


Machine Learning in the Cloud



What is Machine Learning?

Machine Learning (ML) is the utilization of Artificial Intelligence (AI) to analyze data with the goal of discovering pertinent patterns. This data analysis is automated and works on a model-based system, simulating decisions without human intervention. While humans are quite capable of doing this, the raw

amount of data and required iterations makes it far more efficient to use electronic compute. Further, computers can discover patterns that may not be intuitive, while humans can put restraints, or guard rails, on the models to ensure practical results. As computing power – and the need for deeper data analysis – grows, with the improvement of AI being inevitable, ML has already proven useful for dozens of top cloud-based companies but has a brighter future ahead.

How is Machine Learning Utilized?

Almost every major company has its own ML implementation. These companies include Intel, Netflix, Intuit, DoorDash, Uber, Spotify, Etsy, PayPal, Pinterest, and more. All of them use, at least partially, in-house tools designed to improve their ML workflow. The idea is to create, train, and store models which are later utilized for production services, such as predictive feedback. These companies also utilize open source tools through GitHub, being either Python- or Java Virtual Machine (JVM-based). This larger workflow, which has many components, requires low latency and consistent response so that data scientists and engineers can utilize the tools on-the-fly.

These users require an Application Programming Interface (API) and in many cases a Graphical User Interface (GUI) as well. The ability to track model inputs and outputs over time is critical. Models are carefully organized to stay within a set of parameters, being monitored for drift and performance metrics. The ultimate goal is to have compute “learn” useful patterns from the data so that models can be formulated and tested for use in the company’s products. These solutions in turn can be served over the cloud, which often requires low latency response due to the sheer number of queries per second.

Using the companies listed above as a working set, there are at least three primary uses for ML: fraud prevention, recommendations, and logistics. For example, PayPal uses ML to track suspicious online behavior while Intuit uses it to help people with their finances. Netflix, Etsy, Pinterest, and Spotify use it to help make better recommendations for their users, often finding overlapping interests between user subsets that appear unintuitive. DoorDash and Uber utilize ML to improve logistics which includes, for example, meal delivery time predictions. Various other companies use it to improve search, autocomplete, and rankings, internally and externally.

Storage for Machine Learning

As already mentioned, the production and implementation cycles of ML require low latencies, which makes Solid State Drives (SSD) the clear choice over Hard Drives (HDD). It's also important to have consistency of response so that performance issues with modeling can be properly discovered and diagnosed; typical consumer SSDs may be less effective. Also, because every company utilizes in-house designs it is crucial that the storage solution has flexibility – this means custom and adaptive firmware, as SSSTC provides. The SSDs must also be tested for reliability as model storage downtime, for example, could slow the implementation of new features or products. The bottom-line is still Total Cost of Ownership (TCO), but because there are so many different levels of ML implementation it is important to have both SATA and NVMe™ options; companies requiring caching for a larger number of requests per second are better off with the latter, but smaller companies may need less.

Summary

As cloud utilization spreads and data build-up continues, through the pandemic and into the Information Age, the ability to use compute resources for artificial intelligence becomes more attractive. Machine learning, in particular, allows companies to make better sense of their data, which in turn improves their products. The ML workflow needs consistent, flexible, scalable, and responsive storage – SSDs, and particularly NVMe™ SSDs, are the way forward. In-house ML designs need custom storage solutions, as offered by Solid State Storage Technology Corporation, so that they can provide their users with the very best experience possible. Make sure your engineers and data scientists have the tools they need to make the most of your data with our solutions.

Our SSD Solutions



PCIe™ - Our ED1

Series is a powerful, high performance SSD made for edge storage applications. It comes in M.2 and U.2 form factors.



SATA - Our ER2 SSD

Series delivers affordability and performance with superior random read/write speeds of up to 90,000/45,000 IOPS. It comes in M.2 and 2.5" form factors.

Please contact our [Solid State Storage Technology Corp. expert](#) for more information.

*Specifications and features are subject to change without prior notice. Images are samples only, not actual products.

Request Full Specs Sheets



ABOUT US

A subsidiary of KIOXIA Corporation, **Solid State Storage Technology Corporation** is a global leader in the design, development, and manufacturing of digital storage solutions. We offer a comprehensive lineup of high-performance customizable SSDs for the Enterprise, Industrial, and Business Client markets. With various form factors and interfaces, our SSD solutions help businesses simplify their storage infrastructures accelerating variable workloads, improving efficiency, and reducing total cost of ownership.

© 2021 Solid State Storage Technology Corporation. All rights reserved.

Learn more at www.ssstc.com