Machine Learning Model Monitor

Provide essential information to aid decision-makings

Problem

Once a machine learning model is released into the wild, it works with new data. The model performance will inevitably decline over time as data deviates from what the model has seen in training.

Machine learning engineers lack a leading signal to help determine if/when to retrain their models and get the most up-to-date performance statistics in the dev/test/ production environment.

Solution

Users can set up model monitoring jobs at deployment time via config files and commands, which invokes services that calculate model quality metrics and drifts.

Then, users can view all metrics in one customizable dashboard. Also, we provide tools to compare the performance of different model versions.

The detailed features include:

- Model detail: model performance statistics
- Model comparison: compare multiple models' performance on one page
- Alert configuration: customize when & how to be notified of model failure
- Report generation: download and share statistics

Key Features

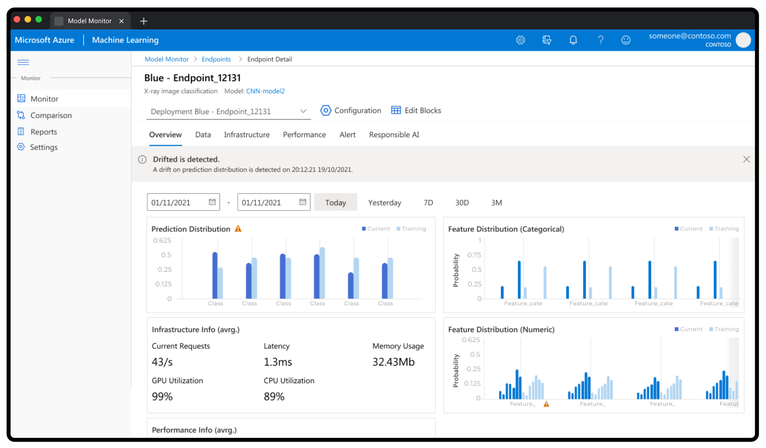

Users can view the endpoint list initially, then select a specific deployment to look at its detail. Users could also edit columns and blocks to customize the dashboard.

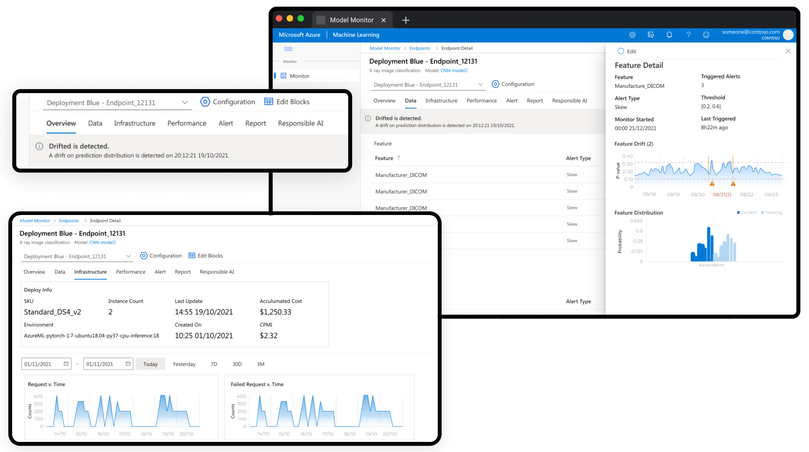

In each deployment, sections of Overview, Data, Infrastructure, Performance, Alert and Report would be provided, along with several Azure built-in features like Responsible AI. Users can see drifts, metrics on those pages.

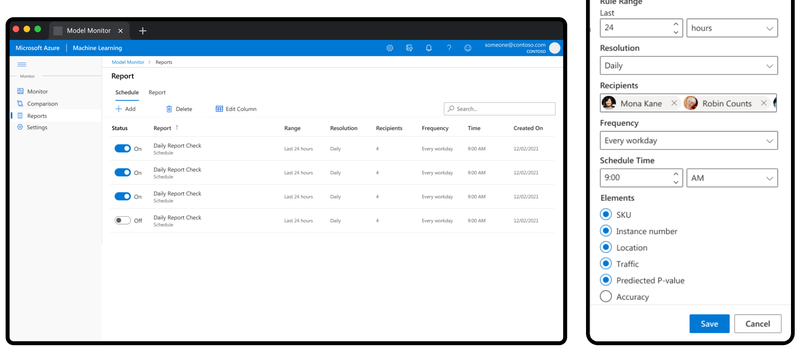

Users could generate reports from selected deployments, and even schedule automatic report generation.

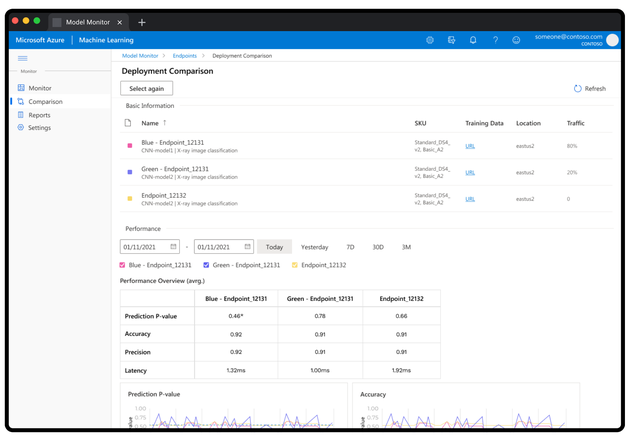

Users are allowed to select different deployments from different endpoints to compare. This feature could aid their decision-making.

Design Process

We followed a user-centered and customer-centered design process. We have two types of key stakeholders, the ML engineers as direct users and the manager in tech companies as the customers.

We tested the design prototype by usability tests and heuristic evaluations with Azure Machine Learning Studio engineers. We also collected advice and needs from customers like a health care team. Based on those feedbacks, we iterated the prototype more than three times.

Software Architecture

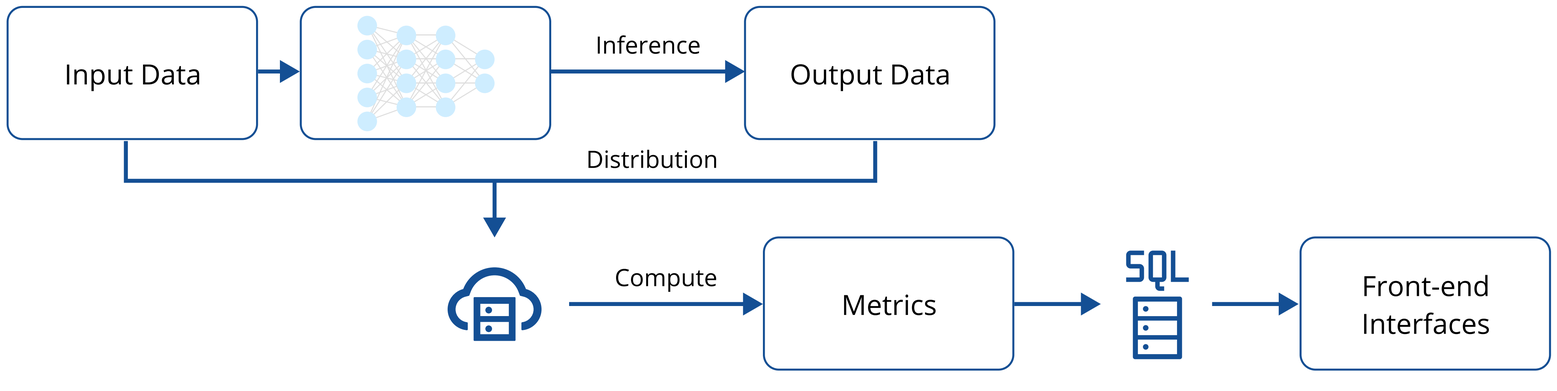

GLOBAL INNOVATION EXCHANGE

Microsoft Azure

Jerry Wang	MSTI	Dec. 2021
Qintian Li	GIX Dual Degree	July 2022
Yanzhang Li	GIX Dual Degree	July 2022
Ziqi Wang	GIX Dual Degree	July 2022