

Operationalizing Machine Learning At Any Scale:

A Look at Companies Succeeding at MLOps

Table of Contents

- 3 Introduction
- 4 SECTION I Doing More Through Automation: How ClearML Helps Daupler Optimize Response Management
- 8 SECTION II How Photomath Uses ClearML to Iterate ML Experiments Quickly
- 13 SECTION III Self-Documenting ML Experiments Slash Wasted Time for Philips
- 17 SECTION IV <u>AgroScout Cultivates New Levels of Efficiency</u> <u>with ClearML Experiment Management Tools</u>
- 21 Summary & Next Steps

Introduction

Machine Learning (ML) is rapidly becoming an integral part of many companies' operations, but at the same time, the deployment and management of ML models can be a complex and time-consuming task.



MLOps, or Machine Learning Operations, is a discipline that can streamline the process of deploying and maintaining ML models.

In this Best Practices Guide, we will explore the benefits of using MLOps in various real-world use cases and introduce examples of companies that have successfully implemented MLOps to unify and streamline their ML workflow. You'll see how four leading companies have implemented ClearML to build, deploy, and maintain their end-to-end machine learning development lifecycles. Read on!

SECTION I

Doing More Through Automation: How ClearML Helps Daupler Optimize Response Management



Client Overview

We recently had a chance to catch up with Heather Grebe, Senior Data Scientist at Daupler, which offers Daupler RMS, a 311 response management system, used by more than 200 cities and service organizations across North America and internationally. This platform helps utilities, public works, and other service organizations coordinate and document response efforts while reducing workload and collecting insights into response operations. With Daupler's technology, organizations improve their response capabilities and strengthen their customers' confidence. The company's goal is to make it easier for clients to focus on their core mission of providing critical services to their customers and communities.

Heather explains the platform handles "both standard and operation-critical 311 service requests that require city or utility response, such as water main breaks, sewer back-ups, traffic signal outages, or general infrastructure damage."



Some of the benefits Daupler provides and the industries they serve

The Challenge

The foundation underlying Daupler RMS is Daupler AI, which can gather and analyze various modes of data relevance, in sequence, based on the last piece of data obtained for a request, apply defined rules configured at the customer level, and initiate and monitor the correct response protocols and resulting data. To do that, Daupler AI receives inputs from multiple sources, such as customer calls, texts, geotags, photos, emails, web forms, and social posts as well as first responder reports and infrastructure alarms. The system analyzes the data using Daupler's proprietary algorithm, which has been trained on data from a pool of millions of request records across many utilities and other service organizations. It then categorizes the data into individual incidents or larger events and routes the information to the appropriate department.

All of that data collection is handled by a machine learning system as part of Heather's realm, where she works with Data Engineer, Carly Matson, to organize millions of data samples, effectively label the most advantageous subset of requests, manage 15 ML models (and growing), and integrate feedback from human-in-the-loop system components. With such a broad remit and lean staff, how does she manage the huge volume of work? The simple answer is automation with ClearML. Heather explains: "It is simply not efficient or effective to spend our limited time manually hacking away at repetitive tasks or writing boilerplate code. That's why I am very interested in automating anything and everything that can be automated, with the proper monitoring."

The Solution

To help do that, Heather is using ClearML Data to create, manage, and version the datasets, ClearML Pipelines to streamline and connect multiple processes, ClearML Orchestrate to handle all our workers, queues, and job scheduling, and ClearML AutoScaler to spin up machines when demand is high (and then spin them back down to reduce costs when no tasks need to be executed).

"ClearML allows us to quickly search for optimal data to improve our models' robustness, label it, get it through the training pipeline, perform detailed analysis with comparisons to our production state, and then stage deployment – all within a day or two."

- Heather Grebe, Senior Data Scientist, Daupler

Automating ML Processes at Daupler

Part of the system Heather works with is a hierarchical text classification system with many different components that work together conditionally. Here's what she had to say about how the automated ML process works:

"Let's say we identify a new concept that needs to be captured in order to serve our customers' diverse needs and we want to build a new model to attempt to extract it from free text that a citizen provides. I'll start by doing some initial sampling and labeling to get a rough MVP concept training dataset in place. Then, once the data has hit the database, it's as simple as a single CLI command to trigger a ClearML pipeline that handles automatic ingestion, analysis, and processing of that new data, ending in an update of our versioned ClearML data set that is then available to any other consuming process.

Once that's updated, it can either automatically trigger training with the current or default hyperparameters or we can actually pull it out and work on it manually through a configuration file to change different parameters – it's all tied in through GitLab CI/CD and the ClearML UI. So we make a change to the configuration file and push the configuration file up through a merge request, then the whole training pipeline will kick off all the way through the evaluation and we'll get a message in our GitLab commit history that says 'Hey, here's a link to your newest model and go in and review it.' If we decide we want to go with it, we can publish it, update a UUID reference in a deployment config, and then push it up through a reviewed merge request to run deployment through GitLab.



An overview of the Daupler AI architecture

We have a robust testing suite that runs in that deployment pipeline to make sure everything is going to have the impact we're expecting, both from a code perspective and from an inference perspective.

Another ClearML feature we use quite heavily is the Hyperparameter Optimization application, which has an incredibly simple UI. To do that, we just take the training step from our pipeline with a simple UUID reference, and specify the method, parameter bounds, run time/iteration bounds, and run-queue and kick it off. We can just pull it out of the pipeline directly and entirely, which I love and run that training step with the queues dedicated to HPT. It doesn't get in the way with any of our other work because I can always manually click and drag jobs in the queue if something has come up that needs to be prioritized before the HPT is complete. In other words, ClearML's queue system enables me to reprioritize tasks and put some manual work before the automatic work to get results out faster. It's very convenient for balancing operation critical work with experimentation and HPT.

Training orchestration is done via the ClearML Autoscaler, so training instances are only up when needed and scaled down when there's nothing to do. The second part of the process, deployment to production environments, is mostly automated and does not usually need any manual wrangling unless alerts are triggered. As long as tests show the model is stable and improving, and our monitoring is silent, everything is fully automatic up until the final push to prod.

In terms of monitoring, we're using ClearML system monitoring for the training phases and Grafana for the bulk of our inference system insights. As for the trigger for deployments, it's written to have a lot of conditional comparisons with different metrics. You can almost think about it like a multi-point decision system that evaluates the model currently in production and the new model on the same testing dataset to determine whether we are seeing the improvements we are looking for without losses to performance in other areas."

Timelines

"If we're building a new concept model, I would say it takes about two weeks to get an MVP to iterate off the ground – between initial data gathering and ontology development, labeling, and then pushing it through the data/training/evaluation pipelines. As far as refreshing the models that are currently in production, it takes about a day; it's just a matter of loading or labeling that active sample, pushing it up, and that's about it," she said.

The Bottom Line

"Using ClearML gives me 99% of what I need, is easy to manage, and once it's set up, we're all set," Heather concluded.

How Photomath Uses ClearML to Iterate ML Experiments Quickly



Client Overview

Photomath is the world's most-used math education app, with more than 300 million learners worldwide. Powered by advanced AI technology, the app scans, solves, and intuitively explains math problems ranging from arithmetic to calculus with step-by-step explanations and animated tutorials to help students learn and practice problem-solving along the way.

The company's mission is to help people understand math, one step at a time. Photomath focuses on being a champion of all learners and making education more accessible. Using the power of AI and a smartphone, Photomath aims to instill a growth mindset in learners, encouraging them to embrace the ups, downs, setbacks, and wins of learning as success. The Photomath app is available in 195 countries and 30+ different languages.



Photomath employees. Source: photomath.com

The Challenge

The main feature of Photomath's app relies on computer vision to read mathematical equations– whether handwritten, from a textbook, or online–analyze them, and provide step-by-step solutions and explanations based on the output of the vision model to help users not only get the solution, but also learn and understand the steps along the way.

"We used Google Sheets to track experiment metadata and performance. But that was error-prone" said Marijan Smetko, Junior Al Engineer at Photomath. "Sometimes we killed each other's experiments due to 'CUDA out of memory' errors or other unpredictable reasons. Quite a few times the wrong data was used." The simple fact was that not using an MLOps platform was lowering Photomath's development velocity.

He noted that there were several instances of the code with slight changes, with little or no tracking. "Plenty of experiments could not theoretically be compared, and that's just a portion of the problems that we knew of. We were inefficient, we wasted time and resources, and there was heavy friction between iterations," he said.

Given those challenges, the team realized they needed an MLOps solution that could track and orchestrate remotely executed experiments and got to work looking for a system that:

- Allowed users to execute experiments remotely on company hardware, without manual sshing
- Enabled AI engineers to track the progress of their experiment training runs and accompanying metrics
- Helped them to compare previous experiments and make more informed decisions



ClearML enables users to get an overview of their research, search for specific data, track experiment performance, and to compare several experiments.

The Solution

"Knowing we had a problem to solve, we were ready to code our own solution, but before we did that, we went online to start searching for more resources and inspiration," said Marijan. "We wanted to see what people had already built that could help us in our endeavors."

Tomislav Filipović, AI Engineer at Photomath, added, "We wanted something that was plug and play. And then, one day, we found ClearML, which does everything we were looking for."

ClearML is an open source MLOps platform that offers data scientists, ML engineers, and DevOps collaborative experiment management, powerful orchestration, easy-to-build data stores, and more – all in one place.

Photomath noted eight distinct advantages about working with ClearML:

1) It's Easy to Set Up

From well-written deployment steps to one Docker Compose invocation, getting started with ClearML was easy.

2) It's Easy to Integrate

Most of Photomath's code was already written as standalone training scripts running on bare metal machines. It was easy for Photomath to write another small script to wrap the experiment running code in a ClearML task and attach some metadata to it.

3) It Offers Experiment Tracking

In order to iterate quickly, it's important to Photomath that their experiments also run quickly. "ClearML's experiment tracking allows us to stop experiments that don't seem to have the opportunity to beat the best one, as well as stop any slowly progressing ones," said Marijan.



One of Photomath's internal experiments related to the LUMEN Data Science competition held by eSTUDENT student organization. One can easily see the model's train and test performance and progress – and can conclude that this example converges too slowly and needs to be changed.

4) It Makes It Simple to Compare Experiments

One of the largest impacts ClearML had on Photomath's AI team is its experiment comparison feature. Previously, when the team used Google Sheets, updating them – finding the right sheet, updating the information, and analyzing the data – was a cumbersome, tedious, and error-prone process. Now with ClearML, that manual process is automated and requires just a few clicks in a browser.

5) It's Easy To Access Experiment Artifacts

ClearML can be configured to store all of the artifacts in the cloud, making it very easy to find a checkpoint from a particular epoch, especially the last and the best-performing epoch. These models are then candidates for deployment. "This enables reusing the artifacts for new experiments, a practice known as warm-start training," said Marijan. "ClearML easily allows our AI engineers to define a model checkpoint from which an experiment can be continued, by downloading it and loading automagically."

6) It Supports Data Management

ClearML will download and run the experiment with the correct data version, and cache the versions and version diffs. Photomath also stores dataset metadata with ClearML Datasets.

7) It Makes It Easy to Report Findings

While the Photomath team uses ClearML to run and evaluate experiments and then choose which experiment was best, they also share those findings with their management team. "Most of our executives come from an engineering background and they are interested in our results. We've had cases when our CEO just opens the ClearML link to see experiments he's interested in," said Tomislav.

8) It's Easy to Use

"We've never had to ask for help with the UI, which means ClearML does its job and it does it well. It solves a very real problem as painlessly as possible. In fact, it gives us everything we need and allows us to work the way we want," concluded Tomislav.



Editor's Note: Parts of this case study were adapted from an article by Photomath, which was originally published here: <u>https://medium.com/photomath-engineering/new-horizons-of-mlops-1579e4d8b45f</u>

Self-Documenting ML Experiments Slash Wasted Time for Philips



Client Overview

Algotec-Philips, part of Philips Radiology Informatic, is a world-renowned pioneer in enterprise software for medical imaging, utilizing cutting-edge technologies in machine learning, computer vision, big data, and natural language processing (NLP). Philips Radiology Informatics provides healthcare facilities with advanced Web-enabled solutions for medical imaging image management, reading, processing, reporting, and distribution. Algotec-Philips' products are the choice of the world's top healthcare institutions, including NIH (The US National Institutes of Health), Johns Hopkins Hospital, Cedars Sinai (Beverly Hills), Institute Curie, Sheba Medical Center and over 2,000 institutions around the globe.

The Challenge

Operating in a highly competitive industry where product quality is of utmost importance, the Philips Radiology Informatics team knew that product delivery and quality were critical, and maintaining a competitive edge would continue to be driven by R&D productivity, processes and workflow. Efficiency in product development is always difficult to measure, but the team knew that they would see results if they could find a tool to drive improvements in the areas where most companies suffer:

Documentation

Like most large development teams, developers were spending substantial time each day manually logging their experiments for review, provenance and root cause analysis. This ongoing record-keeping was often used by both the scientists themselves as they worked, as well as for sharing with colleagues and for archival purposes. Often, in the midst of a complex or intense experiment, this documentation would be superficial, incomplete, or when time was limited ... completely left out of the process.

Collaboration and Transparency

Review/Update meetings were taking much longer than needed; productive debates and brainstorming could occur only after someone provided a long, involved introduction and review of progress. There was no efficient way to prepare and "do one's homework" before coming to the

meeting.

Visualization

Philips' data scientists needed a way to visualize the results of multiple, parallel experiments in order to differentiate, choose the most successful, and further tune the model. Static graphs alone, they knew, weren't enough to provide the whole picture; they had to be able to tie the results and metrics to specific parameters – to dive right down to the underlying data input to spot the precise causes for successes/failures in each case. Most helpful would be the ability to easily jump back and forth between the charts and plots, and the code, hyperparameters and other experiment variables.

• Unification/Standardization

One of the toughest challenges in any development environment where teams need to interact (whether internally, amongst themselves, or with other teams) is sharing one's work using disparate methodologies, terminology and formatting. Even a basic search for a specific configuration file, model, or dataset requires some clear way to express the query, and when the original scientist uses a different syntax, this process can be a challenge.

It was a long, ambitious list, with hurdles familiar to most development organizations. Philips set out to find a tool that would address as many of these challenges as possible.

The Solution

The Philips team's search brought them to ClearML, which ticked all the checkboxes. After installation of the experiment management module, the system began tracking and managing data experiments almost immediately, without additional integration efforts by their algorithm team. Immediately, they felt the difference as efficiency rose in the various areas they had targeted.

First, ClearML's experiment manager "automagically" captures metrics and parameters with the addition of a very few lines of code, so the research basically documents itself. "It was impressive to watch even this one feature begin to operate," says Ohad Silbert, Algorithms Group Leader. "It was like we'd hired a whole team of perfectionist transcriptionists who didn't need to be managed."

Next, data transparency changed the dynamic for team leaders who needed to manage the workflows and plan ahead: Not only do they now come to meetings better prepared and ready to get down to strategic work, but often update meetings were completely eliminated as the information was easily accessed and reviewed.

The ClearML built-in visualization tools elegantly expand the effectiveness of Philips' existing Tensorboards to offer new ways to slice and dice data, for both the developers and their managers. Suddenly, this comprehensive amalgamation of parameters, models, and underlying metrics paint a very concrete, complete picture, on their own or in comparison to related experiments. Ohad explains: "The big leap here is that logging of our Tensorboards happens automatically. Moreover, they are automatically associated with the code and hyperparameters that produced them, so developers can compare experiments tying them all up together."

Finally, ClearML provides an API and UI to manage experiments in a consistent, standardized manner. "We realized that implementing all of our requirements would be too time-consuming," says Ohad. "There were simply too many functions we needed to design, build and continuously tweak. ClearML gives us exactly what we'd been envisioning in that system."

The Results

Overall, performance across the algorithm's team has significantly improved. As Evi Kopelowitz, an Algorithm Developer explained, "ClearML creates an unexpected 'calming' feeling because all the data is just ... right there. I can instantly filter, search for, and archive results, and organize experiments using the Leaderboard. For instance, there was a day I needed to retrieve critical data from a previous experiment, but I simply didn't remember its name. But I did remember the results, so I searched for that, and found exactly what I needed. It sounds simple, but it saved me hours of work. And this kind of thing now happens all the time. Those saved hours add up quickly."

Automatic documentation has also freed up hours per week. Now the team can review past and current experiments, as each is documented cleanly and consistently. Meeting time has been cut down, and the meetings that do occur take less time as they are based on empirical data that's simpler to collect and understand. With new levels of efficiency, the Philips team is now demonstrating measurable boosts in productivity.

Evi explains a challenge that's eliminated by ClearML: "When we run a typical results comparison, we start by comparing all experiments with a specific hyperparameter or metric. Once we identify a few of the best-performing experiments, we dive deeper: We compare metrics and graphs, and finally, we look at hyperparameters to determine the source of those differences. Often, those hyperparameters still don't explain the variance in performance, and we're left with time-consuming code comparison. With ClearML, everything is under one roof; I don't have to go out of the platform, or waste time writing custom one-off code for simply analyzing my experiments. Although we knew we needed this type of organization, the contribution to the efficiency of our work exceeded our expectations."

"When it comes to machine learning, I'll admit it – it's easy to develop bad habits," adds Ohad. "But it's as if ClearML made a list of our inefficiencies, and provided the tools to zap them one by one, helping us invisibly in the background."

"The very nature of workflows and the complex synthesis of code, models, data repositories, and infrastructure means that there are way, way too many moving parts to let us work the way we should," Evi concludes. "Adopting ClearML's solution undoubtedly propelled us forward in tracking, managing, and documenting our research and delivering our algorithms."

SECTION IV

AgroScout Cultivates New Levels of Efficiency with ClearML Experiment Management Tools



Client Overview

AgroScout has created an automated, AI-driven scouting platform for the early detection of pests and disease in vast agricultural areas. With precise, data-driven insight (literally impossible to attain through periodic, manual sampling), a farmer can proactively prevent substantial loss of crops and corresponding revenues as he drives increased productivity across his acreage, all the while reducing his use of pesticides because he can more precisely target problematic areas.

The Challenge

The team needed to focus on their technology, not the infrastructure managing it in the background. AgroScout uses drone and camera footage to gather crop data to identify the subtle indicators of pest and disease. If this were all simply the same Big Data challenge common to companies focused on Al-based imagery analysis (medicine, construction, traffic, etc.), they would be able to handle it by simply building models and throwing plenty of computing resources at them. But working with imagery analysis for agriculture poses a few unique challenges:

- Annotation isn't trivial; even experts can't find all plant defects.
- Analysis requires engaging a team of specialized, global experts to annotate; managing their work and standardizing their terminology needs to be precise.
- High data modality. Each field and even each field's dataset has unique and multi-input features to be learned and merged (e.g., changing lighting or flight angles).
- Scale Just a single image can yield thousands of annotations, and annotation QA is difficult; there are often tens of thousands of these images.

In short, AgroScout's technology involved some technical heavy lifting, and their workflows needed to be optimized to take manual work off their hands. Naturally, they wanted a scalable, self-contained solution that would take little effort to integrate and maintain. The team needed to focus on their technology, not the infrastructure managing it in the background.

The Solution

AgroScout's research led them to ClearML as a way to track research, control their cloud resources, and manage data throughout their development cycles. ClearML helped in six ways:

1) Processing Hi-res Images

As processing massive high-resolution images isn't feasible, AgroScout's approach is to cut them into sub-images before they are fed into the model for training. Naturally, this dramatically increases the number of images to store and track for analysis. ClearML's dataset management feature helped them to create versions for each data set, including one version that had been preprocessed and was ready for training. Facilitating this workflow was already practically sufficient to justify their integration of ClearML, but it was just the beginning.

2) Managing Divergent Annotations from Multiple Agronomists

Every ML model begins with human input to teach the system what to look for in its initial analysis of datasets it is given to work on. When it comes to agriculture, skilled annotators are hard to find; qualified individuals and companies were recruited from around the globe. AgroScout quickly discovered that no matter how directly and clearly they requested standardized annotations, they were receiving data with various label names and terminologies, even from multiple employees within the same vendor. Using dataset statistics in ClearML's UI, they identified mismatches, created unification strategies to map synonymous terms, and finally created alias rules for labels with the same meaning.



Just a single image can yield thousands of annotations. Even for a human, it is hard to detect all the diseases and pests, and to decide how exactly to classify what's been detected.

3) Annotating Incredibly Complex, Multi-faceted Datasets

AgroScout cannot succeed with a single model to cover all agricultural scenarios. Their clients are located around the globe, and naturally, every field has its own unique crop type, weather, age, soil, and other variables.

These factors make the analysis of each particular location different from the analysis of another. All these vectors create a management problem when trying to keep metadata readily available without crawling through multiple folders for every analysis. ClearML's Hyperdatasets allowed AgroScout to organize their data versions and associate them to model training tasks for post-training analysis.

4) Annotation Mass

A single image can sometimes require literally thousands of annotations to record details like ground lift (bumps that indicate that a plant is about to emerge), first emergence, adult plant, and gaps in the field. Not only is this an intense process for the annotators, but the software itself must handle immense quantities of metadata, with full editing features along the way. This type of data management can present a challenge to many data management tools, but ClearML was designed for precisely this type of scalability challenge.

5) Managing Workloads on GPU Instances

Training and testing the type of models used for agricultural analysis – even when optimized – require vast GPU resources. AgroScout built data ingestion pipelines triggered by AWS Lambda functions. Integrating with ClearML Orchestrate, which manages workloads on EC2 instances for training use; executed tasks are then monitored and managed using ClearML's Orchestrate UI. This was a classic example of a heavy, time-consuming, manual task that would have diverted them from core development.

6. Choosing the Winning Models

This particular type of model comparison is exceptionally fluid and dynamic because of the "fuzzy" biological data representing an organic landscape's continuously changing and unpredictable nature. It's rarely a case of "this model succeeded, and this one failed." Conducting the comparisons – deciding which experiments to continue running and which to end – is at the core of ClearML's suite of tools, and proved powerful enough to conduct active, ongoing comparisons to streamline the process of refining the models.

The Results

While overall productivity and accelerated time to market are hard to measure precisely, AgroScout's team feels the improvement across their development lifecycle:

100x

They increase their data volume over 100x without growing the data team

50x

Increased experiment volume 50x with the same team size

50%

Shortened the time to production (time from experimentation start to model in production) by over 50%

Summary and Next Steps

We've seen how companies using ClearML's platform for continuous machine learning streamline the process of building, deploying, and maintaining ML models.

For data scientists, ML engineers, and DevOps professionals who need to develop, orchestrate, and automate ML workflows at scale, ClearML is the only MLOps platform that gives users and customers collaborative experiment management, powerful orchestration, easy-to-build data stores, and one-click model deployment and monitoring.

Unlike alternate solutions, ClearML's end-to-end, open source MLOps platform enables users and customers to focus on developing their ML code and automation, ensuring their work is reproducible and scalable – all in one single platform.

Key Features

The key features of ClearML's unified enterprise offering are:

- **ClearML Experiment,** allowing you to track every part of the ML experimentation process and automate tasks. With it, you can log, share and version all experiments and instantly orchestrate pipelines.
- **ClearML Orchestrate** so that DevOps and data scientists are empowered through autonomy and control over compute resources. The cloud-native solution also enables kubernetes and bare-metal resource scheduling with a simple and unified interface to control costs and workloads.
- **ClearML DataOps**, delivering data store automation. Automate data collection into searchable, accessible, and ML-ready data repositories through cutting-edge MLOps technology.
- **ClearML Hyper-Datasets**, allowing MLOps teams to build data-centric AI workflows. Make the most out of unstructured-data using queryable datasets.
- **ClearML Deploy**, delivering a unifying model repository, custom pipelines, and model serving. This allows MLOps teams to transition from model development to production and gain full workflow visibility with seamless integration to the experiment manager and orchestration.

Every component of ClearML integrates seamlessly with each other, delivering cross-departmental visibility in research, development, and production.

In addition to ClearML for Enterprise, ClearML is available as an entirely modular/a la carte offering. As an open source solution, the software is freely available to all users and can be modified to fit the needs of any specific user. This allows customers to take a scalable approach to their MLOps needs.



ClearML streamlines the process of building and deploying ML models, vastly shortening the time to production and increasing the number of models that make it to production.

How to Get Started

Get started with ClearML by using our free tier servers or by hosting your own. Read our documentation here. You'll find more in-depth tutorials about ClearML on our YouTube channel and we also have a very active Slack channel for anyone that needs help. If you need to scale your ML pipelines and data abstraction or need unmatched performance and control, please request a demo. To learn more about ClearML, please visit: https://clear.ml/.

About ClearML

ClearML is a unified, open source platform for continuous machine learning (ML), trusted by forward-thinking Data Scientists, ML Engineers, DevOps, and decision makers at leading Fortune 500 companies, enterprises, academia, and innovative start-ups worldwide. We enable customers to build continuous ML workflows -- from experiment management and orchestration through data management and scheduling, followed by provisioning and serving -- to achieve the fastest time to ML production, fastest time to value, and increased performance. In this way, ClearML accelerates ML adoption across business units, helping companies reach their revenue potential and materialize their ML investments. With thousands of deployments and a vibrant, engaged community, ClearML is transforming the ML space -- bridging software, machine learning, and automation. To learn more, visit the company's website at https://clear.ml.