Yuta Saito

Address	Ithaca, NY, USA
Email	ys552@cornell.edu
Website	https://usait0.com/en/
Google Scholar	pw4hwS8AAAAJ
GitHub	https://github.com/usaito
Last Update	January 23, 2024

Research Interests

My research lies at the intersection of statistical machine learning and causal inference called *counterfactual learning*. I am interested in the counterfactual nature of logged bandit feedback obtained from interactive systems, and ways of using biased real-world data to achieve safe and efficient automated decision making in the wild. I am also working on fairness in recommender and ranking systems to develop a framework to improve long-term objectives of the platform by guaranteeing fair opportunities among item providers.

Education

2021 - **Cornell University**
Ph.D. Student in Computer Science
Advisor: Thorsten Joachims
Thesis Committee: Jon Kleinberg, Karthik Sridharan
Research Field: Counterfactual Evaluation, Reinforcement Learning, Fairness in RecSys
Expected Graduation: 05/2026

2016 - 2021 **Tokyo Institute of Technology**
B.Eng. in Industrial Engineering and Economics

Publications

International Conference Proceedings (refereed)

According to Google Scholar: h-index is 13 and total citation count is 720+.

1. **Yuta Saito**, Himan Abdollahpouri, Jesse Anderton, Ben Carterette, and Mounia Lalmas. Long-term Off-Policy Evaluation and Learning. In *Proceedings of the ACM Web Conference **(TheWebConf)***, 2024 (to appear).

2. Riku Togashi, Kenshi Abe, and **Yuta Saito**. Scalable and Provably Fair Exposure Control for Large-Scale Recommender Systems. In *Proceedings of the ACM Web Conference **(TheWebConf)***, 2024 (to appear).

3. Haruka Kiyohara, Masahiro Nomura, and **Yuta Saito**. Off-Policy Evaluation of Slate Bandit Policies via Optimizing Abstraction. In *Proceedings of the ACM Web Conference **(TheWebConf)***, 2024 (to appear).

4. Haruka Kiyohara, Ren Kishimoto, Kosuke Kawakami, Ken Kobayashi, Kazuhide Nakata, and **Yuta Saito**. Towards Assessing and Benchmarking Risk-Return Tradeoff of Off-Policy Evaluation. In *Proceedings of the Twelfth International Conference on Learning Representations **(ICLR)***, 2024 (to appear).

5. Haruka Kiyohara, Masatoshi Uehara, Yusuke Narita, Nobuyuki Shimizu, Yasuo Yamamoto, **Yuta Saito**. Off-Policy Evaluation of Ranking Policies under Diverse User Behavior. In *Proceedings of the 29th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining **(KDD)***, 2023.

6. **Yuta Saito**, Qingyang Ren, and Thorsten Joachims. Off-Policy Evaluation for Large Action Spaces via Conjunct Effect Modeling. In *Proceedings of 40th International Conference on Machine Learning **(ICML)***, 2023.

7. Takuma Udagawa, Haruka Kiyohara, Yusuke Narita, **Yuta Saito**, and Kei Tateno. Policy-Adaptive Estimator Selection for Off-Policy Evaluation. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence **(AAAI)***, 2023 (**Oral**).

8. **Yuta Saito** and Thorsten Joachims. Fair Ranking as Fair Division: Impact-Based Individual Fairness in Ranking. In *Proceedings of the 28th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining **(KDD)***, 2022.

9. **Yuta Saito** and Thorsten Joachims. Off-Policy Evaluation for Large Action Spaces via Embeddings. In *Proceedings of 39th International Conference on Machine Learning **(ICML)***, 2022.

10. **Yuta Saito** and Masahiro Nomura. Towards Resolving Propensity Contradiction in Offline Recommender Learning. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence **(IJCAI)***, 2022 (**Long Talk**).

11. Haruka Kiyohara, **Yuta Saito**, Tatsuya Matsuhiro, Yusuke Narita, Nobuyuki Shimizu, and Yasuo Yamamoto. Doubly Robust Off-Policy Evaluation for Ranking Policies under the Cascade Behavior Model. In *Proceedings of International Conference on Web Search and Data Mining **(WSDM)***, 2022 (**Best Paper Runner-Up Award**).

12. Daisuke Moriwaki, Yuta Hayakawa, Isshu Munemasa, **Yuta Saito**, and Akira Matsui. A Real-World Implementation of Unbiased Lift-based Bidding System. In *Proceedings of the 2021 IEEE International Conference on Big Data **(BigData)***, 2021.

13. **Yuta Saito**, Shunsuke Aihara, Megumi Matsutani, and Yusuke Narita. Open Bandit Dataset and Pipeline: Towards Realistic and Reproducible Off-Policy Evaluation. In *Proceedings of the Neural Information Processing Systems **(NeurIPS)** Track on Datasets and Benchmarks*, 2021.

14. Masahiro Nomura* and **Yuta Saito*** (*equal contribution). Efficient Hyperparameter Optimization under Multi-Source Covariate Shift. In *Proceedings of the 30th ACM International Conference on Information and Knowledge Management **(CIKM)***, 2021.

15. **Yuta Saito***, Takuma Udagawa*, Haruka Kiyohara, Kazuki Mogi, Yusuke Narita, and Kei Tateno (*equal contribution). Evaluating the Robustness of Off-Policy Evaluation. In *Proceedings of the 15th ACM Conference on Recommender Systems **(RecSys)***, 2021.

16. Nathan Kallus, **Yuta Saito**, and Masatoshi Uehara. Optimal Off-Policy Evaluation from Multiple Logging Policies. In *Proceedings of 38th International Conference on Machine Learning **(ICML)***, 2021.

17. **Yuta Saito**. Doubly Robust Estimator for Ranking Metrics with Post-Click Conversions. In *Proceedings of the 14th ACM Conference on Recommender Systems **(RecSys)***, 2020.

18. **Yuta Saito**. Unbiased Pairwise Learning from Biased Implicit Feedback. In *Proceedings of 6th ACM SIGIR International Conference on the Theory of Information Retrieval **(ICTIR)***, 2020.

19. **Yuta Saito** and Shota Yasui. Counterfactual Cross-Validation: Stable Model Selection Procedure for Causal Inference Models. In *Proceedings of 37th International Conference on Machine Learning **(ICML)***, 2020.

20. **Yuta Saito**. Asymmetric Tri-training for Debiasing Missing-Not-At-Random Explicit Feedback. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval **(SIGIR)***, 2020.

21. **Yuta Saito**, Gota Morishita, and Shota Yasui. Dual Learning Algorithm for Delayed Conversions. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval **(SIGIR)***, 2020 (short paper).

22. **Yuta Saito**, Hayato Sakata, and Kazuhide Nakata. Cost-Effective and Stable Policy Optimization Algorithm for Uplift Modeling with Multiple Treatments. In *Proceedings of the 2020 SIAM International Conference on Data Mining **(SDM)***, 2020.

23. **Yuta Saito**, Suguru Yaginuma, Yuta Nishino, Hayato Sakata, and Kazuhide Nakata. Unbiased Recommender Learning from Missing-Not-At-Random Implicit Feedback. In *Proceedings of the 13th International Conference on Web Search and Data Mining **(WSDM)***, 2020.

24. **Yuta Saito**, Hayato Sakata, and Kazuhide Nakata. Doubly Robust Prediction and Evaluation Methods Improve Uplift Modeling for Observational Data. In *Proceedings of the 2019 SIAM International Conference on Data Mining **(SDM)***, 2019.

International Conference Workshop Papers (refereed)

1. Haruka Kiyohara, Kosuke Kawakami, and **Yuta Saito**. Accelerating Offline Reinforcement Learning Application in Real-Time Bidding and Recommendation: Potential Use of Simulation. *RecSys 2021 Workshop on Simulation Methods for Recommender Systems (**SimuRec**)*, 2021.

2. **Yuta Saito**, Shunsuke Aihara, Megumi Matsutani, and Yusuke Narita. A Large-scale Open Dataset for Bandit Algorithms. *RecSys 2020 Workshop on Bandit and Reinforcement Learning from User Interactions **(REVEAL)***, 2020 (**Oral Presentation**).

3. **Yuta Saito**, Takuma Udagawa, and Kei Tateno. Data-Driven Off-Policy Estimator Selection: An Application in User Marketing on An Online Content Delivery Service. *RecSys 2020 Workshop on Bandit and Reinforcement Learning from User Interactions **(REVEAL)***, 2020.

4. **Yuta Saito**, Shunsuke Aihara, Megumi Matsutani, and Yusuke Narita. A Large-scale Open Dataset for Bandit Algorithms. *ICML 2020 Workshop on Real World Experiment Design and Active Learning **(RealML)***, 2020.

5. Daisuke Moriwaki, Yuta Hayakawa, Isshu Munemasa, **Yuta Saito**, and Akira Matsui. Unbiased Lift-based Bidding System. In *Proceedings of the 2020 AdKDD&TargetAd Workshop, held in conjunction with the 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining **(AdKDD)***, 2020.

Internships

Summer 2023	**Spotify Research** Developed a statistical method to estimate and optimize the long-term value of policies from logged bandit data and put the method and results into a research paper (under submission)

Conference Tutorials

2022-09	**CONSEQUENCES Workshop (at RecSys2022)** (topic: off-policy evaluation in large action spaces)
2022-08	**KDD2022** (topic: off-policy evaluation and learning)
2021-09	**RecSys2021** (topic: off-policy evaluation and learning)

Invited Talks

2022-12	**Booking.com** (topic: off-policy evaluation in large action spaces)
2022-11	**Amazon** (topic: off-policy evaluation in large action spaces)
2022-11	**Instacart** (topic: off-policy evaluation in large action spaces)
2022-09	**Twitter** (topic: off-policy evaluation in large action spaces)
2022-06	**University of Amsterdam** (topic: impact-based individual item fairness in ranking)
2021-06	**Tokyo Institute of Technology** (topic: python package and open dataset for OPE research)
2021-03	**Royal Melbourne Institute of Technology** (topic: python package and open dataset for OPE research)
2021-03	**California Institute of Technology** (topic: python package and open dataset for OPE research)
2021-02	**Cornell University** (topic: python package and open dataset for OPE research)

Awards

2023	RecSys 2023 Outstanding Reviewer Award
2022	Forbes Japan 30 Under 30 2022
2022	WSDM 2022 Best Paper Runner-Up Award
2021	NeurIPS 2021 Outstanding Reviewer Award

Scholarships

2021 - 2023 **Funai Overseas Scholarship**
Doctoral research fellowship by the Funai Foundation (a private foundation in Japan).
Granted two full years of tuition plus a monthly stipend of $3,000 for living expenses.

Involved Research Projects

- **Open Bandit Project** (https://github.com/st-tech/zr-obp)
  Open Bandit Project is an open-source research project that aims to enable realistic and reproducible experiments on bandit algorithms and their off-policy evaluation. The project consists of a large-scale real-world dataset called *Open Bandit Dataset* and Python software called *Open Bandit Pipeline*. The project was awarded **The Prime Minister's Award for Open Innovation by the Japanese Government**.

Teaching Assistant

Fall 2023 **Advanced Machine Learning (CS6780, Cornell)**

Professional Service

Conference Program Committee

- **ICML** 2021, 2022, 2023
- **NeurIPS** 2021, 2022, 2023
- **KDD** 2022, 2023
- **AAAI** 2023, 2024
- **AISTATS** 2021
- **ICLR** 2022
- **WSDM** 2022, 2023, 2024
- **WWW** 2024
- **SIGIR** 2023
- **RecSys** 2023
- **ECIR** 2024

Workshop Organizer

- RecSys 2023 Workshop on Causality, Counterfactuals and Sequential Decision-Making for Recommender Systems
- RecSys 2022 Workshop on Causality, Counterfactuals, Sequential Decision-Making & Reinforcement Learning

Workshop Program Committee

- NeurIPS 2022 Workshop on Offline Reinforcement Learning
- NeurIPS 2022 Workshop on Reinforcement Learning for Real Life
- NeurIPS 2021 Workshop on Offline Reinforcement Learning
- NeurIPS 2021 Workshop on Causal Inference Challenges in Sequential Decision Making

Journal Reviewer

- IEEE Transactions on Signal Processing
- IEEE Transactions on Knowledge and Data Engineering (TKDE)
- ACM Transactions on Intelligent Systems and Technology (TIST)
- ACM Transactions on Information Systems (TOIS)
- ACM Transactions on Recommender Systems (TORS)
- Transactions on Machine Learning Research (TMLR)

Languages

Japanese (native), English (TOEFL iBT: 105)