CALL FOR PAPERS

ACM Transactions on Probabilistic Machine Learning

Information For Authors

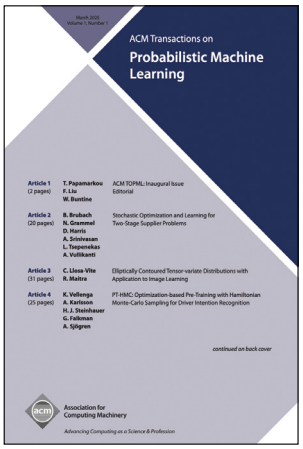

ACM Transactions on Probabilistic Machine Learning (TOPML) is a new Open Access publication from ACM. The journal focuses on probabilistic methods that learn from data to improve performance on decision-making or prediction tasks under uncertainty. Optimization, decision-theoretic or information-theoretic methods are within the remit if they are underpinned by a probabilistic structure. Probabilistic methods may be harnessed to address questions related to statistical inference, uncertainty quantification, predictive calibration, data generation and sampling, causal inference, stability, and scalability. Examples of approaches relevant to the scope include Bayesian modelling and inference, variational inference, Gaussian processes, Monte Carlo sampling, Stein-based methods, and ensemble modelling. Examples of models for which probabilistic approaches are sought include neural networks, kernel-based models, graph-based models, reinforcement learning models, recommender systems, and statistical and stochastic models. Ethical considerations of probabilistic machine learning, such as data privacy and algorithmic fairness, should be addressed in papers where there is a direct ethical connection or context for the work being described.

The journal welcomes theoretical, methodological, and applied contributions. Purely theoretical contributions are of interest if they introduce novel methodology. Methodological and applied contributions are of interest provided that proposed probabilistic approaches are motivated and empirically corroborated by non-trivial examples or applications. Multidisciplinary approaches with a probabilistic structure are within the scope.

Feature Articles:

"ACM TOPML: Inaugural Issue Editorial"
By Co-EICs , Theodore Papamarkou, Fang Liu, and Wray Buntine

"Stochastic Optimization and Learning for Two-Stage Supplier Problems"
By Brian Brubach, Nathaniel Grammel, David G. Harris, Aravind Srinivasan, Leonidas Tsepenekas, and Anil Vullikanti

"Elliptically-Contoured Tensor-variate Distributions with Application to Image Learning"
By Carlos Llosa-Vite and Ranjan Maitra

"PT-HMC: Optimization-based Pre-Training with Hamiltonian Monte-Carlo Sampling for Driver Intention Recognition"
By Koen Vellenga, Alexander Karlsson, H. Joe Steinhauer, Göran Falkman, and Anders Sjögren

"On the Importance of the Execution Schedule for Bayesian Inference"
By Patrick W.A. Wijnings, Martin Roa-Villescas, Sander Stuijk, Bert de Vries and Henk Corporaal

"DRD-GAN: A Novel Distributed Conditional Wasserstein Deep Convolutional Relativistic Discriminator GAN with Improved Convergence"
By Arunava Roy and Dipankar Dasgupta

"Hierarchical Bayesian Data Selection"
By Simon L. Cotter

Association for Computing Machinery

Open Access

Visit topml.acm.org for further information or to submit your manuscript.