

# REVEAL 2020: Bandit and Reinforcement Learning from User Interactions

Thorsten Joachims  
Cornell University  
tj@cs.cornell.edu

Maria Dimakopoulou  
Netflix  
mdimakopoulou@netflix.com

Yves Raimond  
Netflix  
yraimond@netflix.com

Flavian Vasile  
Criteo  
f.vasile@criteo.com

Olivier Koch  
Criteo  
o.koch@criteo.com

Adith Swaminathan  
Microsoft Research  
adswamin@microsoft.com

## ABSTRACT

The REVEAL workshop<sup>1</sup> focuses on framing the recommendation problem as a one of making personalized interventions, e.g. deciding to recommend a particular item to a particular user. Moreover, these interventions sometimes depend on each other, where a stream of interactions occurs between the user and the system, and where each decision to recommend something will have an impact on future steps and long-term rewards. This framing creates a number of challenges we will discuss at the workshop. How can recommender systems be evaluated offline in such a context? How can we learn recommendation policies that are aware of these delayed consequences and outcomes?

## CCS CONCEPTS

- Information systems → Recommender systems.

## KEYWORDS

recommender systems; reinforcement learning; off-policy; offline evaluation; causal inference; multi-armed bandits

## ACM Reference Format:

Thorsten Joachims, Yves Raimond, Olivier Koch, Maria Dimakopoulou, Flavian Vasile, and Adith Swaminathan. 2020. REVEAL 2020: Bandit and Reinforcement Learning from User Interactions. In *Fourteenth ACM Conference on Recommender Systems (RecSys '20), September 22–26, 2020, Virtual Event, Brazil*. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3383313.3411536>

## 1 INTRODUCTION

State-of-the-art recommender systems are notoriously hard to design and improve upon, due to their interactive and dynamic nature, since they involve a multi-step decision-making process, where a stream of interactions occurs between the user and the system. Leveraging reward signals from these interactions and creating a scalable and performant recommendation inference model is a key challenge. Traditionally, to make the problem tractable, the

<sup>1</sup><https://sites.google.com/view/reveal2020/>

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).

RecSys '20, September 22–26, 2020, Virtual Event, Brazil

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-7583-2/20/09.

<https://doi.org/10.1145/3383313.3411536>

interactions are often viewed as independent, but in order to improve recommender systems further, the models will need to take into account the delayed effects of each recommendation and start reasoning/planning for longer-term user satisfaction. To this end, our workshop invites contributions that enable recommender systems to adapt effectively to diverse forms of user feedback and to optimize the quality of each user's long-term experience.

Due to their interactive nature, recommender systems are also notoriously hard to evaluate. When evaluating their systems, practitioners often observe significant differences between a new algorithm's offline and online results, and therefore tend to mostly rely on online methods, such as A/B testing. This is unfortunate, since online evaluation is not always possible and often expensive. Offline evaluation, on the other hand, provides a scalable way of comparing recommender systems and enables the participation of academic research in an industry-relevant problem.

In the past, recommender systems have been evaluated using proxy offline metrics coming from supervised methods, such as regression metrics (mean squared error, log likelihood), classification metrics (area under precision/recall curve) or ranking metrics (precision@k, normalized discounted cumulative gain). Recent research on recommender systems makes the link with counterfactual inference for offline A/B testing that reuses logged interaction data, and as well as the use of simulators that entirely avoid the use of potentially privacy-sensitive user data.

In this context, we believe it is timely to organize a workshop that re-visits the problem of designing and evaluating recommender systems and makes sure the community, spanning academic and industrial backgrounds, is working on the right problem: find for each user, the most impactful recommendation.

This workshop is the follow-up of REVEAL'18 and REVEAL'19, both of which had strong participation at RecSys 2018 in Vancouver and RecSys 2019 in Copenhagen. In this 2020 instance, we will keep pushing forward the boundary of research on the following topics:

- Reinforcement learning and bandits for recommendation
- Robust estimators and counterfactual evaluation
- Using simulation for recommender systems evaluation
- Open datasets and new offline metrics

The benefits of this workshop will be:

- To bridge the gap between academia and industry through new datasets and methods
- To increase the productivity of all researchers and practitioners in their development of recommender systems

## 2 PROGRAM

The workshop will consist of a fast-paced full-day workshop with a focus on high quality paper presentations and four invited keynotes. The workshop will conclude with a fireside chat with the speakers with the aim to provide recommendations for future research.

The workshop will feature a mix of contributed talks, contributed posters, and invited talks by leading researchers from diverse backgrounds working in these areas. We will also have a specific segment of the schedule reserved for the presentation of open problems, and will have plenty of time for discussions where we will explicitly look to spark off collaborations amongst the attendees.

The morning and afternoon breaks will include poster sessions (if this is compatible with the RecSys organization). The workshop will be covered by a dedicated twitter channel. The organizers will put a strong focus on interactivity throughout the day.

## 3 ORGANIZING COMMITTEE

- Thorsten Joachims, Cornell University
- Maria Dimakopoulou, Netflix
- Adith Swaminathan, Microsoft Research
- Yves Raimond, Netflix
- Olivier Koch, Criteo
- Flavian Vasile, Criteo

## 4 PAPER SUBMISSION AND REVIEW PROCESS

This workshop will accept paper submission which will be reviewed through a single-blind process. Accepted papers will be presented as posters or talks. The Program Committee will be the same as the Organization Committee. Accepted papers will be made publicly

available as non-archival reports, allowing future submissions to archival conferences and journals.

## 5 PRIOR WORKSHOPS (MOST RECENT FIRST)

- Workshop on Reinforcement and Robust Estimators for Recommendation, RecSys 2019 Workshop, September 2019 <https://recsys.acm.org/recsys19/reveal/>
- Offline Evaluation of Recommender Systems, RecSys 2018 Workshop, October 2018 <https://recsys.acm.org/recsys18/reveal/>
- Machine Learning for Causal Inference, Counterfactual Prediction, and Autonomous Action ICML 2018 Workshop, July 2018 <https://sites.google.com/site/faim18wscausalml/>
- From 'What If?' To 'What Next?': Causal Inference and Machine Learning for Intelligent Decision Making, NIPS 2017 Workshop, December 2017 <https://sites.google.com/view/causalnips2017>
- What If? Inference and Learning of Hypothetical and Counterfactual Interventions in Complex Systems, NIPS 2016 Workshop, December 2016 <https://sites.google.com/site/whatif2016nips/>
- Controlled Experimentation in Recommendations, Ranking & Response Prediction, in conjunction with RecSys 2014, Silicon Valley <https://recsys.acm.org/recsys14/recsysab/>
- ACM RecSys Workshop on Recommender Systems Evaluation: Dimensions and Design, REDD 2014, Silicon Valley <http://ir.ii.uam.es/redd2014/>
- RecSys'13: Workshop on Reproducibility and Replication in Recommender Systems Evaluation, in conjunction with RecSys 2013, Hong Kong <http://rep sys.project.cwi.nl/>
- BARS 2013: Workshop on Benchmarking Adaptive Retrieval and Recommender Systems, in conjunction with ACM SIGIR, 2013 <http://bars-workshop.org/>
- RUE 2012: Workshop on Recommendation Utility Evaluation: Beyond RMSE, in conjunction with RecSys 2012, Dublin <http://ir.ii.uam.es/rue2012/>