

FAT/MM'19: 1st International Workshop on Fairness, Accountability, and Transparency in MultiMedia

Xavier Alameda-Pineda
Inria, Perception Team
Grenoble, France
xavier.alameda-pineda@inria.fr

Miriam Redi
Wikimedia Foundation
San Francisco, California
mredi@wikimedia.org

Elisa Celis
Yale University
New Haven, Connecticut
elisa.celis@yale.edu

Nicu Sebe
University of Trento
Trento, Italy
niculae.sebe@unitn.it

Shih-Fu Chang
Columbia University
New York, New York
sfchang@cs.columbia.edu

ABSTRACT

The series of FAT* events aim at bringing together researchers and practitioners interested in fairness, accountability, and transparency of computational methods. The FAT/MM workshop focuses on addressing these issues in the Multimedia field. Multimedia computing technologies operate today at an unprecedented scale, with a growing community of scientists interested in multimedia models, tools and applications. Such continued growth has great implications not only for the scientific community, but also for the society as a whole. Typical risks of large-scale computational models include model bias and algorithmic discrimination. These risks become particularly prominent in the multimedia field, which historically has been focusing on user-centered technologies. To ensure a healthy and constructive development of the best multimedia technologies, this workshop offers a space to discuss how to develop fair, unbiased, representative, and transparent multimedia models, bringing together researchers from different areas to present computational solutions to these issues. Related Workshop Proceedings are available in the ACM DL at: <https://dl.acm.org/citation.cfm?id=3347447>.

CCS CONCEPTS

• Information systems; • Security and privacy; • Human-centered computing; • Social and professional topics;

ACM Reference Format:

Xavier Alameda-Pineda, Miriam Redi, Elisa Celis, Nicu Sebe, and Shih-Fu Chang. 2019. FAT/MM'19: 1st International Workshop on Fairness, Accountability, and Transparency in MultiMedia. In *Proceedings of the 27th ACM International Conference on Multimedia (MM '19), October 21–25, 2019, Nice, France*. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3343031.3350555>

1 INTRODUCTION

The computational inclusiveness and transparency of automatic information processing methods is a research topic that exhibited

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

MM '19, October 21–25, 2019, Nice, France

© 2019 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6889-6/19/10.

<https://doi.org/10.1145/3343031.3350555>

growing interest in the past years. In the era of digitized decision-making software where the push for artificial intelligence happens worldwide and at different strata of the socio-economic tissue, the consequences of biased, unexplainable and opaque methods for multimedia analysis and content retrieval, can be dramatic. In this context, the multimedia community must put together the necessary efforts in applying its expertise and know-how and investigate how to transform the current technical tools and methodologies so as to derive computational models that are transparent and inclusive. Information processing is one of the fundamental pillars of multimedia, it does not matter whether data is processed for content, experience or systems applications, the automatic processing of information is used in every corner of our community. This is why it is crucial to start bringing the notion of fairness, accountability and transparency into ACM Multimedia.

2 TOPICS

The workshop aims to foster research around a timely and crucial topic for the present digitized society: the fairness, accountability and transparency of multimedia algorithms. The workshop has a strong scientific link with the FAT/ML workshop and the ACM FAT* conference. Differently from FAT/ML, which is anchored in machine learning, the FAT/MM workshop addresses fairness, accountability and transparency in the core of the multimedia community. We expect submissions covering any topic closely related to the multimedia community AND falling in the following categories:

Models. : Techniques and models for fairness-aware multimedia modeling, retrieval, and recommendation; Interpretable and explainable models in multimedia; Models and frameworks for conducting FAT audits of multimedia systems; Models for addressing inclusion and exclusion in multimedia.

Algorithm evaluation. Qualitative, quantitative, and experimental studies on subjective perceptions of algorithmic bias and unfairness; Experimental results of FAT audits of multimedia systems; Objective metrics for measuring unfairness and bias; Generation of human-readable explanations for multimedia models and algorithmic outputs; Metrics for measuring inclusiveness in multimedia systems.

Data collection and curation. Defining, measuring and mitigating problematic biases in multimedia datasets; Improvement of data

collection processes to be more fair, diverse, and inclusive; Data collection regarding potential unfairness in systems.

3 PROGRAM

The program will follow the next guidelines:

- Invited talk: Prof. Xirong Li
- Oral presentation #1
- Oral presentation #2
- Invited Oral Presentation: Martim Brandao
- Oral presentation #3
- Coffee break
- Invited talk: Prof. Jahna Otterbacher
- Oral presentation #4
- Oral presentation #5

There are five papers accepted at the workshop with the following titles: Learning Facial Recognition Biases through VAE Latent Representations, QoE-Fair Resource Allocation for DASH Video Delivery Systems, Social Multimedia, Diversity, and Global South Cities: A Double Blind Side, A Software Defined Network based research on Fairness in Multimedia, Toward Fairness in Face Matching Algorithms.

The first paper discusses recent concerns about biases embedded in algorithms leveraging face datasets for law enforcement applications. The objective of this work is to implement a variational autoencoder (VAE) that generates low-dimensional representations of faces and then analyze the low-dimensional representations to interpret potential learned biases. While the field of interpreting facial recognition biases is still emerging, previous work has also relied on VAEs to better understand the relationship between images and their latent representations. The authors implement a 10-layer VAE and analyze the encoding of images to a single latent feature and to ten latent features. The experimental results show that protected features such as race (e.g. skin color) play a role in latent representations. This lays the groundwork for interpreting the encodings of individual latent features to address algorithmic biases. Future work could entail finding a set of latent features that more accurately represents said protected characteristics.

The second paper deals with fairness in delivering satisfactory service in the absence of resources. Video delivering to massive audiences is required to provide the users with a satisfactory Quality of Experience (QoE) to keep high engagement and avoid service abandonment. Adaptive BitRate algorithms (ABR) running in video players are designed to dynamically change the video bitrate to provide the best possible QoE given the user device features and the end-to-end network available bandwidth. Well-designed ABR algorithms strive to improve the individual QoE obtained by each user resulting, in the optimal case, in the maximization of the sum of QoE individually perceived by users. However, when resources are scarce, maximizing the sum of the QoE might result in favoring some clients at the expense of others which instead obtain poor QoEs with the possible service abandonment. Thus, the paper argues that video service providers should directly address fairness issues when designing their delivery networks so to gracefully degrade the QoE for all users when resources are scarce. This paper addresses this open issue and shows that the Multi-Commodity Flow Problem (MCFP) optimization framework is a proper methodology to achieve a QoE-fair distribution of the resources. Obtained

results show that the proposed approach is able to achieve its goal and provide a fair level of QoE to heterogeneous clients.

The third paper deals with leveraging multi-modal data from social networks populated with content from the Global South, arguing that this should be a priority of multimedia research. The lack of research and data generates a double blind side problem that is concerning but also potentially harmful if not addressed. The problem is exemplified by studying Ma3Route, a platform that uses mobile and social media to crowdsource and broadcast transit reports in Nairobi, Kenya. Using multimedia data from Ma3Route's Twitter stream, the authors first conduct a descriptive analysis that shows it to be a highly active community generating rich traffic-related reports, and then discover latent topics that identify both regular and ephemeral thematic clusters of reports involving accidents, traffic conditions, and attitudes of citizens towards authorities. In the second place, a deep learning-based analysis of Ma3Route images is conducted to reveal the kind of visual content shared in the platform, and show limitations of using deep neural network models trained with data collected in the developed world, which do not reflect the reality and diversity of other world regions. The paper concludes by presenting a multidisciplinary research agenda for future work in this domain.

The fourth paper is about Multimedia quality of experience (QoE), and how to make sure that users in a multi-user multi-device environment can perceive fairness in terms of experience and resources allocation. The paper discusses how network designs such as software-defined networks can help maintaining application and human-level fairness in networked multimedia applications.

The last paper presented in the workshop deals also with the facial business. In particular, with automated face matching algorithms, since they are used in a wide variety of societal applications ranging from access authentication, to criminal identification, to application customization. If the algorithms work well only for certain racial or gender identities, they would adversely affect others. Recent efforts in algorithmic fairness literature (typically not focused on multimedia or computer vision tasks) have argued for designing algorithms and architectures to tackle such bias via trade-offs between accuracy and fairness. The authors show that adopting an adversarial deep learning-based approach allows for the model to maintain the accuracy at face matching while also reducing demographic disparities compared to a baseline (non-adversarial deep learning) approach at face matching. The results motivate and pave way for more accurate and fair face matching algorithms.

4 ORGANIZERS

Our organization committee is very diverse in terms of expertise, topics, geographic location and seniority.

- Dr. Xavier Alameda-Pineda (Research Scientist, Perception Group, Inria Grenoble)
- Dr. Miriam Redi (Research Scientist, Wikimedia Foundation)
- Prof. Elisa Celis (Assistant Professor, Yale University).
- Prof. Nicu Sebe (Professor and Head of Dept. of Information Engineering and Computer Science, University of Trento)
- Prof. Shih-Fu Chang (Richard Dicker Professor and senior Vice Dean of the School of Engineering and Applied Sciences at Columbia University)