



Self-Documenting ML Experiments Slash Wasted Time for Philips

Philips Radiology Informatics: 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, suffered:

- **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 comparisons 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.

“ClearML’s solution undoubtedly propelled us forward in tracking, managing and documenting our research and delivering our algorithms.”

Evi explains a typical 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 Trains, 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 exceed 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 within Trains 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.”



Allegro AI makes ML and DL researchers more effective by giving them tools to manage their own experiments and data. The company’s open source ClearML platform automates and simplifies developing and managing machine learning solutions for thousands of data science teams all over the world. Allegro AI is trusted by brands such as: NVIDIA, NetApp, Samsung, Hyundai, Bosch, Microsoft, Intel, IBM and Philips.

Contact us to learn how we can help you: info@clear.ml