CloudAngles

mlangles

mlangles Predictive AI

Pharmaceutical Drug Identification using 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.

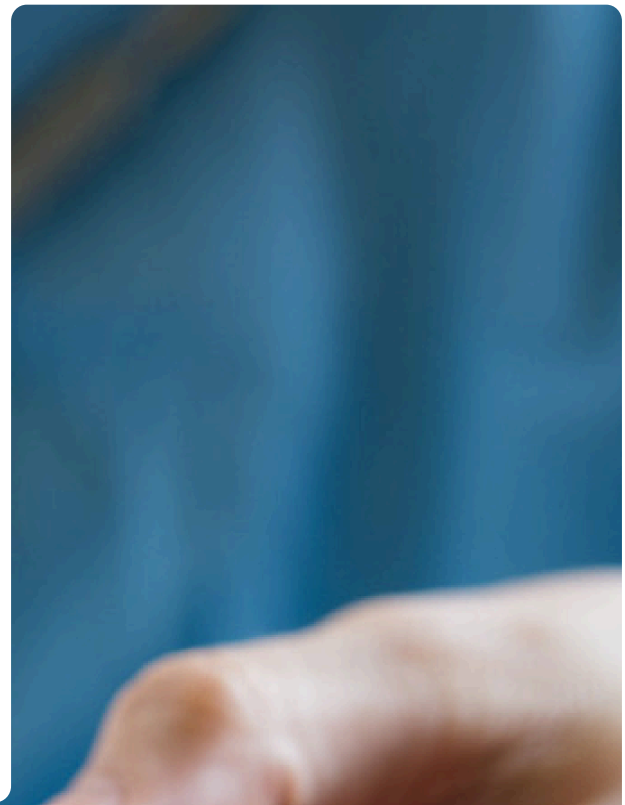

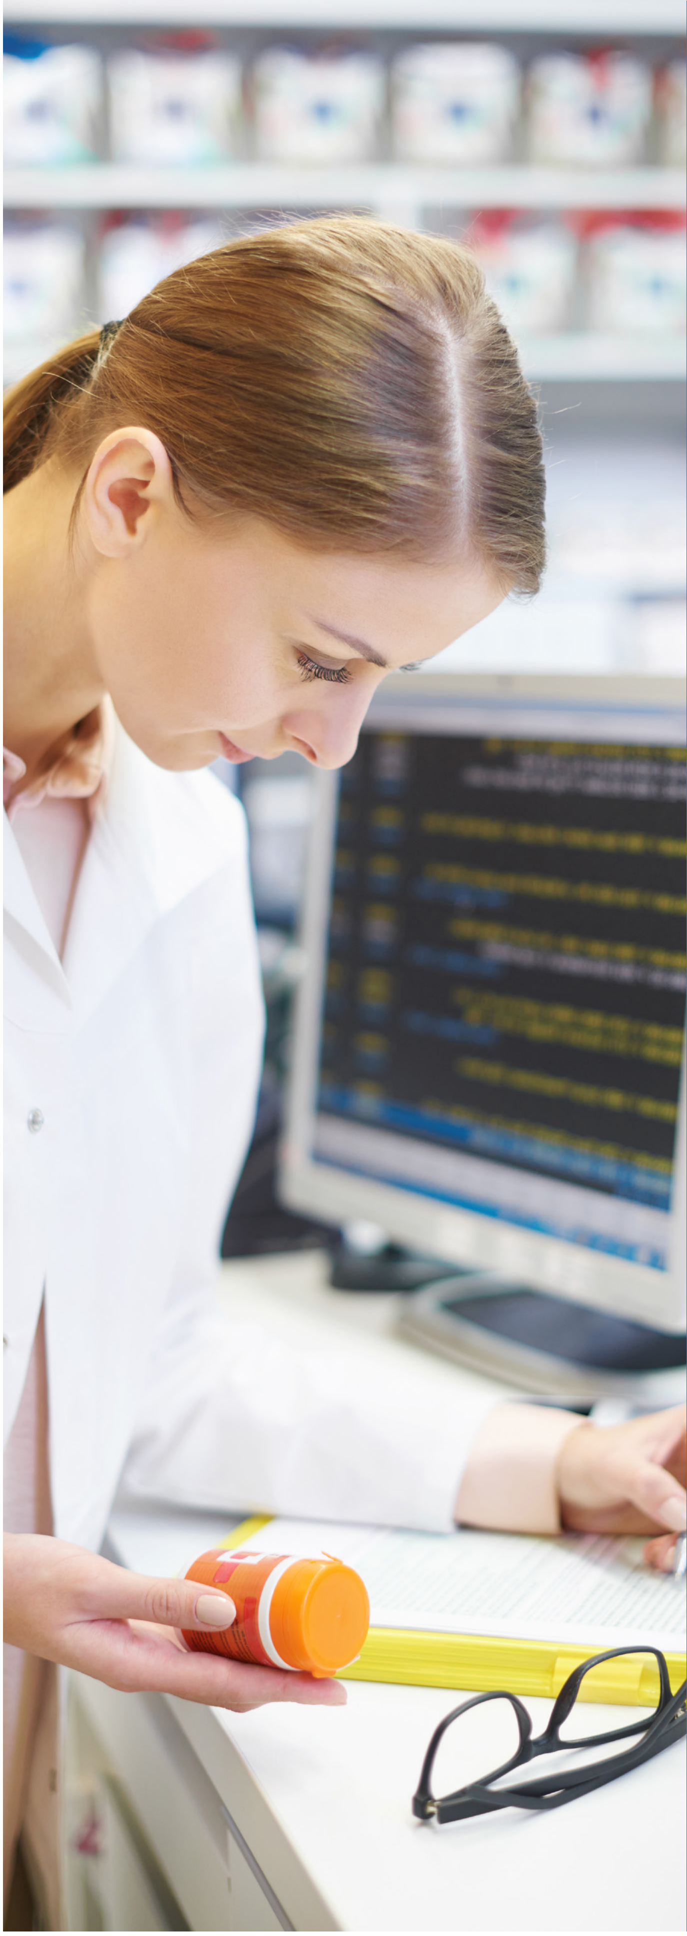

Objective of the Use Case

To develop a deep learning model to accurately identify pharmaceutical drugs, thereby reducing unnecessary mishaps. The dataset features 10 different classes of pharmaceutical drugs. Using classification methods, the developed model predicts the drug type, aiding healthcare professionals in making precise identifications.

Working of the Use Case

- The AI problem will be tackled through a phased approach, starting with the data engineering phase utilizing the pipeline module.
- The modelling process will follow, with the experiment tracking module aiding in the selection of suitable hyperparameters.
- Subsequently, the model will be trained and executed on the provided dataset.
- Predictions generated by the model will be presented using the serving module.
- Continuous monitoring will be implemented to maintain the accuracy and effectiveness of the model over time.

Explanation of the Use Case

In both the USA and South Korea, pharmacists spend an average of 20 hours per month identifying medications at university hospitals. However, approximately 25% of medications remain unidentified using conventional methods. This inefficiency persists even when considering a broader city-wide scale.

Developing a deep learning model to accurately identify pharmaceutical drugs offers significant benefits across the healthcare industry. For healthcare professionals, utilizing such a model ensures precise identification of drugs, thereby reducing errors in medication administration and improving patient safety. This capability can streamline workflows, enhance efficiency in clinical settings, and potentially lower healthcare costs by minimizing medication-related mishaps.

Presently, the dataset comprises ten distinct groups of pill images:

Bactidol: Bactidol is a common antiseptic mouthwash used to kill bacteria and reduce mouth odour. It's often used to help treat sore throats and mouth infections.

Dayzinc: Dayzinc is a vitamin supplement that contains zinc. Zinc is important for our immune system and helps our body fight off illnesses. Dayzinc is often taken to help prevent and treat colds and other infections.

Kremil S: Kremil S is an antacid medication used to relieve heartburn, indigestion, and upset stomach. It works by neutralizing excess stomach acid to provide relief from discomfort.

Fish oil: Bactidol is a common antiseptic mouthwash used to kill bacteria and reduce mouth odor. It's often used to help treat sore throats and mouth infections.

Decolgen: Decolgen is a combination medication used to relieve symptoms of the common cold, flu, allergies, and sinusitis. It typically contains a decongestant to relieve nasal congestion, an antihistamine to reduce allergy symptoms, and a pain reliever to alleviate aches and pains.

Alaxan: Alaxan is a pain reliever medication that contains ibuprofen and paracetamol. It is commonly used to relieve mild to moderate pain, such as headaches, muscle aches, toothaches, and menstrual cramps.

Neozep: Neozep is a cold and allergy medication that contains a combination of an antihistamine and a decongestant. It helps relieve symptoms such as nasal congestion, sneezing, runny nose, and itchy or watery eyes associated with allergies and the common cold.

Biogesic: Biogesic is a brand of paracetamol, which is a common pain reliever and fever reducer. It is used to alleviate mild to moderate pain and reduce fever associated with various conditions, such as headaches, muscle aches, toothaches, and colds.

Bioflu: Bioflu is a cold and flu medication similar to Decolgen. It contains a combination of active ingredients to relieve symptoms like nasal congestion, runny nose, sneezing, sore throat, fever, and body aches.

Medicol: Medicol is a non-steroidal anti-inflammatory drug (NSAID) that contains ibuprofen. It is used to relieve mild to moderate pain and reduce inflammation associated with conditions such as arthritis, menstrual cramps, muscle aches, and headaches.

For image classification, we will use these descriptions to fine-tune a pre-trained model to recognize each drug from its picture. We'll teach the computer to distinguish between the different pills based on their visual features, allowing it to classify new images of the drugs correctly.

Step 1: Data Engineering & Pipeline Creation

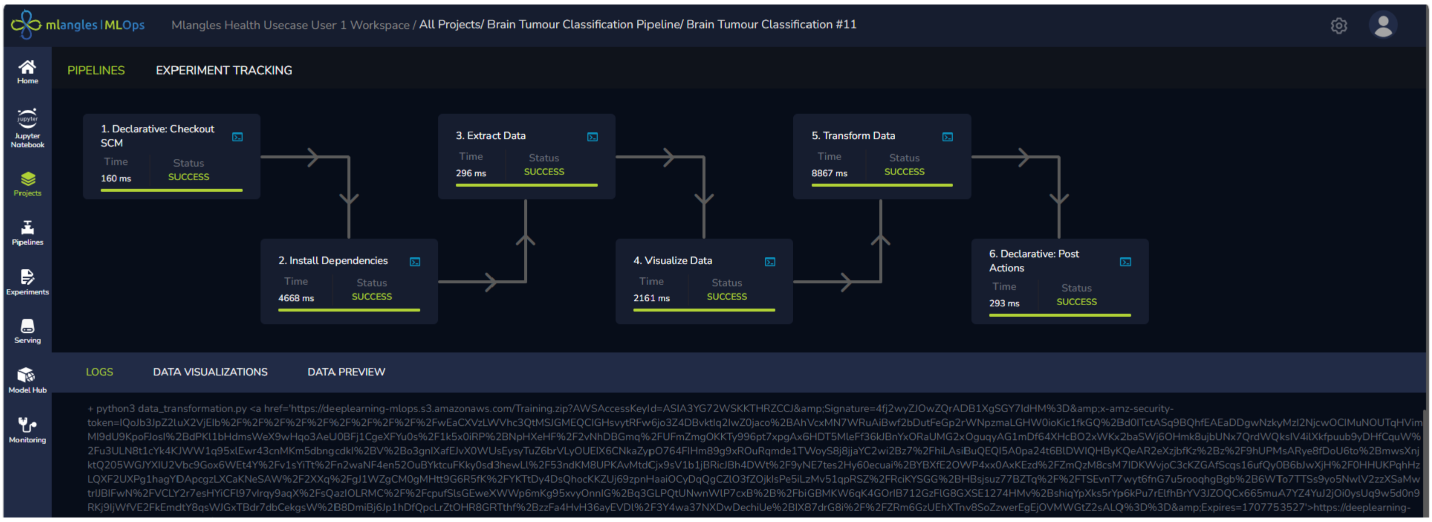

Install Dependencies: Essential packages and libraries have been installed.

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

Data Visualization: One sample image from each of the 10 types of drugs which are bactidol, dayzinc, kermil s, fish oil, decolgen, alaxan, neozep, biogesic, bioflu and medicol are displayed.

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.

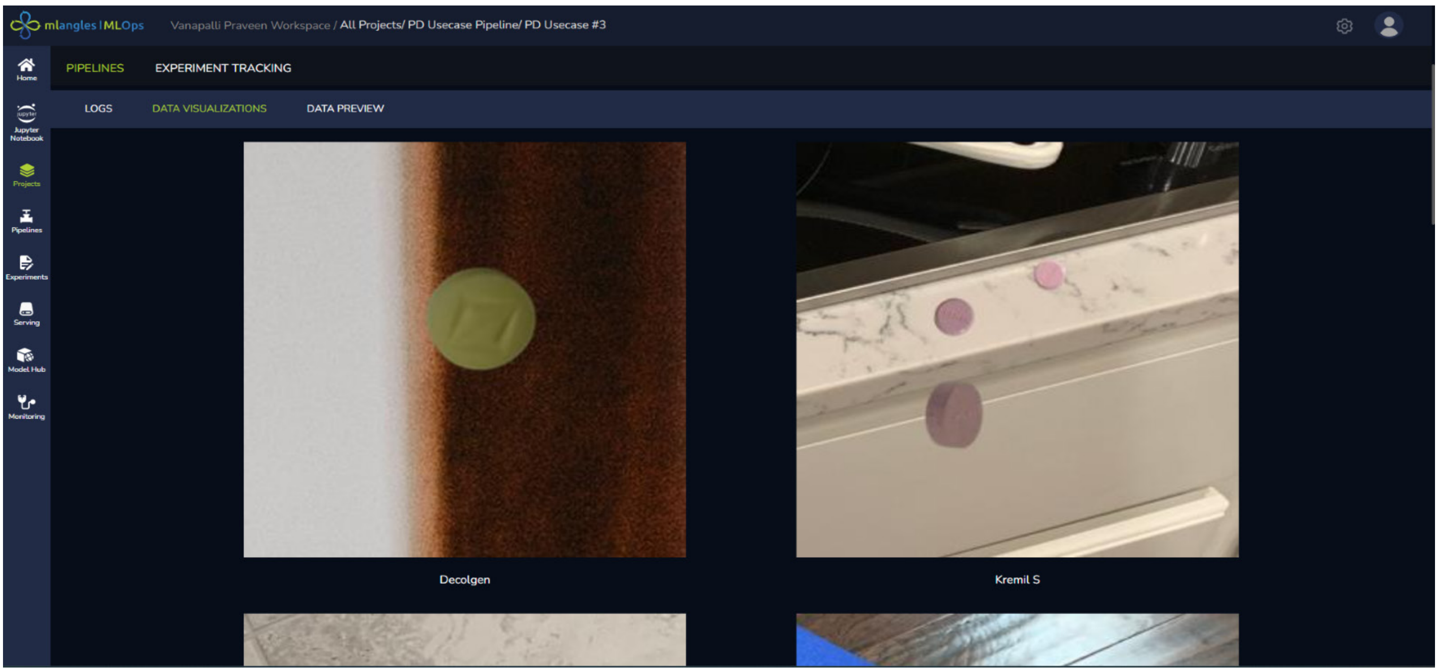

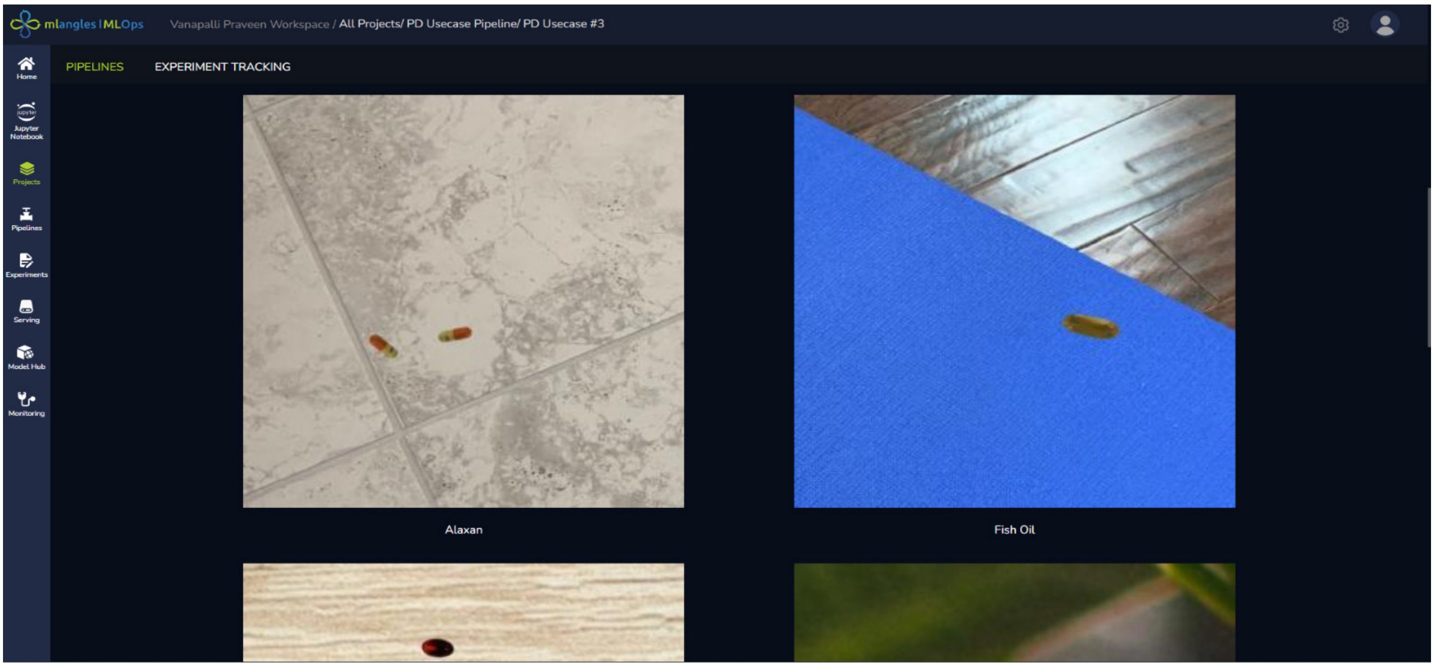

Data Versioning:

- Various processed data versions can be generated through different transformations applied to the same raw dataset, such as deleting columns or applying various transformations on specific columns.
- Throughout the data pipeline, diverse transformations can be executed at each iteration. Consequently, the resulting data at the pipeline's end is systematically versioned.
- Given that each version of the final data is distinct, models trained on these different versions will exhibit varying behaviors.

Step 2: Experiment Tracking – Modelling

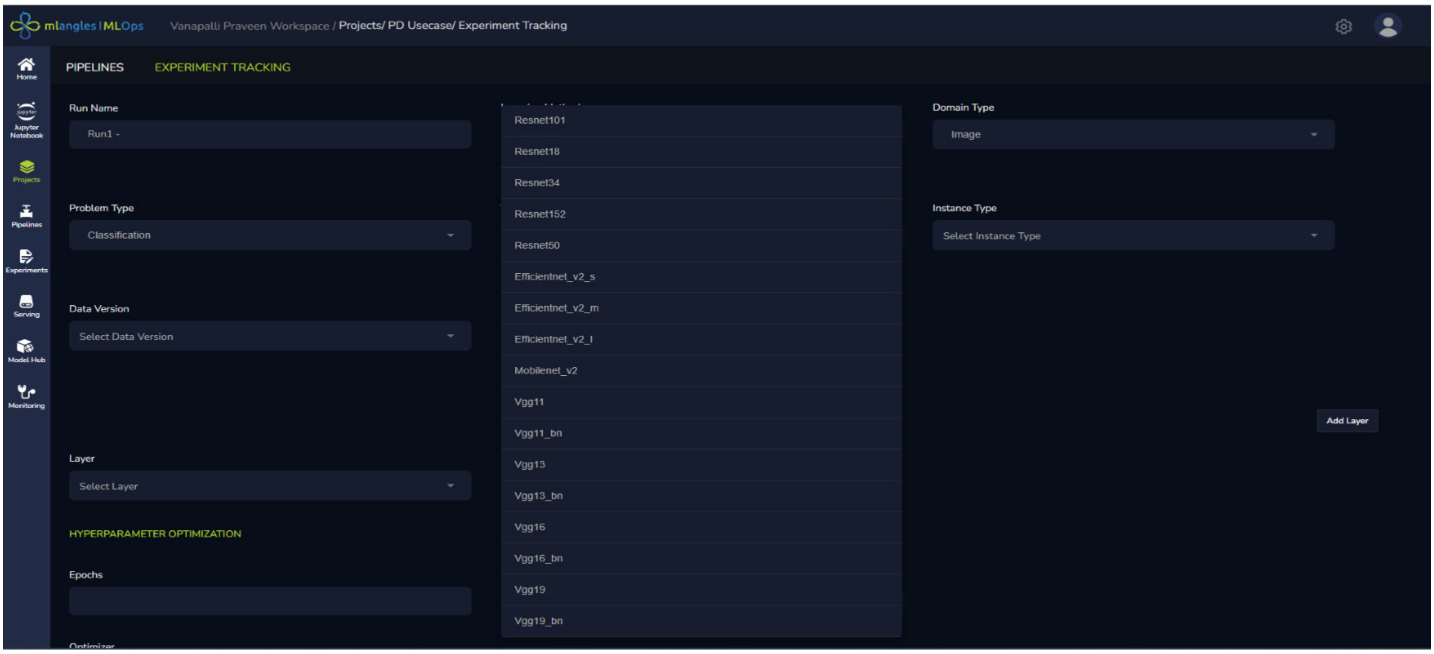

After the dataset is generated through the pipelines, it is utilized to fine-tune any of the various pretrained models such as ResNet, VGG, or EfficientNet. These models have been previously trained on extensive datasets, enabling effective fine-tuning even with a relatively small number of images, resulting in notably good outcomes. Among these models, ResNet stands out as a convolutional neural network (CNN) architecture developed by Microsoft Research. It aims to mitigate the vanishing gradient problem encountered in very deep networks. Noteworthy features of ResNet include residual connections, which allow shortcut connections to skip one or more layers, addressing degradation issues and facilitating the training of exceptionally deep networks.

The ResNet model undergoes training on the current dataset for a specified number of epochs, thereby refining its ability to identify the type of drug depicted in the given image. Once the model has been fine-tuned, the paers utilized for fine-tuning, along with the results of the fine-tuning process, are presented in the form of parameters, metrics, and artifacts, respectively.

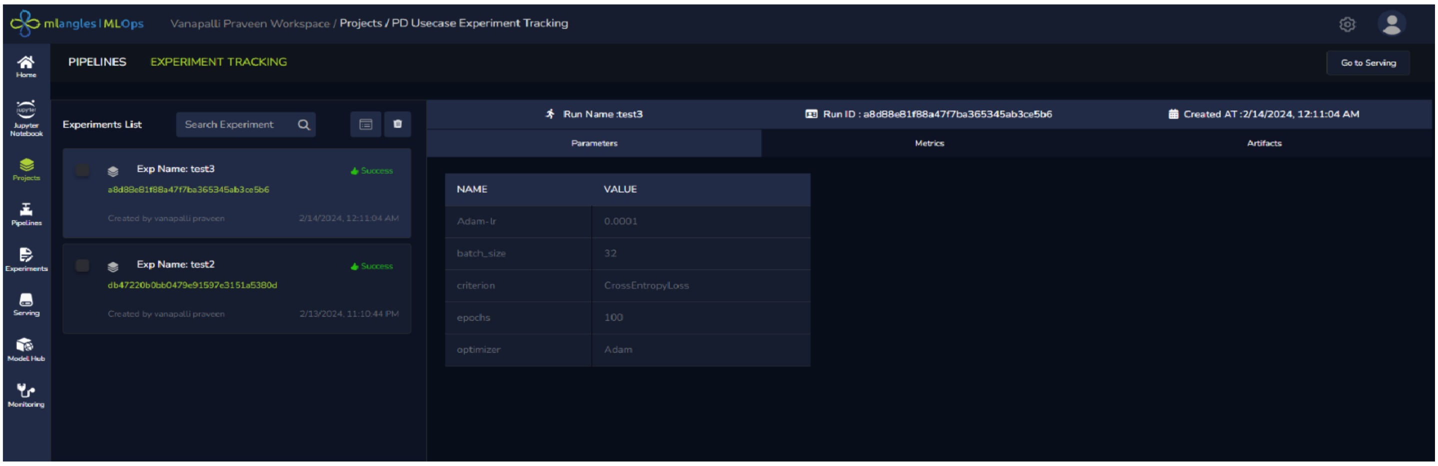

After the experiment ran it gives us detailed information on the parameters used such as batch_size, Criterion, Number of epochs, Optimizer. These parameters play a crucial role in determining the model's performance and behaviour.

Batch Size: The batch size refers to the number of data samples processed in each iteration during training. It affects the model's training speed, memory usage, and generalization ability. Larger batch sizes may lead to faster convergence but require more memory. Smaller batch sizes may provide more stochasticity and regularization but may converge more slowly.

Criterion (Loss Function): The criterion, or loss function, measures the inconsistency between predicted outputs and actual targets. The choice of loss function depends on the problem type (e.g., regression, classification) and the desired behavior of the model.

Number of Epochs: An epoch represents one complete pass through the entire dataset during training.

Optimizer: The optimizer updates the model parameters during training to minimize the loss function. Common optimizers include SGD (Stochastic Gradient Descent), Adam, RMSprop, etc.

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.

Best Validation Accuracy: The best validation accuracy represents the highest accuracy achieved by the model on the validation set during finetuning. Achieving a high validation accuracy implies that the model has learned meaningful patterns from the training data and can successfully generalize to new instances.

Least Validation Loss: The least validation loss refers to the lowest value of the loss function (e.g., Cross-Entropy Loss) recorded on the validation set during finetuning.

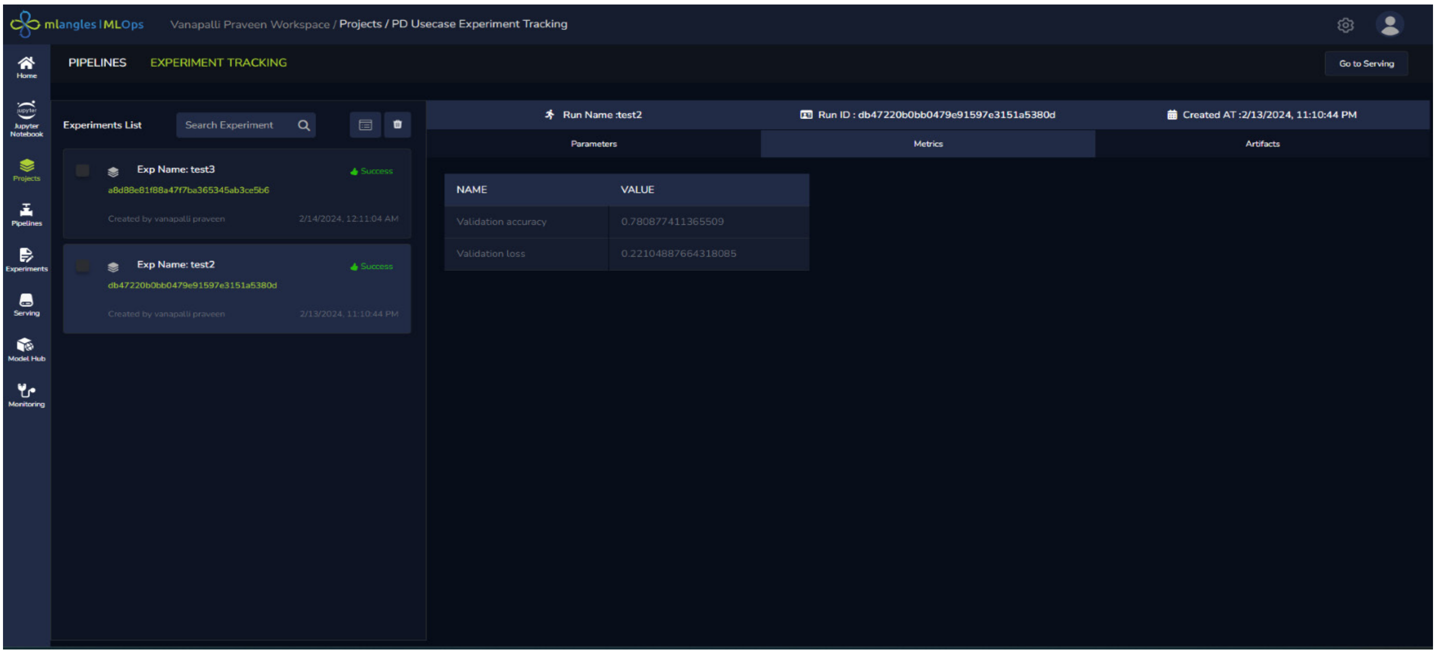

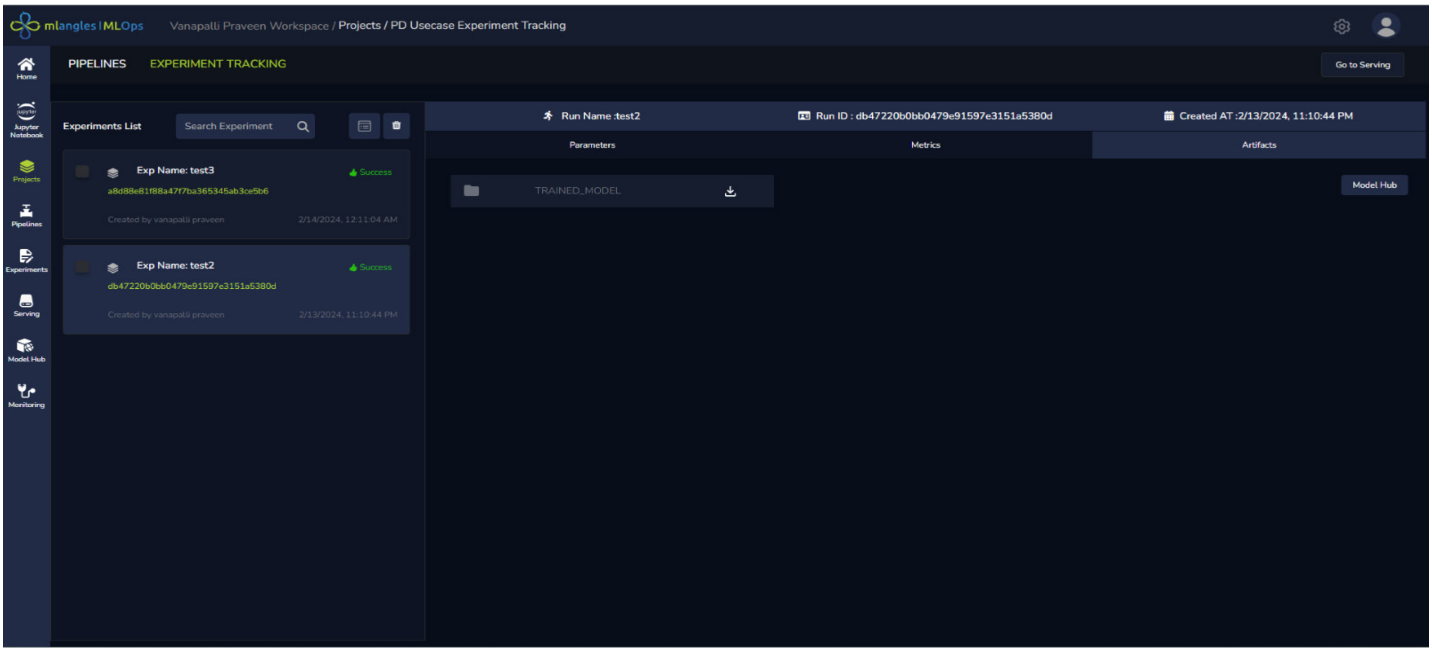

The trained model along with its dependencies could be downloaded from artifacts to perform predictions offline.

Model Versioning:

- Models are sensitive to a plethora of hyperparameters and parameters, including learning rate, loss function, and optimizers.
- Consequently, a model selected for training, with both the model and final data versions remaining constant but changes in parameters, may yield differing performance metrics.
- These diverse model versions can be uploaded to the model hub, facilitating the management of multiple iterations and variations.

Step 3: Prediction/ Serving

The image represents Decolgen medication and is fed into the trained model. The model accurately detects the presence and specific type of the drug, with the resulting identification displayed on the screen.

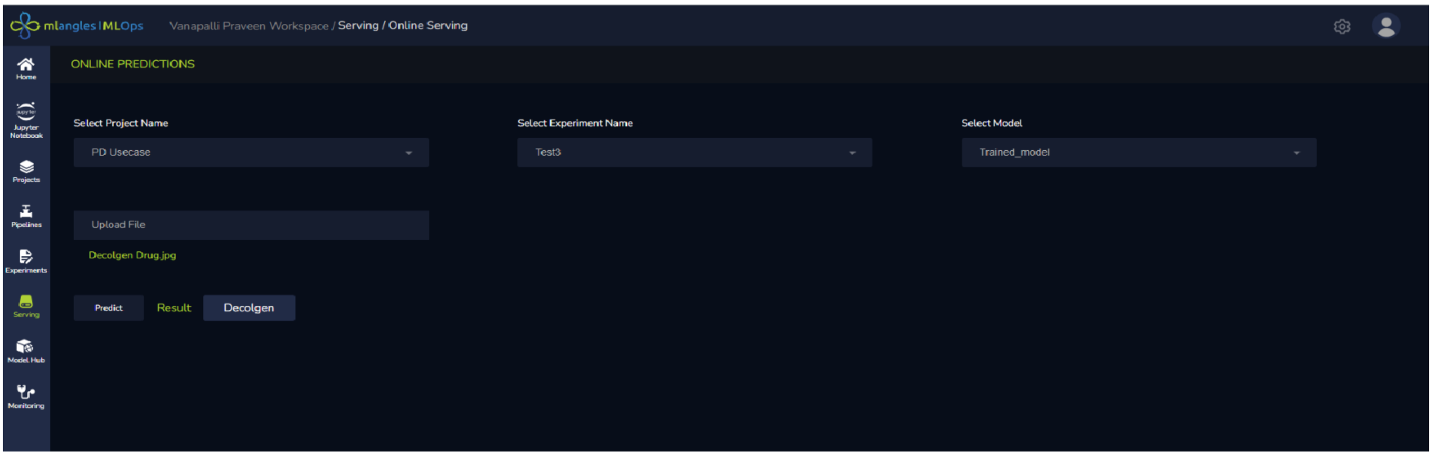

Model Hub:

- Trained models are uploaded to the model hub, whereupon d
- 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.

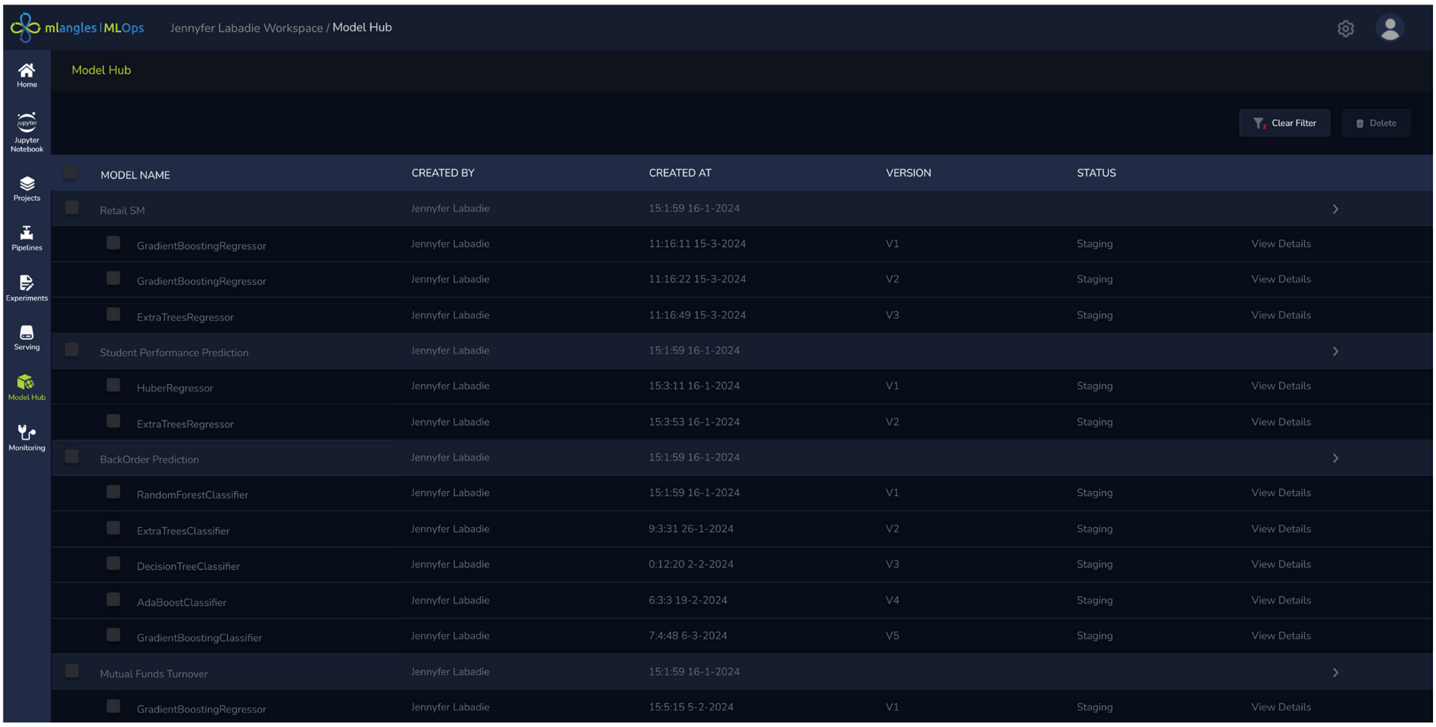

Step 4: Monitoring

IData drift refers to the phenomenon in which the statistical characteristics of images within a dataset change over time, potentially leading to a decline in the performance of machine learning models trained on that dataset. This alteration can stem from various factors, including modifications in imaging equipment, fluctuations in lighting conditions, variances in image acquisition methods, or shifts in the attributes of the target population undergoing imaging.

In medical imaging, for instance, image drift may manifest due to updates or modifications in imaging protocols, alterations in patient demographics, or changes in the hardware or software of imaging devices. When a model is trained on data that does not adequately represent the current image distribution, its performance may suffer when applied to new data, as it may struggle to generalize effectively to the altered distribution.

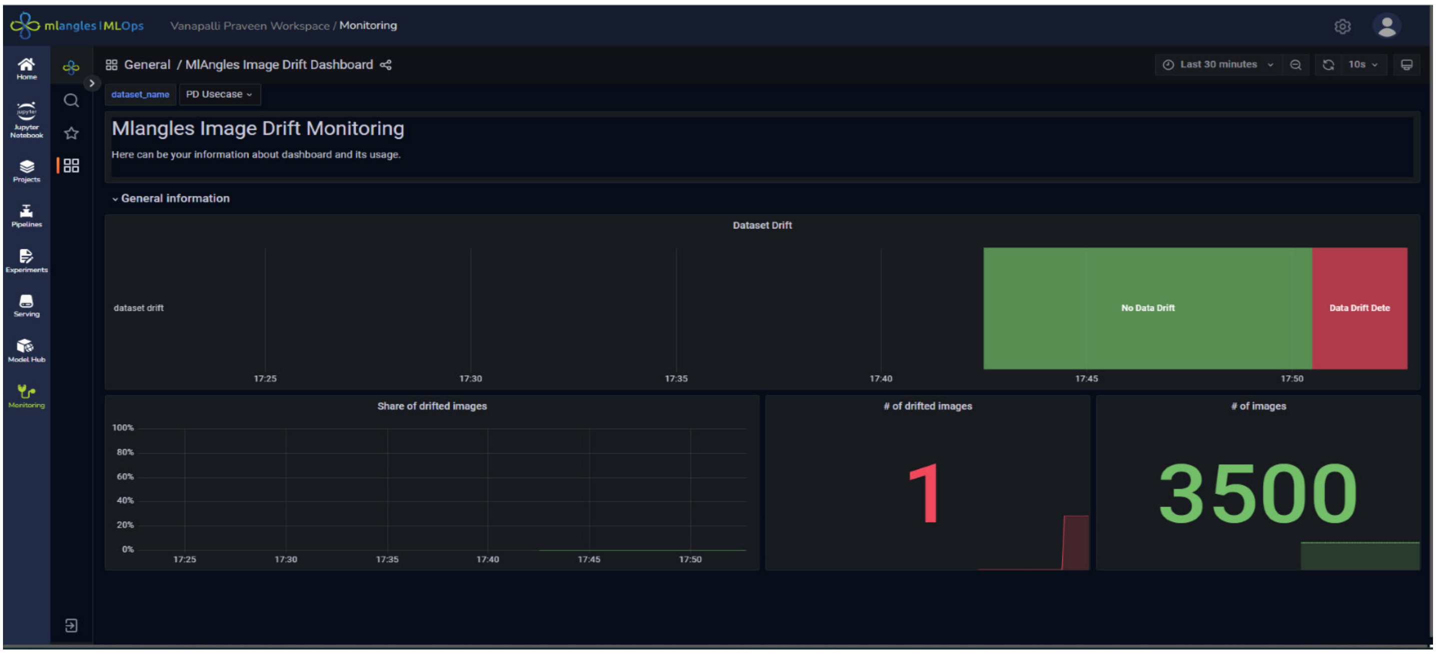

Conclusion

In conclusion, employing pretrained models such as ResNet for fine-tuning on our dataset has yielded promising outcomes in predicting 10 distinct types of drugs: Bactidol, Dayzinc, Kermil S, Fish oil, Decolgen, Alaxan, Neozep, Biogesic, Bioflu, and Medicol, and effectively identifying them. This project highlights the effectiveness of transfer learning in medical image analysis, demonstrating its potential for accurate and efficient drug identification. Our results emphasize the importance of integrating advanced machine learning techniques into healthcare practices, paving the way for enhanced drug capabilities and improved patient outcomes in healthcare.

To Setup a demo

info.mlangles@cloudangles.com

Visit : www.mlangles.ai