



Latent Factor Models for Web Recommender Systems

Bee-Chung Chen
Deepak Agarwal, Pradheep Elango, Raghu Ramakrishnan
Yahoo! Research & Yahoo! Labs

Outline

- Overview of recommender problems at Yahoo!
- Basics of matrix factorization
- Matrix factorization + feature-based regression
- Matrix factorization + topic modeling
- Matrix factorization + fast online learning
- Research problems beyond factor models
 - Explore/exploit (bandit problems)
 - Offline evaluation
 - Multi-objective optimization
 - Whole-page optimization



Web Recommender Systems

Recommend **items** to **users** to maximize some **objective(s)**



Web Search

My Yahoo! | Make Y! your homepage

Sign In | New here? Sign Up | Have something to share? | Page Options ▾

YAHOO! SITES

Edit

- Mail
- Autos
- Chat
- Fantasy Sports
- Finance
- Games
- Horoscopes
- HotJobs
- Maps
- Messenger
- Movies
- omg!
- Personals
- Shopping
- Sports
- Travel
- Updates
- Weather

More Yahoo! Sites

MY FAVORITES

Edit

- eBay
- Facebook
- Twitter

TODAY - July 14, 2010



World Cup octopus could make millions

Paul the octopus is in high demand after a perfect run of predicting soccer game winners. » Possible opportunities

More on the octopus
Cup winners and losers
U.S.'s top moments



Salsa tied to food illness



Octopus could be worth millions



Lottery winner rich in mystery



High schooler's impressive dunk

5 - 8 of 28

NEWS WORLD LOCAL FINANCE

- 9 killed, 10 missing as typhoon lashes Philippines | Photos
- Testing delayed on tighter cap for Gulf oil well | Photos
- W.Va. mine disaster prompts bill to toughen worker safety rules
- Military won't establish 'separate but equal' housing for gays
- Small banks struggling despite gov't bailouts, watchdog reports
- Tiny mushroom blamed for 400 deaths in southwest China
- CHP pursuit ends in two-car crash in San... - S.J. Mercury N...
- Oakland talks break down: layoffs for 80... - S.F. Chronic

TRENDING NOW

1. Kourtney Kardash...
2. Anna Chapman
3. Al Pacino
4. French Toast Rec...
5. Nina Garcia
6. Susan Boyle
7. Job Search
8. Yogi Berra
9. Philippines Typh...
10. Sunscreen

Recommend search queries

Recommend packages:
Image
Title, summary
Links to other pages

Pick 4 out of a pool of K
 $K = 20 \sim 50$
to maximize clicks

Routes traffic other pages

Recommend news article

Recommend applications

Web Recommender Systems

- Goal
 - Recommend **items** to **users** to maximize some **objective(s)**
- A new scientific discipline that involves
 - Machine Learning & Statistics (for learning user-item affinity)
 - Offline Learning
 - Online Learning
 - Collaborative Filtering
 - Explore/Exploit (bandit problems)
 - Multi-Objective Optimization
 - Click-rates (CTR), time-spent, revenue
 - User Understanding
 - User profile construction
 - Content Understanding
 - Topics, “aboutness”, entities, follow-up of something, breaking news,...



YAHOO! SITES ⚙ Edit

- Mail
- Autos
- Chat
- Fantasy Sports
- Finance
- Games
- Horoscopes
- HotJobs
- Maps
- Messenger
- Movies
- omg!
- Personals
- Shopping
- Sports
- Travel
- Updates
- Weather

More Yahoo! Sites

MY FAVORITES ⚙ Edit

- eBay
- Facebook
- Twitter

TODAY - July 14, 2010



World Cup octopus could make millions

Paul the octopus is in high demand after a perfect run of predicting soccer game winners. [» Possible opportunities](#)

More on the octopus
Cup winners and losers
U.S.'s top moments



Salsa tied to food illness



Octopus could be worth millions



Lottery winner rich in mystery



High schooler's impressive dunk

5 - 8 of 28

NEWS WORLD LOCAL FINANCE

- 9 killed, 10 missing as typhoon lashes Philippines | Photos
 - Testing delayed on tighter cap for Gulf oil well | Photos
 - W.Va. mine disaster prompts bill to toughen worker safety rules
 - Military won't establish 'separate but equal' housing for gays
 - Small banks struggling despite gov't bailouts, watchdog reports
 - Tiny mushroom blamed for 400 deaths in southwest China
 - CHP pursuit ends in two-car crash in San... - S.J. Mercury N...
 - Oakland talks break down; layoffs for 80... - S.F. Chronic...
 - Stanford grad student dies in Yosemite... - Mountain Vie...
- NBA · NHL · MLB · Tennis · Golf · Soccer · NASCAR

updated 01:49 am

More: [News](#) | [Popular](#) | [Buzz](#)

TRENDING NOW

1. Kourtney Kardash...
2. Anna Chapman
3. Al Pacino
4. French Toast Rec...
5. Nina Garcia
6. Susan Boyle
7. Job Search
8. Yogi Berra
9. Philippines Typh...
10. Sunscreen

Recommend packages:

Image

Title, summary

Links to other pages

Pick 4 out of a pool of K


$K = 20 \sim 50$

to maximize clicks

Routes traffic other pages

CTR Curves for Two Days on Yahoo! Front Page

Each curve is the CTR of an item in the Today Module on www.yahoo.com over time



Traffic obtained from a controlled randomized experiment (no confounding)

Things to note:

(a) Short lifetimes, (b) temporal effects, (c) often breaking news stories



Problem Definition



Algorithm selects item j with item features \mathbf{x}_j
(keywords, content categories, ...)



User i visits with user features \mathbf{x}_i
(demographics, browse history, geo-location, search history, ...)

(i, j) : response y_{ij}
(click/no-click)

Which item should we select?

- The one with highest predicted CTR **Exploit**
- The one most useful for improving the CTR prediction model **Explore**

Our Strategy

	Most Popular Recommendation	Personalized Recommendation
Offline Learning		Collaborative filtering profile construction [KDD'09, WSDM'10]
Online Learning	Time-series models [WWW'09]	Online regression [NIPS'08]
Intelligent Initialization	Prior estimation	Prior estimation, dimension reduction [KDD'10]
Explore/Exploit	Multi-armed bandits [ICDM'09]	Bandits with covariates [Li, WWW'10]



Model Choices

- Feature-based (or content-based) approach
 - Use features to predict response
 - User features: Age, gender, geo-location, visit pattern, ...
 - Item features: Category, keywords, topics, entities, ...
 - Linear regression, Bayes Net, SVM, tree/forest methods, mixture models, ...
 - Bottleneck: Need predictive features
 - Difficult to capture signals at granular levels: Cannot distinguish between users/items having same feature vectors
- Collaborative filtering (CF)
 - Make recommendation based on past user-item interaction
 - User-user, item-item, matrix factorization, ...
 - See [Adomavicius & Tuzhilin, TKDE, 2005], [Konstan, SIGMOD'08 Tutorial]
 - Good performance for users and items with enough data
 - Does not naturally handle new users and new items (**cold-start**)



Factorization Methods

- Matrix factorization

- Model each user/item as a vector of **factors** (learned from data)

rating that user i
gives item j




$$y_{ij} \sim \sum_k u_{ik} v_{jk} = u'_i v_j \Leftrightarrow \begin{matrix} Y \\ M \times N \end{matrix} \sim \begin{matrix} U \\ M \times K \end{matrix} \begin{matrix} V \\ K \times N \end{matrix}$$

factor vector of user i factor vector of item j

$K \ll M, N$

M = number of users

N = number of items



Factorization Methods

- Matrix factorization

- Model each user/item as a vector of **factors** (learned from data)

rating that user i
gives item j



$$y_{ij} \sim \sum_k u_{ik} v_{jk} = u'_i v_j \quad \Leftrightarrow \quad \begin{matrix} Y \\ M \times N \end{matrix} \sim \begin{matrix} U \\ M \times K \end{matrix} \begin{matrix} V \\ K \times N \end{matrix}$$

factor vector of user i factor vector of item j

$K \ll M, N$

M = number of users

N = number of items

- Better performance than similarity-based methods [Koren, 2009]
 - No factor for new items/users, and expensive to rebuild the model!!
- How to prevent overfitting
 - How to handle cold-start
 - Use **features** (given) to predict the **factor values**



How to Prevent Overfitting

- Loss minimization
- Probabilistic model

$$\ell(\mathbf{u}, \mathbf{v}) =$$

$$\frac{1}{2\sigma^2} \sum_{(i,j)} (y_{ij} - \mathbf{u}'_i \mathbf{v}_j)^2$$

$$+ \frac{1}{2\sigma_u^2} \sum_i \|\mathbf{u}_i\|^2$$

$$+ \frac{1}{2\sigma_v^2} \sum_j \|\mathbf{v}_j\|^2$$

$$y_{ij} \sim N(\mathbf{u}'_i \mathbf{v}_j, \sigma^2)$$

$$\mathbf{u}_i \sim N(0, \sigma_u^2 \mathbf{I})$$

$$\mathbf{v}_j \sim N(0, \sigma_v^2 \mathbf{I})$$

Given $\sigma^2, \sigma_u^2, \sigma_v^2$, find

$$\arg \min_{\mathbf{u}, \mathbf{v}} \ell(\mathbf{u}, \mathbf{v}) \quad \overset{\text{equivalent}}{\longleftrightarrow} \quad \arg \max_{\mathbf{u}, \mathbf{v}} \Pr[\mathbf{u}, \mathbf{v} \mid \mathbf{y}]$$

How to set $\sigma^2, \sigma_u^2, \sigma_v^2$?



Probabilistic Matrix Factorization

- Probabilistic model

$$y_{ij} \sim N(u_i'v_j, \sigma^2)$$

$$u_i \sim N(0, \sigma_u^2 I)$$

$$v_j \sim N(0, \sigma_v^2 I)$$

$$\text{Let } \Theta = (\sigma^2, \sigma_u^2, \sigma_v^2)$$

$$\log \Pr(\mathbf{y}, \mathbf{u}, \mathbf{v} \mid \Theta) = \text{constant}$$

$$-\frac{1}{2\sigma^2} \sum_{(i,j)} (y_{ij} - u_i'v_j)^2 - R \log \sigma^2$$

$$-\frac{1}{2\sigma_u^2} \sum_i \|u_i\|^2 - Mr \log \sigma_u^2$$

$$-\frac{1}{2\sigma_v^2} \sum_j \|v_j\|^2 - Nr \log \sigma_v^2$$

How to determine Θ ?

–Maximum likelihood estimate

$$\arg \max_{\Theta} \Pr(\mathbf{y} \mid \Theta) = \arg \max_{\Theta} \int \Pr(\mathbf{y}, \mathbf{u}, \mathbf{v} \mid \Theta) d\mathbf{u} d\mathbf{v}$$

–Use the EM algorithm



Model Fitting: EM Algorithm

- Find

$$\hat{\Theta} = \arg \max_{\Theta} \Pr(\mathbf{y} | \Theta) = \arg \max_{\Theta} \int \Pr(\mathbf{y}, \mathbf{u}, \mathbf{v} | \Theta) d\mathbf{u} d\mathbf{v}$$

- Iterate between E-step and M-step until convergence

- Let $\hat{\Theta}^{(n)}$ be the current estimate

- E-step: Compute $f(\Theta) = E_{(\mathbf{u}, \mathbf{v} | \mathbf{y}, \hat{\Theta}^{(n)})} [\log \Pr(\mathbf{y}, \mathbf{u}, \mathbf{v} | \Theta)]$

$$-\frac{1}{2\sigma^2} \sum_{(i,j)} E[(y_{ij} - u_i'v_j)^2] - \frac{1}{2\sigma_u^2} \sum_i E \|u_i\|^2 - \frac{1}{2\sigma_v^2} \sum_j E \|v_j\|^2$$

$$-R \log \sigma^2 - Mr \log \sigma_u^2 - Nr \log \sigma_v^2$$

- The expectation is not in closed form
- We draw Gibbs samples and compute the Monte Carlo mean

- M-step: Find $\hat{\Theta}^{(n+1)} = \arg \max_{\Theta} f(\Theta)$



Example: timeSVD++

- Example of matrix factorization in practice
- Part of the winning method of Netflix contest [Koren 2009]


$$y_{ij,t} \sim \mu + \underbrace{b_i(t)}_{\text{user bias}} + \overbrace{b_j(t)}^{\text{item popularity}} + \underbrace{u_i(t)}_{\text{user factors (preference)}}' v_j$$

$$b_i(t) = b_i + \alpha_i \underbrace{\text{dev}_i(t)}_{\text{distance to the middle rating time of } i} + b_{it}$$

$$b_j(t) = b_j + \underbrace{b_{j,\text{bin}(t)}}_{\text{time bin}}$$

$$u_i(t)_k = u_{ik} + \alpha_{ik} \text{dev}_u(t) + u_{ikt}$$

Model parameters: $\mu, b_i, \alpha_i, b_{it}, b_j, b_{jd}, u_{ik}, \alpha_{ik}, u_{ikt}$
for all user i , item j , factor k , time t , time bin d



Subscript:
user i ,
item j
time t



How to Handle Cold Start?

- For new items and new users, their factor values are all 0
- Simple idea
 - Predict their factor values based on features
 - For new user i , predict u_i based on x_i (user feature vector)

$$u_i \sim G x_i$$

x_i : feature vector of user i

- An item may be represented by a bag of words (later)

RLFM: Regression-based Latent Factor Model

- Incorporate features into matrix factorization

- x_i : feature vector of user i
- x_j : feature vector of item j

- Probabilistic model

$$y_{ij} \sim N(u'_i v_j, \sigma^2)$$

$$u_i \sim N(Gx_i, \sigma_u^2 I)$$

$$v_j \sim N(Dx_j, \sigma_v^2 I)$$

$$\text{Let } \Theta = (G, D, \sigma^2, \sigma_u^2, \sigma_v^2)$$

$$\log \Pr(\mathbf{y}, \mathbf{u}, \mathbf{v} | \Theta) = \text{constant}$$

$$- \frac{1}{2\sigma^2} \sum_{(i,j)} (y_{ij} - u'_i v_j)^2 - R \log \sigma^2$$

$$- \frac{1}{2\sigma_u^2} \sum_i \|u_i - Gx_i\|^2 - Mr \log \sigma_u^2$$

$$- \frac{1}{2\sigma_v^2} \sum_j \|v_j - Dx_j\|^2 - Nr \log \sigma_v^2$$

Find

$$\hat{\Theta} = \arg \max_{\Theta} \Pr(\mathbf{y} | \Theta) = \arg \max_{\Theta} \int \Pr(\mathbf{y}, \mathbf{u}, \mathbf{v} | \Theta) d\mathbf{u} d\mathbf{v}$$

Comparison

- Zero-mean factorization

$$y_{ij} \sim N(u_i'v_j, \sigma^2)$$

$$u_i \sim N(0, \sigma_u^2 I)$$

$$v_j \sim N(0, \sigma_v^2 I)$$

- Factorization with features (RLFM)

$$y_{ij} \sim N(u_i'v_j, \sigma^2)$$

$$y_{ij} \sim N(x_i'G'Dx_j + \delta_i'Dx_j + x_i'G'\eta_j + \delta_i'\eta_j, \sigma^2)$$

$$u_i \sim N(Gx_i, \sigma_u^2 I)$$

$$u_i = Gx_i + \delta_i, \quad \delta_i \sim N(0, \sigma_u^2 I)$$



$$v_j \sim N(Dx_j, \sigma_v^2 I)$$

$$v_j = Dx_j + \eta_j, \quad \eta_j \sim N(0, \sigma_v^2 I)$$



- Feature-only model

$$y_{ij} \sim N(x_i'G'Dx_j, \sigma^2)$$

Illustration



**Factorization
with features**



**Factorization
without feature**



Non-linear RLFM

rating that user i gives item j $y_{ij} \sim b(x'_{ij}) + \alpha_i + \beta_j + u'_i v_j$ x_i = feature vector of user i
 x_j = feature vector of item j
 x_{ij} = feature vector of (i, j)

- **Bias of user i :** $\alpha_i = g(x_i) + \varepsilon_i^\alpha$, $\varepsilon_i^\alpha \sim N(0, \sigma_\alpha^2)$
- **Popularity of item j :** $\beta_j = d(x_j) + \varepsilon_j^\beta$, $\varepsilon_j^\beta \sim N(0, \sigma_\beta^2)$
- **Factors of user i :** $u_i = G(x_i) + \varepsilon_i^u$, $\varepsilon_i^u \sim N(0, \sigma_u^2 I)$
- **Factors of item j :** $v_j = D(x_j) + \varepsilon_j^v$, $\varepsilon_j^v \sim N(0, \sigma_v^2 I)$

b, g, d, G, D are regression functions

Any regression model can be used here!!



fLDA: Factorization through LDA Topic Model

- An item is represented by a bag of word
- Model the rating y_{ij} that user i gives to item j as the user's affinity to the topics that the item has

$$y_{ij} = \dots + \sum_k \overset{\text{User } i \text{'s affinity to topic } k}{s_{ik}} \underset{\substack{\text{Pr(item } j \text{ has topic } k) \text{ estimated by averaging} \\ \text{the LDA topic of each word in item } j}}{\bar{z}_{jk}}$$

The topic distribution z_{jk} of a new item i is predicted based on the bag of words in the item

- Unlike regular unsupervised LDA topic modeling, here the LDA topics are learnt in a supervised manner based on past rating data
- These supervised topics are likely to be more useful for the prediction purpose



Supervised Topic Assignment

The topic of the n th word in item j

↓
 $\Pr(z_{jn} = k \mid \text{Rest})$

$$\propto \frac{Z_{kl}^{-jn} + \eta}{Z_k^{-jn} + W\eta} (Z_{jk}^{-jn} + \lambda)$$

Same as unsupervised LDA

Probability of observing y_{ij}
given the model

$$\prod_{i \text{ rated } j} f(y_{ij} \mid z_{jn} = k)$$

Likelihood of observed ratings
by users who rated item j when
 z_{jn} is set to topic k



Experimental Results (MovieLens)

- Task: Predict the rating that a user would give a movie
- Training/test split:
 - Sort observations by time
 - First 75% → Training data
 - Last 25% → Test data
- User cold-start scenario
 - 56% test data with new users
 - 2% new items in test data

Model	Test RMSE
RLFM	0.9363
fLDA	0.9381
Factor-Only	0.9422
FilterBot	0.9517
unsup-LDA	0.9520
MostPopular	0.9726
Feature-Only	1.0906
Constant	1.1190



Summary

- Factorization methods usually have better performance than pure feature-based methods
 - Netflix
 - Our experience
- Metadata (feature vector or bag of words) can be easily incorporated into matrix factorization
- Next step
 - Matrix factorization with social networks
 - Friendship: Address book
 - Communication: Instant messages, emails
 - Multi-application factorization
 - E.g., joint factorization of the (user, news article) matrix and the (user, query) matrix



Fast Online Learning for Time-sensitive Recommendation

- Examples of time-sensitive items
 - News stories, trending queries, tweets, updates, events ...
- Real-time data pipeline that continuously collects new ratings (clicks) on new items
- Modeling requirements:
 - **Fast learning**: Learn good models for new items using **little data**
 - Good initial guess (without ratings on new items)
 - Fast convergence
 - **Fast computation**: Build good models using **little time**
 - Efficient
 - Scalable
 - Parallelizable



FOBFM: Fast Online Bilinear Factor Model

Per-item online model $y_{ij} \sim u_i' \beta_j, \quad \beta_j \sim N(\mu_j, \Sigma)$

Subscript:

user i
item j

Data:

y_{ij} = rating that user i gives item j
 u_i = offline factor vector of user i
 x_j = feature vector of item j

- Feature-based model initialization

$$\beta_j \sim N(\underbrace{Ax_j}_{\text{predicted by features}}, \Sigma) \iff y_{ij} \sim u_i' Ax_j + u_i' v_j$$

$$v_j \sim N(0, \Sigma)$$

- Dimensionality reduction for fast model convergence

$$v_j = B \theta_j$$

$$\theta_j \sim N(0, \sigma_\theta^2 I)$$

B is a $n \times k$ linear projection matrix ($k \ll n$)

project: high $\dim(v_j) \rightarrow$ low $\dim(\theta_j)$

low-rank approx of $\text{Var}[\beta_j]$: $\beta_j \sim N(Ax_j, \sigma_\theta^2 BB')$

$$v_j = B \theta_j$$

Offline training: Determine A, B, σ_θ^2
(once per day)



FOBFM: Fast Online Bilinear Factor Model

Per-item online model $y_{ij} \sim u_i' \beta_j, \quad \beta_j \sim N(\mu_j, \Sigma)$

Subscript:

user i
item j

Data:

y_{ij} = rating that user i gives item j
 u_i = offline factor vector of user i
 x_j = feature vector of item j

- Feature-based model initialization

$$\beta_j \sim N(\underbrace{Ax_j}_{\text{predicted by features}}, \Sigma) \iff y_{ij} \sim u_i' Ax_j + u_i' v_j$$
$$v_j \sim N(0, \Sigma)$$

- Dimensionality reduction for fast model convergence

$$v_j = B \theta_j \quad B \text{ is a } n \times k \text{ linear projection matrix } (k \ll n)$$
$$\theta_j \sim N(0, \sigma_\theta^2 I) \quad \text{project: high dim}(v_j) \rightarrow \text{low dim}(\theta_j)$$

low-rank approx of $\text{Var}[\beta_j]$: $\beta_j \sim N(Ax_j, \sigma_\theta^2 BB')$

- Fast, parallel online learning

$$y_{ij} \sim \underbrace{u_i' Ax_j}_{\text{offset}} + \underbrace{(u_i' B) \theta_j}_{\text{new feature vector (low dimensional)}}, \quad \text{where } \theta_j \text{ is updated in an online manner}$$

- Online selection of dimensionality ($k = \text{dim}(\theta_j)$)
 - Maintain an ensemble of models, one for each candidate dimensionality




Experimental Results: My Yahoo! Dataset (1)

- My Yahoo! is a personalized news reading site
 - Users manually select news/RSS feeds
- ~12M “ratings” from ~3M users to ~13K articles
 - Click = positive
 - View without click = negative



Experimental Results: My Yahoo! Dataset (2)




Methods:

- **No-init:** Standard online regression with ~ 1000 parameters for each item
 - **Offline:** Feature-based model without online update
 - **PCR, PCR+:** Two principal component methods to estimate B
 - **FOBFM:** Our fast online method
- Item-based data split: Every item is new in the test data
 - First 8K articles are in the training data (offline training)
 - Remaining articles are in the test data (online prediction & learning)
 - Our supervised dimensionality reduction (reduced rank regression) significantly outperforms other methods



Experimental Results: My Yahoo! Dataset (3)



- Small number of factors (low dimensionality) is better when the amount of data for online learning is small
- Large number of factors is better when the data for learning becomes large
- The online selection method usually selects the best dimensionality

Experimental Results: MovieLens Dataset

- Training-test data split
 - Time-split: First 75% ratings in training; rest in test
 - Movie-split: 75% randomly selected movies in training; rest in test


Model	RMSE Time-split	RMSE Movie-split
FOBFM	0.8429	0.8549
RLFM	0.9363	1.0858
Online-UU	1.0806	0.9453
Constant	1.1190	1.1162

FOBFM: Our fast online method

RLFM: [Agarwal 2009]


Online-UU: Online version of user-user collaborative filtering

Online-PLSI: [Das 2007]



Experimental Results: Yahoo! Front Page Dataset

- Training-test data split
 - Time-split: First 75% ratings in training; rest in test



- ~2M “ratings” from ~30K frequent users to ~4K articles
 - Click = positive
 - View without click = negative
- Our fast learning method outperforms others

Summary

- Recommending time-sensitive items is challenging
 - Most collaborative filtering methods do not work well in cold start
 - Rebuilding models can incur too much latency when the numbers of items and users are large
- Our approach:
 - Periodically rebuild the offline model that
 - uses feature-based regression to **predict the initial point** for online learning, and
 - **reduces the dimensionality** of online learning
 - Rapidly update online models once new data is received
 - Fast learning: Low dimensional and easily parallelizable
 - Online selection for the best dimensionality

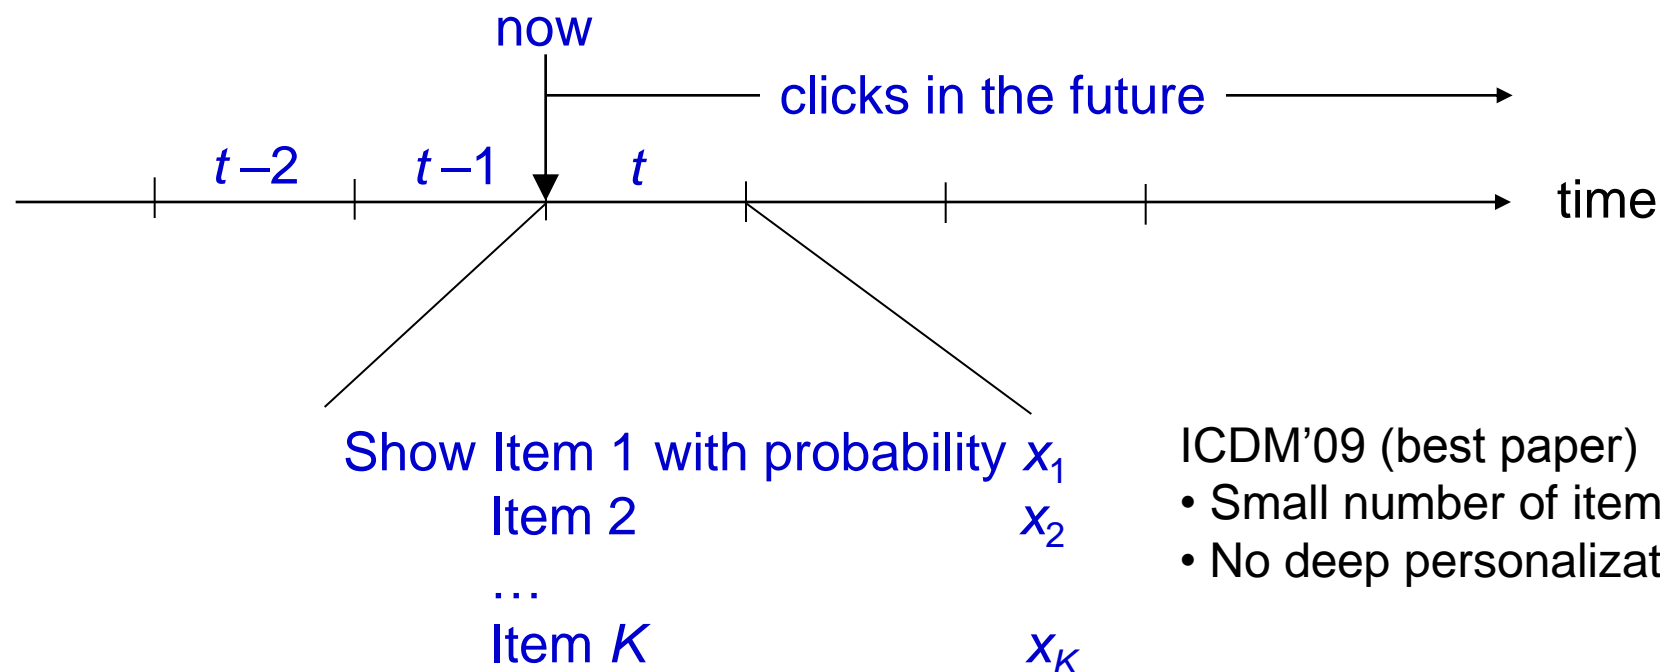


Important Problems Beyond Factor Models

- How to explore/exploit with small traffic, a large item pool, at a fine granularity
- Offline evaluation
- Multi-objective optimization under uncertainty
- Whole page optimization



Explore/Exploit



Determine (x_1, x_2, \dots, x_K) based on clicks and views observed before t in order to maximize the expected total number of clicks in the future

- Challenges**
- Large number of items
 - Small traffic
 - Deep personalization




Offline Evaluation

- Ultimate evaluation: Online bucket test
- Unbiased offline evaluation based on random-bucket data
 - [Lihong Li, WWW'10, WSDM'11]
 - Random bucket: A small user population to which we show each item with equal probability
 - Assumptions:
 - Single recommendation per visit (instead of top- K)
 - All the users respond to the recommended item in an *iid* manner
 - Replay-match methodology
- Challenges
 - How to handle non-random data
 - How to extend to top- K recommendation
 - How to capture users' "non-*iid*" behavior in a session



Multi-Objective Optimization

- Maximize time-spent (or revenue) s.t. click drop < 5%



Challenges:

- Deep personalization
- Optimization in the presence of uncertainty



Whole Page Optimization

The screenshot shows the Yahoo! homepage as of July 14, 2010. The page is divided into several sections:

- Navigation:** Top navigation links for Web, Images, Video, Local, Shopping, and More. A search bar with a "Web Search" button is located below the navigation.
- Left Sidebar (YAHOO! SITES):** A vertical list of site categories including Mail, Autos, Chat, Fantasy Sports, Finance, Games, Horoscopes, HotJobs, Maps, Messenger, Movies, omg!, Personals, Shopping, Sports, Travel, Updates, and Weather. Below this is a "MY FAVORITES" section with links to eBay, Facebook, and Twitter.
- Main Content Area (TODAY - July 14, 2010):** Features a large image of Paul the octopus with the headline "World Cup octopus could make millions". Below the headline is a sub-headline and a short paragraph. To the right of the main text are links for "More on the octopus" and "U.S.'s top moments". Below the main text are four smaller article thumbnails: "Salsa tied to food illness", "Octopus could be worth millions", "Lottery winner rich in mystery", and "High schooler's impressive dunk".
- Right Sidebar (TRENDING NOW):** A list of trending topics: 1. Kourtney Kardash..., 2. Anna Chapman, 3. Al Pacino, 4. French Toast Rec..., 5. Nina Garcia, 6. Susan Boyle, 7. Job Search, 8. Yogi Berra, 9. Philippines Typh..., 10. Sunscreen.
- Below Trending Now:** An advertisement for Chase Ultimate Rewards with the text "Anything you want, you got it with Ultimate Rewards." and a "Click to see your reward" button.
- Below Ad:** A "Must-see music news & features" section with a "Music" icon and four article thumbnails: "Britney lays down the law with kids", "Lady Gaga photo irks Beatles fans", "'Idol' runner-up gets cosmetic work", and "Kellie Pickler's new retro '40s video".
- Bottom Section:** A "DAILY OFFERS" section with a thumbnail and the text "Mortgage rates low as 3.32% APR".

Challenge:
How to jointly optimize all these modules

- Diversity
- Consistency
- Relatedness





Thank You!

Contact me for job/internship opportunities in Yahoo! Labs
beechun@yahoo-inc.com