ResponsibleRecs

Fairness, Diversity and Transparency in Recommender Systems

Kostas Stefanidis

Tampere University, Finland

konstantinos.stefanidis@tuni.fi

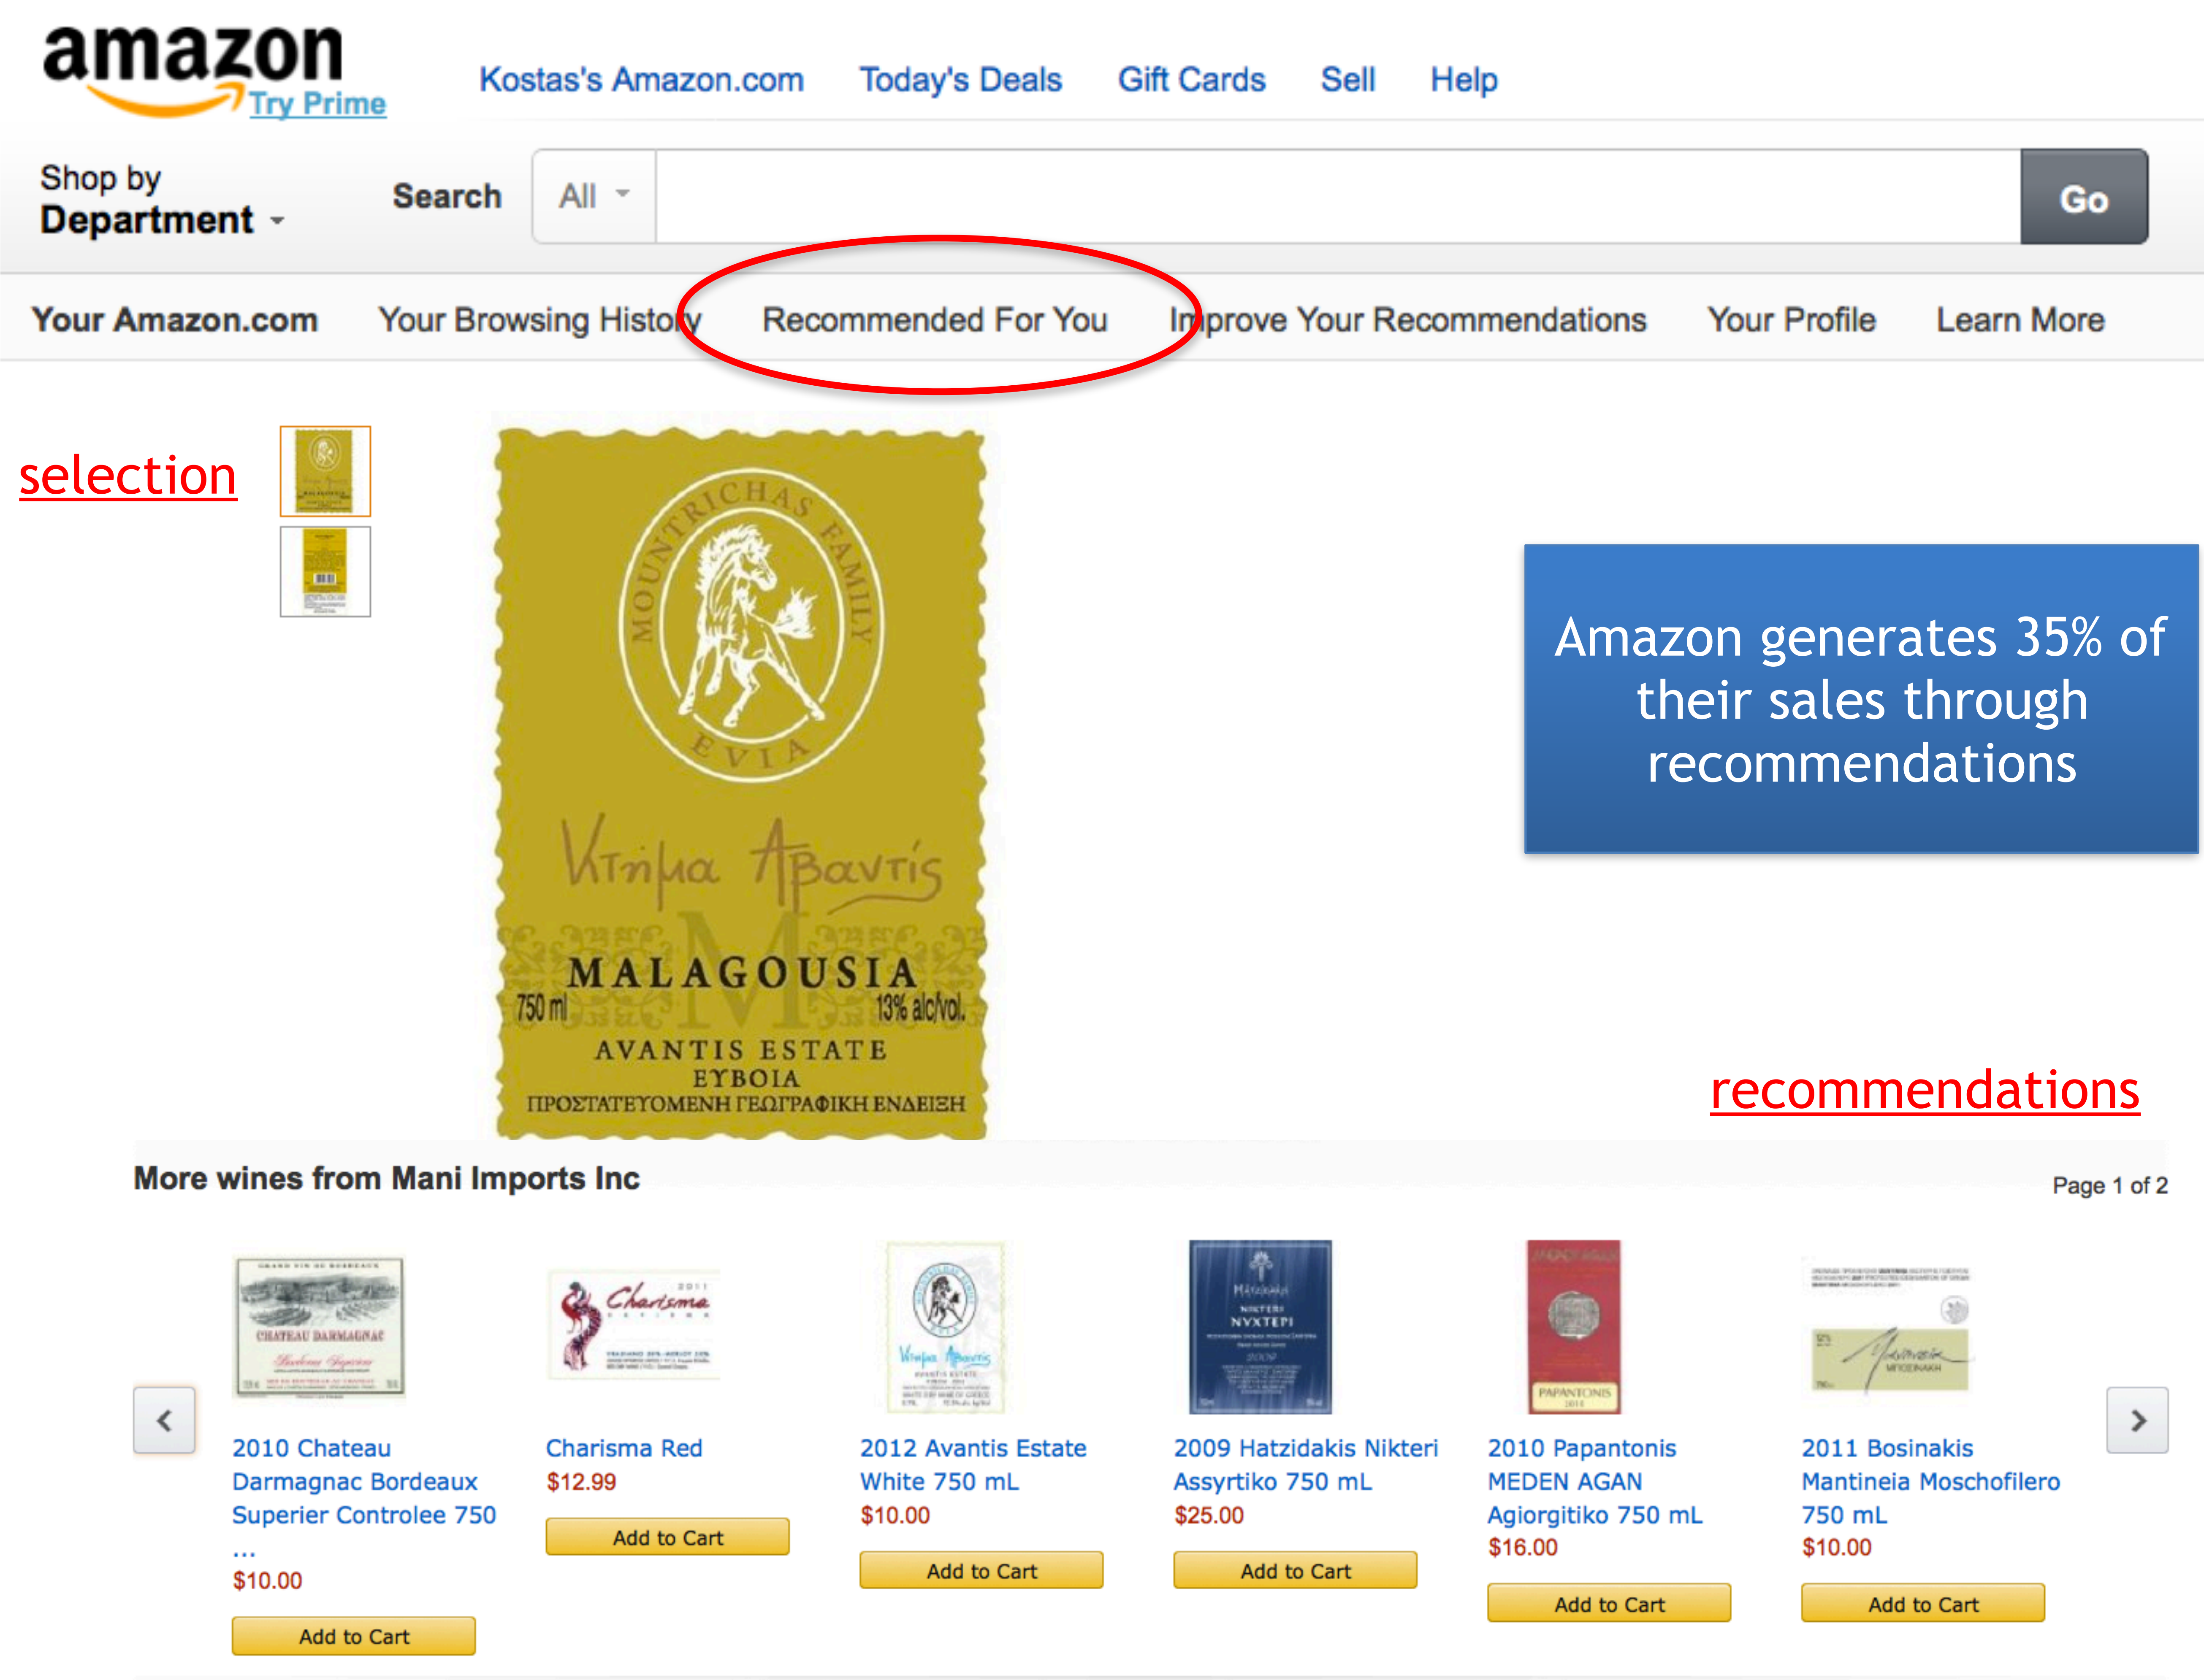

amazon
Try Prime
Kostas's Amazon.com
Today's Deals
Gift Cards
Sell
Help
Shop by
Department
Search
All
Go
Your Amazon.com
Your Browsing History
Recommended For You
Improve Your Recommendations
Your Profile
Learn More
selection
MOUNTRICHAS FAMILY
EVIA
MALAGOUSIA
750 ml
13% alc/vol.
AVANTIS ESTATE
ΕΥΒΟΙΑ
ΠΡΟΣΤΑΤΕΥΟΜΕΝΗ ΓΕΩΓΡΑΦΙΚΗ ΕΝΔΕΙΞΗ
Amazon generates 35% of their sales through recommendations
recommendations
More wines from Mani Imports Inc
Page 1 of 2
2010 Chateau Darmagnac Bordeaux Superier Controlee 750 ...
$10.00
Add to Cart
Charisma Red
$12.99
Add to Cart
2012 Avantis Estate White 750 mL
$10.00
Add to Cart
2009 Hatzidakis Nikteri Assyrtiko 750 mL
$25.00
Add to Cart
2010 Papantonis MEDEN AGAN Agiorgitiko 750 mL
$16.00
Add to Cart
2011 Bosinakis Mantineia Moschofilero 750 mL
$10.00
Add to Cart

Booking.com

Aquila Atlantis Hotel ★★★★★

2, Igias Str., Heraklio Town, 71202, Greece – Show map Share

Book now

Available rooms Facilities House rules The fine print See all reviews

selection

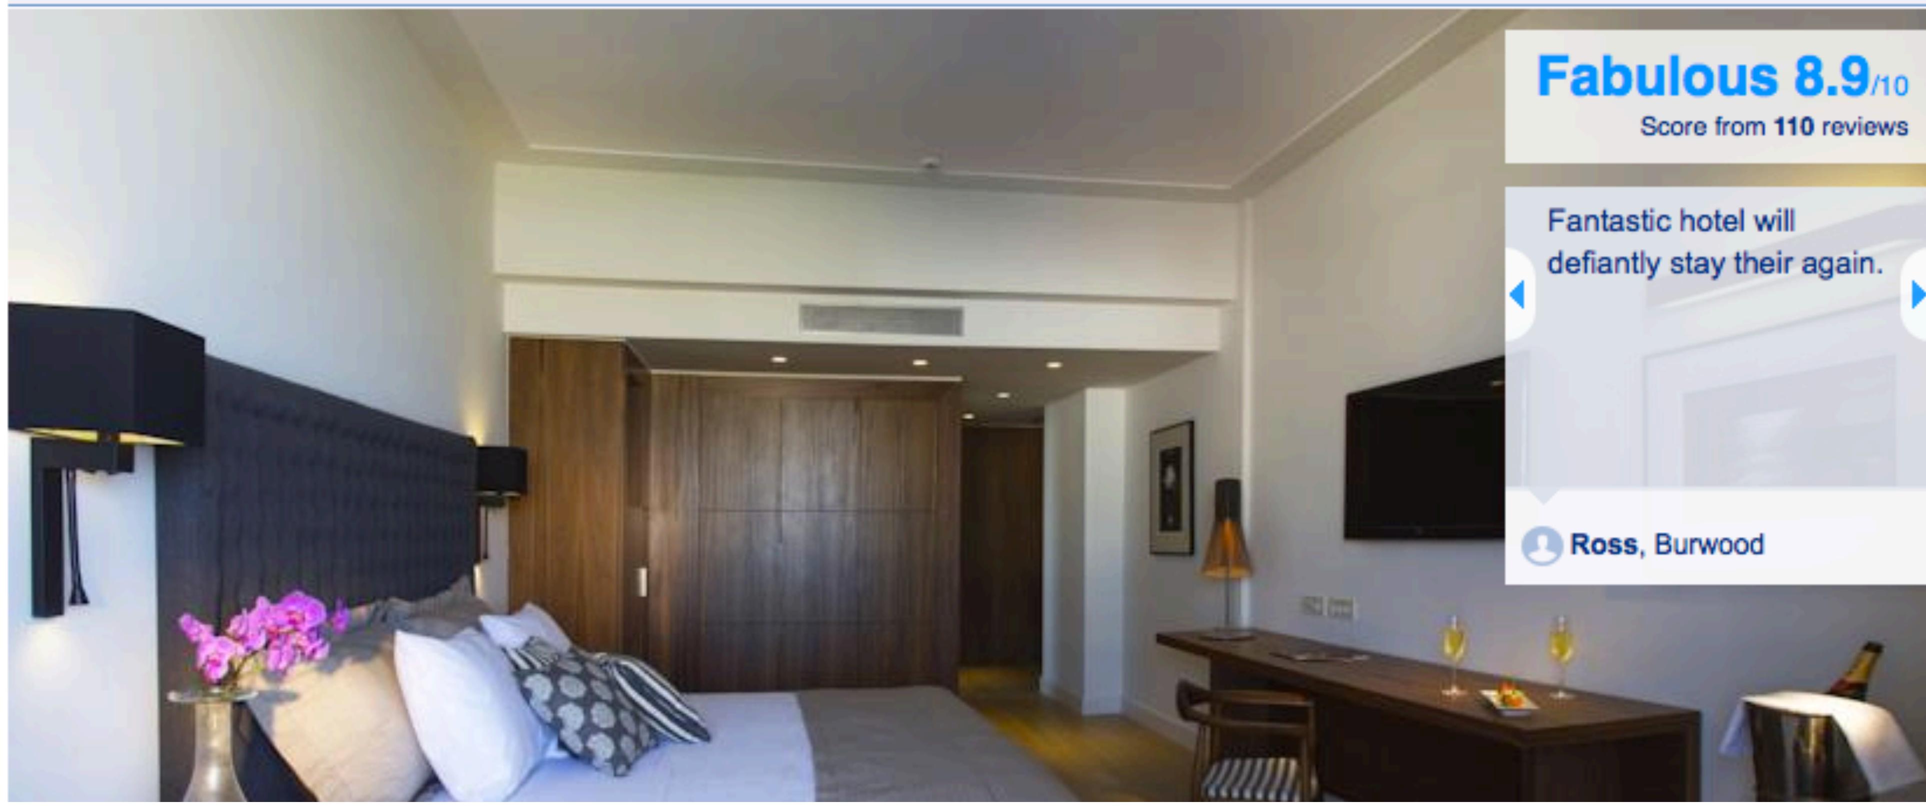

recommendations

Kostas, we found more properties like Aquila Atlantis Hotel

Galaxy Iraklio Hotel ★★★★★

Located in Heraklion's elegant district, this 5-star hotel offers 2 gourmet restaurants, a free wellness centre and a large freshwater pool. Luxurious rooms feature balconies with pool and city views....

Most recent booking for this hotel was today at 00:40

Score from 450 reviews
Fabulous - 8.9/10

Total price from:
€ 119

Book now

Capsis Astoria Hotel ★★★★

This well-known central hotel in Heraklion is located next to the Archaeological Museum.

Most recent booking for this hotel was today at 05:57

Score from 252 reviews
Very good - 8.2/10

Total price from:
€ 75.60

Book now

Atrion Hotel ★★★

Just a short walk from Heraklion centre and the sandy beach, Atrion offers elegant accommodation with free Internet access.

Most recent booking for this hotel was today at 12:51

Score from 458 reviews
Very good - 8.5/10

Total price from:
€ 73

Book now

Lato Boutique Hotel ★★★

Situated opposite the old city harbour, Lato Boutique Hotel features a rooftop restaurant-bar overlooking Heraklion's Venetian Fortress.

Most recent booking for this hotel was today at 11:38

Score from 799 reviews
Fabulous - 8.7/10

Total price from:
€ 82

Book now

LinkedIn

Groups You May Like

Computer Vision and Pattern Recognition
Join - Professional Group

AMMA (Amsterdam Master's in Medical Anthropology) Alumnae
Join - Alumni Group

Agile Research
Join - Professional Group

Feedback | See more »

Companies You May Want To Follow

Feedback | See more »

Jobs You May Be Interested In

Senior Web Developer
Linakis Digital - Greece

Information Security Consultant
Trustwave - Athens, Greece

Enterprise Propositions & Pricing...
Vodafone - Athens

Feedback | See more »

Recommender Systems

Recommender systems aim at suggesting to users data items of potential interest to them

Two main steps:

- <u>Estimate</u> a rating for each item and user
- <u>Recommend</u> to the user the item(s) with the highest rating(s)

Data for and about people

The data is here!

Data for and about people

The potential benefits of recommenders are well-accepted

- The importance of using such techniques in a fair, diverse and transparent manner is only recently considered

Our objective: *create new algorithms for generating responsible recommendations for individual users and groups*

- Ensure **fairness**, **diversity** and **transparency**

Design and develop algorithms for both offline and online scenarios:

- All data items are available before any selections have to be made
- Not all data items are available at once; classify each individual item as presented, into the selected ones or not for recommendations

Apply to many domains, from traditional *products recommendations* to *health-related recommendations*

The promise of Big Data

With the growing complexity of the available online information, users find themselves overwhelmed by the mass of choices available

- Recommender systems provide suggestions on data items of potential interest to the users

Big data technology comes with the promise to improve people's lives towards this direction by enhancing the discovery of interesting information, and provide results tailored to users' profiles

- Remember the power!: Enormous datasets, Enormous computational power, Massively parallel processing

However!

The same technology, if not used responsibly, may lead to discrimination, amplify biases in the original data, restrict transparency and strengthen unfairness

- Recommendations may play an important role in guiding users' decisions and forming their opinions

E.g., consider scenarios in which models based on biased data can abet violence, increase diversity issues, have an impact on economy policies, or even amplify discrimination in the justice system

Fairness

By fairness, we typically mean lack of bias

Is it correct to assume that insights achieved via computations on data are unbiased simply because data was collected automatically or processing was performed algorithmically?

- Bias may come from the algorithm, reflecting, e.g., commercial or other preferences of its designers, or even from the actual data, e.g., if a survey contains biased questions

So far: recommenders consider the notion of fairness *indirectly*

- No explicit models and algorithms
- For group recommendations, there are approaches that consider disagreements [1] and social relationships [2] among group members, or use fairness to measure the quality of a set of items for a group [3]

Our goal is to model and formally define fairness, as well as to introduce algorithms that directly optimize it

Diversity

Diversity ensures that different kinds of data items are represented in the output of an algorithmic process

- E.g., in a news recommender, suggest news from several political parties and not only from the user's favourite one

There is considerable work on search result diversification:

- Diverse keyword database search using preferences [4] / Assess the topical diversity of recommendation lists using order-independent intra-list similarity measures [5] / …

Our goal here is to consider a family of diversity constraints that can express coverage-based, in addition to distance-based (relying on the pairwise similarity between items), diversity.

Re-define the measures of fairness

Fairness is related to diversification: when considering that a fair set of data items is likely to include items that represent different, or even all, categories of data items, or when considering groups, that satisfy different users

Transparency

Users many times want to know and control both what is being recorded about them, and how this piece of information is being used, e.g., to recommend content or for target advertising

<u>We skip privacy for now</u>

A transparent data analysis framework requires suggestions that can be easily understood by the users

- Can we explain the output of the algorithm?

Our goal is to suggest explainable items along with explanations — <u>Provide the WHY!</u>

- A user-based explanation is based on similar users
- An item-based explanation presents the items that had the highest impact on the recommender's decision

We aim at an integrated approach that considers explanations in the recommendation process rather than separating the explanation from the recommendation process

- <u>Recommendations along with their associated explanations will form graph-based summaries that include data items that ensure fairness and diversity</u>

Online Recommendations

But, all data items are available for evaluation at once?

- Many times, items appear one at a time, with a decision to be made on the specific item instantaneously

Online scenario / Streaming data / Sequential recommendations, i.e., suggestions in rounds, by exploiting newly added data items

The idea: Process items *incrementally*, maintaining a valuable recommendations set at any point in time

- [6] considers a fixed window of recent items, posing a problem for items not generated at a fixed rate
- [7] proposes algorithms to diversify a stream of results using a jumping window approach

Our goal: Provide a fair, diverse and transparent set of items considering the whole item set, rather than a fixed number of recent items

- Allow actively withdrawing items from the recommendations set, instead of simply dropping items as they leave the window

To our knowledge, combing fairness, diversity and transparency, especially in an online setting has not been considered

Research Agenda

The focus is on <u>set-based selections</u>: Both fairness and diversity are set-based concepts

- It makes no sense to talk about an individual item as being diverse
- Currently, most algorithmic decision-making approaches are based on individual items, where a utility score is associated with each item typically computed with respect to the values of the item

<u>Group Recommendations</u>: Even more so than in traditional recommendations for individual users, identifying data items of high relevance to a group is challenging, especially for cases where group members disagree on their dearest items

- We focus on developing novel data analysis methods that ensure **fairness**, **diversity** and **transparency** in *set selection for recommendations*
- We aim to design and develop algorithms that generate recommendations for **individual users** and **groups**, for both when data items are available before any selections have to be made and the **online case**

Generating responsible recommendations is timely due to the available big data technologies and the current discussions on bias and discrimination in algorithmic processing and decision making, yet is not enough supported by existing models and algorithms

Objective - Fairness & Diversity

Previous work focused *separately* on fairness or diversity in query processing and recommender systems

- This is a useful basis but must be significantly extended to bring in both fairness and diversity

Preliminary: *Fairness can be defined as the proportional representation of the values of attributes of particular concern, and diversity as the existence of such values*

How fairness and diversity can be combined with respect to the users preferences as expressed by the users' ratings?

Objective - Transparency & Explainability

Recommend items based on their explainability: Integrate recommendation and explanation to improve transparency

NOT! Data employed for producing recommendations can be different from the data used in generating the explanations, leading to explanation generation modules that are separate from the recommender system

E.g., consider “black boxes” in machine learning

Provide explainable recommendations through summaries: produce an abridged version of results

- We do not focus on providing the data items with the maximum utility score, but on summaries consisting of explainable items that exhibit fairness and diversity
- We aim to handle efficiently exploratory operations, like zoom-in and zoom-out, providing granular information access to the user

Objective - Individual User & Group Recommendations

Special focus on group recommendations!

<u>New definitions for fairness and diversity applicable to groups</u>

Fairness for groups via *envy-freeness*: A user considers a set of items fair for him/her, if there are items for which the user does not feel jealous, i.e., the presented items have utility scores within the range of scores of the best items for him/her

Diversity for groups via *coverage*: Rely on the existence of a number of items for all group members

+ Integration of items' *explainability*

Objective - Static & Online Processing

The static case solves the problem making the assumption that all data items are available at processing time

- Fairness, diversity and transparency constraints will direct the process; return the items with the highest utility computed with respect to our constraints

In the online case, not all items are available at once

- Classify each individual item, as presented, into the selected or not selected ones based on the fairness, diversity and transparency constraints
- How? Extend the K-choice Secretary Problem [8] to design online methods for picking items, presented in random order - independently to their utility, subject to fairness, diversity and transparency constraints

References

1. Sihem Amer-Yahia, Senjuti Basu Roy, Ashish Chawla, Gautam Das, and Cong Yu. 2009. Group Recommendation: Semantics and E ciency. PVLDB 2, 1 (2009), 754–765
2. Keqian Li, Wei Lu, Smriti Bhagat, Laks V.S. Lakshmanan, and Cong Yu. 2014. On social event organization. In KDD
3. Shuyao Qi, Nikos Mamoulis, Evaggelia Pitoura, and Panayiotis Tsaparas. 2016. Recommending Packages to Groups. In ICDM
4. Kostas Stefanidis, Marina Drosou, and Evaggelia Pitoura. 2010. PerK: personalized keyword search in relational databases through preferences. In EDBT
5. Cai-Nicolas Ziegler, Sean M. McNee, Joseph A. Konstan, and Georg Lausen. 2005. Improving recommendation lists through topic diversification. In WWW
6. Marina Drosou, Kostas Stefanidis, and Evaggelia Pitoura. 2009. Preference- aware publish/subscribe delivery with diversity. In DEBS
7. Marina Drosou and Evaggelia Pitoura. 2009. Diversity over Continuous Data. IEEE Data Eng. Bull. 32, 4 (2009), 49–56
8. Moshe Babaio, Nicole Immorlica, and Robert Kleinberg. 2007. Matroids, secretary problems, and online mechanisms. In SODA