The Emerging Practice of Operational ML
USENIX OpML Conference

NISHA TALAGALA AND JOEL YOUNG

Nisha Talagala is the CEO/Founder of Pyxeda. Previously, Nisha co-founded ParallelM and pioneered MLOps, acquired by DataRobot. Nisha also drove the USENIX Conference on Operational Machine Learning, the first conference on production ML. Nisha has 20 years of experience in software, distributed systems, technical strategy, and product leadership. Her prior roles include Fellow at SanDisk and Fellow/Lead Architect at Fusion-io, at Intel as Server Flash/Persistent Memory Lead and at Gear6 as CTO. Nisha earned her PhD at UC Berkeley, holds 69 patents, is a frequent speaker at industry/academic events, and is a contributing writer to several online publications.
nishatalagala@gmail.com

Joel leads LinkedIn's Assessments team and co-led LinkedIn's Productive Machine Learning effort. Prior to LinkedIn, he was an Assistant Professor at the Naval Postgraduate School and also at the Air Force Institute of Technology. In addition to academia and over 20 years in the Air Force, he was an Eagle Scout, a hospital janitor, a carnival operator, and a ditch digger. Joel earned his PhD in computer science from Brown University in 2005. His research areas include digital forensics, data mining, and machine learning. He has also published in web search, computational biology, and temporal modeling. joel.d.young@gmail.com

At OpML '19, the first USENIX Conference on Operational Machine Learning, we learned many useful lessons. Moving forward, we expect the same will hold true for the second conference, coming this May 2020. In this article, we discuss some of the pragmatic practices that came out of the first conference.

Machine learning (ML) and its variants such as deep learning (DL) and reinforcement learning are starting to impact every commercial industry. In recognition of the growing need to drive ML into production, and the unique technical challenges therein, USENIX launched OpML in 2019 (Conference on Operational Machine Learning). The first conference dedicated to the operational aspects of machine learning and its variants, OpML is focused on the full life cycle of deploying and managing ML into production [1]. OpML '19 was an energetic gathering of experts, practitioners, and researchers who came together for one day in Santa Clara, CA, to talk about the problems, practices, new tools, and cutting-edge research on production machine learning in industries ranging from finance, insurance, health care, security, web scale, manufacturing, and others [2].

While there were many great presentations, papers, panels, and posters (too many to talk about individually—check out all the details here [2]), there were several emergent trends and themes (previously described here [11]). We expect each of these will expand and become even more prominent over the next several years as more organizations push ML into production and adopt machine learning ops practices to scale ML in production.

Agile Methodologies Meet Machine Learning

Many practitioners emphasized the importance of iteration and continuous improvement to achieving production ML success. Much like software, machine learning improves through iteration and regular production releases. Those who have ML running at scale make it a point to recommend that projects should start with either no Machine Learning or simple Machine Learning to establish a baseline. As one practitioner put it, you don't want to spend a year investing in a complex deep learning solution, only to find out after deployment that a simpler non-ML method can outperform it [3].

Bringing agility to ML also requires that the infrastructure be optimized to support agile rollouts (and rollbacks!). This means that successful production ML infrastructure includes automated deployment, modularity, use of microservices, and also avoiding fine-grained optimization early on [3].

ML-Specific Production Diagnostics because ML Bugs Differ from Software Bugs

Various presentations provided memorable examples of how ML errors not only bypass conventional production checks but can actually look like better production performance. For example—an ML model that fails and generates a default output can actually cause a performance boost!

Detecting ML bugs in production requires specialized techniques like Model Performance Predictors [4], comparisons with non-ML baselines, visual debugging tools [5], and metric-driven design of the operational ML infrastructure. Facebook, Uber, and other organizations experienced with large-scale production machine learning ops, emphasized the importance of ML-specific production metrics that range from health checks to ML-specific (such as GPU) resource utilization metrics [6].

Rich Open Source Ecosystem for All Aspects of Machine Learning Ops

The rich open source ecosystem for model development (with TensorFlow, Scikit-learn, Spark, PyTorch, R, etc.) is well known. OpML showcased how the open source ecosystem for machine learning ops is growing rapidly, with powerful publicly available tooling used by large and small companies alike. Examples include Apache Atlas for governance and compliance, Kubeflow for machine learning ops on Kubernetes, MLflow for life-cycle management, and Tensorflow tracing for monitoring. Classic enterprise vendors are starting to integrate these open source packages to fill solutions (see, e.g., Cisco's support of Kubeflow). Furthermore, web-scale companies are open sourcing the core infrastructure that drives their production ML, such as the ML orchestration tool TonY from LinkedIn [7].

As these tools become more prominent, full end-to-end use cases are also being documented by practitioners, creating design patterns that can be used as best practices by others.

Cloud-Based Services and SaaS Make Production ML Easier

For a team trying to deploy ML in production for the first few times, the process can be daunting, even with open source tools available for each stage of the process. The cloud offers an alternative because the resource management aspects (such as machine provisioning, auto-scaling, elasticity, etc.) are handled by the cloud back end. When accelerators (GPUs, TPUs, etc.) are used, production resource management is challenging. Using cloud services is a way to get started by leveraging the investments made by cloud providers to optimize accelerator usage. Find out more in Ananthanarayanan et al.'s slides at [8].

Cloud deployment can also create a ramp-up path for an IT organization to try ML deployment without a large in-house infrastructure roll out. As discussed by Wenzel and Maurice [9], even on-premise enterprise deployments are moving to self-service production ML models similar to cloud services, enabling the IT organization to serve the production ML needs of multiple teams and business units.

Leverage Expertise from At-Scale Web-Based ML Operations for Enterprise

At-scale experts like LinkedIn, Facebook, Google, Airbnb, Uber, and others, who were the first ML adopters, had to build from scratch all of the infrastructure and practices needed to extract monetary value out of ML. Now these experts are sharing not only their code but also their experiences and hard-won knowledge, which can be adopted for the benefits of enterprise. As the Experts Panel at OpML pointed out [3], the best practices that these organizations follow for ML infrastructure (from team composition and reliability engineering to resource management) contain powerful insights that enterprises can benefit from as they seek to expand their production ML footprint. Experiences from scale ML deployments at Microsoft and others [2] can show enterprises how to deliver performant machine learning into their business applications.

Other end-to-end experiences from at-scale companies [2] showed how business metrics can be translated into ML solutions and the consequent ML solution iteratively improved for business benefit. Finally, organizations facing the unique challenges that edge deployment places on machine learning ops can benefit from learning of scale deployments already in place.

Moving Forward: OpML '20

The goal of the OpML conference is to help develop robust practices for scaling the management of models (i.e., artifacts of learning from big data) throughout their life cycle. Through such practices, we can help organizations transition from manual hand-holding to automated management of ML models in production—the ML version of the move in server operations from "pets to cattle" [12]. Production ML is still a nascent field, and OpML '19 showcased some emerging best practices as described above. New challenges emerge every day, however, such as regulatory concerns brought on by GDPR and CCPA, migrating from legacy infrastructure to cloud, and security attacks on ML systems, just to name a few. OpML '20, to be held in May in Santa Clara, CA, USA, will continue the example set by OpML '19 and be a venue for experts, practitioners, and researchers to discuss, debate, and share the state of the art in Operational ML.

Summary

A great op-ed piece by Michael Jordan in Medium—"Artificial Intelligence: The Revolution Hasn't Happened Yet"—highlighted the importance of an engineering practice for AI [9]. OpML '19, the first Machine Learning Ops conference, illustrated how the ML and AI industry is maturing in this direction, with more and more organizations either struggling with the operational and life-cycle management aspects of machine learning in production, or pushing to scale ML operations and develop operational best practices. This is great news for the AI industry since it is a step further towards generating real ROI from AI investments. OpML '20, following last year's success, will continue to support and bring together the Operational ML community and help realize the long-awaited potential of AI business value. Please join us at OpML '20! [13].

References

[1] OpML '19: https://www.usenix.org/conference/opml19.

[2] OpML '19 Program: https://www.usenix.org/conference/opml19/program.

[3] J. Young, "How the Experts Do It: Production ML at Scale," June 7, 2019: https://www.linkedin.com/pulse/how-experts-do-production-ml-scale-joel-young/.

[4] S. Ghanta, S. Subramanian, L. Khermosh, H. Shah, Y. Goldberg, S. Sundararaman, D. Roselli, N. Talagala, "MPP: Model Performance Predictor" (slides): https://www.usenix.org/sites/default/files/conference/protected-files/opml19_slides_ghantapdf.pdf.

[5] L. Li, Y. Bai, Y. Wang, "Manifold: Model-Agnostic Visual Debugging Tool for ML" : https://www.usenix.org/sites/default/files/conference/protected-files/opml19_slides_li-lezhi.pdf.

[6] Y. Yan, Z. Zweiger, "The Power of Metrics: How to Monitor and Improve ML Efficiency": https://www.usenix.org/sites/default/files/conference/protected-files/opml19_slides_yan.pdf.

[7] J. Hung, "TonY: An Orchestrator for Distributed Machine Learning Jobs" (slides): https://www.usenix.org/sites/default/files/conference/protected-files/opml19_slides_hsu.pdf.

[8] R. Ananthanarayanan, P. Brandt, M. Joshi, M. Sathiamoorthy, "Opportunities and Challenges of Machine Learning Accelerators in Production" (slides): https://www.usenix.org/sites/default/files/conference/protected-files/opml19_slides_ananthanarayanan.pdf.

[9] T. Wenzel and V. Maurice, "Machine Learning Models as a Service" (slides): https://www.usenix.org/sites/default/files/conference/protected-files/opml19_slides_wenzel.pdf.

[10] https://medium.com/@mijordan3/artificial-intelligence-the-revolution-hasnt-happened-yet-5e1d5812e1e7.

[11] https://www.forbes.com/sites/cognitiveworld/2019/07/01/five-trends-in-machine-learning-ops-takeaways-from-the-first-operational-ml-conference/364421147132.

[12] Impact of Data Regulations and Bias on Operational ML panel: https://www.usenix.org/conference/opml19/presentation/panel-uttamchandani.

[13] OpML '20: https://www.usenix.org/conference/opml20.