The Art of Building ML as a Product: Key Learnings and Strategies

Shashank Shashikant Rao
Intuit

ABSTRACT

Despite the promise of Artificial Intelligence (AI) and Machine Learning (ML), many organizations' efforts in the field are falling short with most only scratching the surface of what's possible [2]. The pace of innovation at many organizations has fallen short of expectations. This is in part due to a lack of understanding of how to build ML as a product, which requires a different approach than optimizing for the last mile of performance that ML practitioners are used to in research settings.

To build ML as a product, technical leadership must wear multiple hats across the product lifecycle, including Product Management, Design, Engineering, Content, and User Experience. This requires a shift in thinking for research teams, who must focus on providing value to the business rather than just optimizing for model performance.

Oftentimes, ML based mission teams are tasked with launching models and experiences where a baseline model or data driven approach does not exist. There is an art to organizing teams for moving with speed continuing to deliver value while continuing to optimize for both business and model performance.

The talk will share key learnings from the author's work at Intuit in progressively building towards a larger platform goal of being an AI Driven Expert Platform [3]. It will discuss data-driven approaches to machine learning management such as identifying metrics for measuring impact in nascent projects such as customer engagement or revenue generation. It also involves lessons learnt in crafting A/B tests and designing launch experiments for new features, which can help to validate the effectiveness of ML models and algorithms before they are deployed in production.

ACM Reference Format:
Shashank Shashikant Rao. 2023. The Art of Building ML as a Product: Key Learnings and Strategies. In Proceedings of The Workshop on Applied Machine Learning Management at The 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD WAMLM '23). ACM, New York, NY, USA, 1 page. https://doi.org/XXXXXXX. XXXXXXX.

DISCUSSION POINTS

This talk discusses some of the principles, key learnings and example case studies of building an ML driven product. As examples, the talk will cover relevant areas of interest to an ML practitioner about the processes of managing applied research from the industry.

- Booking.com's lessons learned in experiment design sophistication and evaluation of models [1]
- Learnings from Spotify engineering culture on the concept of decoupled releases and small frequent releases [4]
- Uber's approach on organizing Machine Learning Teams For Model Developer Velocity [5]

KDD WAMLM '23, Aug. 6–10, 2023, Long Beach, CA.

https://doi.org/XXXXXXX. XXXXXXX

- Principles of Incrementality and data driven sufficiency to add value to business at Doordash [6]

The case studies will weave through the following discussion points:

Stakeholder Management

- Influencing business direction with machine learning.
- Building and designing products to be AI First.
- Stewarding business decisions with ML at the core.
- Architectural oversight from ML teams.

Measuring business impact

- Standardized metrics to evaluate the performance and impact of ML models.
- ML in industries with long product development cycles.
- Difficulty in quantifying the benefits of ML in terms of cost savings, revenue generation, or other business metrics.
- Complexity of ML models and their interactions with other systems, making it difficult to isolate their impact.

Importance of A/B Tests

- Reducing the risk of deploying ML models that may not perform as expected.
- Make data-driven decisions about which ML models to deploy and how to optimize them.
- Resolve cross team alignment and pick the "winner" using data.

Responsible AI

- Ensuring the product is compliant with relevant regulations and ethical considerations.
- Moving fast without breaking things.
- Limited understanding of the impact of ML on customer experience and satisfaction.

RELEVANCE TO WORKSHOP

This talk is highly relevant to academic and industry ML practitioners in understanding the nuances of building ML as a product. The goal is to share knowledge and ideas between different organizations on management of ML projects and measuring business impact.

MAIN PRESENTER BIO

Shashank is a Senior Data Science Manager at Intuit, part of the AI, Data and Analytics group. As part of Intuit AI, Shashank is currently focused on demand forecasting, capacity planning and optimization for tax expert hiring, training, scheduling and operations. Shashank holds a Master's degree from Columbia University, New York specializing in Data Science. At Intuit Shashank has worked on deploying AI solutions in the Quickbooks suite of products to forecast business cashflow and to identify cash flow insights and recommendations. Previously at PayPal and Citrix, Shashank has 7+ years of experience in the field leading high performing data scientists and machine learning engineers in mission based teams. Shashank has received several internal awards with 8 granted patents in the field.

COMPANY PORTRAIT

Intuit helps consumers and small businesses prosper by delivering financial management and compliance products and services. We also provide specialized tax products to accounting professionals, who are key partners that help us serve small business customers. Our global financial technology platform, which includes TurboTax, Credit Karma, QuickBooks, and Mailchimp, is designed to help consumers and small businesses manage their finances, save money, pay off debt and do their taxes. For those customers who run small businesses, we are also focused on helping them

find and keep customers, get paid faster, pay their employees, manage and get access to capital, and ensure that their books are done right. ProSeries and Lacerte are our leading tax preparation offerings for professional accountants. Incorporated in 1984 and headquartered in Mountain View, California, we sell our products and services primarily in the United States.

REFERENCES

[1] Lucas Bernardi, Themistoklis Mavridis, and Pablo Estevez. 2019. 150 Successful Machine Learning Models: 6 Lessons Learned at Booking.com. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '19). Association for Computing Machinery, New York, NY, USA, 1743–1751. https://doi.org/10.1145/3292500.3330744

[2] Building the AI-Powered Organization. *Harvard Business Review*. Retrieved June 2, 2023 from https://hbr.org/2019/07/building-the-ai-powered-organization

[3] Ashok N. Srivastava. Paving a path to virtual financial advice at scale. *Intuit Blog*. Retrieved June 2, 2023 from https://www.intuit.com/blog/innovative-thinking/taxpayer-empowerment/paving-a-path-to-virtual-financial-advice-at-scale/

[4] Spotify engineering culture (part 1) - Spotify Engineering : Spotify . Retrieved June 2, 2023 from https://engineering.atspotify.com/2014/03/spotify-engineering-culture-part-1/

[5] Scaling Michelangelo. *Uber Blog*. Retrieved June 2, 2023 from https://www.uber.com/blog/michelangelo-machine-learning-platform/

[6] Alok Gupta. Organizing Machine Learning. *Doordash Engineering Blog*. Retrieved June 5, 2023 from https://doordash.engineering/2020/02/12/organizing-machine-learning-every-flavor-welcome/