Case Study
Neural Guard

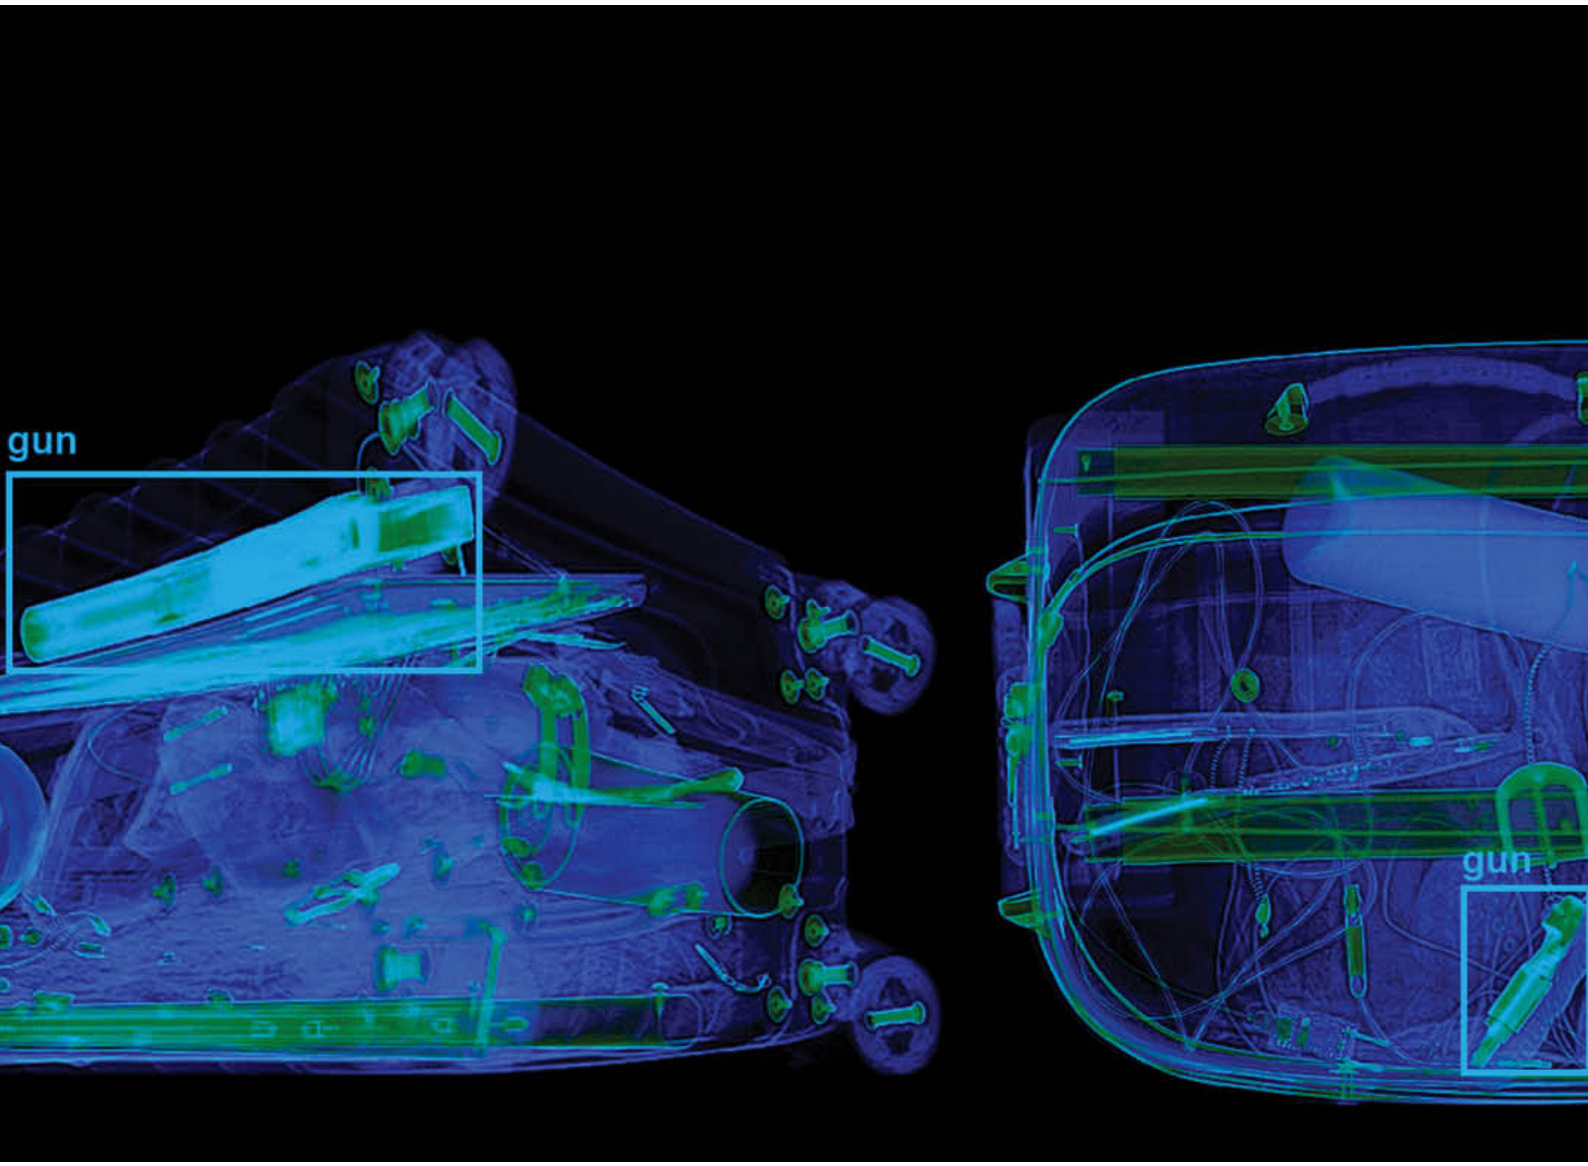

Neural Guard's AI Workflow: The Image of Efficiency

Neural Guard built a state-of-the-art, production-grade data pipeline for building, maintaining and serving multiple object detection models ... all on top of ClearML.

Neural Guard: Client Overview

With the expansion of global trends like urbanization, aviation, mass transportation, and global trade, the associated security and commercial challenges have become ever more crucial. Neural Guard produces artificial intelligence-based auto-detection solutions for the security screening market. Neural Guard technology detects specific, high-risk items in CT and X-Ray imagery by leveraging cutting-edge artificial intelligence algorithms to analyze a security scanner's output.

The Challenge

The team at Neural Guard face the challenge of building, optimizing and maintaining deep learning (DL) models that recognize multiple unique objects found in high resolution X-Ray imagery. As opposed to processing instantly readable, text-based data, imagery analysis requires incredibly large data sets and models that take into account unusually extensive potential variations in each image as it is identified, tagged and fed into the model. Additionally, each X-ray or other detection machine has unique, subtle deviations in its output that require attention to achieve high-quality detection results. Even more challenging is the fact that new samples – uniquely shaped knives, rare gun models, home-made weapons – are constantly added into even "mature" models, refining them even further. In short, Neural Guard's solution requires an ongoing, data-heavy experiment process.

From a business value standpoint, Neural Guard had two key objectives for its detection system:

- Deliver the highest quality detection capabilities for security screening machines.
- Create as automated a system as possible to handle the vast complexity of the multiple environments it will be installed on, while saving on costs.

Naturally, managing this matrix of data sets and models demanded a powerful management platform. It had to be scalable enough to handle the growing

data, the vast array of thousands of machines it is installed on, and the always-expanding collection of models as they move through the pipeline to be documented, reproduced, compared, shared, stored and easily searched. And all this, ideally, without requiring substantial DevOps effort, or a data scientist's own hands-on involvement to manage the logistics of these processes.

While creating object detection models was important, it was vividly clear that the most important piece of the puzzle was overcoming the AI data management challenge: To effectively and accurately process huge datasets and prepare the highest quality datasets for continuous training of thousands upon thousands of ever-changing object detection models.

Neural Guard clearly needed a best-of-breed, scalable solution to manage the pipeline processes key to efficient DL development. As they began to explore options, they quickly discovered that there were few platforms scalable enough, comprehensive enough, and designed to easily integrate into customer-specific workflows.

The Solution

As the first step, Neural Guard designed a plan for an automated pipeline that would be able to:

- Receive images generated by its x-ray machines
- Analyze the images to identify the key images that add value to the models
- Annotate / label them
- QA / QC the labeled data
- Create datasets by defining subsets of the total data, for training and testing of custom models built, for each product line
- Run training jobs
- Evaluate results of created models
- Compare results of older versions to new ones
- Deploy best models to the cloud and edge devices

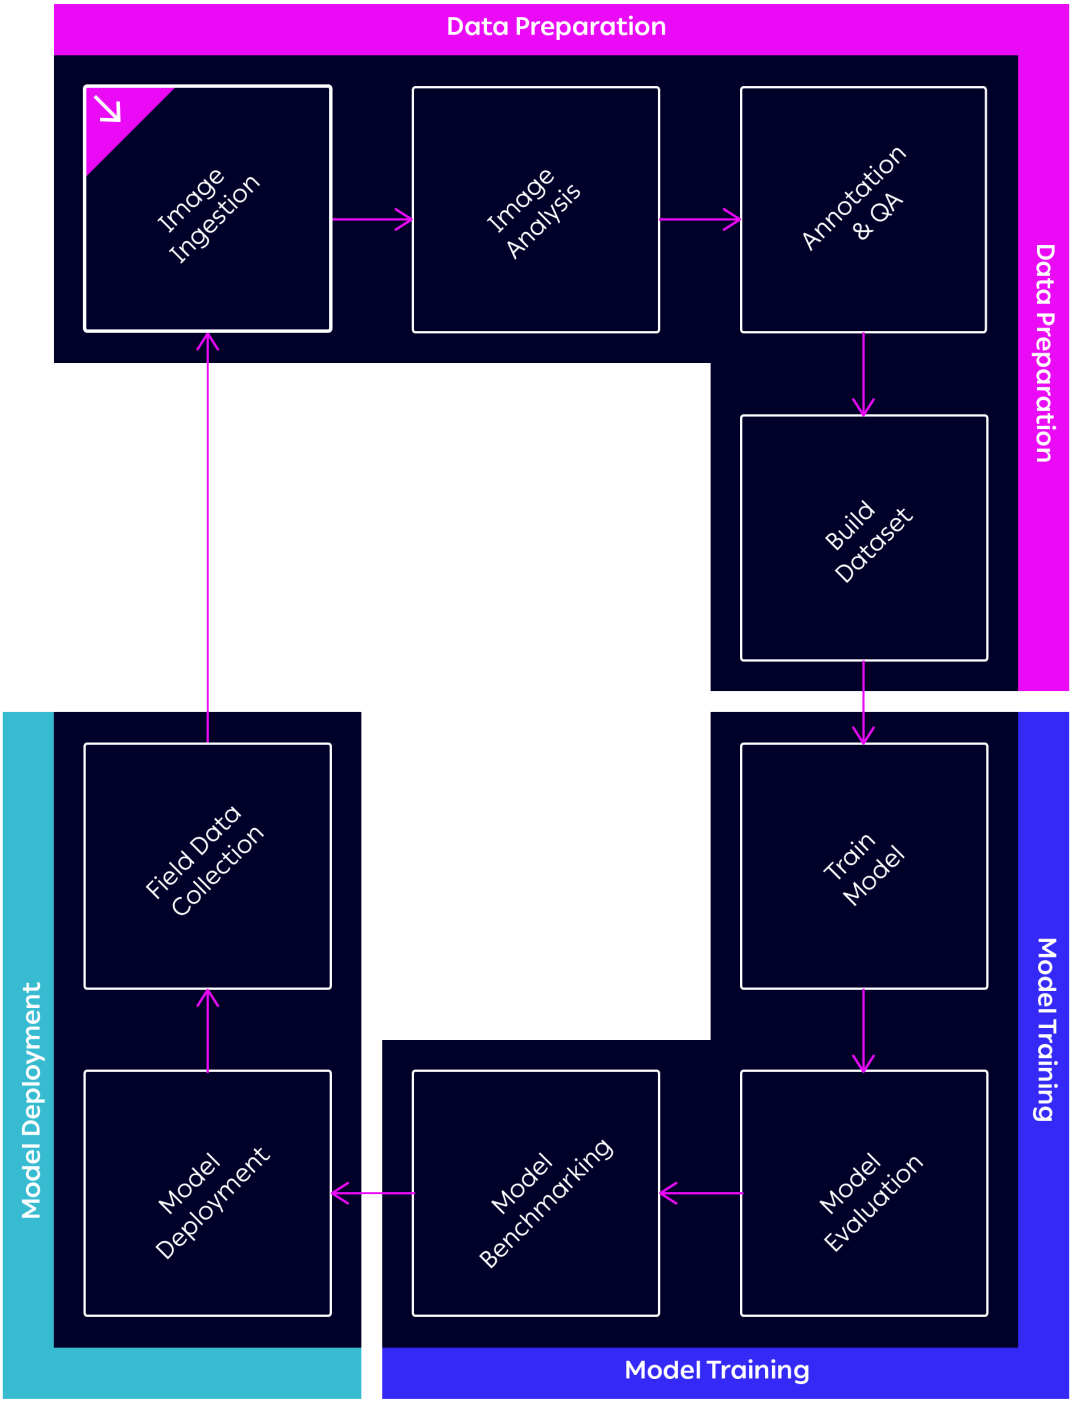

Diagram of Neural Guard's automated pipeline

At the heart of this challenge lay the ability to create a system that would enable Neural Guard to "take ownership" of its data at a very granular level, both in terms of being able to analyze the data at hand, and also manipulate it. Effectively, they sought to set up a debiased, optimized training data set for each machine and object.

As their solution architecture began taking shape, Neural Guard realized that they would need to rely a lot on specific software development to build this pipeline, including -- among other things -- building their own human-labeling management system. But they also realized that building the core data management piece is a gargantuan task. Luckily, they knew about ClearML, the experiment management, ML-Ops and data management platform. It was the only commercial platform they found that would be able to deliver the data management capabilities they were looking for. With all this in mind, Neural Guard set out and built this state-of-the-art, production-grade data pipeline for building, maintaining and serving multiple object detection models, all on top of ClearML.

"Using ClearML data management features proved to be an invaluable tool for us," explains Raviv Pavel, CTO for Neural Guard, "With it, we were able to understand our data and data requirements on a much more profound level. One major factor we were able to accurately measure, for instance, was how much data do we actually need. It turns out that when we can track and compare multiple experiments easily, including what data went in, we actually did not need as much data as we thought. This was a huge cost saver for us."

"Using ClearML ... we were able to understand our data and data requirements on a much more profound level."

Another huge benefit was the ability to build a true scalable, continuous learning pipeline. With ClearML taking care of version logging and fetching the data, Neural Guard focused only on bringing more data in and evaluating its benefits. "The ability to truly track what each dataset contributed to the model performance was very powerful," says Pavel. "We were able to focus solely on analyzing the results, and not spending time on building an infrastructure that would support the process."

> "The most fundamental key to success is building an automated, high-quality scalable data pipeline. ClearML catapulted us in what we have been able to achieve in both the time and resources needed."

Another highly appreciated benefit was ClearML flexible, robust and easy-to-integrate with SDK and APIs. "Using ClearML's SDK made tedious work like updating metadata on images into a simple task," explains Pavel. "It was another important component of our system that we didn't have to design and then to build on our own – and it became indistinguishable from other native parts of our system."

ClearML was integrated and customized on other fronts as well. For example, the experiment manager was customized to add custom metadata to experiments, and, using the ClearML REST API, to extend the existing dashboards to provide particularly relevant metrics and graphs for Neural Guard.

This integration was especially helpful when using ClearML's dataset management with the experiment manager to access required data; Neural guard could leverage the same codebase, and use the UI to quickly change the data used. Add to that the fact that ClearML also manages local data caching (including prefetching the data and ensuring that the latest version is always present, without any manual work), and it confirmed for the team the wisdom of the decision to choose ClearML.

A final aspect that Neural Guard had to take into consideration was model deployment. The world of security presents its own rigorous regulations, limitations, and restrictions, so it was clear to Neural Guard that model deployment had to remain safely on-premise at their customers' sites. This reality introduced a host of specific challenges, including permissions and versioning. Neural Guard leveraged ClearML's model management to easily identify the best performing model, then fetch and distribute it to their custom-built model deployment solution. To achieve this, the team built an identity and permission management system, complete with a deployment and update pipeline, custom-tailored to their service's workflow.

The Results

Neural Guard's benefits in leveraging ClearML can be divided into three categories:

- Saving on cost and shortening time-to-market by not building their own data management solution. Their estimate is decades of man hours.
- Ongoing saving related to training, maintaining and deploying multiple DL models. Much of these savings are due to eliminating the need to hire additional data scientists, as well as the need for a dedicated large ML-engineering and data-engineering team, including supporting DevOps staffers.
- Delivering a best-in-class solution to its customers.

All of this simply as a result of using ClearML as a core component in its training and deployment pipeline.

"For a company whose key value proposition to its customers is the quality of its AI detection algorithms at a competitive price point, the most fundamental key to success is building an automated, high-quality scalable data pipeline. ClearML catapulted us in what we have been able to achieve in both the time and resources needed," concludes Pavel. "There is currently no comparable commercial solution to ClearML out there."