

AdKDD 2021

Abraham Bagherjeiran
eBay, Inc.
San Jose, CA, USA
abagherjeiran@ebay.com

Kuang-chih Lee
Alibaba Group
Beijing, China
leekc307@gmail.com

Nemanja Djuric
Aurora Innovation
Pittsburgh, PA, USA
nemanja@temple.edu

Kun Liu
Amazon
Palo Alto, CA, USA
kunliu1@gmail.com

Suju Rajan
LinkedIn
Sunnyvale, CA, USA
srajan@linkedin.com

Mihajlo Grbovic
AirBnB
San Francisco, CA, USA
mihajlo.grbovic@airbnb.com

Vladan Radosavljevic
Spotify
New York City, NY, USA
vladan@temple.edu

ABSTRACT

The digital advertising field has always had challenging ML problems, learning from petabytes of data that is highly imbalanced, reactivity times in the milliseconds and more recently compounded with the complex user's path to purchase across devices, across platforms and even online/real-world behavior. The AdKDD workshop continues to be a forum for researchers in advertising, during and after KDD. Our website which hosts slides and abstracts receives approximately 2,000 monthly visits. In surveys during AdKDD 2019 and 2020, over 60% agreed that AdKDD is the reason they attended KDD and over 90% indicated they would attend next year. The 2021 edition is particularly timely because of ongoing developments in ad tracking. We will aim to discuss notions of privacy and tracking enforced by GDPR and through company policies. In addition, we will seek papers that discuss fairness in the context of advertising, to what extent does hyper-personalization work, and on whether the ad industry as a whole needs to think through more effective business models such as incrementality. Ad tech is in an interesting place of evolution/maturity now and we would like to use the AdKDD forum to get the researchers to think not only about the ML aspects but also spark conversations about the societal ones.

CCS CONCEPTS

• **Information systems** → **Online advertising**; • **Applied computing** → **Electronic commerce**.

KEYWORDS

Computational advertising, Ad targeting, User modeling

ACM Reference Format:

Abraham Bagherjeiran, Nemanja Djuric, Mihajlo Grbovic, Kuang-chih Lee, Kun Liu, Vladan Radosavljevic, and Suju Rajan. 2021. AdKDD 2021. In

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).

KDD '21, August 14–18, 2021, Virtual Event, Singapore.

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8332-5/21/08.

<https://doi.org/10.1145/3447548.3469486>

Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '21), August 14–18, 2021, Virtual Event, Singapore. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3447548.3469486>

1 INTRODUCTION

An average consumer spends 8+ hours a day across all devices interacting with online content almost entirely sponsored by advertisements. At over \$300B global market size in 2021 and expected to pass \$1T by 2027, online advertising has already surpassed traditional ads in global spend. Moreover, computational advertising in particular is perhaps the most visible and ubiquitous application of machine learning and one that interacts directly with consumers. When done right, ads help us enrich our lives and creep us out when done badly. Looking at the published literature over the last few years, many researchers might consider computational advertising as a mature field. Yet, the opposite is true. The field is evolving, however, from ads controlled by monolithic publishers and randomly rotating banner ads to highly personalized content experiences in news feeds on mobile devices and even on TV—all utilizing data amassed from petabytes of stored user data. Ads are far from done.

2 WORKSHOP TOPICS

Evolution of computational advertising: Online advertising has progressed beyond the notion of traditional desktop ads to ads that are native, social, mobile, and contextual. In tandem, the rise of new mechanisms, such as header bidding, complex ad exchanges, repeated auctions, ad blockers, viewability trackers, and others, challenge the traditional notions of advertising. There also continue to exist controversial issues in advertising such as privacy, security, fraud, ethics, and economic attribution. We invite papers that are focused on some of the above aspects.

Large-scale and novel ad targeting: Recent advances in real-time, big data systems, and easier accessibility to different types of data make it possible to design more personalized and efficient ad targeting systems. We invite papers that advance the state-of-the-art in related areas of ad targeting.

Deployed systems & battle scars: We particularly encourage papers that highlight experience in deploying real-time ad targeting

systems, data and audience insights, as well as position papers on the future of online advertising.

3 ML CRITEO CHALLENGE

The Online Advertising industry is seeing a major shift today in its operational constraints with a global movement towards more privacy. Popular techniques for privacy-compliant advertising such as aggregation and differential privacy mechanisms were shown to match high privacy standards but also raise concerns about the possibility to learn relevant machine learning models for ad placement.

We propose in this challenge to explore the trade-off between privacy level and prediction performance, on data donated by Criteo - an industry leader that already released several open data sets for research purposes. To anchor the competition in reality, the challenge design is inspired by (and as close as possible/convenient to) current propositions in the Privacy Sandbox discussed in the Improving Web Advertising forum at W3C.

4 PAST WORKSHOPS

There have been a total of fourteen AdKDD and TargetAd workshops to date, organized every year since 2007, which focused on highlighting state-of-the-art advances in computational advertising. All the workshops were well attended, often with standing room only, and very well received both by the academic community and the advertising industry. For the 2020 edition virtually in San Diego, USA we had 100-150 Zoom attendees, 500+ playbacks of the YouTube live stream, and 1500+ views of recorded presentations.

In 2018 and 2017 we organized workshops at KDD under the name of "AdKDD & TargetAd" while in 2019 we organized the workshop under the name of AdKDD. The reason for the name change was that the organizers of the two successful workshops (i.e., AdKDD and TargetAd) came together to organize an even stronger venue for high-quality publications. As a result in 2019 we fully merged the two workshops into one, given shared area and values of the two workshops, and named the new joint workshop simply "AdKDD" to simplify branding and marketing of the event for the time to come. This year we are continuing with the same organization and branding.

5 ORGANIZERS

Abraham Bagherjeiran, eBay Abraham is a Senior Director of Applied Research at eBay. His team drives the search ranking and monetization optimization for eBay Search businesses at the intersection of e-commerce and advertising. He received his PhD in Computer Science at the University of Houston, Texas. He has authored over 25 papers and patents in the field of computational advertising.

Nemanja Djuric, Aurora Technologies Nemanja is a Staff Autonomy Engineer and Technical Lead Manager at Aurora Technologies, while prior to his current position he worked in the same role at Uber ATG. Previously he was a Research Scientist at Yahoo Labs working on computational advertising. Dr. Djuric published more than 50 peer-reviewed publications at the leading Machine Learning, Data Mining, Robotics, Computer Vision, and Web Science conferences and journals, in addition to 4 granted patents and

10+ pending patent applications. Nemanja received his PhD degree in Computer Science from Temple University in 2013. His work was featured in Market Watch, VentureBeat, IEEE Innovation at Work, and other news outlets across the world.

Mihajlo Grbovic, Airbnb Mihajlo is a Principal Machine Learning Scientist on the Search Team at Airbnb working on Ranking and Recommendation Systems. Prior to that, he was a Senior Research Manager at Yahoo Labs working on Advertising Sciences. He has more than 10 years of technical experience in applied Machine Learning, acting as a Science Lead in a portfolio of advertising technology projects at Yahoo and Tumblr. Dr. Grbovic published more than 40 peer-reviewed publications at top Machine Learning and Web Science conferences and co-authored more than 10 pending patents. His work was featured in Wall Street Journal, Scientific American, MIT Technology Review, Popular Science, and Market Watch.

Kuang-chih Lee, Alibaba Kuang-chih is the head of Marketplace Governance in AliExpress.com. He manages all aspects of research and development for real-time personalized e-commerce marketplace. He has 30+ research papers published in top conferences (CVPR, NIPS, AAAI, CIKM, and KDD) and journals (PAMI, and CVIU), as well as 20+ patents. Prior to joining Alibaba Inc, Kuang-chih Lee was the principal scientist and research director at Yahoo Inc. Before that he held various research and development leadership positions at Turn Inc, Flashfoto, DigitalPersona, and like.com. Kuang-chih Lee received CS PhD degree in UIUC in 2005.

Kun Liu, Amazon Kun is a Senior Manager of Machine Learning and AI at Amazon. His team is building machine-learning infrastructure and predictive models to support Amazon's hyper growth Sponsored Product business. Prior to joining Amazon, Kun was a Senior Manager leading the Ads Relevance team at LinkedIn, where he worked on the full spectrum of computational advertising problems: ads CTR and conversion prediction, bidding optimization, lookalike modeling, data-driven reserve price, bid landscaping, campaign insights and optimization recommendation. Kun received his Ph.D. degree in Computer Science from University of Maryland Baltimore County in 2007.

Vladan Radosavljevic, Spotify Vladan is a Machine Learning Chapter Lead at Spotify. Vladan's teams build recommendation systems for Spotify's home page personalization. Previously, Vladan was a Head of Data Science at OLX Group where his team built solutions for two-sided marketplace platform. Before OLX, he was a Senior Scientist at Uber ATG working on systems for autonomous driving. Prior to Uber, he was a Research Scientist at Yahoo Labs where he worked on computational advertising problems. Vladan received his PhD from Temple University in Philadelphia in 2011. His work was featured in Market Watch, VentureBeat, IEEE Innovation at Work, and other news outlets across the world.

Suju Rajan, LinkedIn Suju is a Sr. Director at LinkedIn where she heads machine learning teams that build solutions for LinkedIn's Enterprise offerings. Prior to LinkedIn, she was a SVP at Criteo, where she headed the Criteo AI Labs to advance both state-of-the-art and business impact in the field of computational advertising. Before Criteo, she was the Director of Personalization Sciences at Yahoo Research where her team worked on personalized recommendations for several Yahoo products. She received her PhD from the University of Texas at Austin.