UMAP 2019 Fairness in User Modeling, Adaptation and Personalization (FairUMAP 2019) Chairs' Welcome & Overview

It is our great pleasure to welcome you to the Second *FairUMAP* workshop at *UMAP 2019*. This full-day workshop brings together researchers working at the intersection of user modeling, adaptation, and personalization on one hand, and bias, fairness and transparency in algorithmic systems on the other hand.

The workshop was motivated by the observation that these two fields increasingly impact one another. Personalization has become a ubiquitous and essential part of systems that help users find relevant information in today's highly complex, information-rich online environments. Machine learning techniques applied to big data, as done by recommender systems, and user modeling in general, are key enabling technologies that allow intelligent systems to learn from users and adapt their output to users' needs and preferences. However, there has been a growing recognition that these underlying technologies raise novel ethical, legal, and policy challenges. It has become apparent that a single-minded focus on user characteristics has obscured other important and beneficial outcomes such systems must be able to deliver. System properties such as fairness, transparency, balance, and other social welfare considerations are not captured by typical metrics based on which data-driven personalized models are optimized. Indeed, widely-used personalization systems in popular sites such as Facebook, Google News and YouTube have been heavily criticized for personalizing information delivery too heavily at the cost of these other objectives.

Bias, fairness, and transparency in algorithmic systems are topics of considerable recent research interest. However, more work is needed to expand and extend this field of research into algorithmic and modeling approaches where personalization is of primary importance. The goal of this workshop was to bring together a growing community of experts from academia and industry to discuss ethical, social, and legal concerns related to personalization, and specifically to explore a variety of mechanisms and modeling approaches that help mitigate bias and achieve fairness in personalized systems.

We invited papers on a range of topics including Bias and discrimination in user modeling, personalization and recommendation; Computational techniques and algorithms for fairness-aware personalization; Definitions, metrics and criteria for optimizing and evaluating fairness-related aspects of personalized systems; Data preprocessing and transformation methods to address bias in training data; User modeling approaches that take fairness and bias into account; User studies to evaluate impact of personalization on fairness, balance, diversity, and other social welfare criteria; Balancing needs of multiple stakeholders in recommender systems and other personalized systems; and 'Filter bubble' or 'balkanization' effects of personalization; Transparent and accurate explanations for recommendations and other personalization outcomes.

These topics were addressed by five contributed papers as well as other speakers. More details about these papers and speakers can be found on the workshop Website: <https://fairumap.wordpress.com/>. Below, we provide a brief summary of the accepted contributions.

Three of the papers focus on various issues surrounding *fairness in recommendation systems*.

Zhang, Ghane, and Sabouri present a first application of multi-stakeholder recommendation systems (MSRS) in the context of a personalized learning system, which recommends datasets for student projects. They develop a mechanism through which the utilities of various stakeholders (students and instructors) are balanced, recommending the optimal solution in the Pareto optimal set.

Sonboli and Burke introduce the concept of local fairness, which, in contrast to notions of group or individual fairness, is more suitable for use in recommendation systems, where what is considered fair depends on the local context. They present a case study of the crowdfunding site, Kiva.org, which aims to equalize access to capital across all types of small businesses in developing countries. The analysis demonstrates that, rather than defining protected/unprotected businesses a priori (i.e., across all loan-seekers at Kiva), it is necessary to define protected groups relative to local conditions, before fairness-awareness algorithms are applied.

In a position paper, *Herder* argues in favor of researchers taking a multi-stakeholder perspective, in which commercial needs are considered, alongside those of end-users. After a brief comparison of the typical approaches taken in the UMAP and RecSys communities, he arrives at the conclusion that to truly realize multi-sided fair recommendations, researchers must give monetary benefits/incentives full consideration.

One of the contributions focuses on *fairness in machine learning* and specifically, classification.

The paper by *Cummings, Gupta, Kimpara and Morgenstern* demonstrates the existence of a classifier that simultaneously archives privacy and fairness goals; individuals' privacy in the training data is protected, while Equality of False Positives/Negatives is ensured for two populations. In particular, the authors motivate and develop a private version of the Fair No-Regret Dynamics (FairNR) algorithm.

Finally, one contribution considers the *reasonableness of machine inferences*.

Inspired by recent calls for legal measures to ensure citizens' "Right to Reasonable Inferences," *Engelmann and Grossklags* present a case study of the Clarifai image tagging service and its treatments of people images, which explores the question of what is "reasonable." Drawing from philosophical notions of reasoning, they argue that a tagging algorithms' accuracy is independent of the reasonableness of its inferences.

FairUMAP 2019 Organization and Program

Workshop Chairs: Bettina Berendt (*KU Leuven, Belgium*)
Veronika Bogina (*University of Haifa, Israel*)
Robin Burke (*University of Colorado, USA*)
Michael Ekstrand (*Boise State University, USA*)
Alan Hartman (*University of Haifa, Israel*)
Styliani Kleanthous (*Open University of Cyprus, Cyprus*)
Tsvi Kuflik (*University of Haifa, Israel*)
Bamshad Mobasher (*DePaul University, USA*)
Jahna Otterbacher (*Open University of Cyprus, Cyprus*)

Program Committee: Solon Barocas (*Microsoft Research, Cornell University, USA*)
Pablo Castells (*Universidad Autónoma de Madrid, Spain*)
David Chin (*University of Hawaii, USA*)
Jean Garcia-Gathright (*Spotify, USA*)
Krishna Gummadi (*MPI-SWS, Germany*)
Sara Hajian (*Eurecat, Technology Center of Catalonia, Spain*)
Eelco Herder (*Radboud Univeristy, The Netherlands*)
Toshihiro Kamishima (*AIST, Japan*)
Daniel Kluver (*University of Minnesota / Macalester College, USA*)
Bart Knijnenburg (*Clemson University, USA*)
Pradeep Kumar Murukannaiah (*Rochester Institute of Technology, USA*)
Yong Zheng (*Illinois Institute of Technology, USA*)

Accepted Papers:

- Rachel Cummings, Varun Gupta, Dhamma Kimpara and Jamie Morgenstern
On the Compatibility of Privacy and Fairness
- Severin Engelmann and Jens Grossklags
Setting the Stage: Towards Principles for Reasonable Image Inferences
- Eelco Herder
The Need For Identifying Ways To Monetize Personalization and Recommendation
- Nasim Sonboli and Robin Burke
Localized Fairness in Recommender Systems
- Yong Zheng, Nastaran Ghane and Milad Sabouri
Personalized Learning with Multi-Stakeholder Optimizations