

AdKDD 2024

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

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

Kuang-Chih Lee
Walmart
Sunnyvale, CA, USA
leekc307@gmail.com

Linsey Pang
Salesforce
San Francisco, CA, USA
xpang@salesforce.com

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

Suju Rajan
Amazon
Palo Alto, CA, USA
rsuju@amazon.com

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 and 1,800 active users during the KDD 2021. 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 2024 edition is particularly timely because of the increasing application of Graph-based NN and Generative AI models in advertising. Coupled with privacy-preserving initiatives enforced by GDPR, CCPA the future of computational advertising is at an interesting crossroads. For this edition, we plan to solicit papers that span the spectrum of deep user understanding while remaining privacy-preserving. In addition, we will seek papers that discuss fairness in the context of advertising, to what extent does hyper-personalization work, and whether the ad industry as a whole needs to think through more effective business models such as incrementality. We have hosted several academic and industry luminaries as keynote speakers and have found our invited speaker series hosting expert practitioners to be an audience favorite. We will continue fielding a diverse set of keynote speakers and invited talks for this edition as well. As with past editions, we hope to motivate researchers in this space to think not only about the ML aspects but also to spark conversations about the societal impact of online advertising.

CCS CONCEPTS

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

KEYWORDS

Computational advertising, Ad targeting, User modeling

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 '24, August 25–29, 2024, Barcelona, Spain
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0490-1/24/08.
<https://doi.org/10.1145/3637528.3671476>

ACM Reference Format:

Abraham Bagherjeiran, Nemanja Djuric, Kuang-Chih Lee, Linsey Pang, Vladan Radosavljevic, and Suju Rajan. 2024. AdKDD 2024. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24)*, August 25–29, 2024, Barcelona, Spain. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3637528.3671476>

1 INTRODUCTION

In 2023, the average worldwide internet user spent on average 6.5 hours daily across all devices interacting with online content almost entirely sponsored by advertisements. At almost \$700B global market size in 2024, and expected to pass \$830B by 2026, digital advertising has already surpassed traditional ads in global spend and continues to grow despite economic headwinds. Digital advertising and in particular computational advertising is perhaps the most visible and ubiquitous application of machine learning and one that interacts directly with consumers. When done right, ads connect us to opportunities to enrich our lives and creep us out when done badly. Recently at the forefront of political battles between governments, large multinational corporations, and consumers, digital advertising remains a dynamic industry and research area.

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. Computational advertising is evolving from simple rule-based ads controlled by monolithic publishers and randomly rotating banner ads to highly personalized content experiences within native, video, and display formats on mobile devices, connected TV, and audio – all utilizing data amassed from petabytes of stored user data collected increasingly through inferred identity providers. The increasing use of sequence models and the rise of Generative AI is spawning new directions such as the auto-generation of ad creatives or pay-for-use shopping assistants. Ads are far from done.

The AdKDD workshops held in conjunction with KDD conferences in the past years (2017-2023) continue to generate interest from academia and industry as one of the top venues specifically for advertising research. We believe this is a unique forum for folks interested in aspects of digital advertising to get together, exchange notes and get a pulse for the state of the art, especially in the industry. We have also had several inquiries on whether we would repeat this workshop for 2024 and are confident we will have an engaged and productive workshop session. Although we have considered creating a separate conference dedicated to advertising, we believe KDD is still the best venue for our attendees.

2 WORKSHOP TOPICS

The workshop focuses on three main aspects of computational advertising, discussed below.

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, as well as new technologies such as generative AI, challenge the traditional notions of advertising. Controversial issues in advertising continue to exist 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 PAST WORKSHOPS

There have been a total of seventeen 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 2023 edition held in Long Beach, CA we had 100+ attendees, as well as 100 YouTube views on average per each recorded presentation.

4 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 Innovation Nemanja is a Principal Technical Lead Manager at Aurora Innovation, 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, Computer Vision, Robotics, and Web Science conferences and journals, in addition to 10 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.

Kuang-chih Lee, Walmart Kuang-chih is a Senior Director of AdTech Data Science at Walmart. Prior to Walmart, he worked at Alibaba as the head of Marketplace Governance in two different organizations, Alimama.com and AliExpress.com. He managed all aspects of research and development for a 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 a CS PhD degree in UIUC in 2005.

Linsey Pang, Salesforce Linsey is Principal Applied Machine learning scientist at Salesforce, working on personalization, recommendation related projects. Prior to this, she was Principal Data Scientist and Data science Manager II at Walmart Lab. She was leading the Pricing, One demand data science projects in the Merchant Technology organization. Prior to joining Walmart Lab, she was working as an applied scientist at eBay Inc. She has co-authored top conference papers including KDD, NIPS, ICDM, etc. She got her Ph.D degree in Data Mining from University of Sydney in 2015 and her research interests include data mining, machine learning, high performance computing, etc.

Vladan Radosavljevic, Spotify Vladan is a Machine Learning Chapter Lead at Spotify. Vladan is leading teams that build foundational recommendation models at Spotify. Previously, Vladan was a Head of Data Science at OLX Group where his team built solutions for two-sided marketplace platforms. 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, Amazon Suju Rajan is a Director of Applied Science for Sponsored Products at Amazon. Prior to Amazon, she was a Sr. Director at LinkedIn where she headed machine learning teams that built 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.