VBPR: Visual Bayesian Personalized Ranking from Implicit Feedback.

R. He & J. McAuley. AAAI 2016.

발표자 김형준

Introduction

- **Recommender Systems(RSs) provide personalized suggestions by learning from historical feedback.**
  - Historical feedback : explicit (star ratings, etc.), implicit (purchase history, book marks, etc.)

- **Matrix Factorization(MF) approaches have been proposed to uncover the most relevant latent dimensions (both in explicit / implicit feedback settings).**
  - MF approaches suffer from **cold-start issues** dues to the sparsity of real-world dataset.

$\Longrightarrow$ incorporate visual appearance of the items into the preference predictor.

VBPR: Visual Bayesian Personalized Ranking

- **Objective**
  - For each user u, we want to generate a **personalized ranking** of the items which the user **haven't provided feedback yet**. (i.e., $\boldsymbol{I} \backslash \boldsymbol{I}_{\boldsymbol{u}}^{+}$)

Notation	Explanation
$\mathcal{U}, \mathcal{I}$	user set, item set
$\mathcal{I}_u^+$	positive item set of user u
$\widehat{x}_{u,i}$	predicted 'score' user u gives to item i
K	dimension of latent factors
D	dimension of visual factors
F	dimension of Deep CNN features
α	global offset (scalar)
β_u, β_i	user u's bias, item i's bias (scalar)
γ_u, γ_i	latent factors of user u, item i ($K \times 1$)
θ_u, θ_i	visual factors of user u, item i ($D \times 1$)
f_i	Deep CNN visual features of item i ($F \times 1$)
$\mathbf{E}$	$D \times F$ embedding matrix
β'	visual bias vector (visual bias = $\beta'^T f_i$)

Notations

VBPR: Visual Bayesian Personalized Ranking

- **Preference predictor**
  - Built on top of Matrix Factorization (MF)

$$\widehat{x}_{u,i} = \alpha + \beta_u + \beta_i + \gamma_u^T \gamma_i,$$

VBPR: Visual Bayesian Personalized Ranking

- **Preference predictor**
  - Built on top of Matrix Factorization (MF)

$$\widehat{x}_{u,i} = \underbrace{\boxed{\alpha}}_{\text{global offset}} + \underbrace{\boxed{\beta_u}}_{\text{user bias}} + \underbrace{\boxed{\beta_i}}_{\text{item bias}} + \underbrace{\boxed{\gamma_u^T}}_{\substack{\text{latent factor of}\\ \text{user } u\\ \text{(K-dim)}}} \underbrace{\boxed{\gamma_i}}_{\substack{\text{latent factor of}\\ \text{item } i\\ \text{(K-dim)}}},$$

  - Problem : existence of **'cold'** or **'cool'** items.
    - Too few associated observations to estimate the latent dimensions.

$\Longrightarrow$ Solve by using visual factors.

VBPR: Visual Bayesian Personalized Ranking

- **Preference predictor**
  - Use the inner product of the visual factors → **visual interaction** between user u and item i .

$$\widehat{x}_{u,i} = \alpha + \beta_u + \beta_i + \gamma_u^T \gamma_i + \theta_u^T \theta_i,$$

VBPR: Visual Bayesian Personalized Ranking

- **Preference predictor**
  - Use the inner product of the visual factors → **visual interaction** between user u and item i .

$$\widehat{x}_{u,i} = \alpha + \beta_u + \beta_i + \gamma_u^T \gamma_i + \boxed{\theta_u^T}\boxed{\theta_i},$$

visual factor of user u
(D-dim)

visual factor of item i
(D-dim)

$$\theta_i = \boxed{\mathbf{E}}\boxed{f_i}$$

Embedding matrix
(D×F-dim)

Deep CNN visual features
(F×1-dim)

Final prediction model :

$$\widehat{x}_{u,i} = \alpha + \beta_u + \beta_i + \gamma_u^T \gamma_i + \boxed{\theta_u^T}\boxed{(\mathbf{E} f_i)} + \boxed{\beta'^T}\boxed{f_i}.$$

Visual bias vector

VBPR: Visual Bayesian Personalized Ranking

- **Preference predictor**
  - Use the inner product of the visual factors → **visual interaction** between user u and item i .

$$\widehat{x}_{u,i} = \alpha + \beta_u + \beta_i + \gamma_u^T \gamma_i + \theta_u^T (\boxed{\mathbf{E}} \boxed{f_i}) + \beta'^T f_i.$$

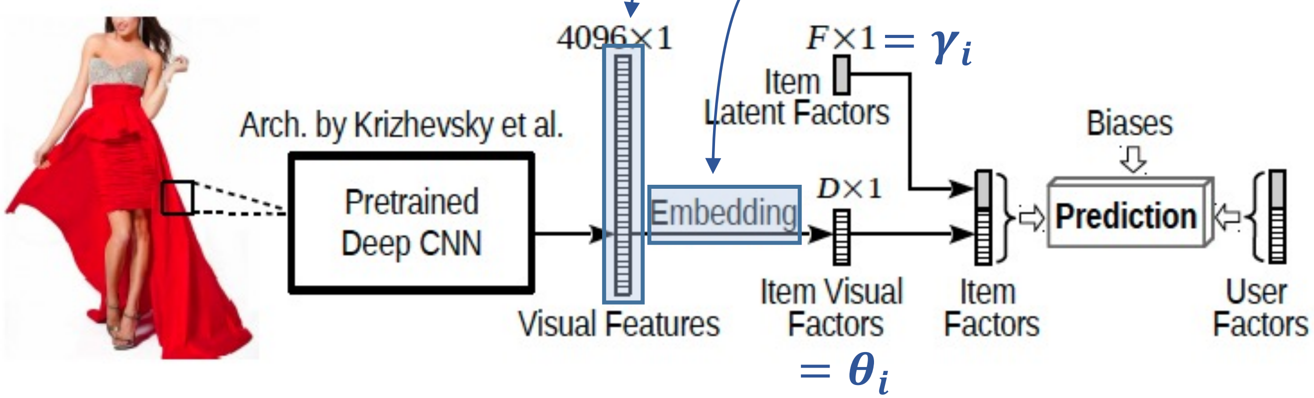

Visual Features

- **AlexNet**
  - 5 Convolution layers followed by 3 fully-connected layers.
  - Pre-trained on 1.2million ImageNet images
  - Take the output of the second fully-connected layer
    - F = 4096 dimensional visual feature vector f_i

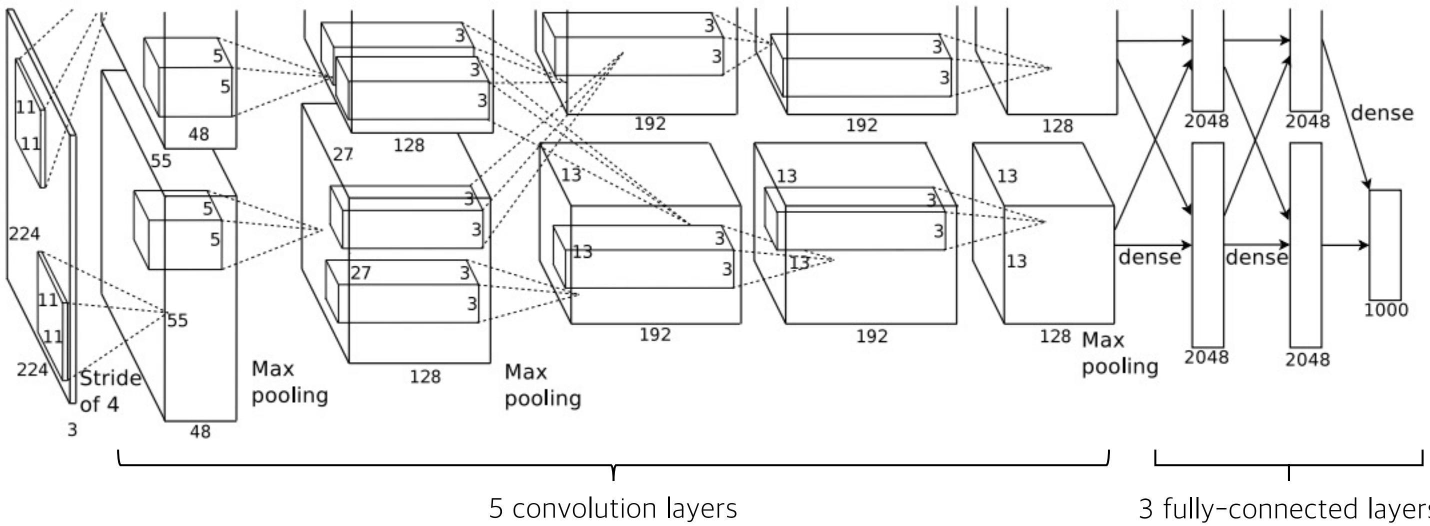

Reference :
Imagenet classification with deep convolutional neural networks. Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. NIPS 2012.

Visual Features

- **AlexNet**

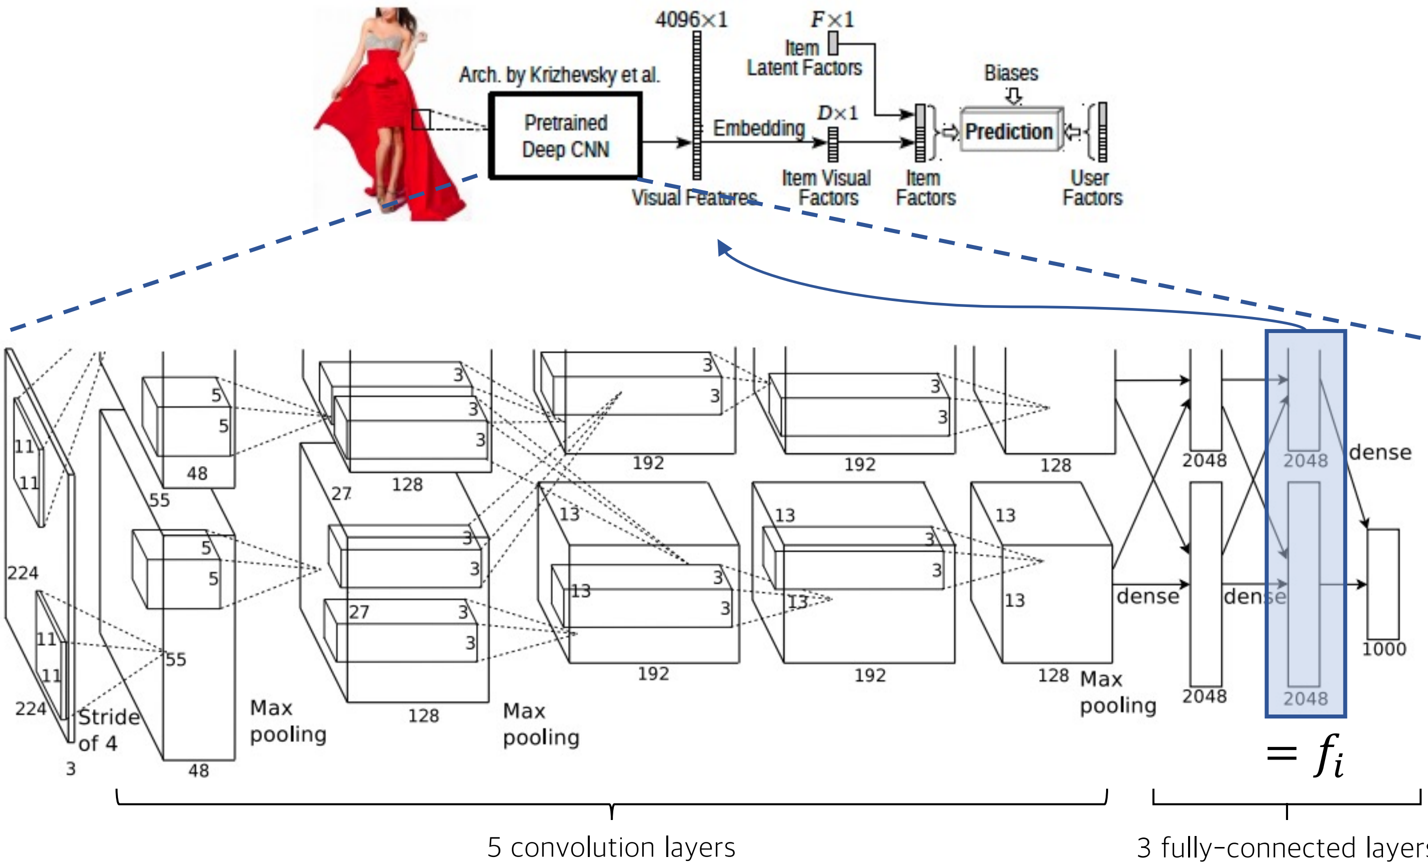

Reference :
Imagenet classification with deep convolutional neural networks. Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. NIPS 2012.

Model Learning Using BPR

- **Model learning using Bayesian Personalized Ranking (BPR)**
  - Training set D_s consist of triples (u, i, j)

$$D_S = \{(u, i, j) | u \in \mathcal{U} \wedge i \in \mathcal{I}_u^+ \wedge j \in \mathcal{I} \setminus \mathcal{I}_u^+\}.$$

    - u : user
    - i : item that user u expressed positive feedback
    - j : non-observed item

Model Learning Using BPR

- **Model learning using Bayesian Personalized Ranking (BPR)**
  - Following optimization criterion is used for **personalized ranking (BPR-OPT)**

$$BPR - OPT = \sum_{(u,i,j) \in D_S} ln\sigma(\widehat{x}_{uij}) - \lambda_\Theta ||\Theta||^2$$

  - Where $\hat{x}_{uij}$ is defined as : $\widehat{x}_{uij} = \widehat{x}_{u,i} - \widehat{x}_{u,j}$, [1]

For more details :
BPR Paper : BPR: Bayesian Personalized Ranking from Implicit Feedback. Rendle et al. (2009) (https://arxiv.org/pdf/1205.2618.pdf)
Blog post : https://leehyejin91.github.io/post-bpr/

Model Learning Using BPR

- **Model learning using Bayesian Personalized Ranking (BPR)**
  - Following optimization criterion is used for **personalized ranking (BPR-OPT)**

$$BPR - OPT = \sum_{(u,i,j)\in D_S} \underbrace{\boxed{ln\sigma}}_{\text{Logistic sigmoid function}}(\hat{x}_{uij}) - \underbrace{\boxed{\lambda_\Theta}}_{\text{Regularization term}}||\underbrace{\boxed{\Theta}}_{\text{Parameterize relationship between } (u,i,j)}||^2$$

  - Where $\hat{x}_{uij}$ is defined as : $\hat{x}_{uij} = \hat{x}_{u,i} - \hat{x}_{u,j}$, [1]

For more details :
BPR Paper : BPR: Bayesian Personalized Ranking from Implicit Feedback. Rendle et al. (2009) (https://arxiv.org/pdf/1205.2618.pdf)
Blog post : https://leehyejin91.github.io/post-bpr/

Model Learning Using BPR

- **Model learning using Bayesian Personalized Ranking (BPR)**
  - Following optimization criterion is used for **personalized ranking (BPR-OPT)**

$$BPR - OPT = \sum_{(u,i,j)\in D_S} \boxed{ln\sigma(\widehat{x}_{uij})} - \lambda_\Theta ||\Theta||^2$$

Derivative of $\partial ln\sigma(\hat{x}_{uij})$

$$\frac{\partial ln\sigma(\hat{x}_{uij})}{\partial \ominus} = \frac{1}{\sigma(\hat{x}_{uij})}\frac{\partial\sigma(\hat{x}_{uij})}{\partial\hat{x}_{uij}}\frac{\partial\hat{x}_{uij}}{\partial \ominus}$$

$$= \frac{1}{\sigma(\hat{x}_{uij})}\sigma(\hat{x}_{uij})\left(1-\sigma(\hat{x}_{uij})\right)\frac{\partial\hat{x}_{uij}}{\partial \ominus}$$

$$= \left(1-\sigma(\hat{x}_{uij})\right)\frac{\partial\hat{x}_{uij}}{\partial \ominus} = \boldsymbol{\sigma(-\hat{x}_{uij})\frac{\partial\hat{x}_{uij}}{\partial \ominus}}$$

  - Where $\hat{x}_{uij}$ is defined as : $\widehat{x}_{uij} = \widehat{x}_{u,i} - \widehat{x}_{u,j}$, [1]

  - Updates the parameters in the following way:

$$\Theta \leftarrow \Theta + \eta \cdot \left(\sigma(-\widehat{x}_{uij})\frac{\partial\widehat{x}_{uij}}{\partial\Theta} - \lambda_\Theta\Theta\right),$$

For more details :
BPR Paper : BPR: Bayesian Personalized Ranking from Implicit Feedback. Rendle et al. (2009) (https://arxiv.org/pdf/1205.2618.pdf)
Blog post : https://leehyejin91.github.io/post-bpr/

Model Learning Using BPR

- **Model learning using Visual Bayesian Personalized Ranking (VBPR)**
  - For VBPR, $\hat{x}_{u,i}$ and $\hat{x}_{u,i}$ considers newly-introduced visual parameters as well.

$$\widehat{x}_{u,i} = \alpha + \beta_u + \beta_i + \gamma_u^T \gamma_i + \boxed{\theta_u^T}(\boxed{\mathbf{E}} f_i) + \boxed{\beta'}^T f_i.$$

Instead of $\widehat{x}_{u,i} = \alpha + \beta_u + \beta_i + \gamma_u^T \gamma_i,$

  - Use the same optimization criterion :

$$BPR - OPT = \sum_{(u,i,j) \in D_S} ln\sigma(\widehat{x}_{uij}) - \lambda_\Theta ||\Theta||^2$$

  - Updates the parameters in the following way:

$$\begin{aligned}
\boxed{\theta_u} &\leftarrow \theta_u + \eta \cdot (\sigma(-\widehat{x}_{uij})\mathbf{E}(f_i - f_j) - \lambda_\Theta \theta_u), \\
\boxed{\beta'} &\leftarrow \beta' + \eta \cdot (\sigma(-\widehat{x}_{uij})(f_i - f_j) - \lambda_\beta \beta'), \\
\boxed{\mathbf{E}} &\leftarrow \mathbf{E} + \eta \cdot (\sigma(-\widehat{x}_{uij})\theta_u (f_i - f_j)^T - \lambda_{\mathbf{E}} \mathbf{E}).
\end{aligned}$$

Gradient for each parameters θ_u, β', E

$$\frac{\partial ln\sigma(\hat{x}_{uij})}{\partial \theta_u} = \frac{1}{\sigma(\hat{x}_{uij})} \frac{\partial \sigma(\hat{x}_{uij})}{\partial \hat{x}_{uij}} \frac{\partial \hat{x}_{uij}}{\partial \theta_u} = \sigma(-\hat{x}_{uij}) \frac{\partial \hat{x}_{uij}}{\partial \theta_u} = \sigma(-\hat{x}_{uij}) \left(\frac{\partial \hat{x}_{u,i}}{\partial \theta_u} - \frac{\partial \hat{x}_{u,j}}{\partial \theta_u} \right) = \sigma(-\hat{x}_{uij})(Ef_i - Ef_j) = \boldsymbol{\sigma}(-\boldsymbol{\hat{x}_{uij}})\boldsymbol{E}(\boldsymbol{f_i} - \boldsymbol{f_i})$$

$$\frac{\partial ln\sigma(\hat{x}_{uij})}{\partial \beta'} = \frac{1}{\sigma(\hat{x}_{uij})} \frac{\partial \sigma(\hat{x}_{uij})}{\partial \hat{x}_{uij}} \frac{\partial \hat{x}_{uij}}{\partial \beta'} = \sigma(-\hat{x}_{uij}) \frac{\partial \hat{x}_{uij}}{\partial \beta'} = \sigma(-\hat{x}_{uij}) \left(\frac{\partial \hat{x}_{u,i}}{\partial \beta'} - \frac{\partial \hat{x}_{u,j}}{\partial \beta'} \right) = \sigma(-\hat{x}_{uij})(f_i - f_j) = \boldsymbol{\sigma}(-\boldsymbol{\hat{x}_{uij}})(\boldsymbol{f_i} - \boldsymbol{f_i})$$

$$\frac{\partial ln\sigma(\hat{x}_{uij})}{\partial E} = \frac{1}{\sigma(\hat{x}_{uij})} \frac{\partial \sigma(\hat{x}_{uij})}{\partial \hat{x}_{uij}} \frac{\partial \hat{x}_{uij}}{\partial E} = \sigma(-\hat{x}_{uij}) \frac{\partial \hat{x}_{uij}}{\partial E} = \sigma(-\hat{x}_{uij}) \left(\frac{\partial \hat{x}_{u,i}}{\partial E} - \frac{\partial \hat{x}_{u,j}}{\partial E} \right) = \sigma(-\hat{x}_{uij})(\theta_u f_i^T - \theta_u f_j^T) = \boldsymbol{\sigma}(-\boldsymbol{\hat{x}_{uij}})\theta_u(\boldsymbol{f_i} - \boldsymbol{f_i})^T$$

Scalability

- **Efficiency of VBPR**

BPR → $\Theta \leftarrow \Theta + \eta \cdot (\sigma(-\widehat{x}_{uij})\frac{\partial \widehat{x}_{uij}}{\partial \Theta} - \lambda_\Theta \Theta),$ → $O(K)$

VBPR →

$\theta_u \leftarrow \theta_u + \eta \cdot (\sigma(-\widehat{x}_{uij})\mathbf{E}(f_i - f_j) - \lambda_\Theta \theta_u),$ → $O(D \times F) = O(D)$

$\beta' \leftarrow \beta' + \eta \cdot (\sigma(-\widehat{x}_{uij})(f_i - f_j) - \lambda_\beta \beta'),$ → $O(F)$

$\mathbf{E} \leftarrow \mathbf{E} + \eta \cdot (\sigma(-\widehat{x}_{uij})\theta_u(f_i - f_j)^T - \lambda_\mathbf{E} \mathbf{E}).$ → $O(D \times F) = O(D)$

- $\boldsymbol{F}$ is a constant (fixed dimension of Deep CNN features)
- $\boldsymbol{K}$ is the dimension of latent factors.
- $\boldsymbol{D}$ is the dimension of visual factors.

- Overall time complexity of updating each triple (u, i, j) :

$$O(K) + O(D \times F) = O(K) + O(D) = \boldsymbol{O(K + D)}$$

Linear in the number of dimensions!

Notation	Explanation
$\mathcal{U}, \mathcal{I}$	user set, item set
$\mathcal{I}_u^+$	positive item set of user u
$\widehat{x}_{u,i}$	predicted 'score' user u gives to item i
K	dimension of latent factors
D	dimension of visual factors
F	dimension of Deep CNN features
α	global offset (scalar)
β_u, β_i	user u's bias, item i's bias (scalar)
γ_u, γ_i	latent factors of user u, item i ($K \times 1$)
θ_u, θ_i	visual factors of user u, item i ($D \times 1$)
f_i	Deep CNN visual features of item i ($F \times 1$)
$\mathbf{E}$	$D \times F$ embedding matrix
β'	visual bias vector (visual bias = $\beta'^T f_i$)

notations

Datasets

- **Amazon.com**
  - Women's Clothing
  - Men's Clothing
  - Cell Phones (expect visual characteristics to play a smaller role than clothing)
  - Use user's review histories as user feedback.

- **Tradesy.com**
  - Use user's purchase history and "thumbs-up" as user feedback.
  - Marketplace for woman to buy and sell pre-owned clothing items.
    - Requires **cold-start** predictions : **"one-off"** trading characteristics.
      - Probably sell one item to one user, resulting in cold start environment.

  - "cold start" training set : items that had fewer than 5 positive feedbacks.

Dataset	#users	#items	#feedback
Amazon Women	99,748	331,173	854,211
Amazon Men	34,212	100,654	260,352
Amazon Phones	113,900	192,085	964,477
Tradsy.com	19,823	166,526	410,186
Total	267,683	790,438	2,489,226

Experiments

- **Evaluation Methods : AUC**
  - For each user u, random item is used for validation V_u, another item is used to testing T_u , remaining data is used for training.
  - Predicted **ranking** is evaluated on T_u with the metric **AUC** (area under the ROC curve).

$$AUC = \frac{1}{|\mathcal{U}|}\sum_{u}\frac{1}{|E(u)|}\sum_{(i,j)\in E(u)}\delta(\widehat{x}_{u,i} > \widehat{x}_{u,j})$$

  - $\delta(b)$: indicator function that returns 1 if b is true.
  - Report the performance on the test set T for the hyperparameters with the best performance on the validation set V.

The AUC per user is defined as : $\mathrm{AUC}(u) := \frac{1}{|I_u^+|\,|I \setminus I_u^+|}\sum_{i\in I_u^+}\sum_{j\in|I\setminus I_u^+|}\delta(\hat{x}_{uij} > 0)$

The average AUC for all users : $\mathrm{AUC} := \frac{1}{|U|}\sum_{u\in U} AUC(u) = \frac{1}{|\mathcal{U}|}\sum_{u}\frac{1}{|E(u)|}\sum_{(i,j)\in E(u)}\delta(\widehat{x}_{u,i} > \widehat{x}_{u,j})$

Reference :
BPR: Bayesian Personalized Ranking from Implicit Feedback. Rendle et al. UAI 2009.

Experiments

- **Baselines**
  - Random (**RAND**) : rank items randomly
  - Most Popular (**MP**) : rank items according to their popularity
  - Matric Factorization methods : **MM-MF, BPR-MF**
  - Image-based Recommendation (**IBR**) : learns a visual space and retrieves similar items to a query image using nearest-neighbor search.
  - **WRMF** : point-wise method

  - Used the same total number of dimensions for all MF-based methods.
    - For VBPR, visual and non-visual dimensions are fixed to a 50-50 spilt.

Experiments

Table 3: AUC on the test set $\mathcal{T}$ (#factors = 20). The best performing method on each dataset is boldfaced.

Dataset	Setting	(a) RAND	(b) MP	(c) IBR	(d) MM-MF	(e) BPR-MF	(f) VBPR	improvement f vs. best	improvement f vs. e
Amazon Women	All Items	0.4997	0.5772	0.7163	0.7127	0.7020	**0.7834**	9.4%	11.6%
	Cold Start	0.5031	0.3159	0.6673	0.5489	0.5281	**0.6813**	2.1%	29.0%
Amazon Men	All Items	0.4992	0.5726	0.7185	0.7179	0.7100	**0.7841**	9.1%	10.4%
	Cold Start	0.4986	0.3214	0.6787	0.5666	0.5512	**0.6898**	1.6%	25.1%
Amazon Phones	All Items	0.5063	0.7163	0.7397	0.7956	0.7918	**0.8052**	1.2%	1.7%
	Cold Start	0.5014	0.3393	**0.6319**	0.5570	0.5346	0.6056	-4.2%	13.3%
Tradesy.com	All Items	0.5003	0.5085	N/A	0.6097	0.6198	**0.7829**	26.3%	26.3%
	Cold Start	0.4972	0.3721	N/A	0.5172	0.5241	**0.7594**	44.9%	44.9%

- **VBPR (built on top of BPR-MF) on average improves on BPR-MF by over 12% + over 28% on cold start setting.**
  - Shows the importance of CNN features (visual features).

- **IBR (Image-based recommendation) outperforms MF-based methods in cold start settings ↔ IBR loses to MF-based methods in warm start settings.**
  - Not trained on historical user feedback.

Experiments

Table 3: AUC on the test set $\mathcal{T}$ (#factors = 20). The best performing method on each dataset is boldfaced.

Dataset	Setting	(a) RAND	(b) MP	(c) IBR	(d) MM-MF	(e) BPR-MF	(f) VBPR	improvement f vs. best	improvement f vs. e
Amazon Women	All Items	0.4997	0.5772	0.7163	0.7127	0.7020	**0.7834**	9.4%	11.6%
	Cold Start	0.5031	0.3159	0.6673	0.5489	0.5281	**0.6813**	2.1%	29.0%
Amazon Men	All Items	0.4992	0.5726	0.7185	0.7179	0.7100	**0.7841**	9.1%	10.4%
	Cold Start	0.4986	0.3214	0.6787	0.5666	0.5512	**0.6898**	1.6%	25.1%
Amazon Phones	All Items	0.5063	0.7163	0.7397	0.7956	0.7918	**0.8052**	1.2%	1.7%
	Cold Start	0.5014	0.3393	**0.6319**	0.5570	0.5346	0.6056	-4.2%	13.3%
Tradesy.com	All Items	0.5003	0.5085	N/A	0.6097	0.6198	**0.7829**	26.3%	26.3%
	Cold Start	0.4972	0.3721	N/A	0.5172	0.5241	**0.7594**	44.9%	44.9%

- **VBPR (built on top of BPR-MF) on average improves on BPR-MF by over 12% + over 28% on cold start setting.**
  - Shows the importance of CNN features (visual features).

- **IBR (Image-based recommendation) outperforms MF-based methods in cold start settings ↔ IBR loses to MF-based methods in warm start settings.**
  - Not trained on historical user feedback.

Experiments

Table 3: AUC on the test set $\mathcal{T}$ (#factors = 20). The best performing method on each dataset is boldfaced.

Dataset	Setting	(a) RAND	(b) MP	(c) IBR	(d) MM-MF	(e) BPR-MF	(f) VBPR	improvement f vs. best	improvement f vs. e
Amazon Women	All Items	0.4997	0.5772	0.7163	0.7127	0.7020	**0.7834**	9.4%	11.6%
	Cold Start	0.5031	0.3159	0.6673	0.5489	0.5281	**0.6813**	2.1%	29.0%
Amazon Men	All Items	0.4992	0.5726	0.7185	0.7179	0.7100	**0.7841**	9.1%	10.4%
	Cold Start	0.4986	0.3214	0.6787	0.5666	0.5512	**0.6898**	1.6%	25.1%
Amazon Phones	All Items	0.5063	0.7163	0.7397	0.7956	0.7918	**0.8052**	1.2%	1.7%
	Cold Start	0.5014	0.3393	**0.6319**	0.5570	0.5346	0.6056	-4.2%	13.3%
Tradesy.com	All Items	0.5003	0.5085	N/A	0.6097	0.6198	**0.7829**	26.3%	26.3%
	Cold Start	0.4972	0.3721	N/A	0.5172	0.5241	**0.7594**	44.9%	44.9%

- **VBPR outperforms all the baselines in most cases.**
- **More improvements in Tradesy.com settings. (cold start dataset)**
- **Visual features show greater benefits on clothing than cellphone datasets.**
  - Visual factors play a smaller role when selecting cellphones.
- **Cold start items are ‘unpopular’, popularity-based methods were ineffective**

Experiments

Table 3: AUC on the test set $\mathcal{T}$ (#factors = 20). The best performing method on each dataset is boldfaced.

Dataset	Setting	(a) RAND	(b) MP	(c) IBR	(d) MM-MF	(e) BPR-MF	(f) VBPR	improvement f vs. best	improvement f vs. e
Amazon Women	All Items	0.4997	0.5772	0.7163	0.7127	0.7020	**0.7834**	9.4%	11.6%
	Cold Start	0.5031	0.3159	0.6673	0.5489	0.5281	**0.6813**	2.1%	29.0%
Amazon Men	All Items	0.4992	0.5726	0.7185	0.7179	0.7100	**0.7841**	9.1%	10.4%
	Cold Start	0.4986	0.3214	0.6787	0.5666	0.5512	**0.6898**	1.6%	25.1%
Amazon Phones	All Items	0.5063	0.7163	0.7397	0.7956	0.7918	**0.8052**	1.2%	1.7%
	Cold Start	0.5014	0.3393	**0.6319**	0.5570	0.5346	0.6056	-4.2%	13.3%
Tradesy.com	All Items	0.5003	0.5085	N/A	0.6097	0.6198	**0.7829**	26.3%	26.3%
	Cold Start	0.4972	0.3721	N/A	0.5172	0.5241	**0.7594**	44.9%	44.9%

- **VBPR outperforms all the baselines in most cases.**
- **More improvements in Tradesy.com settings. (cold start dataset)**
- **Visual features show greater benefits on clothing than cellphone datasets.**
  - Visual factors play a smaller role when selecting cellphones.
- **Cold start items are 'unpopular', popularity-based methods were ineffective**

Experiments

Table 3: AUC on the test set $\mathcal{T}$ (#factors = 20). The best performing method on each dataset is boldfaced.

Dataset	Setting	(a) RAND	(b) MP	(c) IBR	(d) MM-MF	(e) BPR-MF	(f) VBPR	improvement f vs. best	improvement f vs. e
Amazon Women	All Items	0.4997	0.5772	0.7163	0.7127	0.7020	**0.7834**	9.4%	11.6%
	Cold Start	0.5031	0.3159	0.6673	0.5489	0.5281	**0.6813**	2.1%	29.0%
Amazon Men	All Items	0.4992	0.5726	0.7185	0.7179	0.7100	**0.7841**	9.1%	10.4%
	Cold Start	0.4986	0.3214	0.6787	0.5666	0.5512	**0.6898**	1.6%	25.1%
Amazon Phones	All Items	0.5063	0.7163	0.7397	0.7956	0.7918	**0.8052**	1.2%	1.7%
	Cold Start	0.5014	0.3393	**0.6319**	0.5570	0.5346	0.6056	-4.2%	13.3%
Tradesy.com	All Items	0.5003	0.5085	N/A	0.6097	0.6198	**0.7829**	26.3%	26.3%
	Cold Start	0.4972	0.3721	N/A	0.5172	0.5241	**0.7594**	44.9%	44.9%

- **VBPR outperforms all the baselines in most cases.**
- **More improvements in Tradesy.com settings. (cold start dataset)**
- **Visual features show greater benefits on clothing than cellphone datasets.**
  - Visual factors play a smaller role when selecting cellphones.
- **Cold start items are 'unpopular', popularity-based methods were ineffective**

Experiments

Table 3: AUC on the test set $\mathcal{T}$ (#factors = 20). The best performing method on each dataset is boldfaced.

Dataset	Setting	(a) RAND	(b) MP	(c) IBR	(d) MM-MF	(e) BPR-MF	(f) VBPR	improvement f vs. best	improvement f vs. e
Amazon Women	All Items	0.4997	0.5772	0.7163	0.7127	0.7020	**0.7834**	9.4%	11.6%
	Cold Start	0.5031	0.3159	0.6673	0.5489	0.5281	**0.6813**	2.1%	29.0%
Amazon Men	All Items	0.4992	0.5726	0.7185	0.7179	0.7100	**0.7841**	9.1%	10.4%
	Cold Start	0.4986	0.3214	0.6787	0.5666	0.5512	**0.6898**	1.6%	25.1%
Amazon Phones	All Items	0.5063	0.7163	0.7397	0.7956	0.7918	**0.8052**	1.2%	1.7%
	Cold Start	0.5014	0.3393	**0.6319**	0.5570	0.5346	0.6056	-4.2%	13.3%
Tradesy.com	All Items	0.5003	0.5085	N/A	0.6097	0.6198	**0.7829**	26.3%	26.3%
	Cold Start	0.4972	0.3721	N/A	0.5172	0.5241	**0.7594**	44.9%	44.9%

- **VBPR outperforms all the baselines in most cases.**
- **More improvements in Tradesy.com settings. (cold start dataset)**
- **Visual features show greater benefits on clothing than cellphone datasets.**
  - Visual factors play a smaller role when selecting cellphones.
- **Cold start items are 'unpopular', popularity-based methods were ineffective**

Experiments

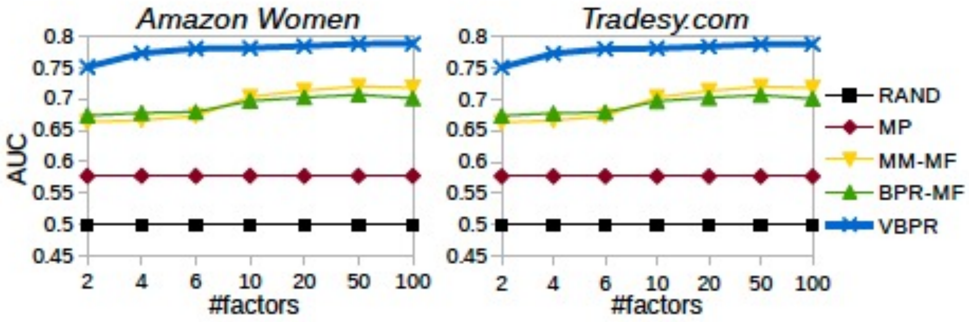

Figure 2: AUC with varying dimensions.

- **For MF-based methods perform better as the number of factors increases.**

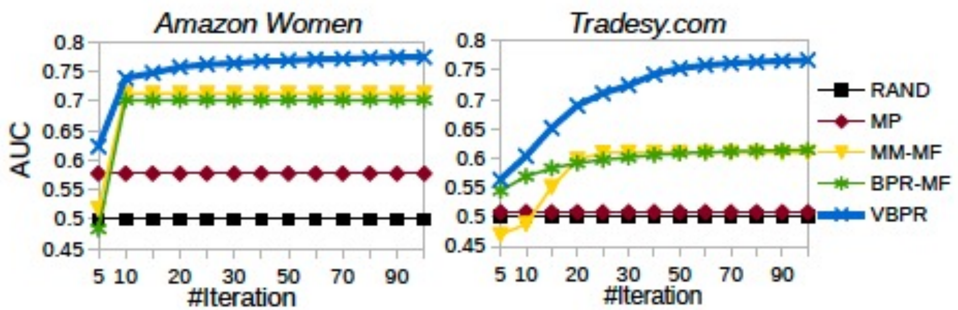

Figure 3: AUC with training iterations (#factors=20).

- **VBPR takes longer to converge than other MF-based methods. (about 3.5 hours)**

Experiments

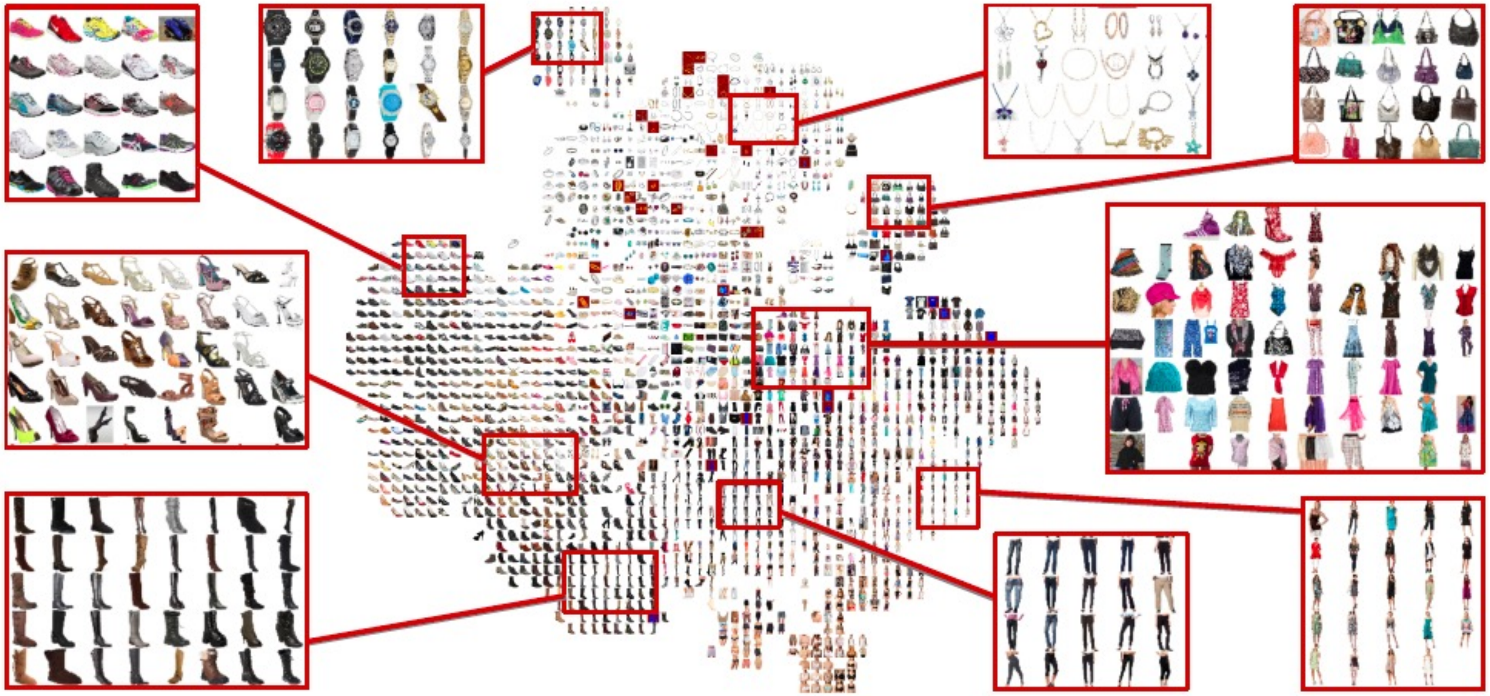

Figure 4: 2-D visualization (with t-SNE (?)) of the 10-D visual space learned from *Amazon Women*. All images are from the test set. For clarity, the space is discretized into a grid and for each grid cell one image is randomly selected among overlapping instances.

- **Although visual features were extracted from a CNN pre-trained on a different dataset, they were able to learn a visual transition across different categories.**

Conclusion

- **VBPR : methods that incorporates visual features for personalized ranking tasks.**
  - Trained on Bayesian Personalized Ranking (BPR) using stochastic gradient ascent.

- **VBPR outperforms state-of-the-art ranking techniques and alleviate cold start issues.**

INTELLIGENT DATA SYSTEMS LABORATORY

Thank you