - Very challenging
  - Pairwise comparison more suited than absolute rating → higher cost
  - Expert-sourced manual labeling

Lesson 6: Measure Everything

1. **Implicit** user feedback
   - E.g. impressions, immediate actions (clicks), secondary actions
   - Understand and optimize flows, e.g.,
     - Impact of ‘See more’ link and landing page
     - Conversion rate of a job view into a job application from various channels
       - E.g. Homepage vs email vs mobile

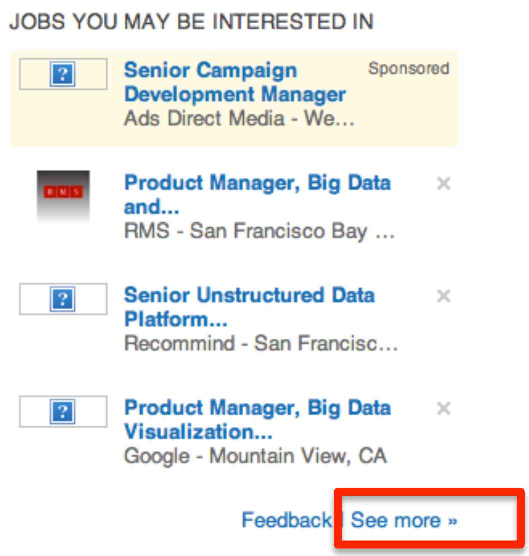

Lesson 6: Measure Everything (cont.)

2. **Explicit** user feedback

- Understand usefulness of the recommendation with unequal depth
  - Text based A/B test!
- Help prioritize future improvements
- Reveal unexpected complexities
  - E.g. ‘Local’ means different things for different members
- Prevent misinterpretation of implicit user feedback!

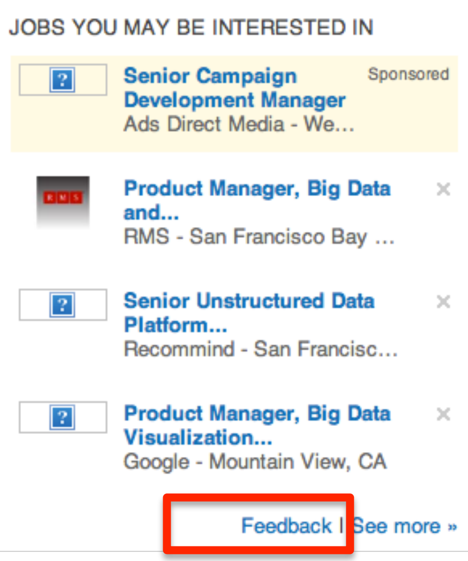

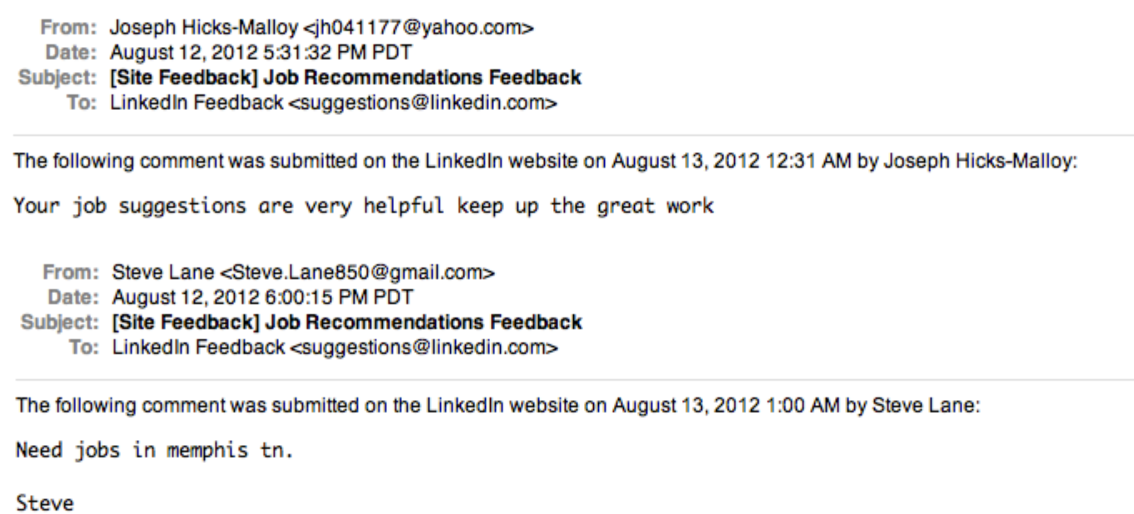

Lesson 7: Beware of some A/B testing pitfalls

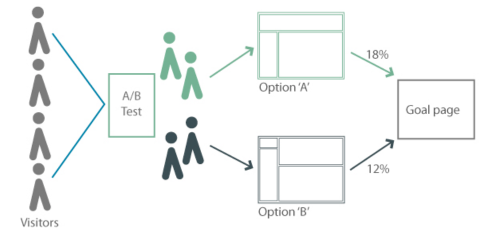

1. Novelty effect
   - E.g., new job recommendation algorithms have week-long novelty effect that shows lifts twice the stationary (real) one

Phase 3		
AlgorithmId in production	Explanation	PRODUCTION bucket
4	4001 + MinMax	0-34
5006	5005 + TBS-0.3	35-44
TBD	5005 + TBS-0.2	45-54
TBD	4001 + TBS-0.2	55-64
TBD	4001 + TBS-0.3	65-74
TBD	4001 + TBS-0.1	75-84
4	4001 + MinMax	85-99

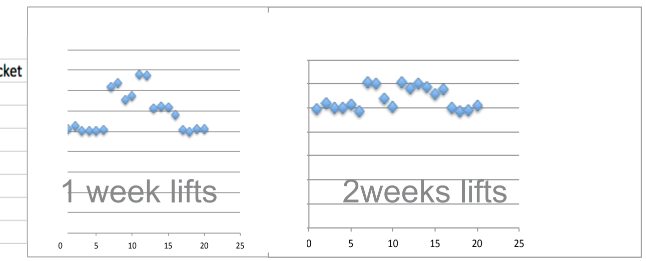

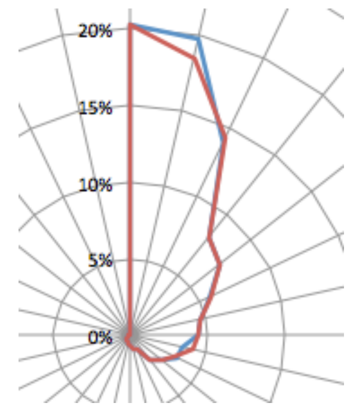

2. Cannibalization
   - Zero-sum game or real lift?

3. Random sampling destroys network effect

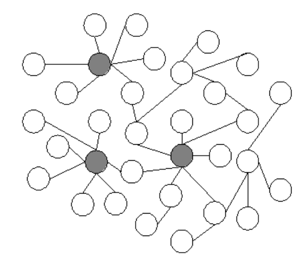

Lesson 8: One Unified Platform, One API

- Scaling innovation
  - Cross-leverage improvements between products
  - Shared knowledgebase

- Maintainability
  - Production serving and tracking
  - Infrastructure for complete upgrades

- Performance
  - Billions of sub second computations

Open Source Technologies

Zoie

Bobo

Kafka

Voldemort

http://sna-projects.com/sna/

Conclusion

- Social/professional networks are a new frontier for recommender systems
- Still many open questions:
  - How do we define and measure engagement?
  - What is the utility of a recommendation for the member?
  - What is the value of a recommendation for the network?
  - How do we reconcile utility and value when they conflict?
  - How do we network A/B test without tears?
  - …
- Learn as you go
  - Track everything and invest in forensics analytics
    - Breadth – understand holistic impact
    - Depth – understand flows

References

- M. Rodriguez, C. Posse and E. Zhang. 2012. Multiple Objective Optimization in Recommendation Systems. To appear in *Proceedings of the Sixth ACM conference on Recommender systems* (RecSys '12)
- M. Amin, B. Yan, S. Sriram, A. Bhasin and C. Posse. 2012. Social Referral : Using Network Connections to Deliver Recommendations. To appear in *Proceedings of the Sixth ACM conference on Recommender systems* (RecSys '12)
- A. Reda, Y. Park, M. Tiwari, C. Posse and S. Shah. 2012. Metaphor: Related Search Recommendations on a Social Network. 2012. To appear in *Proceedings of the 21st International Conference on Information and Knowledge Management* (CIKM '12)
- B. Yan, I. Bajaj and A. Bhasin. 2011. Entity Resolution Using Social Graphs for Business Applications. In *Proceedings of the International Conference on Advances in Social Network Analysis and Mining* (ASONAM '11)