Correcting Popularity Bias by Enhancing Recommendation Neutrality

Toshihiro Kamishima[*], Shotaro Akaho[*], Hideki Asoh[*], and Jun Sakuma[**]

[*]National Institute of Advanced Industrial Science and Technology (AIST), Japan

[**]University of Tsukuba, Japan

The 8th ACM Conference on Recommender Systems, Poster

@ Silicon Valley (Foster City), U.S.A., Oct. 8, 2014

Outline

Overview

Recommendation neutrality

- viewpoint feature, recommendation neutrality

Experiments

- popularity bias, histogram of predicted ratings, accuracy vs neutrality

Applications

- excluding information unwanted by a user, fair treatment of content providers, adherence of laws or regulations

Information-neutral Recommender System

- formalization, neutrality term

Recommendation neutrality vs recommendation diversity

- Recommendation diversity, differences between neutrality and diversity

Conclusion

Overview

Providing neutral information is important in recommendation

- excluding information unwanted by a user
- fair treatment of content suppliers or item providers
- adherence to laws and regulations in recommendation

Information-neutral Recommender System

The absolutely neutral recommendation is intrinsically infeasible, because recommendation is always biased in a sense that it is arranged for a specific user

This system makes recommendation so as to enhance neutrality with respect to a viewpoint feature

Viewpoint Feature

As in a case of standard recommendation, we use random variables
X: a user, Y: an item, and R: a rating

We adopt an additional variable for recommendation neutrality

V : **viewpoint feature**

- It is specified by a user depending on his/her purpose
- Recommendation results are neutral with respect to this viewpoint
- Its value is determined depending on a user and an item

Ex. viewpoint = movie's popularity / user's gender

In this presentation, a viewpoint feature is restricted to a binary type

Recommendation Neutrality

[Kamishima 12, Kamishima 13]

Recommendation Neutrality

- Recommendation results are neutral if no information about a viewpoint feature influences the results
- The status of the viewpoint feature is explicitly excluded from the inference of the recommendation results

the statistical independence
between a result, R, and a viewpoint feature, V

$$\Pr[R|V] = \Pr[R]$$

$$|||$$

$$R \perp\!\!\!\perp V$$

Ratings are predicted
under this constraint of recommendation neutrality

Experimental Results

Popularity Bias

[Celma 08]

Popularity Bias
the tendency for popular items to be recommended more frequently

Flixster data [Jamali+ 10]

The degree popularity of an item is measured
by the number of users who rated the item

short-head (top 1%)
share in ratings: 47.2%
mean rating: **3.71**

⟷

long-tail (bottom 99%)
share in ratings: 52.8%
mean rating: **3.53**

Short-head items are frequently and highly rated

⬇

By specifying this popularity as a viewpoint feature,
enhancing recommendation neutrality corrects this unwanted bias

Histograms of Predicted Ratings

standard recommender

two distributions are largely diverged

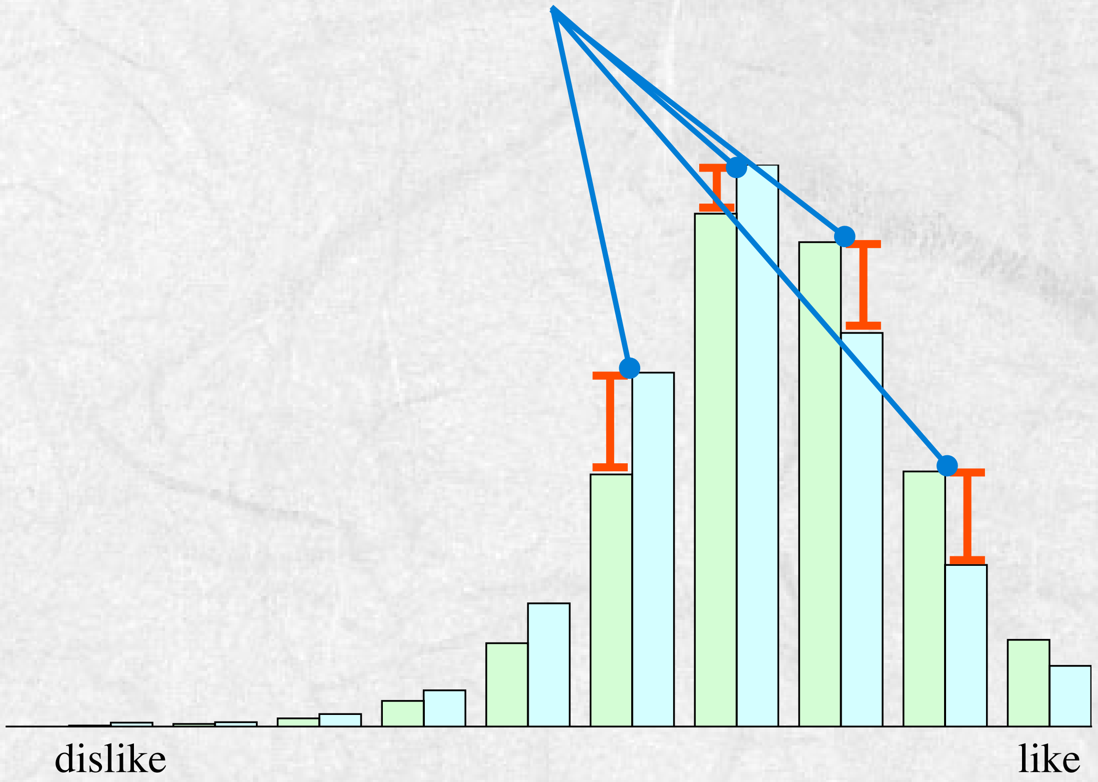

neutrality enhanced

distributions become close by enhancing neutrality

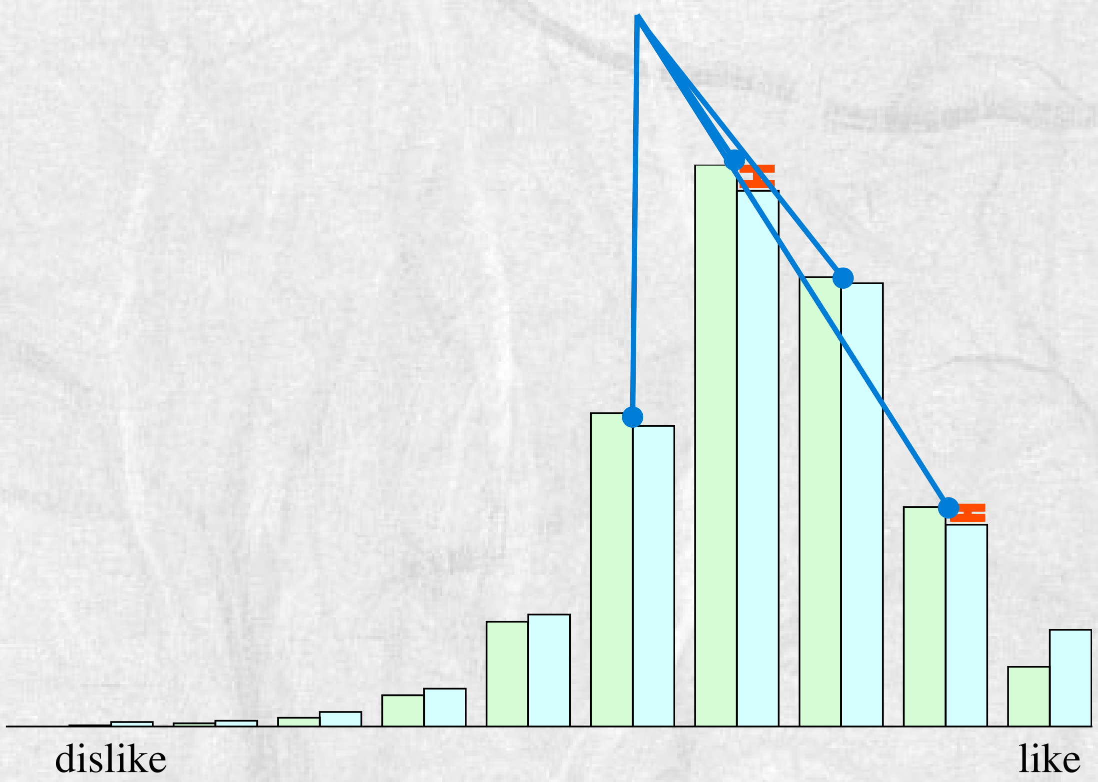

✤each bin of histograms of short-head and long-tail data are arranged

Unwanted information about popularity was successfully excluded by enhancing recommendation neutrality

Accuracy vs Neutrality

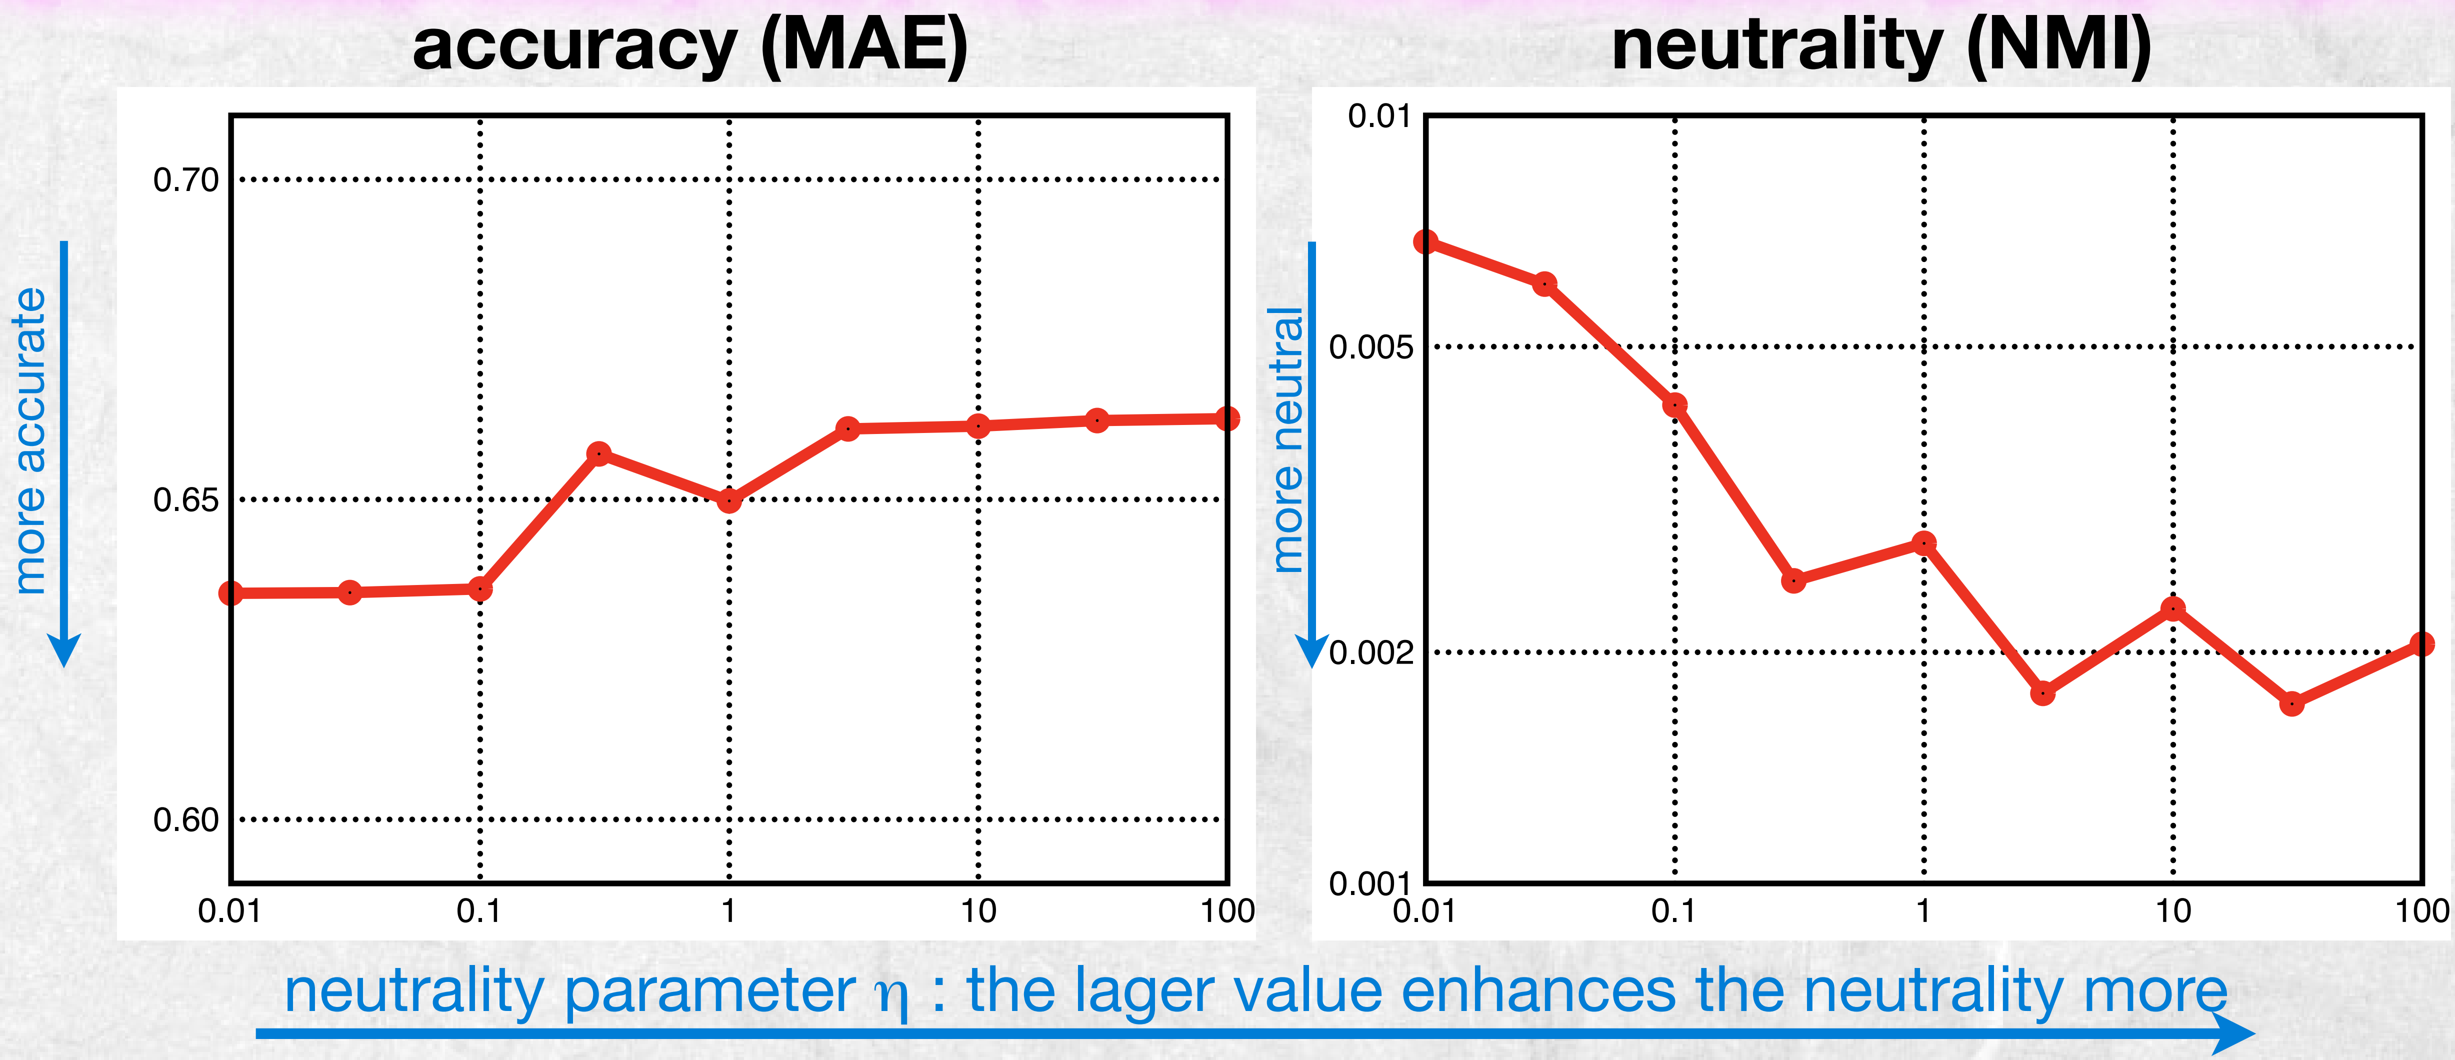

As the increase of a neutrality parameter η,

- accuracy was slightly worsened, but
- neutrality was successfully improved

Neutrality was enhanced, and thus a popularity bias was corrected, without sacrificing recommendation accuracy

Applications of Recommendation Neutrality

Application

Excluding Unwanted Information

[TED Talk by Eli Pariser, http://www.filterbubble.com/]

Information unwanted by a user is excluded from recommendation

Filter Bubble: To fit for Pariser's preference, conservative people are eliminated from his friend recommendation list in FaceBook

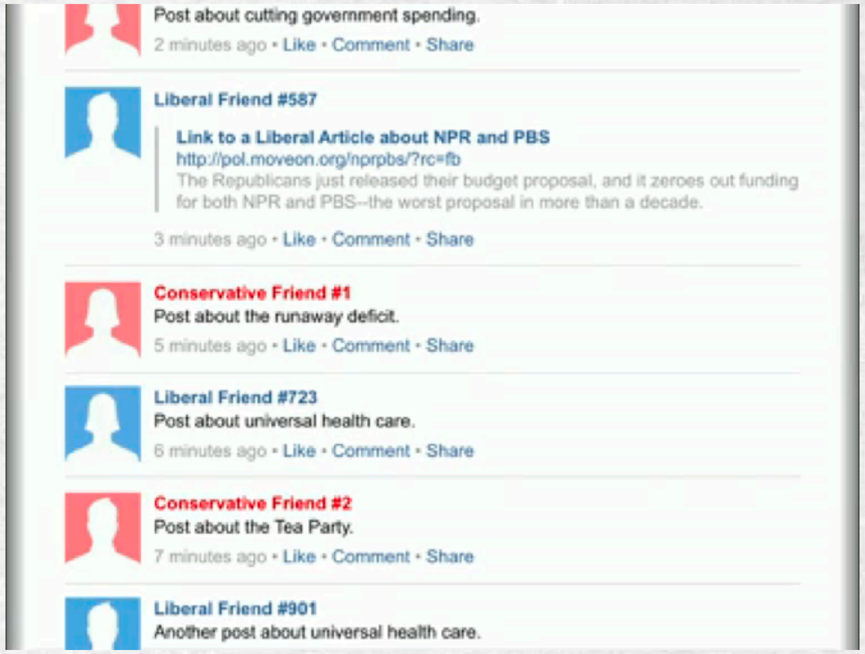

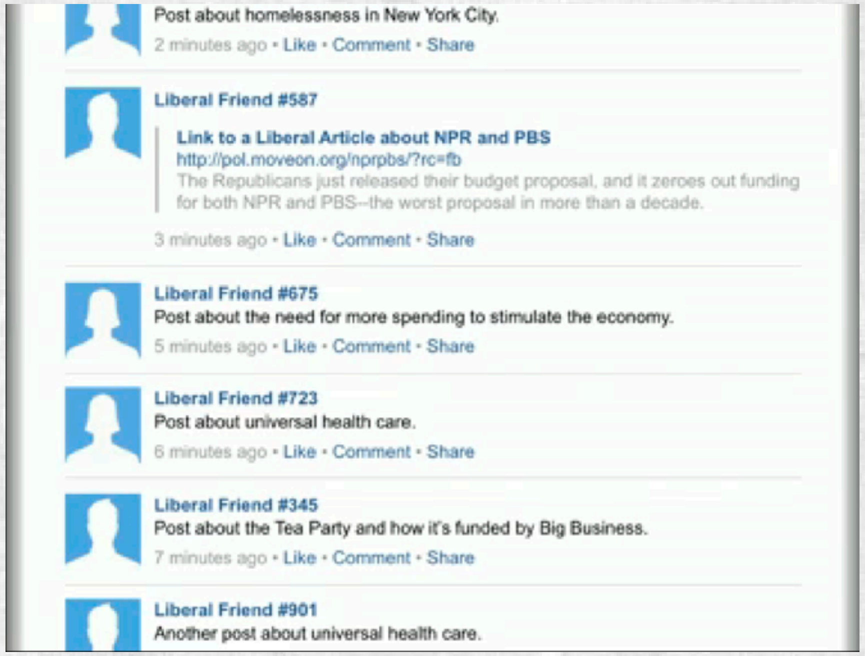

viewpoint = a political conviction of a friend candidate

⬇

Information about a candidate is conservative or progressive does not influence whether he/she is included in a friend list or not

Application
Fair Treatment of Content Providers

System managers should fairly treat their content providers

Ranking in a list retrieved by search engines

[Bloomberg]

The US FTC has been investigating Google to determine whether the search engine ranks its own services higher than those of competitors

Content providers are managers' customers

For marketplace sites, their tenants are customers, and these tenants must be treated fairly when recommending the tenants' products

viewpoint = a content provider of a candidate item

⬇

Information about who provides a candidate item is ignored, and providers are treated fairly

Application
Adherence to Laws and Regulations

[Sweeney 13]

Recommendation services must be managed while adhering to laws and regulations

suspicious placement keyword-matching advertisement

Advertisements indicating arrest records were more frequently displayed for names that are more popular among individuals of African descent than those of European descent

⬇

Socially discriminative treatments must be avoided

viewpoint = users’ socially sensitive demographic information

⬇

Legally or socially sensitive information
can be excluded from the inference process of recommendation

Information-neutral Recommender System

Information-neutral PMF Model

information-neutral version of a PMF model

- **adjust ratings according to the state of a viewpoint**
  incorporate dependency on a viewpoint variable
- **enhance the neutrality of a score from a viewpoint**
  add a neutrality function as a constraint term

adjust ratings according to the state of a viewpoint

viewpoint feature

$$\hat{r}(x,y,v) = \mu^{(v)} + b_x^{(v)} + c_y^{(v)} + \mathbf{p}_x^{(v)} {\mathbf{q}_y^{(v)}}^\top$$

- Multiple models are built separately, and each of these models corresponds to the each value of a viewpoint feature
- When predicting ratings, a model is selected according to the value of viewpoint feature

Neutrality Term and Objective Function

enhance the neutrality of a rating from a viewpoint feature

neutrality term, $\mathrm{neutral}(R, V)$: quantify the degree of neutrality

- It depends on both ratings and view point features
- The larger value of the neutrality term indicates that the higher level of the neutrality

Objective Function of an Information-neutral PMF Model

neutrality parameter to control the balance between the neutrality and accuracy

regularization parameter

$$\underbrace{\textstyle\sum_{\mathcal{D}}(r_i - \hat{r}(x_i, y_i, v_i))^2}_{\text{squared loss function}} - \eta \, \underbrace{\mathrm{neutral}(R, V)}_{\text{neutrality term}} + \lambda \, \underbrace{\|\mathbf{\Theta}\|_2^2}_{L_2 \text{ regularizer}}$$

Parameters are learned by minimizing this objective function

Neutrality Term

Calders&Verwer's Score (CV Score)

make two distributions of R given $V=0$ and 1 similar

$$-\| \Pr[R|V=0] - \Pr[R|V=1] \|$$

mean match method

$$-\left(\mathrm{Mean}_{\mathcal{D}^{(0)}}[\hat{r}] - \mathrm{Mean}_{\mathcal{D}^{(1)}}[\hat{r}]\right)^2$$

Matching means of predicted ratings
for each data set where V=1 and V=0

it matches only the first moment, but it empirically works well

analytically differentiable and efficient in optimization

Recommendation Neutrality
vs
Recommendation Diversity

Recommendation Diversity

[Ziegler+ 05, Zhang+ 08, Latha+ 09, Adomavicius+ 12]

Recommendation Diversity

Similar items are not recommended in a single list, to a single user, to all users, or in a temporally successive lists

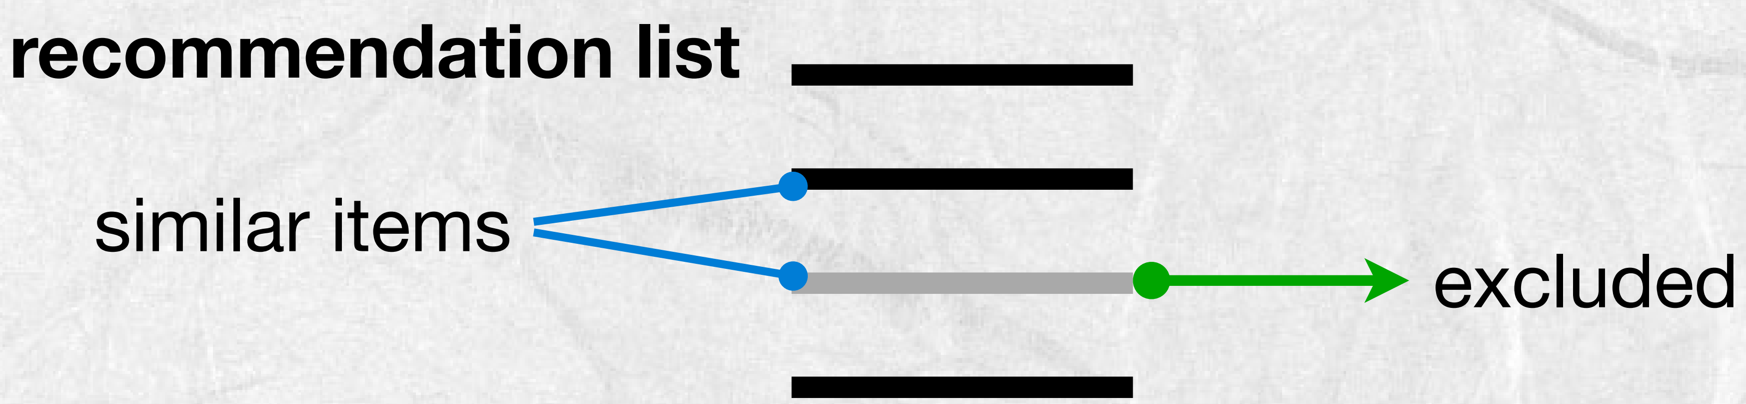

Diversity

Items that are similar in a specified metric are excluded from recommendation results

The mutual relations among results

Neutrality

Information about a viewpoint feature is excluded from recommendation results

The relations between results and viewpoints

Neutrality vs Diversity

Diversity	Neutrality
Depending on the definition of similarity measures	Depending on the specification of viewpoint
↓	↓
Similarity	**Viewpoint**
A function of **two items**	A function of a pair of **an item and a user**

Because a viewpoint depends on a user, neutrality can be applicable for coping with users' factor, such as, users' gender or age, which cannot be straightforwardly dealt by using diversity

Neutrality vs Diversity

standard recommendation

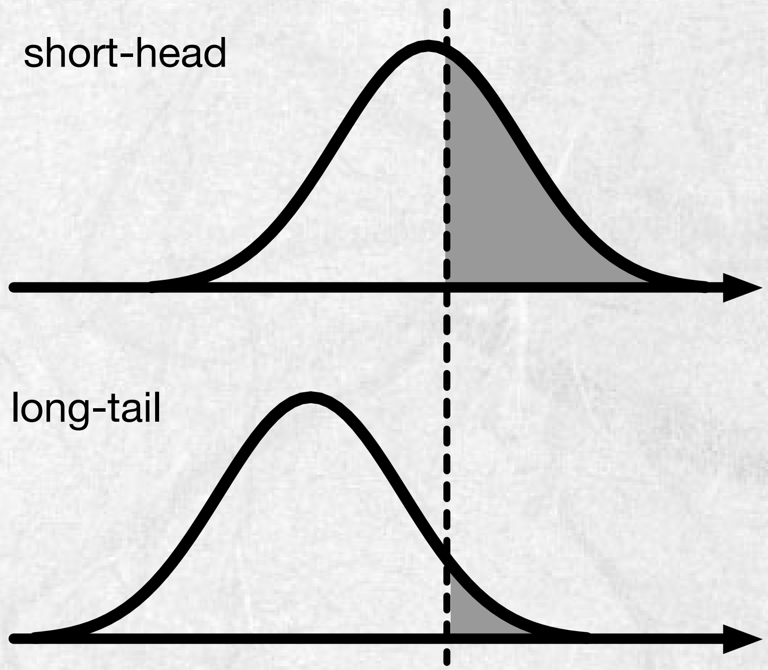

diversified recommendation

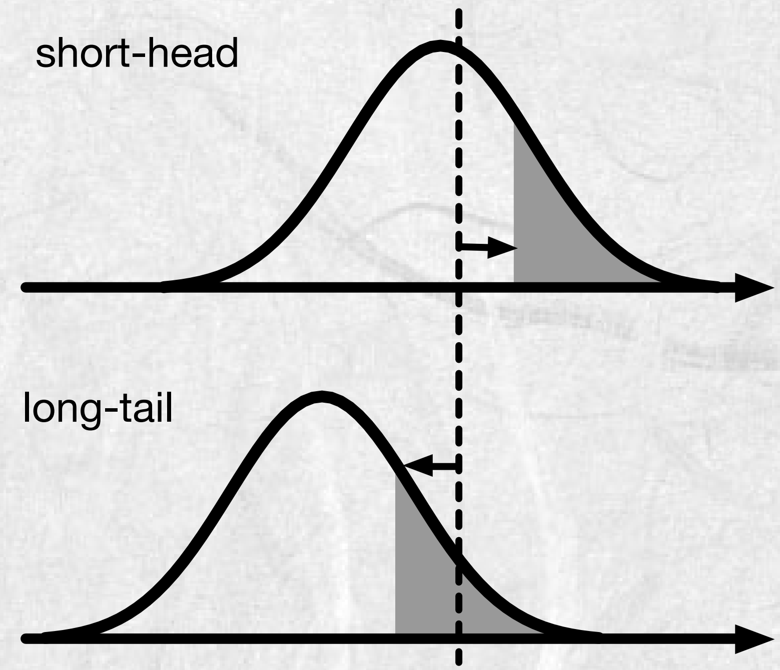

Because a set of recommendations are diversified by abandoning short-head items, **predicted scores are still biased**

Prediction scores themselves are unbiased by enhancing neutrality

Conclusion

Our Contributions

- We formulate the recommendation neutrality from a specified viewpoint feature
- We developed a recommendation technique that can enhance the recommendation neutrality
- We applied this technique to correct popularity bias, and the effectiveness of this is empirically shown

Future Work

- Developing a neutrality term that can more precisely approximate distributions without losing its computational efficiency

Program codes: http://www.kamishima.net/inrs/

Acknowledgements

We would like to thank for providing a data set for Dr. Mohsen Jamali

This work is supported by MEXT/JSPS KAKENHI Grant Number 16700157, 21500154, 24500194, and 25540094