

Masatoshi Uehara.

mu223@cornell.edu <https://masatoshiuehara.com/>

Education

2013-2017 Department of applied mathematics and information science, The University of Tokyo

2017- 2020 Department of Statistics, PHD program, Harvard University

2020- Department of Computer Science, PhD program, Cornell University

Research Interests

Reinforcement learning, Causal inference, Online learning

Publications (* indicates I am the corresponding author or co-first author)

Off-Policy Evaluation and Learning for External Validity under a Covariate Shift

Uehara, M. Kato, T. Yasui, Shota.

In Thirty-fourth Conference on Neural Information Processing Systems (NeurIPS 2020)

Doubly Robust Off-Policy Value and Gradient Estimation for Deterministic Policies

Kallus, N. Uehara, M*.

In Thirty-fourth Conference on Neural Information Processing Systems (NeurIPS 2020)

Statistically Efficient Off-Policy Policy Gradients

Kallus, N. Uehara, M*.

In Proceedings of the 33rd International Conference on International Conference (ICML 2020)

Minimax Weight and Q-Function Learning for Off-Policy Evaluation

Uehara, M. , Huang, J. Jiang, N.

In Proceedings of the 33rd International Conference on International Conference (ICML 2020)

Double Reinforcement Learning for Efficient Off-Policy Evaluation in Markov Decision Processes

Kallus, N. Uehara, M*.

In Proceedings of the 33rd International Conference on International Conference (ICML 2020)

Journal of Machine Learning Research(Longer Version)

Imputation Estimators for Unnormalized Models with Missing Data

Uehara, M. Matsuda, T. Kim, J.W.

In 23nd International Conference on Artificial Intelligence and Statistics (AISTATS 2020)

Unified Estimation Framework for Unnormalized Models with Statistical Efficiency

Uehara, M. Kanamori, T. Takenouchi, T. Matsuda, T.

In 23nd International Conference on Artificial Intelligence and Statistics (AISTATS 2020)

Intrinsically Efficient, Stable, and Bounded Off-Policy Evaluation for Reinforcement Learning

Kallus, N. Uehara, M*.

In Thirty-third Conference on Neural Information Processing Systems (NeurIPS 2019)

Preprints

Efficient Evaluation of Natural Stochastic Policies in Offline Reinforcement Learning

Kallus, N. Uehara, M*. <https://arxiv.org/abs/2006.03886>

Localized Debiased Machine Learning: Efficient Estimation of Quantile Treatment Effects, Conditional Value at Risk, and Beyond

Kallus, N. Mao, X. Uehara, M. <https://arxiv.org/abs/1912.12945>

Efficiently Breaking the Curse of Horizon: Double Reinforcement Learning in Infinite-Horizon Processes

Major revision in *Operations research*

Kallus, N. Uehara, M*. <https://arxiv.org/abs/1909.05850.pdf>

Information Criteria for Non-normalized Models

Matsuda, T. Uehara, M. Hyvarinen, A. <https://arxiv.org/pdf/1905.05976.pdf>

Semiparametric Response Model with Nonignorable Nonresponse

Uehara, M. Kim, J.W. <https://arxiv.org/pdf/1810.12519.pdf>

Analysis of Noise Contrastive Estimation from the Perspective of Asymptotic Variance

Uehara, M. Matsuda, T. Komaki, F. <https://arxiv.org/abs/1808.07983.pdf>

Generative Adversarial Nets from a Density Ratio Estimation Perspective,

Uehara, M. Sato, I. Suzuki, M. Nakayama, K. Matsuo, Y. <https://arxiv.org/abs/1610.02920.pdf>

Coding skills

Python, C++, R