

Towards ML Engineering with TensorFlow Extended (TFX)

Konstantinos Katsiapis & Kevin Haas
Google
katsiapis@google.com & khaas@google.com

ABSTRACT

The discipline of Software Engineering has evolved over the past 5+ decades to good levels of maturity. This maturity is in fact both a blessing and a necessity, since the modern world largely depends on it.

At the same time, the popularity of Machine Learning (ML) has been steadily increasing over the past 2+ decades, and over the last decade ML is being increasingly used for both experimentation and production workloads. It is no longer uncommon for ML to power widely used applications and products that are integral parts of our life. Much like what was the case for Software Engineering, the proliferation of use of ML technology necessitates the evolution of the ML discipline from “Coding” to “Engineering”.

Gus Katsiapis offers a view from the trenches of using and building end-to-end ML platforms, and shares collective knowledge and experience, gathered over more than a decade of applied ML at Google. We hope this helps pave the way towards a world of ML Engineering.

Kevin Haas offers an overview of TensorFlow Extended (TFX), the end-to-end machine learning platform for TensorFlow that powers products across all of Alphabet (and beyond). TFX helps effectively manage the end-to-end training and production workflow including model management, versioning, and serving, thereby helping one realize aspects of ML Engineering.

BIOGRAPHY

Gus: Konstantinos (Gus) is the über tech lead of TensorFlow Extended (TFX), an end-to-end machine learning platform based on TensorFlow (tensorflow.org/tfx). Before that he worked on Sibyl, a massive scale machine learning system (precursor to TensorFlow) widely used at Google. Prior to being a builder of

machine learning infrastructure he was an avid user of it, while leading the Mobile Display Ads Quality team at Google. Prior to Google, Gus gathered knowledge and experience at Amazon, Calian, Ontario Ministry of Finance, Independent Electricity System Operator, and Computron. Gus earned a master’s degree in computer science with a specialization in artificial intelligence from Stanford University and before that a bachelor’s degree in mathematics, majoring in computer science and minoring in economics, from the University of Waterloo.



**Konstantinos (Gus)
Katsiapis**



Kevin Haas

Kevin: Kevin Haas is a senior engineering manager at Google Research, driving the open source adoption of Tensorflow Extended (<https://github.com/tensorflow/tfx>), one of Google’s production ML platforms. Kevin previously served as an engineering leader for multiple machine learning and infrastructure efforts in Google Cloud, Research, and Infrastructure teams. Prior to Google, Kevin led knowledge and search infrastructure efforts in multiple Internet and software companies, including IBM, Microsoft, and Yahoo!. Kevin received his MS from Stanford University in Computer Science in dual specializations of systems.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

KDD '19, August 4–8, 2019, Anchorage, AK, USA.

© 2019 Copyright is held by the owner/author(s).

ACM ISBN 978-1-4503-6201-6/19/08.

DOI: <https://doi.org/10.1145/3292500.3340408>