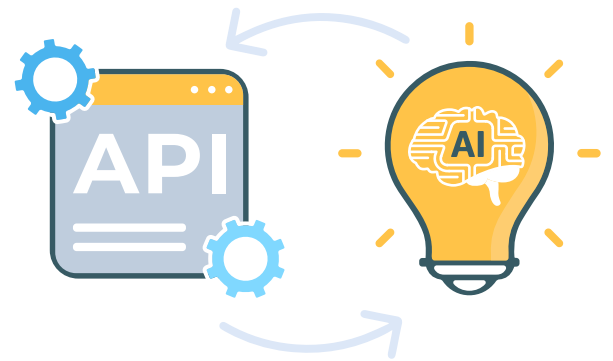# CASE STUDY

How an IT startup facilitated and enhanced ML model management and experiment tracking

EXECUTIVE SUMMARY

A Machine Learning company (The Company) with a team of 83 engineers improved ML model training performance while reducing cloud costs by 27% within two months of usage by leveraging Hystax OptScale, an MLOps and FinOps open source platform with advanced MLOps capabilities.

THE GOAL

The Company aimed to gain complete visibility into the ML model training process and output metrics, therefore, improving the efficiency of their Machine Learning operations, with a specific focus on simplifying model training and hyperparameter tuning processes.
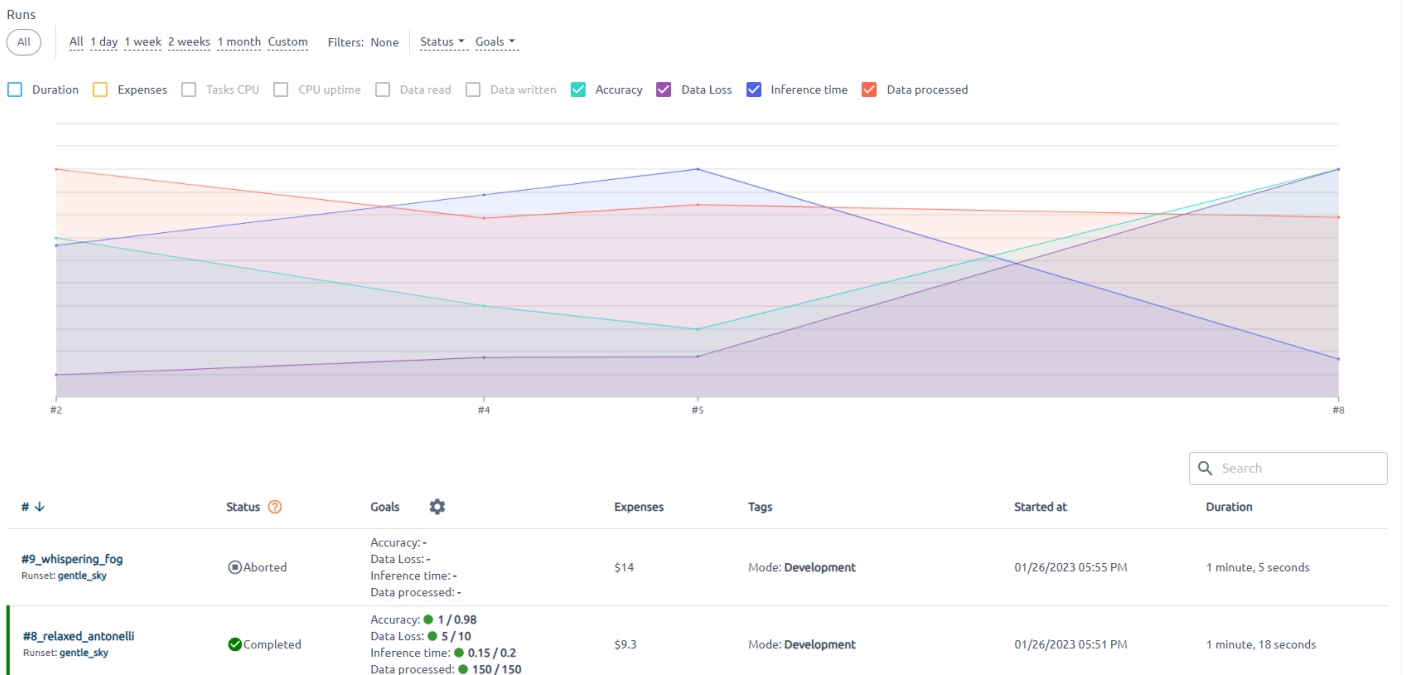
MAIN CHALLENGES

The team needed help comprehending their progress on the ML model training due to the absence of a shared dashboard that displays model training results. Also, they grappled with constantly escalating AWS cloud costs because of the intensive computing power required for ML model training. Additionally, the team needed help managing, monitoring, and optimizing their cloud resources due to the lack of detailed insights into individual ML training performance metrics and overall ML/AI operations. Moreover, hyperparameter tuning could have been streamlined for greater efficiency, and improper budget management often led to exceeded cloud spending limits.
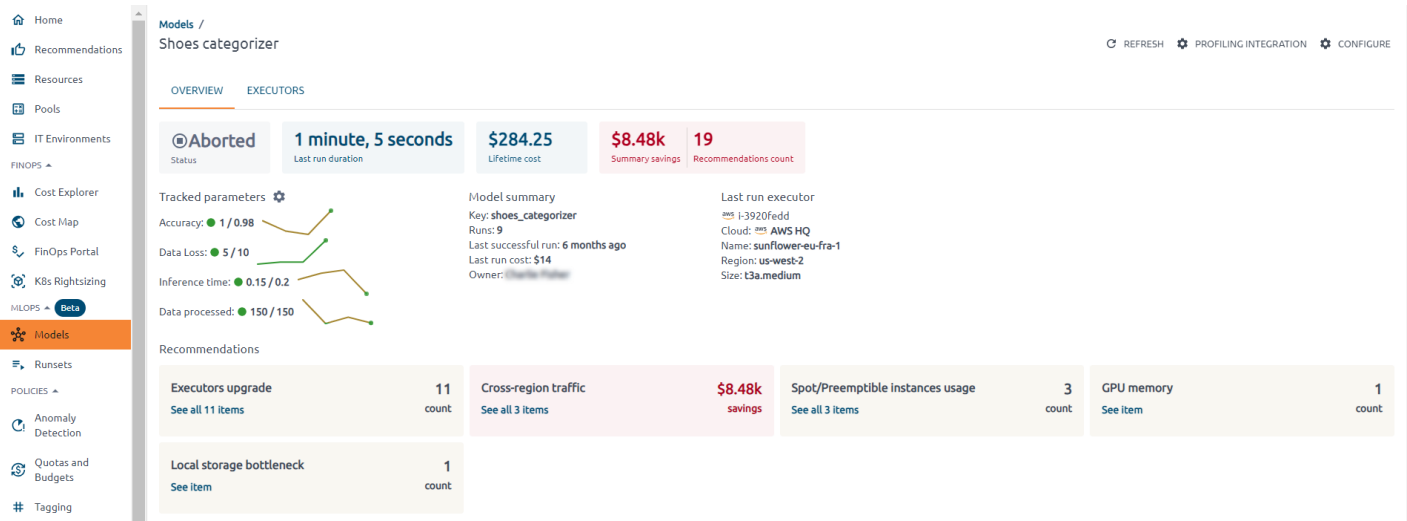
THE SOLUTION

The Company adopted the OptScale SaaS version, taking advantage of its features targeted at cost management (FinOps) and MLOps. OptScale's ability to provide detailed cost information for each cloud resource and its cost optimization recommendation engine helped The Company efficiently manage and optimize its costs.

1 **ML model training instrumentation.** This feature of OptScale facilitated efficient ML model training and utilized resource tracking in the cloud.

2 **Model dashboard/leaderboard.** The OptScale's dashboard allowed a comprehensive view of various training metrics for each ML model. This option helped the team gain insights into model performance, make informed decisions, and adjust operations as needed.

3 **Model training performance insights.** The developers gained valuable insights from the performance metrics gathered for every stage of each ML model training session and implemented code improvements to reduce training time.

4 **Runsets for hyperparameter tuning.** Leveraging this feature, the team set the framework and templates for running hyperparameter tuning sessions using Spot Instances. This approach also enhanced the efficiency of hyperparameters tuning and helped control costs by setting maximum budgets and durations for model training tasks.

5 **Recommendations for optimal cloud usage.** OptScale provided valuable recommendations for the cloud capacity team used for ML training.

6 **Resource allocation.** The OptScale detailed insights into resource utilization enabled the team to monitor and adjust resource allocation effectively for ML training sessions, leading to substantial cost savings and improved ML operations.

Runs

All | 1 day 1 week 2 weeks 1 month Custom | Filters: None | Status ▾ Goals ▾

☐ Duration ☑ Expenses ☐ Tasks CPU ☐ CPU uptime ☐ Data read ☐ Data written ☑ Accuracy ☑ Data Loss ☑ Inference time ☑ Data processed

# ↓	Status ⓘ	Goals ⚙	Expenses	Tags	Started at	Duration
#9_whispering_fog Runset: **gentle_sky**	◉ Aborted	Accuracy: - Data Loss: - Inference time: - Data processed: -	$14	Mode: **Development**	01/26/2023 05:55 PM	1 minute, 5 seconds
#8_relaxed_antonelli Runset: **gentle_sky**	✓ Completed	Accuracy: ● 1 / 0.98 Data Loss: ● 5 / 10 Inference time: ● 0.15 / 0.2 Data processed: ● 150 / 150	$9.3	Mode: **Development**	01/26/2023 05:51 PM	1 minute, 18 seconds

hx hystax

Models /
Shoes categorizer

REFRESH   PROFILING INTEGRATION   CONFIGURE

OVERVIEW   EXECUTORS

⊘ Aborted	1 minute, 5 seconds	$284.25	$8.48k	19
Status	Last run duration	Lifetime cost	Summary savings	Recommendations count

Tracked parameters ⚙

Accuracy: ● 1 / 0.98

Data Loss: ● 5 / 10

Inference time: ● 0.15 / 0.2

Data processed: ● 150 / 150

Model summary
Key: shoes_categorizer
Runs: 9
Last successful run: 6 months ago
Last run cost: $14
Owner:

Last run executor
aws i-3920fedd
Cloud: aws AWS HQ
Name: sunflower-eu-fra-1
Region: us-west-2
Size: t3a.medium

Recommendations

Executors upgrade	11
See all 11 items	count

Cross-region traffic	$8.48k
See all 3 items	savings

Spot/Preemptible instances usage	3
See all 3 items	count

GPU memory	1
See item	count

Local storage bottleneck	1
See item	count

RESULTS

With OptScale, The Company significantly boosted its MLOps efficiency, notably in ML model training, performance and experiment tracking, and hyperparameter tuning.
In addition, they achieved a 27% reduction in cloud costs in the first quarter of usage. This enhancement empowered the team to concentrate more on innovation and delivering high-quality ML solutions for their customers, resulting in increased productivity and client satisfaction.

ABOUT HYSTAX

Hystax develops OptScale, an MLOps & FinOps open source platform that optimizes performance and IT infrastructure cost by analyzing cloud usage, profiling and instrumentation of applications, ML/AI tasks, and cloud PaaS services, and delivering tangible optimization recommendations. The tool aims to find performance bottlenecks, optimize cloud spend and give a complete picture of utilized cloud resources and their usage details.
The platform can be used as a SaaS or deployed from source code; it is optimized for ML/AI teams but works with any workload.

Github project: https://github.com/hystax/optscale

OptScale live demo: https://my.optscale.com/live-demo

hx hystax