mlangles

CloudAngles

mlangles Predictive AI

Real-Time Construction Site Safety Monitoring with Predictive AI

About mlangles Predictive AI

mlangles is a comprehensive AI platform designed to manage the lifecycle of data and models, offering streamlined solutions for every stage of the process.

Through its Predictive AI component, mlangles provides a suite of tools to navigate efficiently through each phase of AI project development, encompassing data engineering, development, deployment, and monitoring. It facilitates continuous integration, continuous deployment, continuous training, continuous monitoring (CI-CD-CT-CM), enabling enterprises to effectively manage their AI initiatives.

Objective of the Use Case

This use case explores the development of a deep learning model for real-time detection and classification of construction site safety equipment usage.

The model accurately identifies and categorizes Personal Protective Equipment (PPE) worn by workers, such as hardhats, masks, safety vests, etc. which will contribute to a safer work environment by enabling real-time monitoring and potential violation alerts.

Explanation of the use case

This use case focuses on developing a deep learning model that can automatically identify Personal Protective Equipment (PPE) worn by construction workers in real-time.

This dataset is a valuable collection of images annotated with specific safety equipment and conditions, such as 'Hardhat', 'Mask', 'NO-Hardhat', 'NO-Mask', 'NO-Safety Vest', 'Person', 'Safety Cone', 'Safety Vest', 'machinery', and 'vehicle'. The dataset comprises 512 training images, 114 validation images, and 82 test images. The detailed labeling, particularly distinguishing between 'Hardhat' and 'NO-Hardhat', is crucial for tracking and monitoring applications aimed at ensuring safety compliance on construction sites.

In practical applications, leveraging this dataset with machine learning models can facilitate real-time detection of safety equipment usage, aid in accident prevention, and streamline compliance monitoring processes on construction sites. The availability of such a well-annotated dataset enhances the effectiveness and reliability of safety-focused AI solutions in construction industry applications.

Working of the use case:

STEP 1: Data Engineering and Pipeline Creation

Install Dependencies: The required packages and libraries are installed.

Data Extraction: After loading the dataset from the S3 bucket, we will use it for data preprocessing and transformation steps.

Data Visualization: A sample image from the dataset displays various classes, such as 'Hardhat', 'Mask', 'NO-Hardhat', 'NO-Mask', 'NO-Safety Vest', 'Person', 'Safety Cone', 'Safety Vest', 'machinery', and 'vehicle'. This visualization helps in understanding the different objects the model needs to detect and classify, aiding in fine-tuning and evaluating the model's performance.

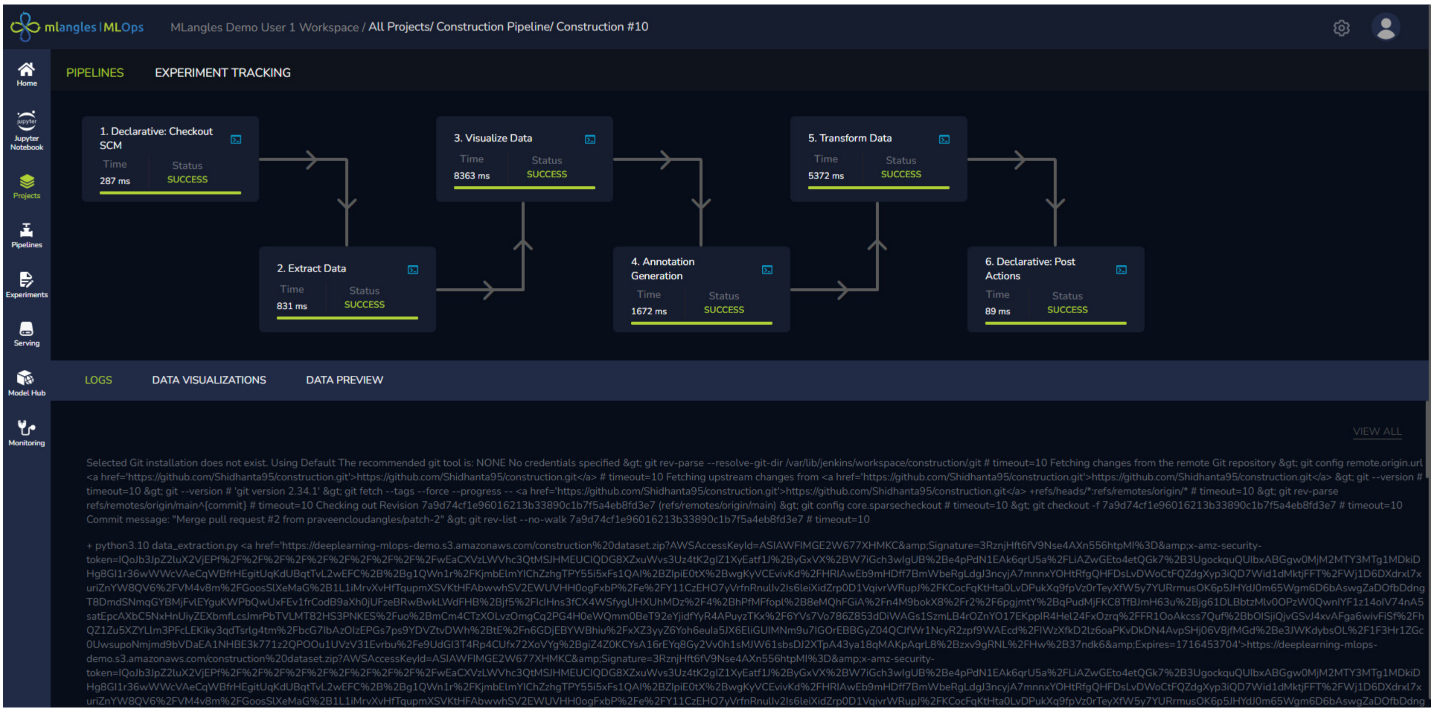

Data Transformation: Data transformation helps improve model performance by ensuring consistency and optimal compatibility with the deep learning model's requirements. The input images undergo a sequence of transformations as part of the preprocessing pipeline:

Resizing: The images are resized to 256x256 pixels to standardize the input dimensions.

Cropping: The center of each image is then cropped to 224x224 pixels to focus on the most relevant part of the image.

Conversion to Tensors: The images are converted into PyTorch tensors, enabling efficient processing by the model.

Normalization: The tensors are normalized using predefined mean and standard deviation values to standardize the pixel intensity values.

These transformed images collectively form the dataset used for training, ensuring that the data fed into the model is consistent and well-prepared, thereby enhancing the model's performance.

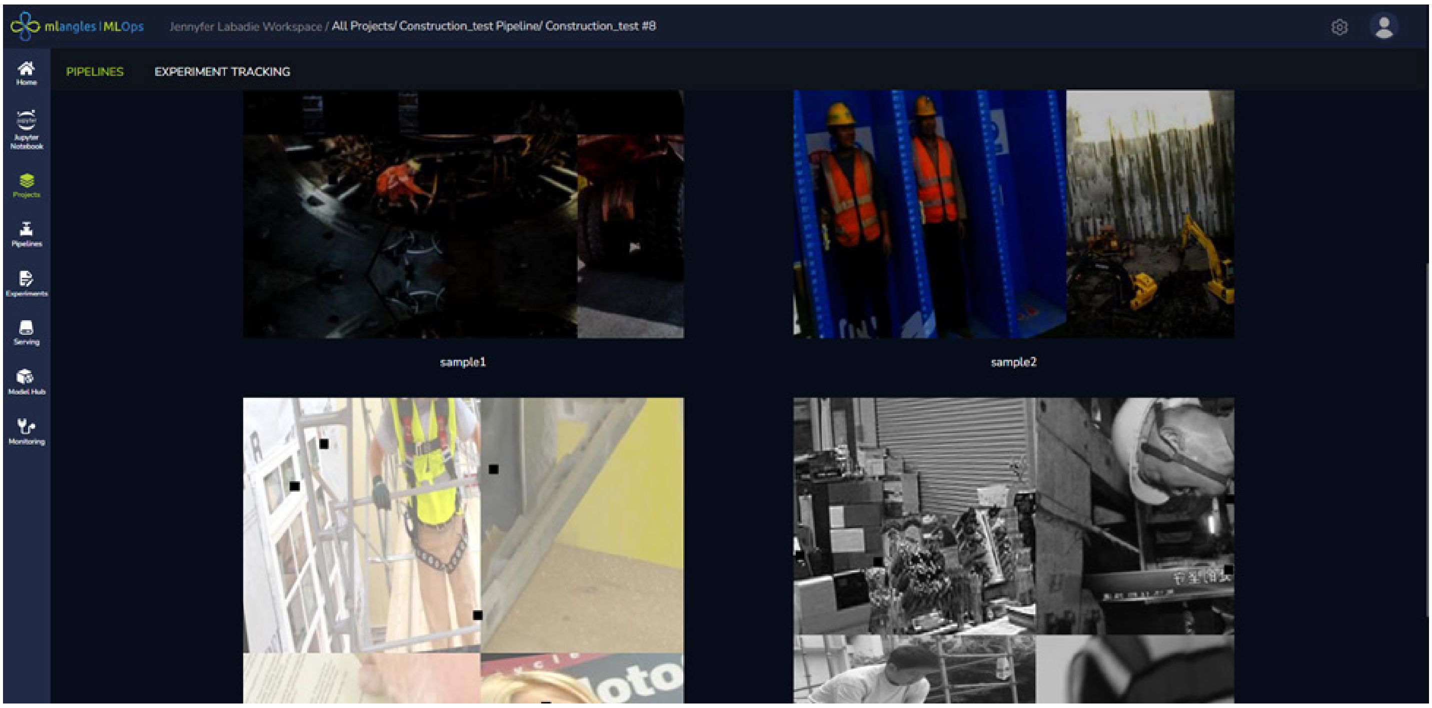

STEP 2: Experiment Tracking – modelling

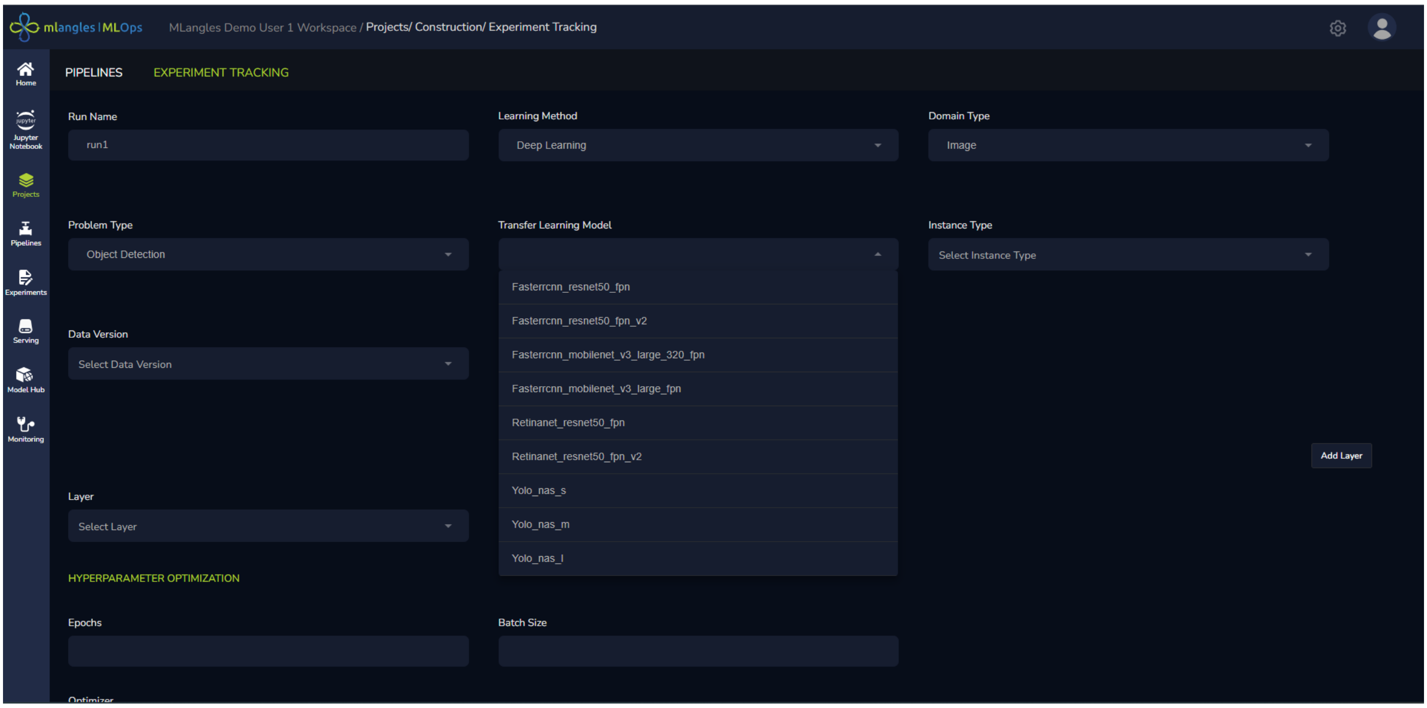

Once the dataset is prepared from the preprocessing pipelines, it is utilized to fine-tune the YOLO (You Only Look Once) NAS (Neural Architecture Search) model for construction equipment labelling. The process involves leveraging transfer learning by adapting a YOLO model that has been pretrained on a large-scale dataset, such as COCO (Common Objects in Context), to a smaller dataset containing images of construction equipment labelling categorized into two classes: Hardhat, Mask, NO-Hardhat, NO-Mask, NO-Safety Vest, Person, Safety Cone, Safety Vest, machinery and vehicle.

Transfer learning with YOLO NAS involves leveraging the learned features from the large-scale dataset to quickly adapt the model to the specific task of construction equipment labelling. This approach can lead to improved accuracy and efficiency compared to training the model from scratch, especially when working with limited labeled data and computational resources.

During fine-tuning, key parameters such as the number of epochs, batch size, optimizer (e.g., SGD, Adam), and loss function (e.g., cross-entropy) are configured to optimize the model's performance on the construction equipment labeling task. The results of fine-tuning the YOLO NAS model include model parameters, evaluation metrics (e.g., accuracy, precision, recall), and visualizations of predicted construction equipment classifications, demonstrating the model's effectiveness in identifying and categorizing different types of safety equipment's.

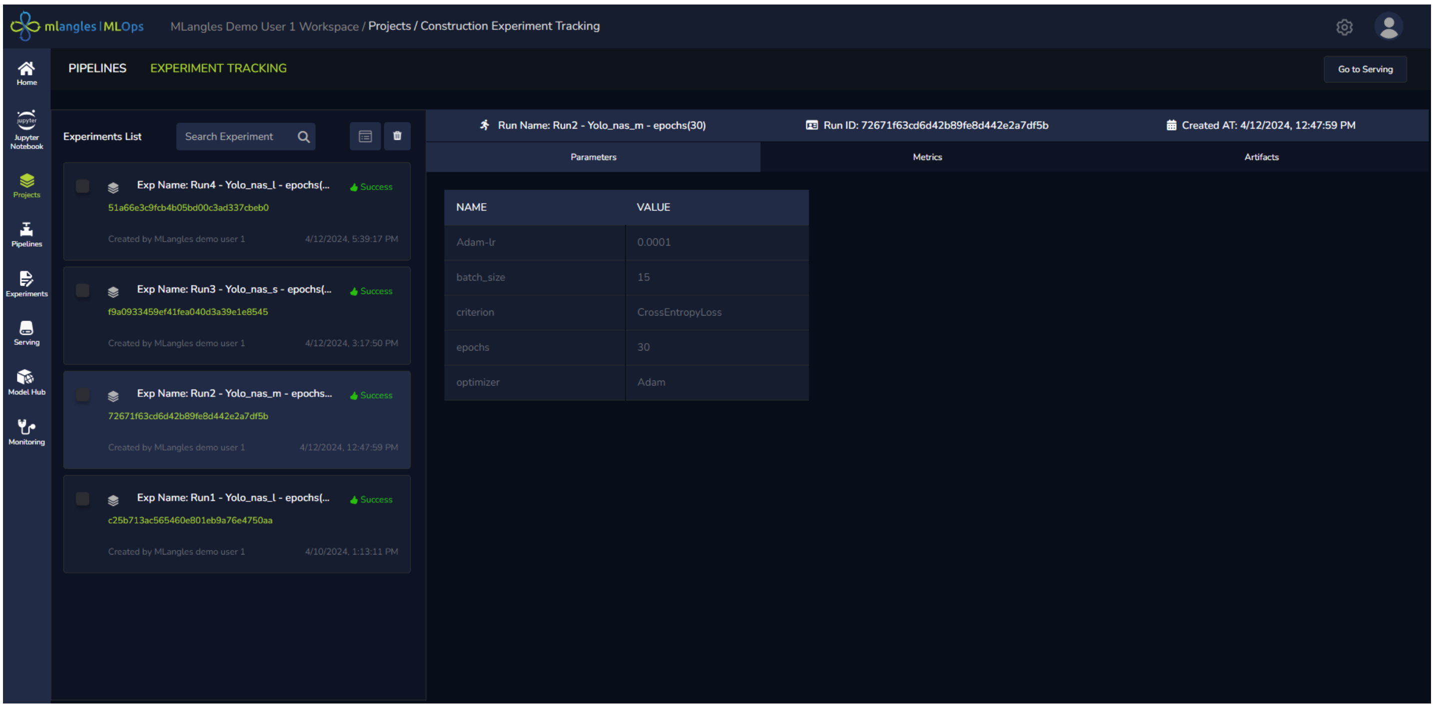

The metrics include the best validation accuracy and the least validation loss recorded during finetuning which measure how well the model generalizes to unseen data, often referred to as the validation set, which is distinct from the training data.

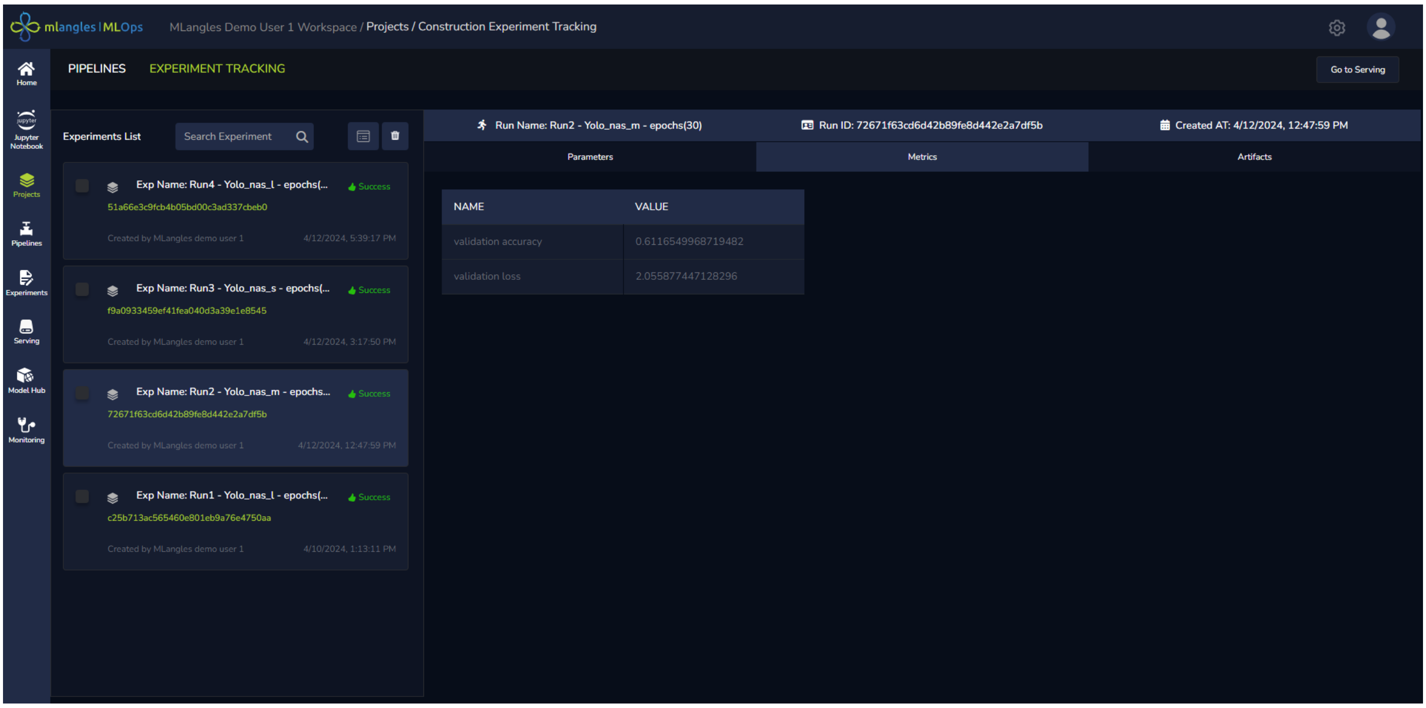

The trained model along with its dependencies could be downloaded from artifacts to perform predictions offline.

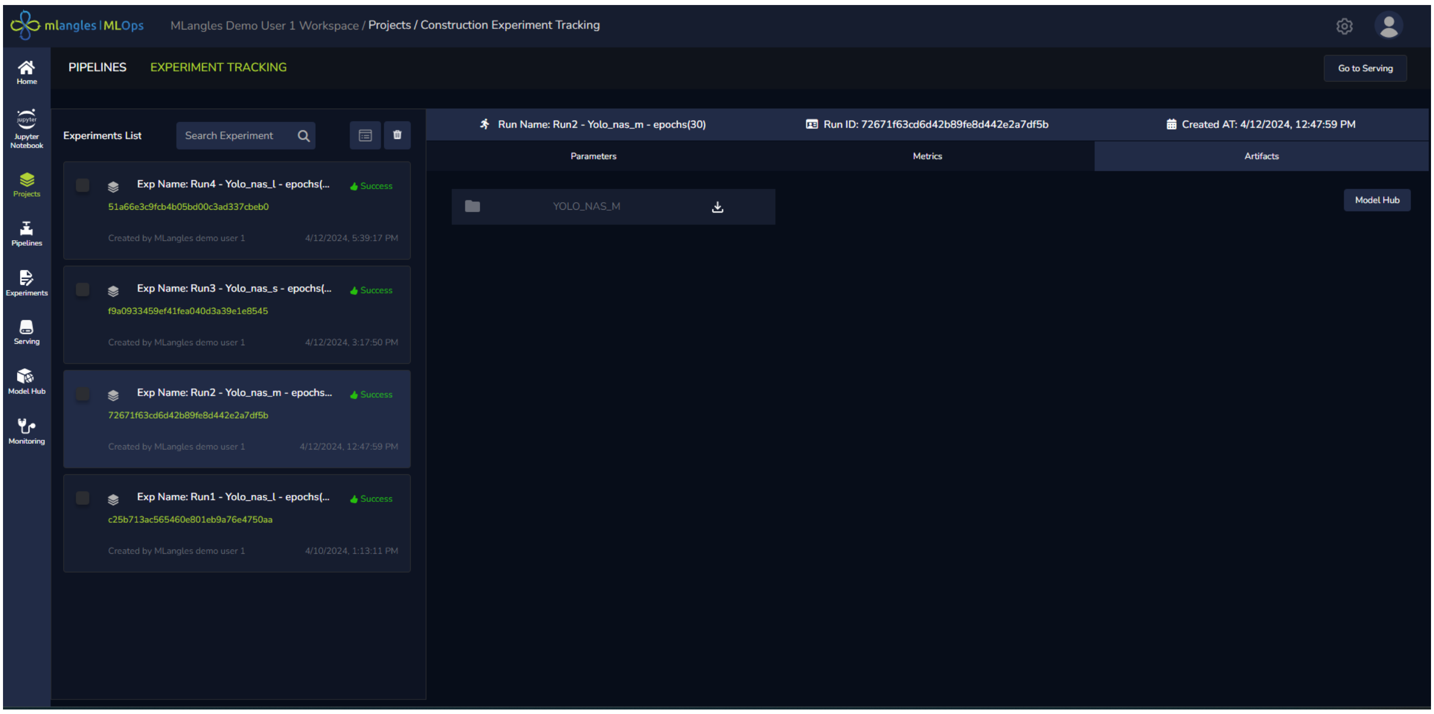

STEP 3: Experiment Tracking – modelling

The image is a photograph of a construction image, and it is provided as input to the trained construction detection model. The model accurately identifies the presence of the safety and classifies its type (e.g., Hardhat, Mask, NO-Hardhat, NO-Mask, NO-Safety Vest, Person, Safety Cone, Safety Vest, machinery and vehicle), displaying the result on the screen. This demonstrates the model's capability to analyze car images and detect specific types of safety types.

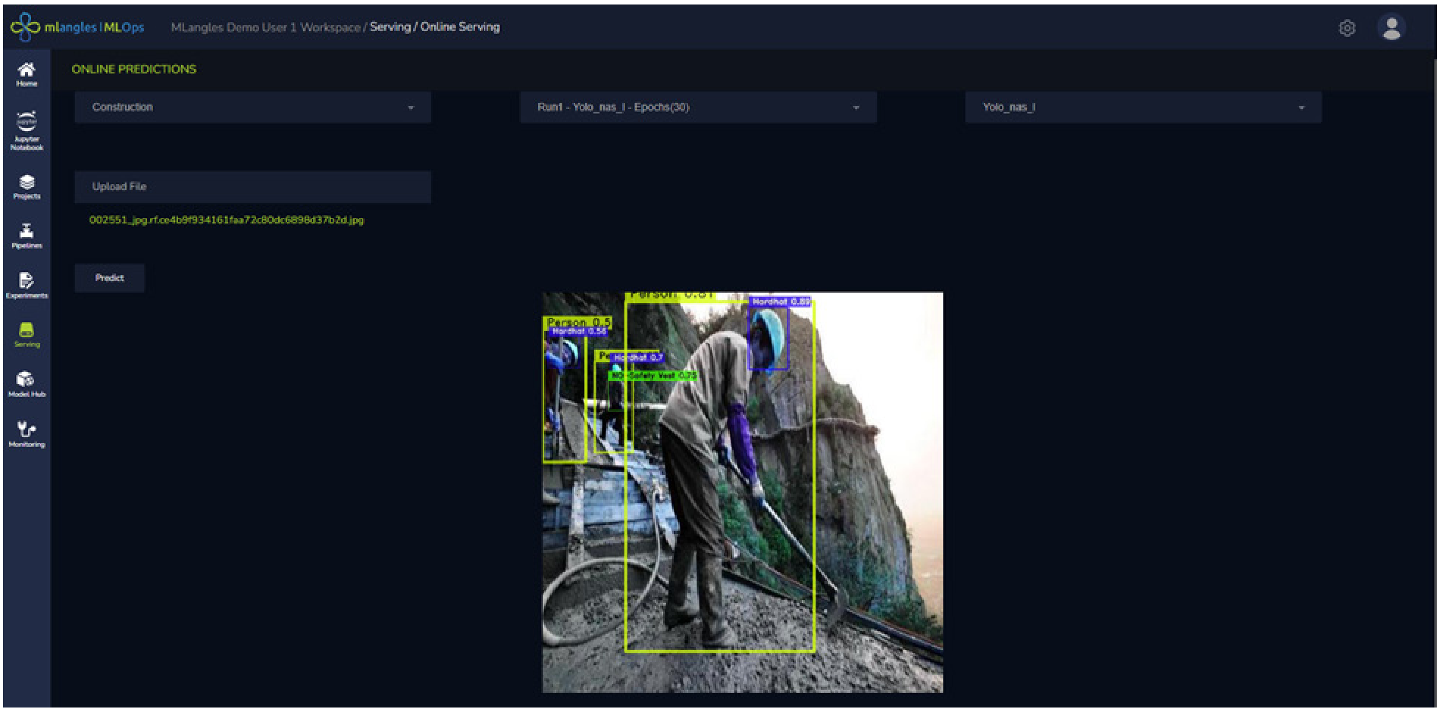

Model Hub:

- Trained models are uploaded to the model hub, whereupon deployment, a REST API endpoint is automatically generated.
- Data is transmitted to this endpoint as a request, triggering the model to execute a prediction and return the output as the response to the request.

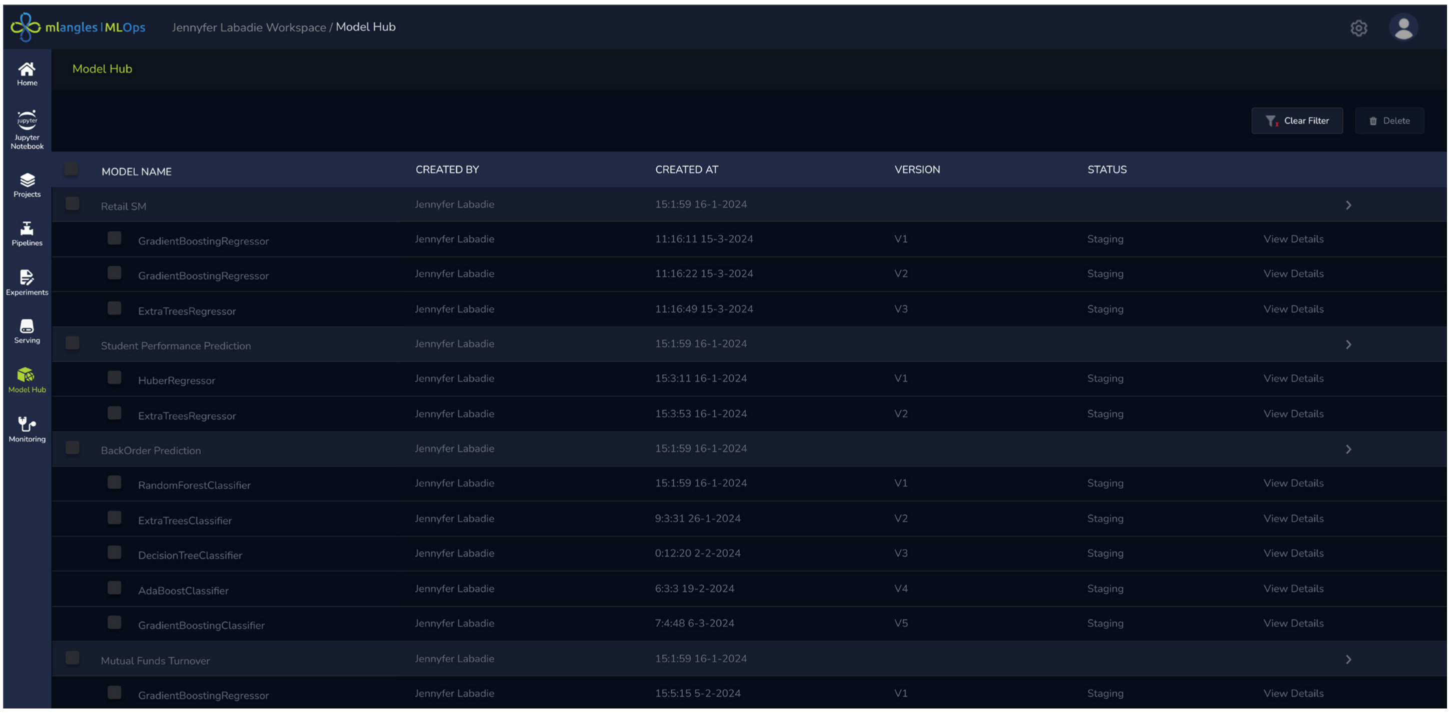

STEP 4: Monitoring

Data drift refers to the phenomenon where the statistical properties of images in a dataset change over time, potentially leading to a decrease in the performance of machine learning models trained on that dataset. This change can occur due to various factors such as changes in imaging equipment, variations in lighting conditions, differences in image acquisition techniques, or shifts in the characteristics of the target population being imaged.

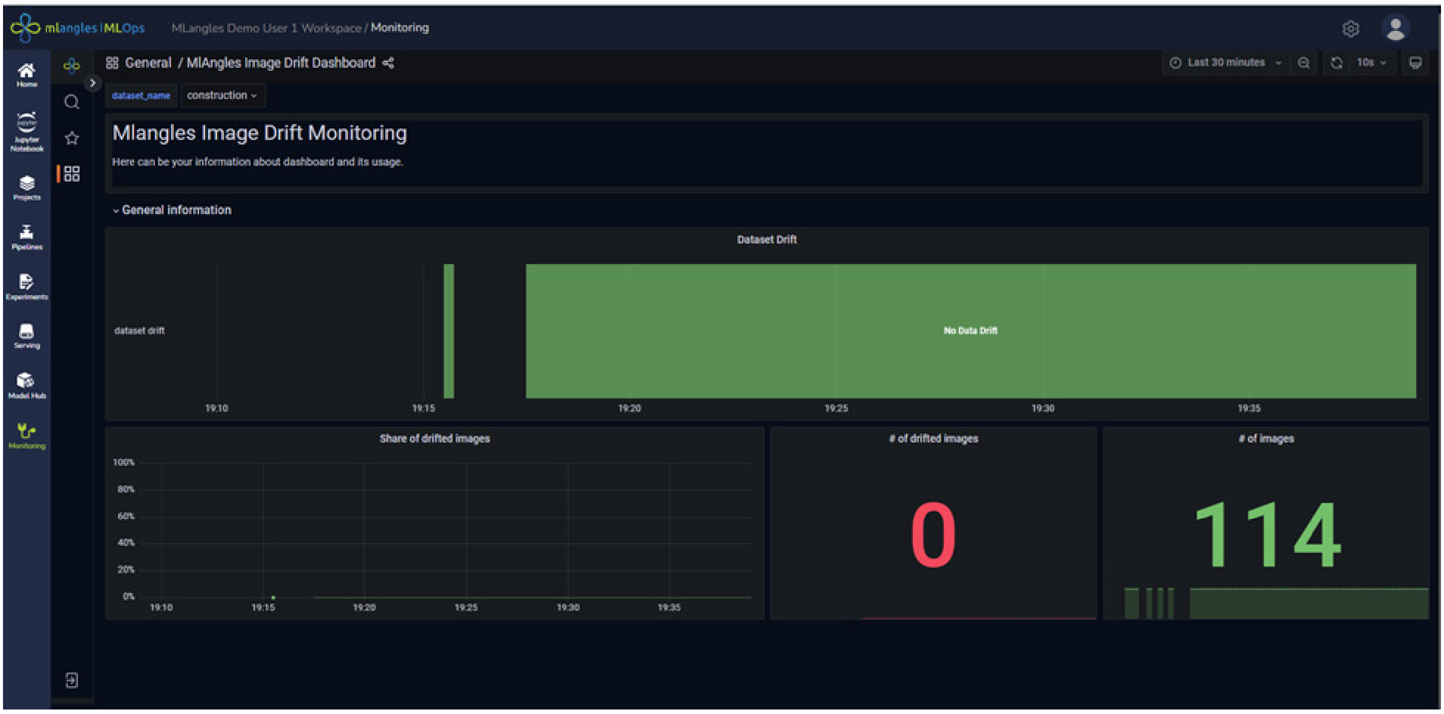

Conclusion

In conclusion, this use case demonstrates the significant potential of deep learning models for improving safety in the construction industry. By leveraging pretrained models like YOLO NAS through transfer learning, these models can efficiently identify and classify types of safety equipment's, such as Hardhat, Mask, NO-Hardhat, NO-Mask, NO-Safety Vest, Person, Safety Cone, Safety Vest, machinery and vehicle, with high accuracy and efficiency. This approach reduces the need for extensive labeled data and computational resources while improving the speed and accuracy of safety equipment assessment.

To setup a demo

Info.mlangles@cloudangles.com

Visit: www.mlangles.ai