

OpML '20: 2020 USENIX Conference on Operational Machine Learning

July 30, 2020, Santa Clara, CA, USA

Sponsored by USENIX, the Advanced Computing Systems Association



Important Dates

Paper Presentations

- Submissions due: Tuesday, February 25, 2020, 5:00 pm PST
- Notification to authors: Friday, March 20, 2020
- Final papers due: Thursday, May 14, 2020

Talk (Non-Paper) Presentations

- Submissions due: Tuesday, February 25, 2020, 5:00 pm PST
- Notification to presenters: Friday, March 20, 2020

Conference Organizers

Program Co-Chairs

Nisha Talagala, *Pyxeda AI*
Joel Young, *LinkedIn*

Program Committee

Fei Chen, *LinkedIn*
Sindhu Ghanta, *ParallelM*
Prassana Padmanaban, *Netflix*
Neoklis Polyzotis, *Google*
Suresh Raman, *Intuit*
Bharath Ramsundar, *Stealth Mode*
Marius Seritan, *Argo AI*
Faisal Siddiqi, *Netflix*
Swaminathan Sundararaman, *Pyxeda*
Eno Thereska, *Amazon*
Boris Tvaroska, *Lenovo*
Todd Underwood, *Google*
Sandeep Uttamchandani, *Intuit*
Manasi Vartak, *Verta AI*
Jesse Ward, *LinkedIn*

Steering Committee

Nitin Agrawal, *Samsung*
Eli Collins, *Accel Partners/Samba Nova*
Casey Henderson, *USENIX Association*
Robert Ober, *Nvidia*
Bharath Ramsundar, *Stealth Mode*
Jairam Ranganathan, *Uber*
D. Sculley, *Google*
Tal Shaked, *Google*
Swaminathan Sundararaman
Nisha Talagala, *Pyxeda AI*
Sandeep Uttamchandani, *Intuit*
Joel Young, *LinkedIn*

Overview

Machine learning (ML), with its variants (deep learning (DL), reinforcement learning, etc.) is impacting every commercial industry. The first Conference on Operational Machine Learning, OpML 2019, was an energetic gathering of practitioners, industry experts and researchers who came together to discuss the critical challenges associated with bringing Machine Learning to Production. The 2020 USENIX Conference on Operational Machine Learning (OpML '20), continues this focus on operational machine learning and its variants, and will focus on the full lifecycle of deploying and managing ML into production. OpML '20 will bring together machine learning and systems researchers and practitioners such as data scientists, data engineers, infrastructure engineers, reliability engineers, sysadmins, and DevOps specialists to develop and bring to practice impactful research advances and cutting-edge solutions to the pervasive challenges of ML production lifecycle management.

OpML is complementary to existing conferences that focus on algorithmic advances and systems design for improved ML methods, performance, and scale. OpML focuses on deployment, automation, orchestration, monitoring, diagnostics, compliance, governance, production scale training and re-training,



and the challenges of safely operating and optimizing production systems running ML/DL/Advanced algorithms on live data. Production ML lifecycle is a necessity for wide-scale adoption and deployment of ML/DL across industries and for businesses to benefit from the core ML algorithms and research advances.

Participation Overview

OpML provides a forum for both researchers and industry practitioners to exchange and debate innovations, learnings, experiences, and problems in Operational Machine Learning.

The Conference seeks submissions in the form of short papers and presentations.

- Accepted presentations will be included in the Presentations Track.
- Accepted short papers will be published in the conference proceedings and will also be presented in the Presentations Track.

Submissions will be judged on their applicability to the problems of Operational ML, originality, technical merit, topical relevance, and the likelihood of leading to insightful discussions that will influence practices of ML and its variants in production and benefit attendees.

Suggested Topics

- Use of state of the art or new/emerging analytic engines (e.g., extensions to Spark, TensorFlow, PyTorch, Ray) in production
- Challenges of production training and re-training including scale, transitioning from experimental models to production models, etc.
- Use of new/emerging ML acceleration hardware and associated production challenges
- Systems for orchestrating, diagnosing, monitoring, and managing ML in production
- Applying existing DevOps and SDLC tools and practices to the ML operational lifecycle
- Diagnostics of ML algorithms operating on live data (e.g., drift detection)
- New model introduction into production (e.g., staging, A/B test)
- Governance of ML models and deployment processes, model risk management, adaptations for new regulations.
- Systems, Tools, Approaches for Model/Dataset version control, and lineage
- Data governance approaches as they relate to production ML
- New visualizations to accommodate diverse users (engineers, data scientists, ops, etc.)
- Experiences with advanced techniques (such as Reinforcement Learning) in production
- Technical advances for addressing ML regulatory requirements (e.g., GDPR)
- Privacy and security challenges in production ML
- Bringing research on Explainable ML into production use
- Experiences with bringing ML techniques to production and scaling ML in production
- Industry-specific best practices for ML production (e.g., edge computing/IoT, healthcare)
- Containerized ML workflows (e.g., Kubernetes for scale in ML)
- Use of Cloud APIs and Cloud based ML services in Production, experience and best practices.

What You Can Expect to Get from OpML

- **Machine Learning Researchers:** Understand the practical challenges, learn about advances in ML explainability, diagnostics, etc. Gain exposure to how hardware advances are used in production.
- **Systems Researchers:** Share innovative techniques for managing ML/DL in production, and learn from the experiences of real-life production challenges.
- **Data Scientists/Data Engineers:** Gain exposure to best practices in production ML in the contexts of compliance, risk management, and legal requirements.
- **System/IT, Reliability Engineers and DevOps:** Learn best practices for deploying ML algorithms and artifacts. Explore the latest software and hardware advances available for production ML lifecycle management.

Conference Details

OpML '20 will be a one-day conference. At least one author of each accepted paper or talk must attend the conference to present. Presentation details and guidelines will be communicated to the authors of the accepted paper or talk.

Submission Instructions

- Paper submissions must be no longer than two pages excluding references and should be submitted electronically via the papers submission form linked from the Call for Participation page. Accepted papers will be published by USENIX and included in the Conference proceedings. Accepted papers will also be presented at the Conference in the Presentations Track (20-minute presentation and 5 minutes for questions).
- Talk submissions must include an outline of the presentation's content as well as identified take-aways for the attendees, and must be submitted via the talks submission form linked from the Call for Participation page. Accepted presentations will be part of the Presentations Track (25 minutes presentation and 5 minutes for questions).

Submitting Papers

Paper submissions should be PDF documents that are viewable by standard tools. Submissions must follow the USENIX formatting guidelines: 10-point type on 12-point (single-spaced) leading, with the text block being no more than 7" wide by 9" deep. See the detailed formatting requirements at www.usenix.org/conference/opml20/requirements-authors.

Submissions may not be under consideration for any other venue. Note that arxiv is not considered a different venue. Submissions of work previously included in arxiv are encouraged if the content is appropriate for the topics listed above and if the presentation will advance the field of Operational ML and generate discussion at the conference. Questions? Contact your program co-chairs, opml20chairs@usenix.org, or the USENIX office, submissionspolicy@usenix.org.

The review process is not blind. The names and affiliations of the authors should be included on the first page. The names of the reviewers, however, will remain anonymous. Papers accompanied by nondisclosure agreement forms will not be considered. Accepted submissions will be treated as confidential prior to publication on the OpML '20 website; rejected submissions will be permanently treated as confidential. All accepted papers will be available online to registered attendees before the conference. If your accepted paper should not be published prior to the event, please notify production@usenix.org. The papers will be available online to everyone on the day of the conference.