User Modeling on the World Wide Web

Andrew Tomkins

Google, Inc

Scope note

- All opinions are mine, not Google's
- Data described in this presentation is all from the published literature.

Privacy

- This talk could focus entirely on privacy
- …But it won’t
- …But maybe it should

“If I have 3 million customers on the web, I should have 3 million stores on the web.”

Jeff Bezos

Ceo, Amazon.com

Goal Revelation

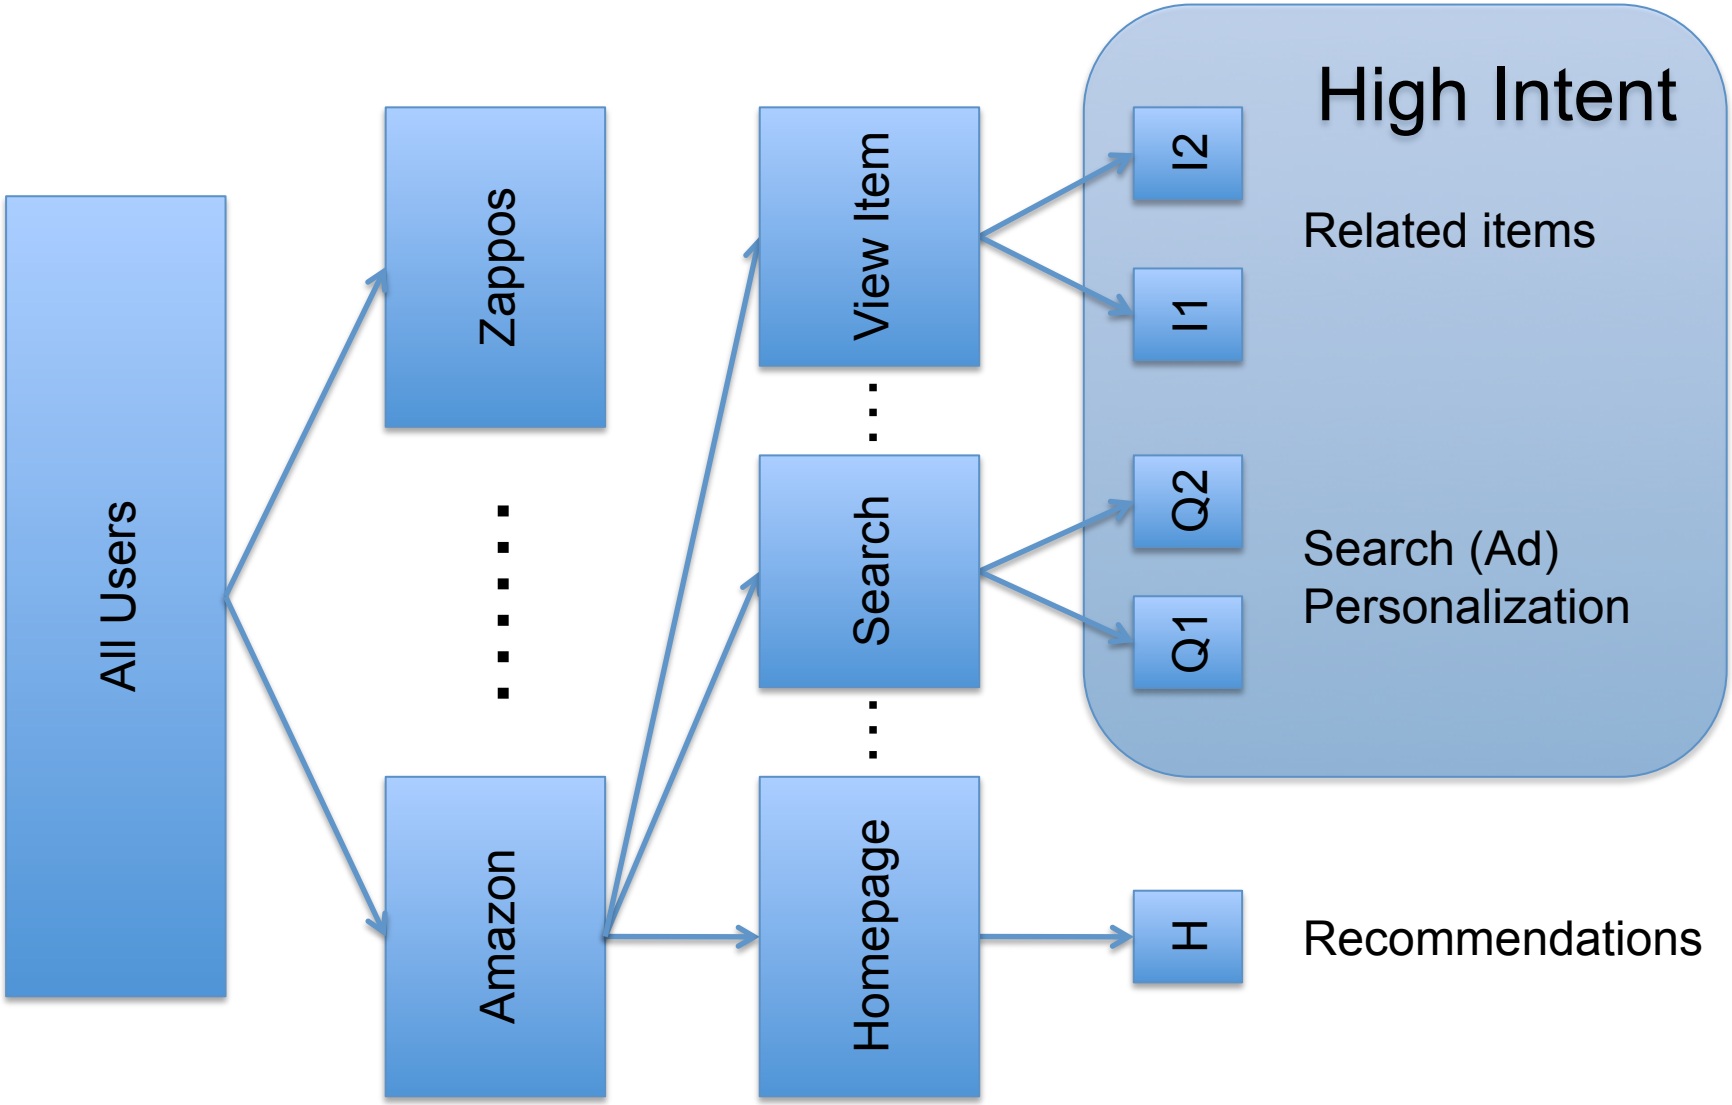

Personalization and Queries

- Basic flow
  - User logs onto computer (initial entropy H_0)
  - User visits site (destination entropy H_d)
  - User expresses intent (intention entropy H_i)
  - User selects result (final entropy $H_f = 0$)
- Search vs Display:
  - Display 16x volume
  - Search 2x revenue
  - H_i much smaller for search than display

Taxonomy of Personalization

Profile Personalization

- Idea: User-settable features significantly determine site behavior
- Examples:
  - "Themes" on many websites like Yahoo, Google, many mobile devices, etc
  - Explicit interests, as in e.g. Quora
  - Portfolio list in finance sites
  - Location

Data Personalization

- Idea: user id significantly determines site behavior
- Examples:
  - Email providers
  - Social update streams
    - facebook, myspace, twitter, etc

Model-Based Personalization

- Idea: rich user model informs presentation of page
- Examples
  - Amazon.com recommendations
  - pandora.com music genome recommendations
  - facebook.com “best view” news feed

Summary of Page Types

- No personalization: same page for every user
- Profile personalization: User-settable features significantly determine site behavior
- Data personalization: user id significantly determines site behavior
- Model personalization: rich user model informs presentation of page

Categories of pageviews

Main category	Sub-category	Count
Communications		
	Social	24.3%
	Mail	9.4%
	Forum	1.4%
	Blog	0.4%
	Total	35.5%
Content		
	Game	6.2%
	MM	5.4%
	Portal	5.4%
	Head Listings	3.4%
	News	3.4%
	Other Vertical	28.1%
	Total	52.0%
Search		
	Main	6.2%
	MM	1.4%
	Item	1.4%
	Total	9.0%
Unknown	Total	3.4%

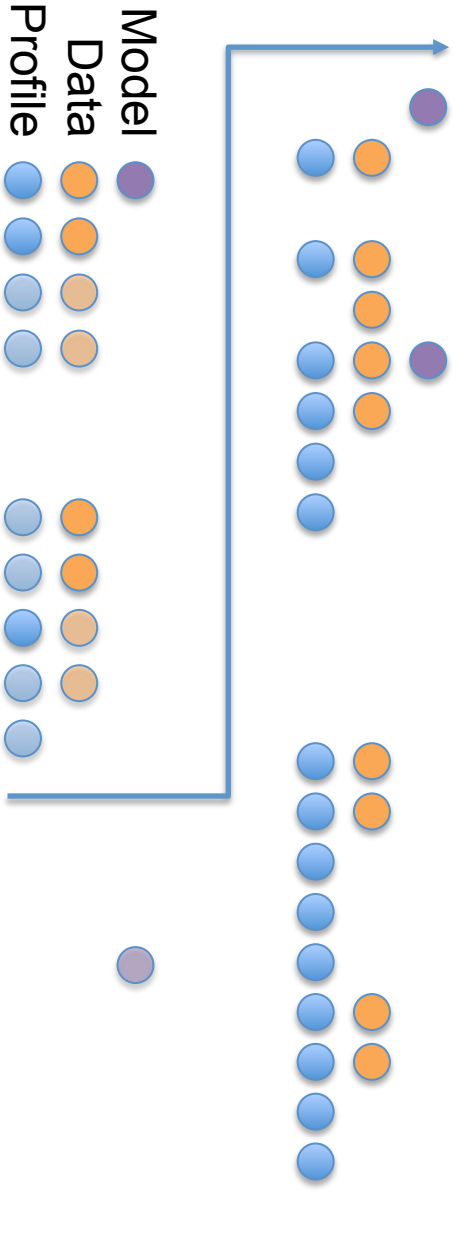

Vertical	Count
Retail	3.4%
Travel	1.8%
Other	1.6%
Finance	1.4%
Education	1.2%
Personals	1.0%
Jobs	1.0%
Services	1.0%
B2B	1.0%
Social	0.8%
Entertainment	0.8%
Mobile	0.8%
Reference	0.8%
Sports	0.8%
Real estate	0.6%
Movies	0.6%
Auto	0.6%
TV	0.6%
Local	0.6%
Radio	0.4%
Food	0.4%
Health	0.4%
Government	0.4%

Breakdown (rough estimate)

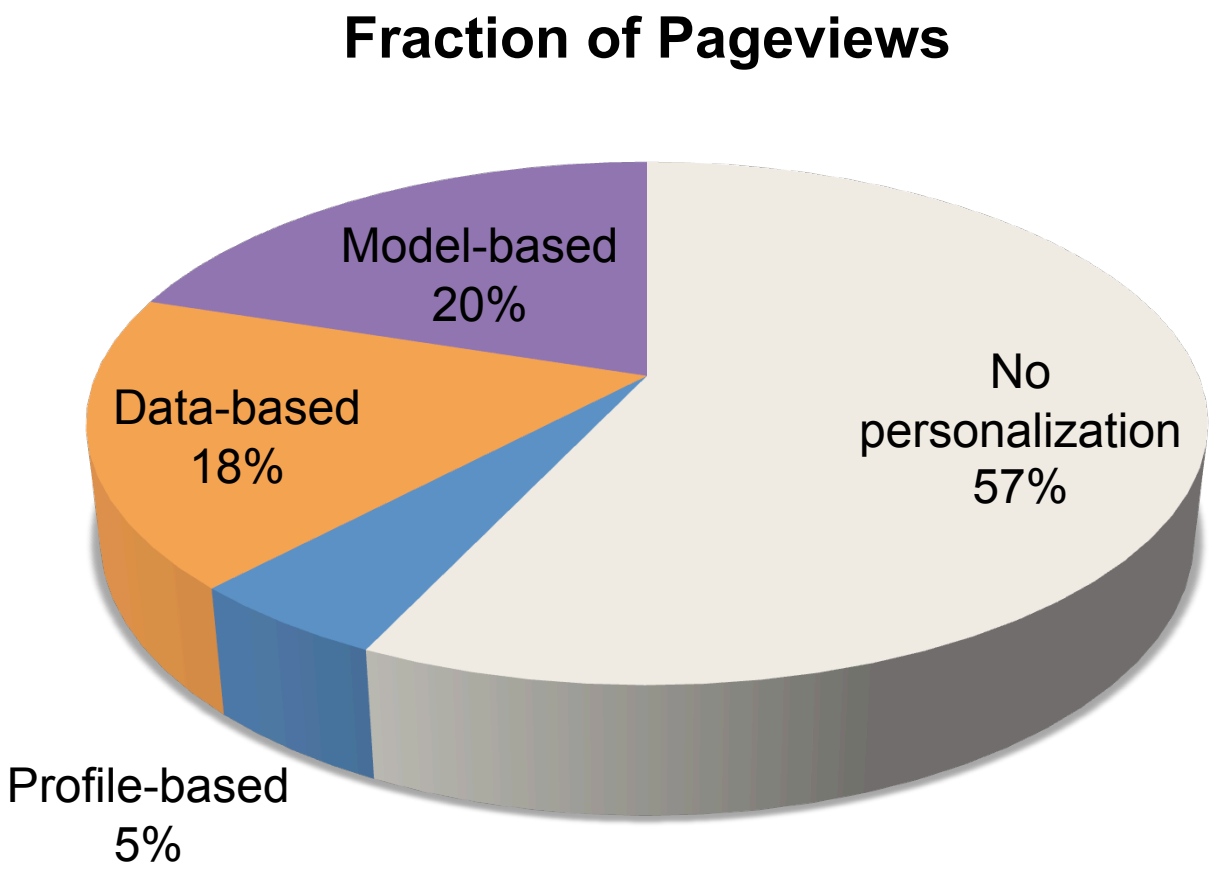

Potential for Personalization

Trends

- More walled gardens and cleaner content – more ability to provide model-based personalization
- More value to logged-in users?

Value Assessment

- Pick any webpage. Could a web-savvy "personal butler" knowing everything about you improve the page?
  - If so, this page is a candidate for model-based personalization.

Personalization by Problem Domain

- Advertising
- Content Optimization
- Search
- Recommendations

Advertising personalization

- [Chen et al] KDD09:
  - Traffic-based behavioral targeting for display ads
  - 20% lift in CTR
- [Yan et al] WWW09:
  - Query-based behavioral targeting
  - 6.7x lift in CTR

Content Optimization

- 40% CTR lift (Agarwal et al, NIPS08)
  - Click dynamics much more important than personalization
  - Other work showed a 13% max lift due to segment personalization in this setting

Search Personalization

- Teevan, Dumais and Horwitz show that committing to a ranking for even 6 people is 46% worse than customizing the ranking for each (under NDCG)
- And search is “high intent,” so difficult to personalize

Recommender Systems

- Koren (NetFlix Grand Prize writeup) reports these RMSE values:
  - Baseline item and user features: 0.96
  - Temporal item and user features: 0.92
  - Full model: 0.86
- [Jahrer et al] KDD10
  - Training: ~2hrs -> ~160hrs
  - Performance: improve RMSE in 4th significant digit
- Chen et al, CHI2010, “Short and Tweet”:
  - 33% interesting -> 72% interesting

Item-item Systems

- Not necessarily model-based personalization, but uses rich intent information
  - Eg, amazon related items, youtube related videos

Obvious Omissions

- Mobile personalization
- Social network advertising
- Social network content recommendation

Research landscape

Problems in Personalization

- Item Recommendation
- Task completion tools
- Information delivery
- State management

Item Recommendation Breakdown

- Serendipitous Content Discovery:
  - Input: user, discovery constraints
  - Output: Interesting stuff
  - Examples: portals, social networks, multimedia
- Task-Specific Recommendations
  - Input: user, task
  - Output: recommendations
  - Examples: search, shopping, local

Item Recommendation Problems

- Many problems are hybrids of serendipitous content discovery and task completion
  - Shopping item recommendation
  - Advertising (search, contextual, display)
- Uniform representation for all these different item types?

Rich Task Completion

- Examples
  - Travel planning, events, considered purchase
- Challenge:
  - Most domains are small (see previous slide)
  - Requires significant domain customization
  - Difficult to justify investment

Information Delivery

- [Broder], many venues

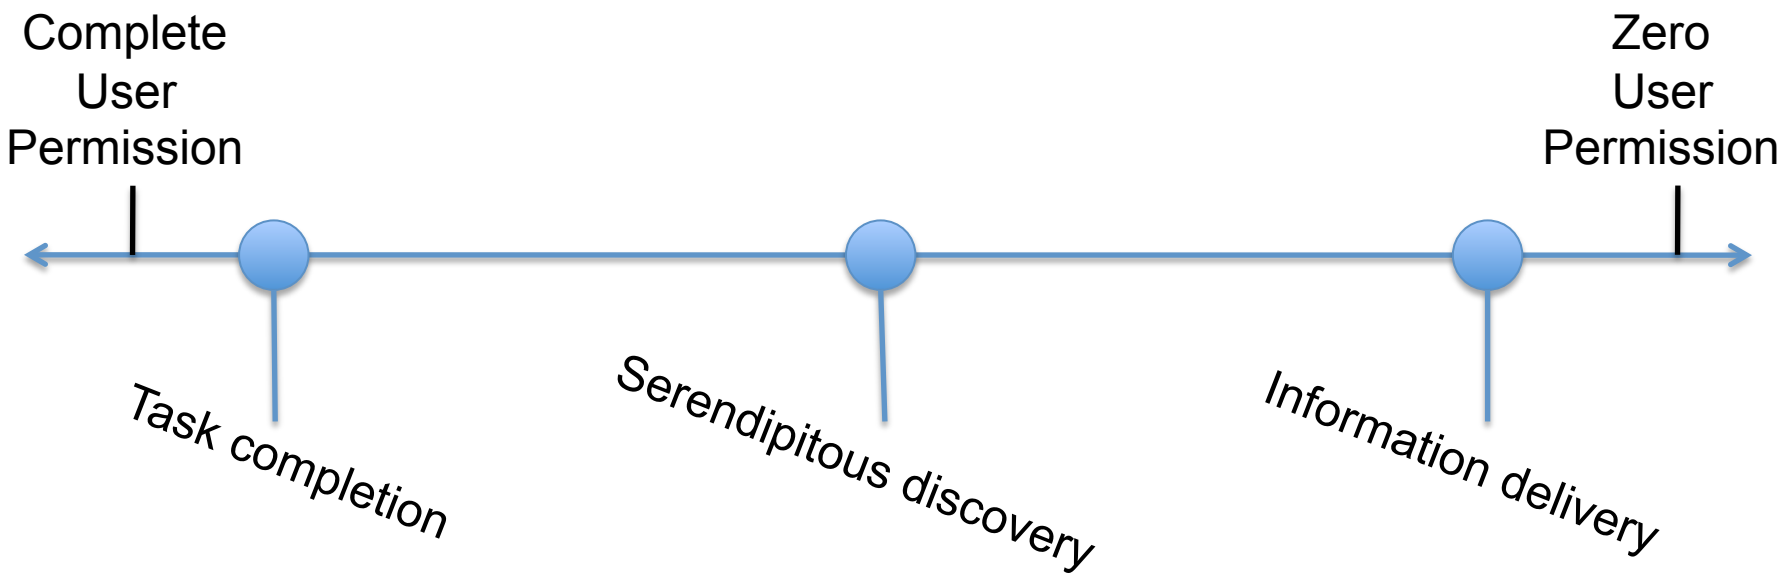

End States

- State 1: local state. User behavior is aggregated separately by each provider, leading to:
  - Fragmentation
  - Oligopoly in each segment due to a “rich get richer” pattern of customized experience
- State 2: shared state. User behavior is combined and sent to a user-specified profile service, made available to all sites

Shared State Personalization

- Develop the right feature spaces:
  - Rich enough to capture niche interests/prefs
  - Structured enough to reason about
  - Examples
    - Raw clicks, queries, visits
    - Categories, topics, entities, normalized representations
    - Preferences for language, multimedia, sophistication, style…
- Securely store and deliver user behavioral/ profile data
- Give user comprehensible control of delivery

Modeling Internal User State

- Work / home
- Research / entertainment
- Moods
- provider architecture, incentive schemes. Start with profiles, move to behavior?

Bigger Questions

- Metrics: CTR, Happiness, Social welfare, User education / capability / knowledge
  - User should choose
  - We don’t know how to do this yet
- Longitudinal optimization
  - Current focus is largely single-page optimization
  - How much do we leave on the table?

The End

Questions and discussion