Multistakeholder Recommendation

Part 1

Objectives

- You should come away from this part of the session
  - Knowing what multistakeholder recommendation is
  - Being able to identify multistakeholder issues that might arise in an application
  - Understanding a range of approaches for implementing and evaluating MSR

Who am I

- Recommender systems researcher since the mid-90s (before “recommender systems” was the accepted term)
- Currently: Professor in the Department of Information Science, University of Colorado, Boulder
  - Director of That Recommender Systems Lab (that-recsys-lab.net)
- Formerly: DePaul University in Chicago (2002-2018)
  - Co-led the Web Intelligence Lab with Bamshad Mobasher
- Current chair of the Steering Committee for the RecSys conference
- Organizing the Recommendation in Multistakeholder Environments worksh (RMSE 2019)
  - next week

We are looking for new PhD students!

There's also a diversity-focused post-doc at CU Boulder

Raise your hand if

- Have prior exposure to the concept of multistakeholder recommendation?
- Have prior exposure to the concept of fairness-aware recommendation?
- Have used a multistakeholder recommender system?
- Have used the Amazon.com web site?
- Have used Facebook?

Outline of these sessions

- Session I: Multistakeholder recommendation
  - Definitions
  - Challenges
  - Evaluation
  - Algorithms
- Session II: Fairness-aware recommendation
  - Definitions
  - Challenges
  - Evaluation
  - Algorithms

Stakeholder (definition)

- Comes from the literature on business management
  - A stakeholder in an organization is (by definition) any group or individual who can affect or is affected by the achievement of the organization's objectives.
    - (Freeman, 2010)
- For recommender systems, my definition
  - A recommendation stakeholder is any group or individual who can affected or is affected by the delivery of recommendations to users.
- Normally in recommender systems research
  - We consider only the user as a stakeholder
  - Optimize recommendations for ”user satisfaction”

Multistakeholder recommendation environment

- An environment / application where the requirements for recommendation generation
  - Include the perspectives of multiple parties
    - Not just the user
  - Example: computational advertising
    - User ~~wants~~ might respond to ads meeting their interests
    - Advertisers want users within an audience segment
    - Publishers want to maximize ad revenues

Isn’t this bad?

- “Recommendation should be all about the user”
- Two answers
  - This is already the case in many e-commerce settings
    - Filter out products that are out-of-stock
    - Optimize for time on site (do users really want that?)
    - Promote “house brands”
    - Promote new sellers / new items to overcome cold-start
    - Explore-exploit
    - Better to be transparent about the considerations
  - Without recognition of the multistakeholder nature of business
    - We can get into trouble: unfairness, bias, filter bubbles
    - Why not make these constraints explicit in our systems
      - Not something tacked on after the fact

Business Roundtable

- Recent report from the US Business Roundtable organization re-defining the role of a corporation:
- *"While each of our individual companies serves its own corporate purpose, we share a fundamental commitment to **all of our stakeholders**. We commit to:*
  - *Delivering value to our customers...*
  - *Investing in our employees...*
  - *Dealing fairly and ethically with our suppliers...*
  - *Supporting the communities in which we work...*
  - *Generating long-term value for shareholders...*
- *Each of our stakeholders is essential. **We commit to deliver value to all of them**, for the future success of our companies, our communities and our country."*
- Signed by 180 CEOs
  - Including Amazon, Apple, Oracle, SAP, etc.

Multistakeholder recommendation

- A multistakeholder recommender system is one in which the objectives of multiple parties, in addition to objectives attributed to the user, are considered in the computation of recommendations,
    - Especially a system in which such parties lie on different sides of the recommendation interaction.

Multisided platforms (MSPs)

- *Especially a system in which such parties lie on different sides of the recommendation interaction.*
- "Multisided platforms are technologies, products or services that create value primarily by enabling direct interactions between two or more customer or participant groups." (Hagiu, 2014)

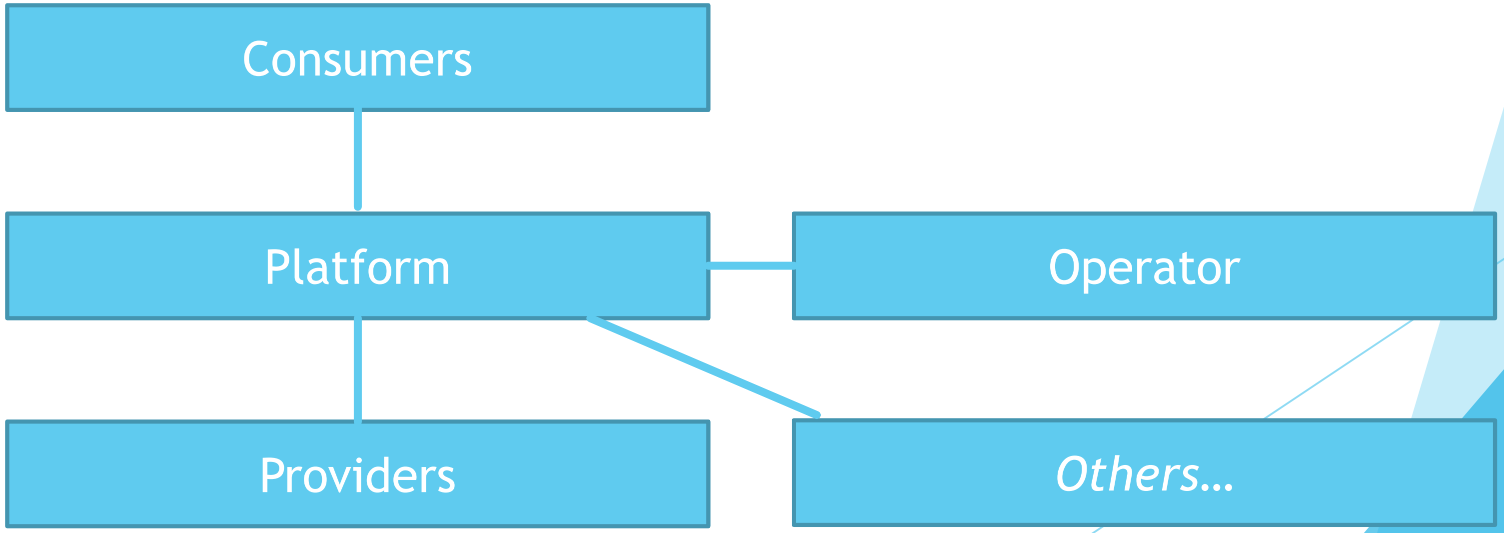

What questions so far?

Example

- Out of school time activities
- Wide range of options
  - entire summer
  - regular after-school
  - one-time events
- Wide range of providers
  - schools
  - non-profits
  - museums
  - sports clubs
- Audience: teens

Exercise

- Form groups of three or four
- Consider a system that recommends such activities
- List stakeholder groups
  - 15 minutes
- Be prepared to discuss your stakeholder groups

Possible confusion

- Isn’t it just the system designer / owner?
- People who are building the system decide what objectives to optimize
  - aren't they the only stakeholders?
- The idea of a stakeholder is about impact, not about control

Distinction

- Evaluation
  - A recommender system implementer can apply a metric that evaluates the impact of its recommendations on different stakeholders
  - To understand how different stakeholders are affected
- Optimization
  - An implementer can incorporate objectives related to different stakeholders as part of system optimization
  - Tuning the system to achieve particular stakeholder impact
- Involvement
  - An implementer can incorporate multiple stakeholders in the design of the system itself and the tradeoffs between different objectives

Related areas

- Group recommendation
  - Long history in recommender systems
  - Multiple parties receive the same recommendation
- Multistakeholder analysis of businesses
- Multisided platforms
- Matching markets
  - Reciprocal recommendation
- Computational advertising
- Non-accuracy methods in recommendation
  - long-tail
  - diversity

Key stakeholders

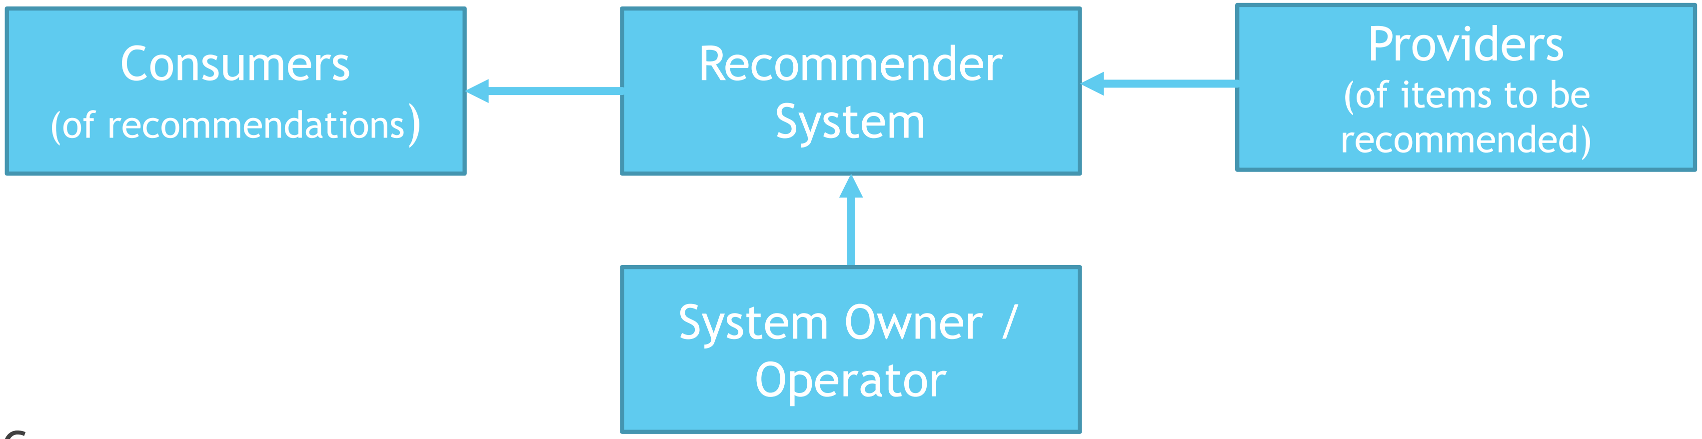

- Consumers
  - Individuals who get recommendations
- Providers
  - Entities who supply items that the system recommends
    - could be multiple parties on this side, depending on type of item
- System
  - Entity that operates the recommender system
- May also be "side stakeholders"
  - Example: different delivery services for items

Provider considerations

- Entities whose items are being recommended
  - maybe this is the system itself, not often not
- Types of objectives
  - Neutral: no related objective
    - We don't care how providers fare
  - Personalized: specific objectives for providers
    - Who are the "good" consumers?
- Types of interactions
  - Passive interaction: implicit feedback
    - Who does the provider accept?
  - Active interaction: provider specifies
    - Who does the provider want?

Examples

- AirBnB hosts
  - A host can decide whether or not to accept a potential guest
  - The system could learn from that as part of matching with users
  - passive
- On-line dating
  - The user specifies the type of match they are seeking
  - The system uses this information to match with users
  - active

Consumer considerations

- Objective will be personalized
  - otherwise not really a recommender system
  - "popular items"
- Types of interactions
  - Passive
    - "Welcome back! Here are some recommendations"
  - Active
    - User specifies a context, query or other information
    - Recommendations tailored to the input
    - Example:
      - search for Spanish restaurants, but the list is ordered in a personalized manner

Aside: interaction semantics

- Technical distinction between recommendation and personalized search
  - not much
  - many fielded systems have both characteristics
- Difference largely in user's mental model
  - what does the user think they are getting
    - answers to a question
    - suggestions based on history
- Much of this discussion applies regardless

System

- How does the system gain from recommendation interactions?
  - and does that depend on specific interactions?
- Neutral
  - the system doesn't care what is recommended and to whom as long as users are satisfied
  - Think MovieLens
- Aggregate
  - the system specifically gains from recommendations in some aggregate way
  - for example, a commission on sales
- Targeted
  - the system has its own objectives about what is recommended and to whom
    - and those objectives might not be shared by other participants
  - Example: fairness (more later)

Some notation

Interaction:
- = passive
+ = active
* = either

$$<C_p^+, P_n^-, S_n>$$

Stakeholder utility:
n = neutral
p = personalized
* = either

System utility:
n = neutral
a = aggregate
t = targeted
a,t = either aggregate or targeted
* = any

Design space

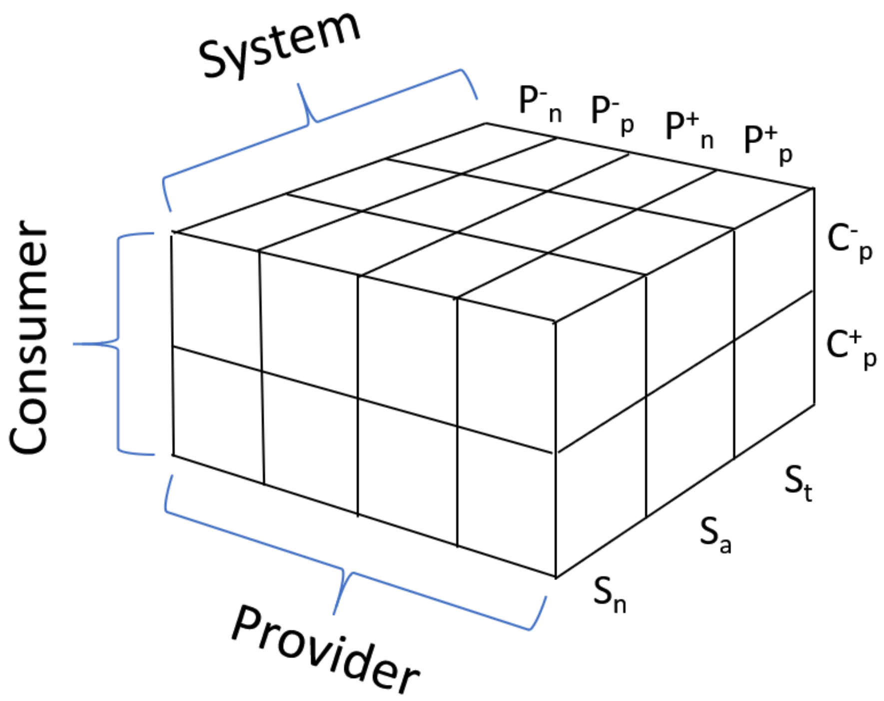

Abdollahpouri, et al. Research Directions in Multistakeholder Recommendation. 2019.

Examples

- On-line dating (reciprocal recommendation)
  - $<C^{+}_{p}, P^{+}_{p}, S_{n}>$
- Display advertising
  - $<C^{-}_{p}, P^{+}_{p}, S_{a}>$
- Social network recommendation (ala Daley et al. 2010)
  - $<C^{-}_{p}, P^{-}_{p}, S_{t}>$

The point

- There are a lot of different multistakeholder configurations
  - Not all solutions are applicable to every configuration
- Example
  - If you can model system utility as an aggregate of provider utility
    - via commission, for example
  - Then you don't have to worry about separate system objectives

What questions so far?

Exercise

- Consider the configuration
- $<C^{+}_{p}, P^{-}_{n}, S_{t}>$
- How does this map to the out-of-school activity recommender?
- Discuss

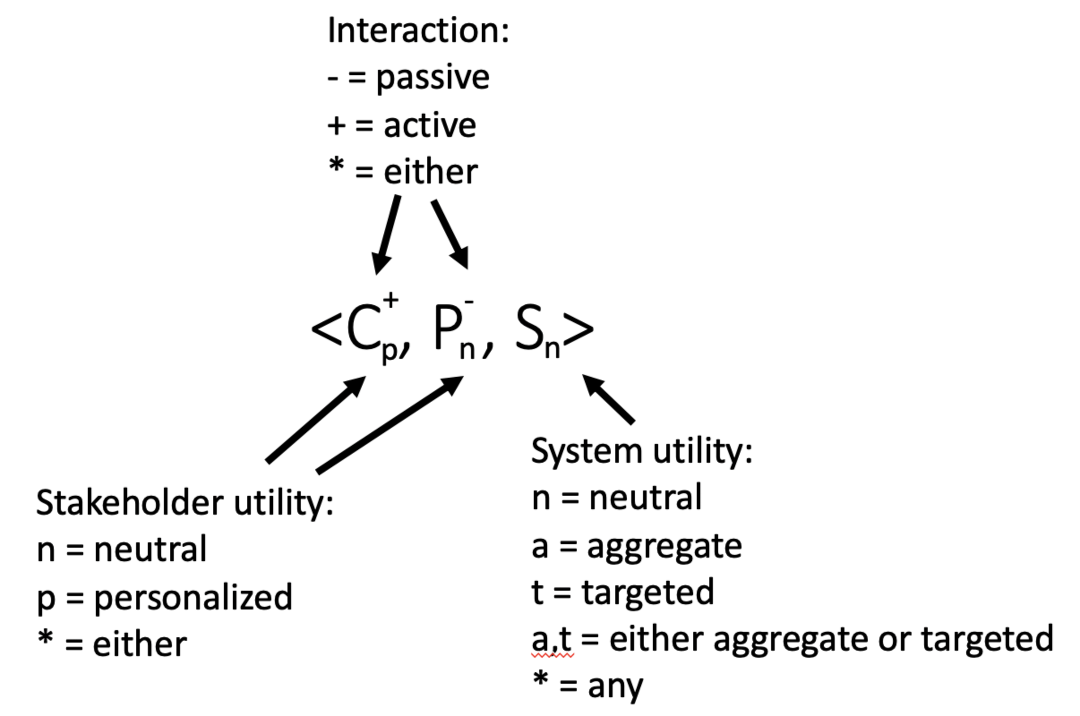

Implementations

- Multi-criteria optimization methods
- Re-ranking
- Note
  - Implementations often similar to techniques used for other non-accuracy metrics
    - diversity, coverage, etc.

Multi-criteria methods

- Combined optimization objective
  - Example: loss = α obj1 + (1 – α) obj2
- Sequential optimization
  - S_1 = opt(obj1)
  - S_2 = opt(obj2) but bound loss on obj1 (1% for example)

Combined objective

- Many examples
  - Recent one
  - Mehrotra et al. "Towards a Fair Marketplace", CIKM 2018
- Application: Playlist recommendation in Spotify
  - Stakeholders are users, artists
  - Users want accuracy recommendations
  - Artists want to be recommended

"Fair Marketplace", cont'd

- Algorithm: contextual bandit
  - learns to maximize reward $s_u^* = argmax_{s \in S_u} \left((1-\beta)\,\phi(u,s) + \beta\,\psi(s)\right)$
  - where ϕ is the relevance and ψ is the fairness
  - β controls the tradeoff

Regularization

- A combined objective
  - where the non-accuracy multistakeholder objective is treated as a regularization over the accuracy objective
  - We'll see an example when we talk about fairness

Sequential optimization

- Derive a solution for the accuracy objective
  - then solve for a second objective constraining the loss on the first
- Example
  - Agarwal, et al. "Click shaping to optimize multiple objectives", KDD 2011
- Application: content for the Yahoo! Front Page
  - Two different KPIs
    - click-through rate
    - time on site (stickiness)
  - These are both about the user (maybe?)
  - However, paper discusses other system stakeholder objectives such as site revenue

"Click shaping"

- Algorithm: Bayesian estimation of CTR
  - followed by constrained optimization of the time-spent metric
  - using linear programming

Multi-criteria methods

- Requires resolving some tricky tuning issues
- Combined objectives
  - require setting a weight on the outcomes for different stakeholders
- Regularization
  - distorts the optimization space, can cause significant accuracy loss
- Sequential optimization
  - have to decide what is an appropriate average accuracy loss
    - distributional control might be better

Re-ranking

- Produce recommendation lists in the usual way
  - optimized for user stakeholders
- Then re-rank to balance original ranking vs other stakeholders' objectives
- Example
  - Sürer, Özge, Robin Burke, and Edward C. Malthouse. "Multistakeholder recommendation with provider constraints". RecSys 2018.
- Application: Recommendation in a multi-supplier marketplace
  - User stakeholders
  - Suppliers want a share of recommendations delivered

"Multi-supplier"

- Algorithm: agnostic to initial algorithm
  - user-based algorithm computes all recommendation lists for all users
  - define desired optimal provider exposure as an integer programming problem
  - use a Lagraingian relaxation of IP to achieve scalability

Provider-side metrics

- Lots of literature on measuring outcomes for users
- How do we measure outcomes for provider?
  - application-specific
  - what matters to these folks?
- Several ideas
  - exposure: people see their products / items
  - audience: who sees their items
  - quality: prediction outcomes

Exposure

- Count the number of recommendations of the provider's items across some set of recommendation lists
  - doesn't matter whether the user is interested

$$\sum_{L_i \in \mathcal{L}} \sum_{j \in L_i} \mathbb{1}(j \in I_p)$$

- Count the number of **Hits** (that is recommendations matching the test data)
  - doesn't work so well for cold-start providers

$$\sum_{L_i \in \mathcal{L}} \sum_{j \in L_i} \mathbb{1}(j \in I_p \wedge r_{ij} \in T)$$

- Could normalize by
  - the number of lists
  - the size of provider's catalog
- Can also take rank into account

Audience

- Count how many users are reached by the provider's items; they see at least one item
  - Quality of match not included

$$\sum_{L_i \in \mathcal{L}} \mathbb{1}(|I_p \cap L_i| > 0)$$

- Count how many users in some targeted group g see the recommendations

$$\sum_{L_i \in \mathcal{L}} \mathbb{1}(|I_p \cap L_i| > 0 \wedge g_p(i))$$

- Variants
  - Use hits instead of just counts

Accuracy

- How accurate are the recommendations of the provider's items when they are given
  - Could use any metric: RMSE, nDCG, etc.
  - Might be deceptive if $|T_p|$ is small

$$[\sum_{r_{ij} \in T_p} m(r_{ij}, \hat{r}_i j)]/|T_p|$$

What questions so far?

Exercise

- Think back to our recommender for out of school activities
- What metric(s) would be appropriate for providers?
- If you think multiple metrics are appropriate,
  - how should tradeoffs between them be managed?

Conclusion

- Multistakeholder recommendation
  - the objectives of multiple parties are considered in making recommendations,
    - when the parties lie on different sides of the recommendation interaction.
- Multistakeholder recommendation is a necessity in many applications
  - It is not "bad" to incorporate the stakeholders other than the user
  - If a system has benefits for users, system viability is in their interest
- Multistakeholder recommendation approaches
  - Multi-criteria optimization
    - incl. regularization
  - Re-ranking

Open questions

- Transparency
  - How to surface / explain multistakeholder aspects of recommendation?
  - Esp. to consumers who might resist such aspects
- Tradeoffs
  - How to define and evaluate tradeoffs?
- Multistakeholder UX
  - How do different stakeholders interact with the system?

References

- Agarwal, et al. Click shaping to optimize multiple objectives. KDD 2011
- Business Roundtable, Statement on the Purpose of a Corporation. https://opportunity.businessroundtable.org/wp-content/uploads/2019/08/BRT-Statement-on-the-Purpose-of-a-Corporation-with-Signatures.pdf, August 2019
- Daly, E., Werner Geyer, and David R. Millen. The network effects of recommending social connections. RecSys 2010.
- Freeman, R. E. *Strategic management: A stakeholder approach*. Cambridge university press, 2010.
- Hagiu, A. *Strategic decisions for multisided platforms*. MIT, 2014.
- Mehrotra et al. Towards a Fair Marketplace: Counterfactual Evaluation of the trade-off between Relevance, Fairness & Satisfaction in Recommendation Systems. CIKM 2018
- Sürer, Ö et al. Multistakeholder recommendation with provider constraints. RecSys 2018.

Break