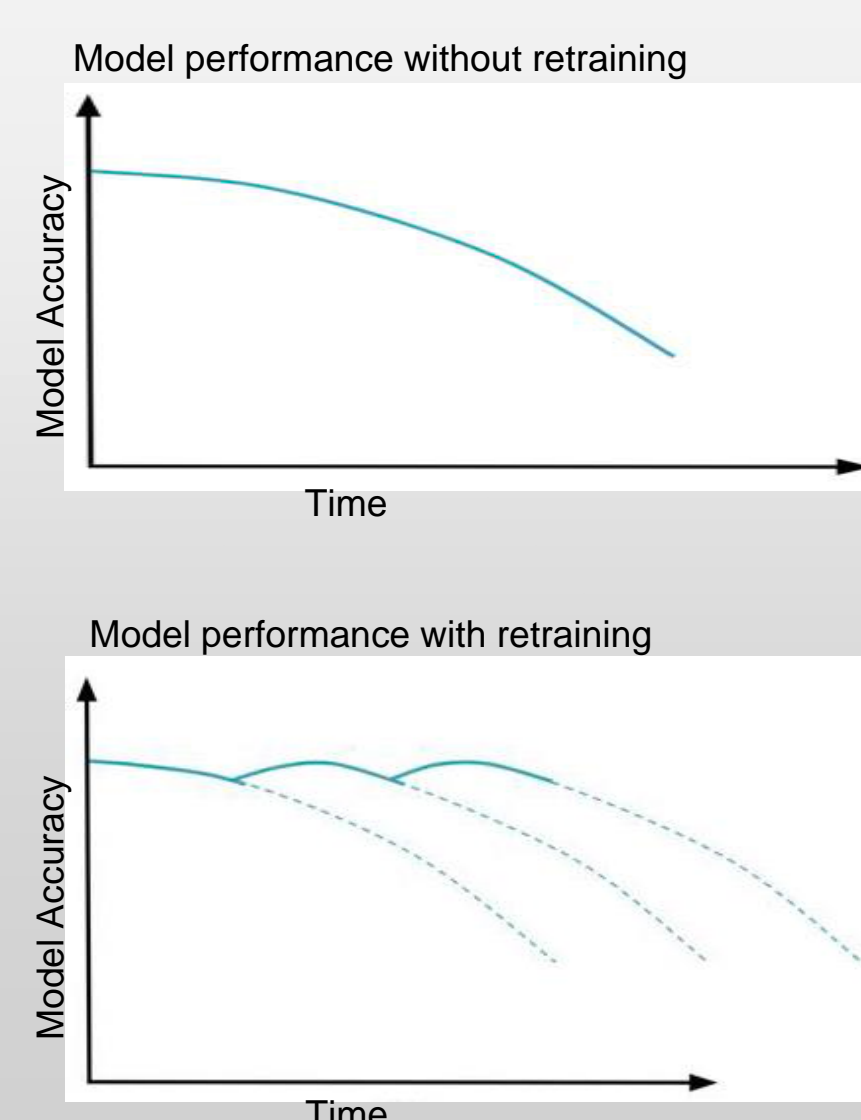


Problem Statement

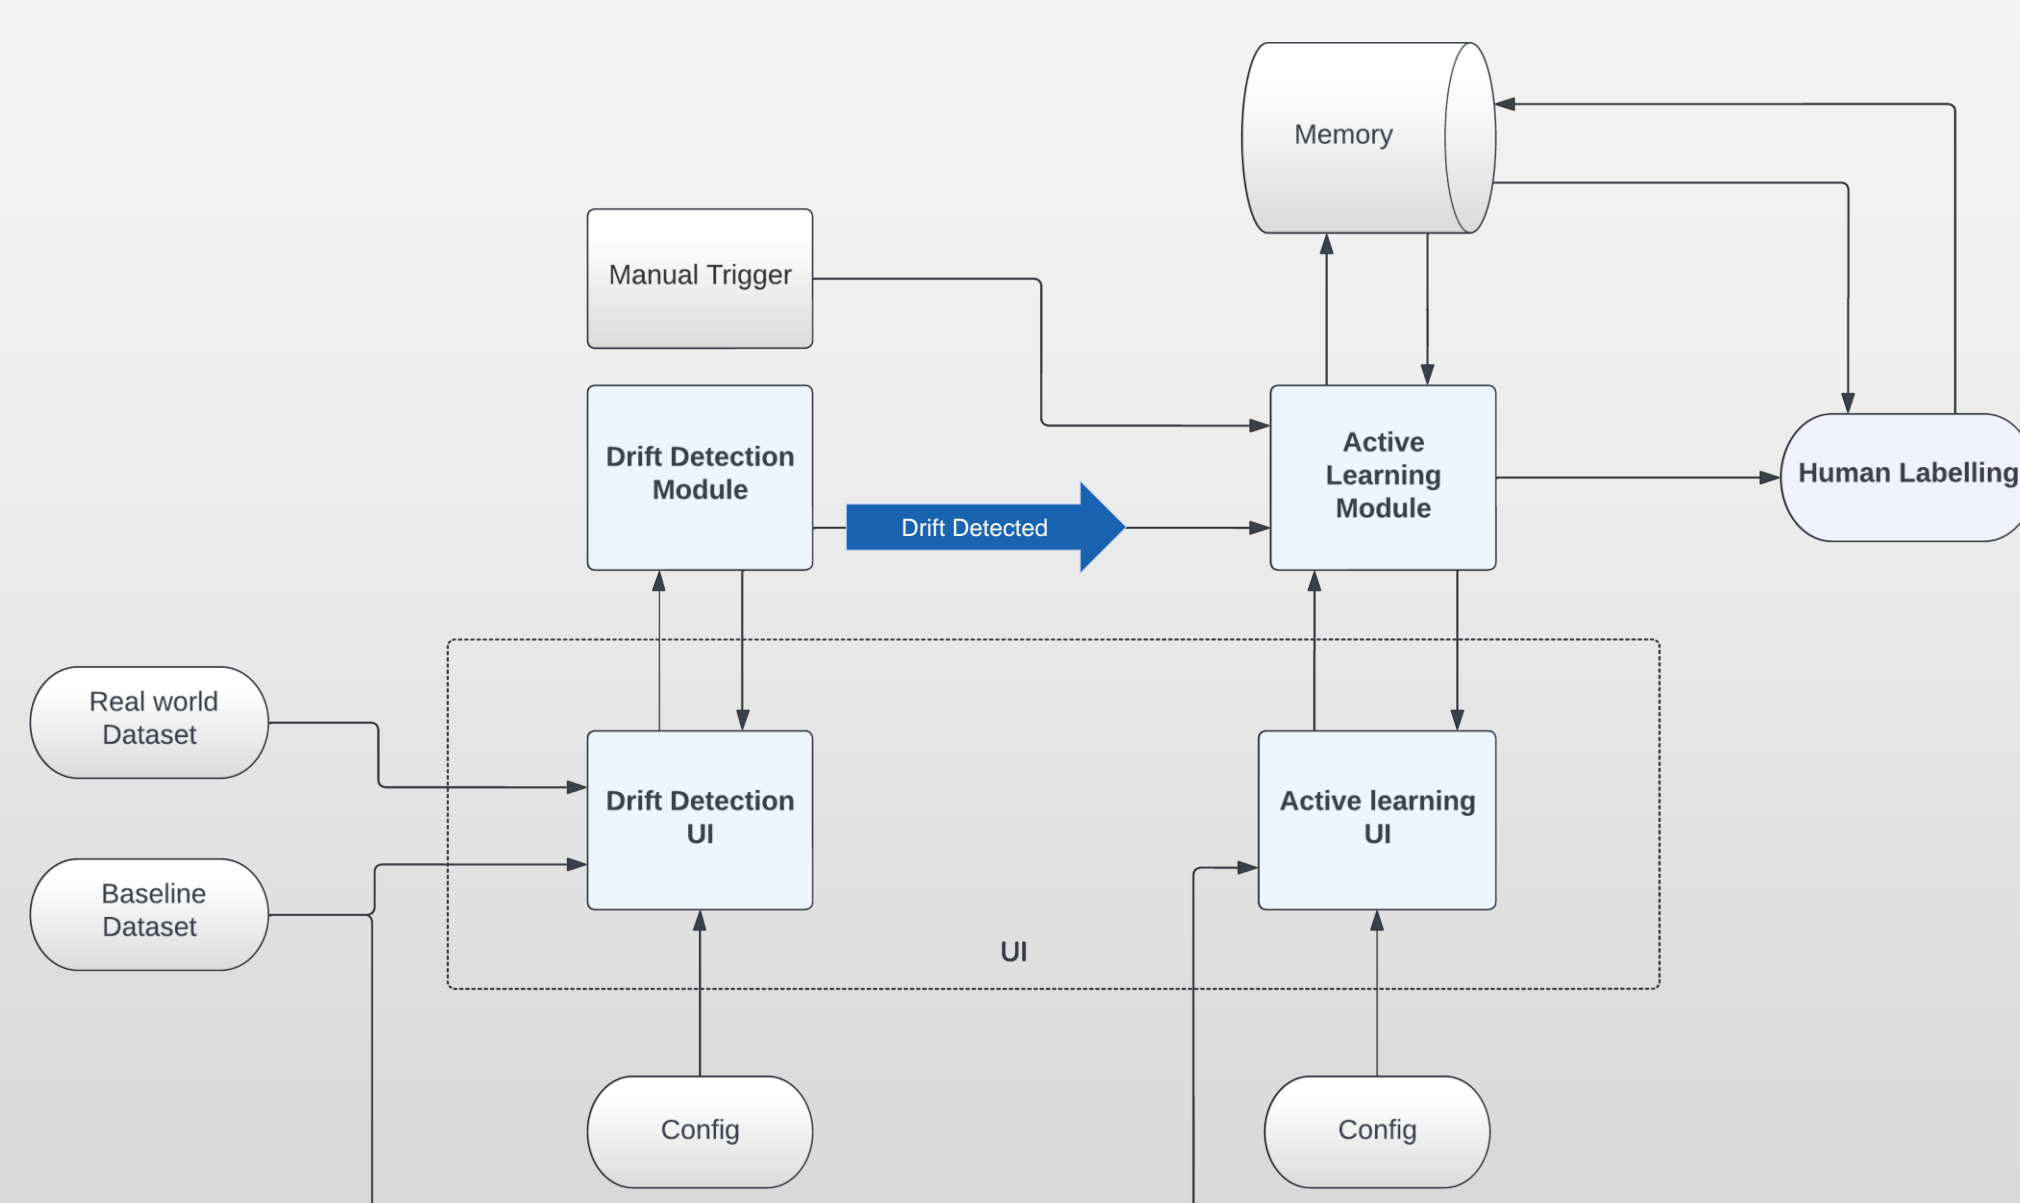
- Data drift is the major reason for the decrease in the accuracy of machine learning models.
- There might be variations in the data distribution during inference time, changes in user behavior compared to the baseline data, or additional factors in real-world interactions that might cause biased predictions.
- Identifying such drifts and automating model retraining ensures that the model remains relevant in production and gives unbiased predictions over time.

Objectives

- To develop a drift detection module (i.e., detecting reduction in the predictive power of an ML model)
- To develop a retraining module for automatically updating the model based on the conditions in the production environment.
- To build a UI that monitors model performance and shows baseline performance, current deviation, measurement criteria, retraining policies, and new model candidates.



Prototype elements



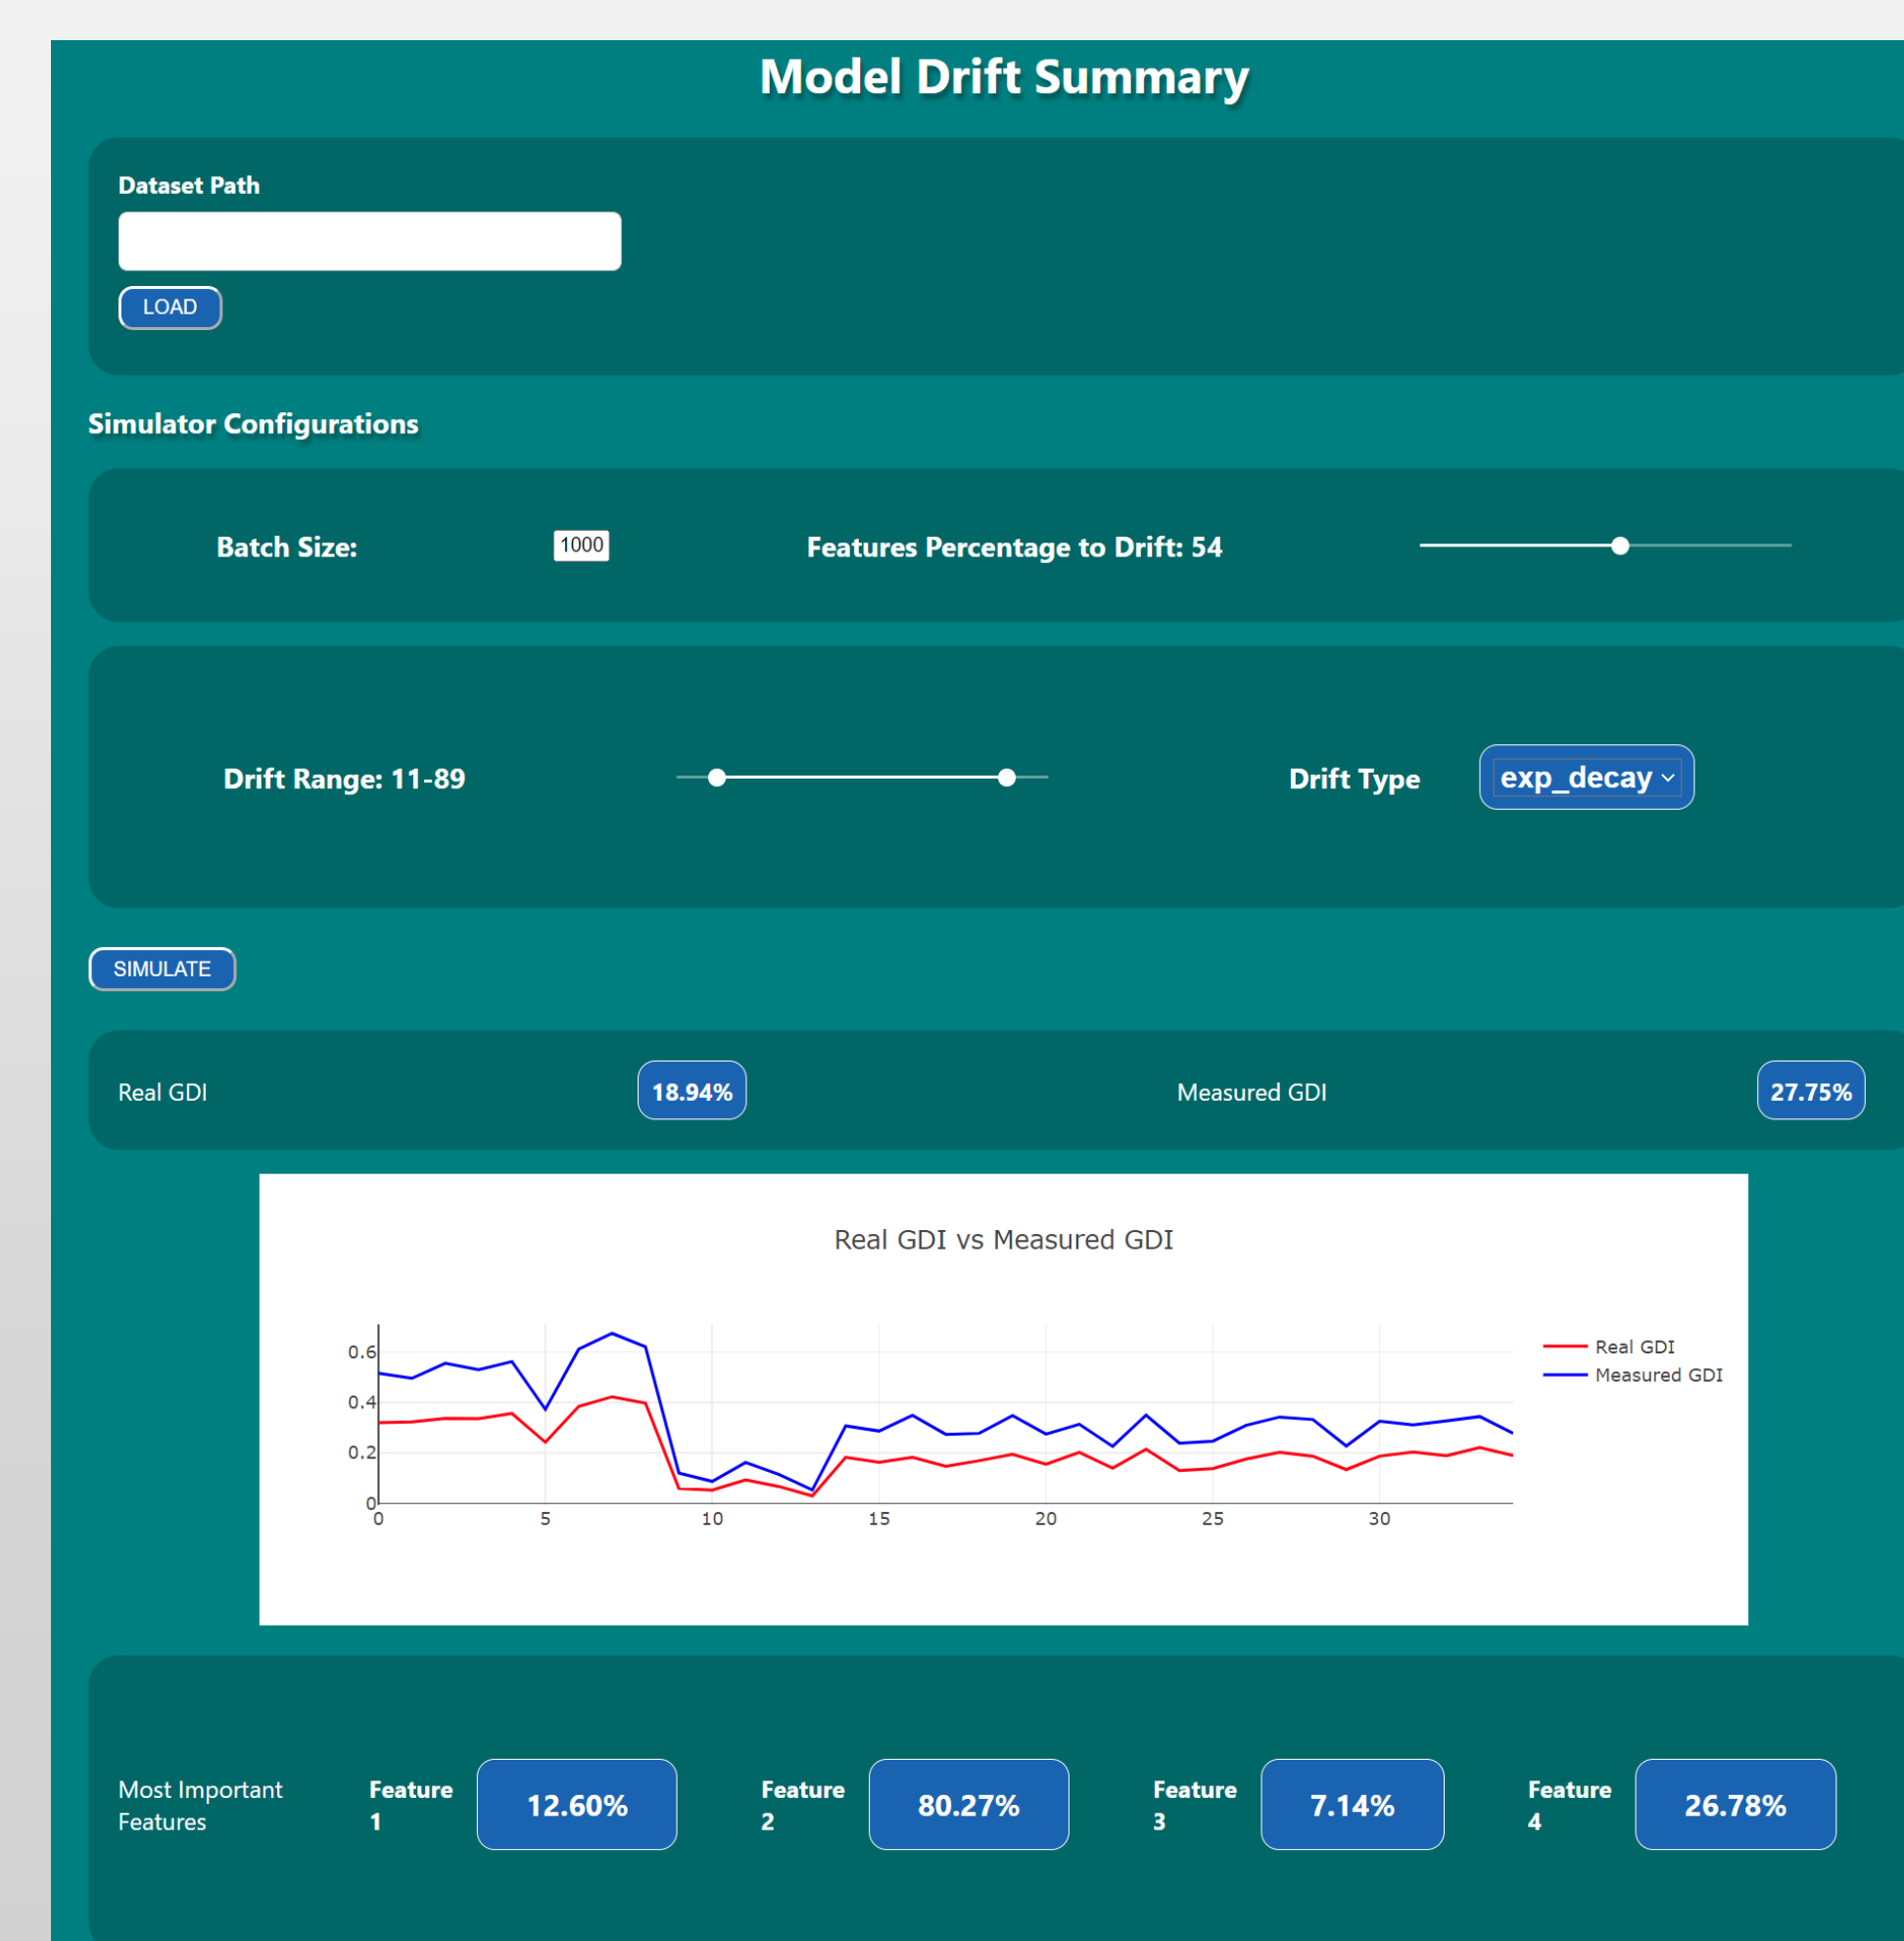
- The drift detection and active learning modules communicate with the frontend and backend services using REST APIs.
- The drift detection module triggers the retraining process by sending the drifted data to the active learning module.
- A human annotator asynchronously labels the most interesting drifted datapoints.

Global Drift Index (GDI)

- The Global Drift Index (GDI) is a novel aspect of the project and is a single number that denotes the percentage drift in the dataset.
- The steps to calculate the proposed Global Drift Index (GDI) are as follows:
 - Find the drift index for every feature using the Chi-square test.
 - Find the feature importance (weights) using Principal Component Analysis (PCA).
 - $GDI = \text{Weighted Average of the drift indices for each feature in the dataset.}$

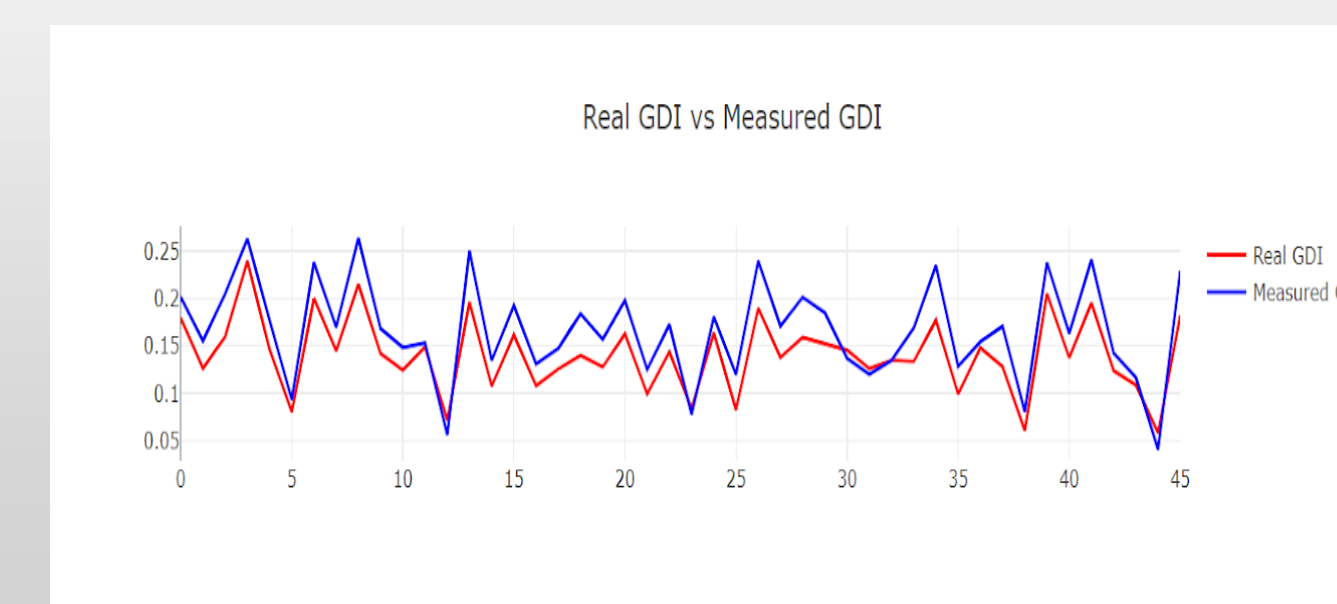
Drift detection

- Based on the GDI threshold set by the user, drift is detected, and the active learning module is triggered.



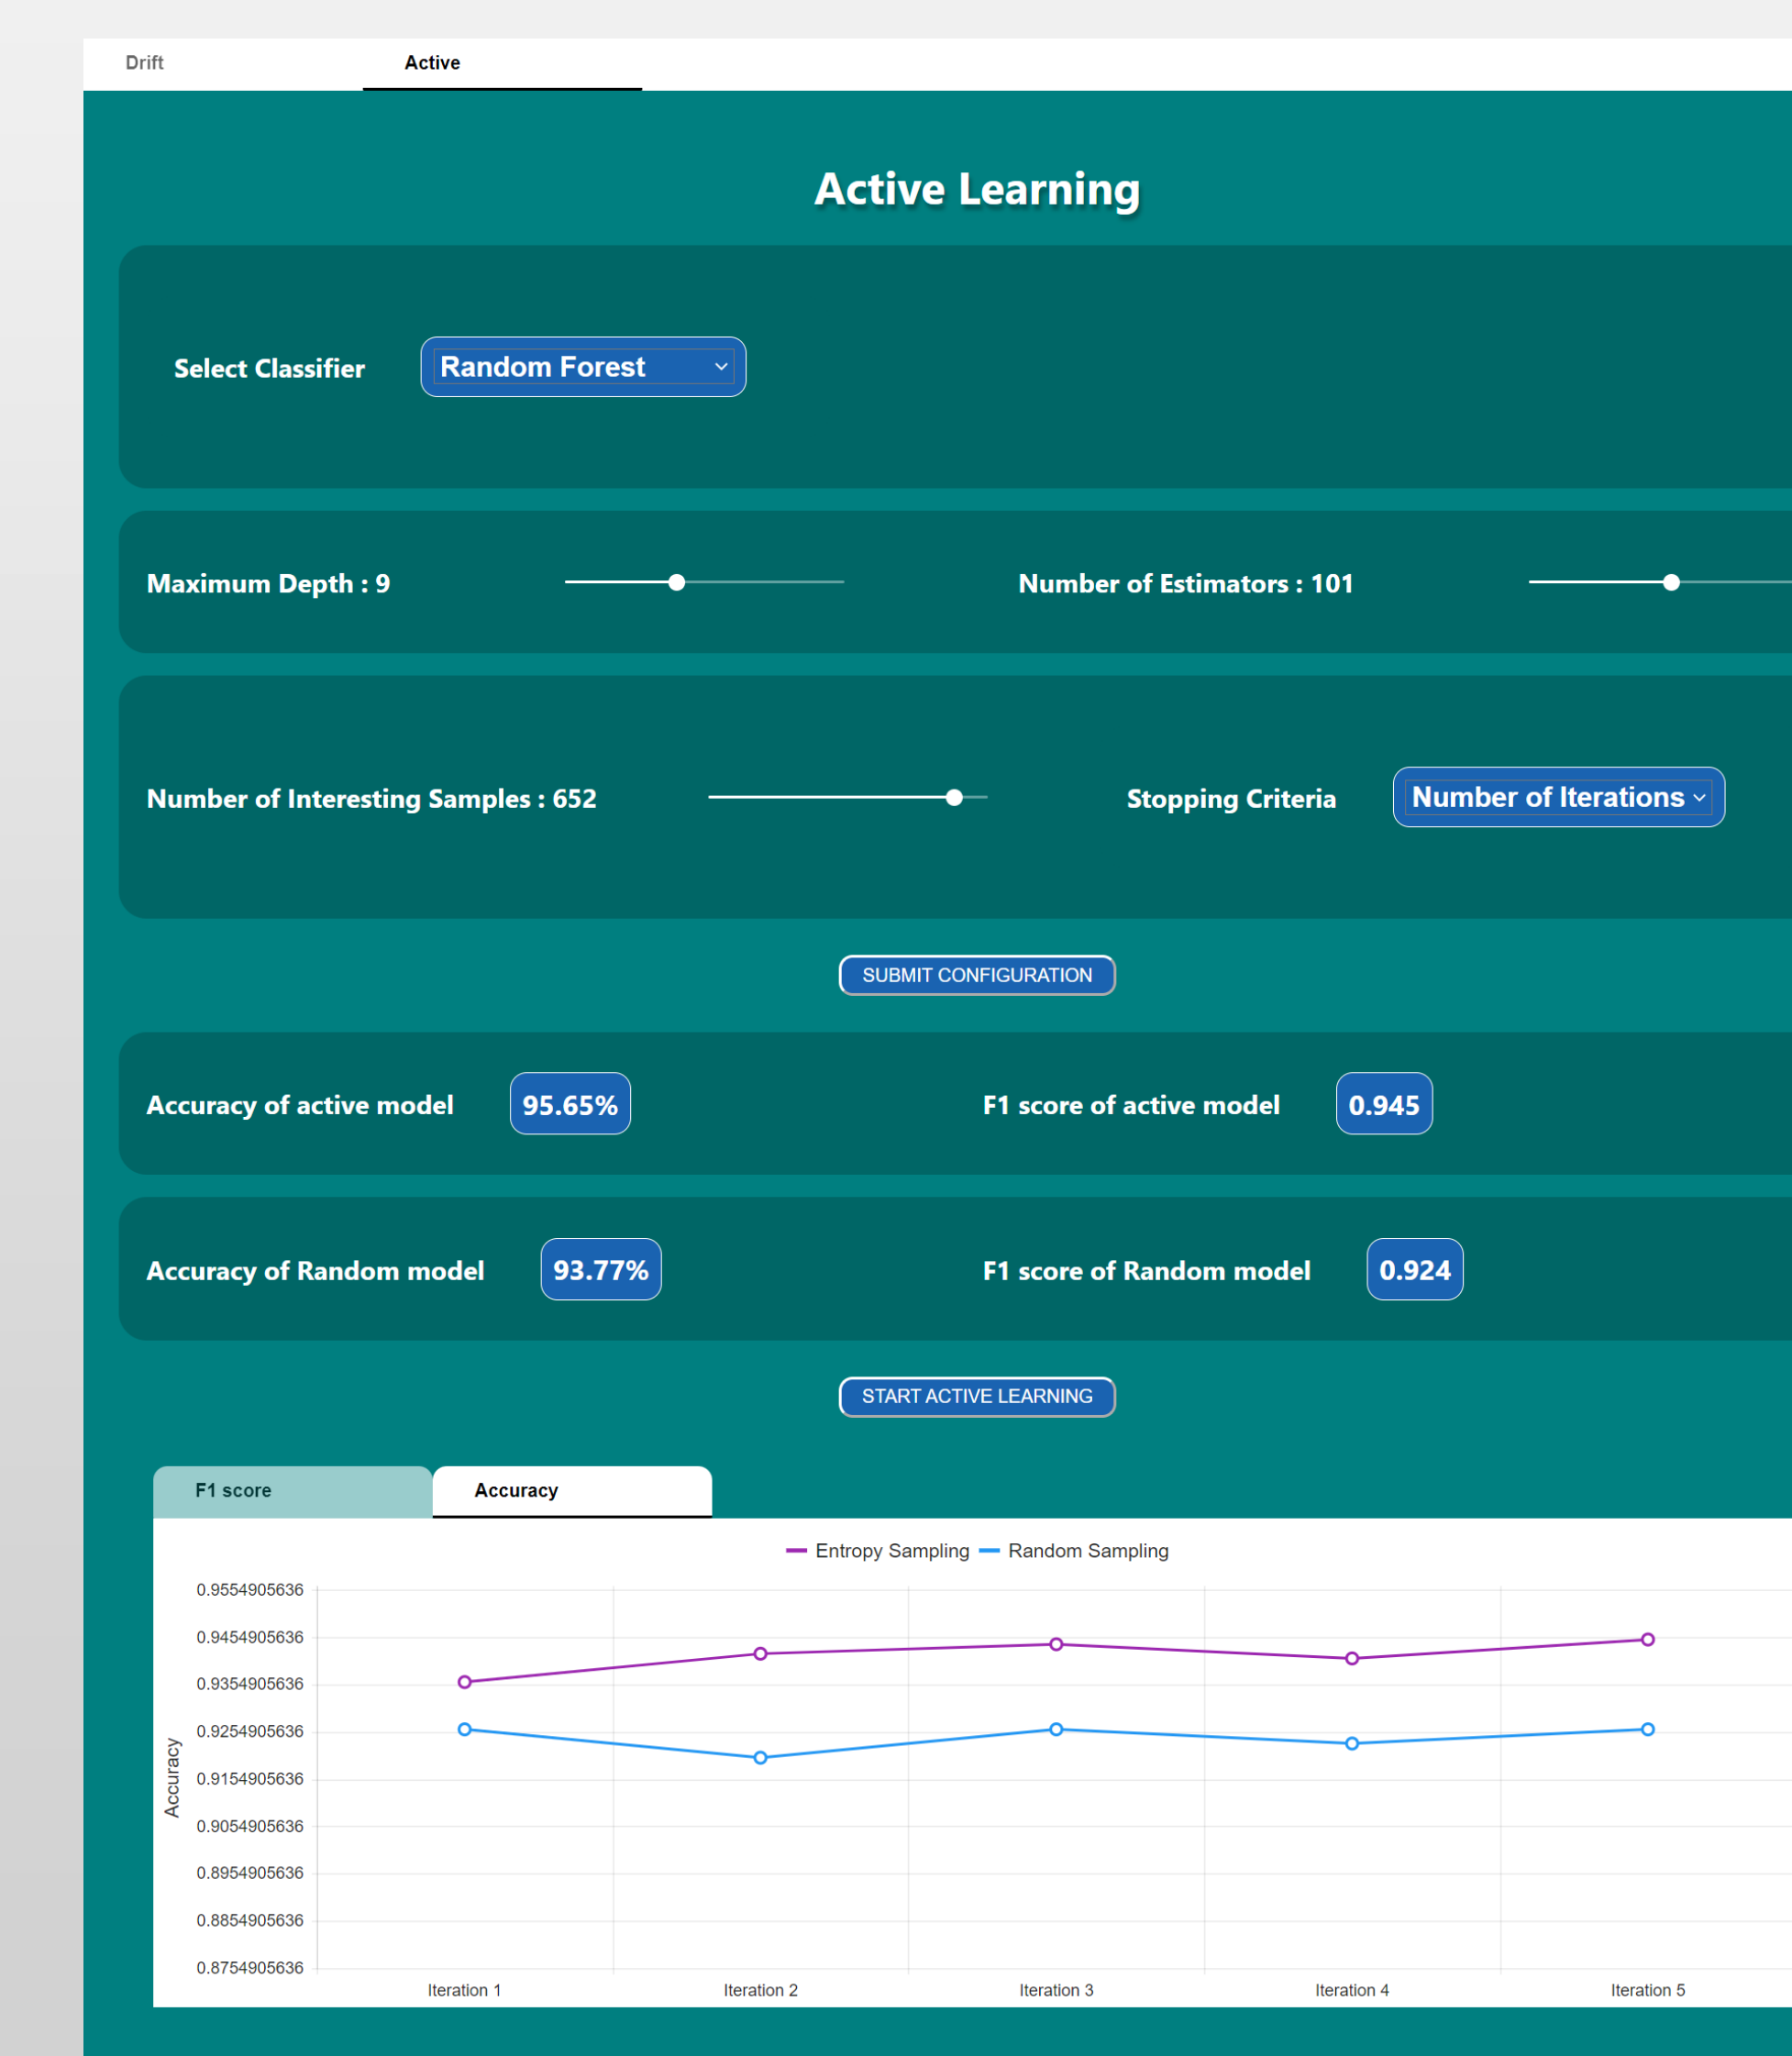
Real-world simulator

- The drift detection module is tested using a simulation module.
- The simulator sends batches of data with known values of drift to the drift detector.
- The measured values of drift (from the detector) are then compared to the real values (from the simulator).



Active Learning

- The fundamental idea behind the active learning concept is that an ML algorithm could potentially reach better results while using a smaller number of training labels if it were allowed to choose the data it wants to learn from.
- The “interesting” samples are chosen based on the entropy measure of the classifier output probabilities.



Conclusion and Future Work

Our work has been able to successfully:

- Detect reductions in the predictive power of an ML model and automatically retrain the model.
- Measure and represent drift using a single number called Global Drift Index (GDI).
- Reduce the data labeling costs by selectively adding the most interesting drifted samples to the training data.

The next steps for the project will be:

- Add modules to perform drift detection for prediction drift and active learning for regression and anomaly detection tasks.
- Detect drift due to changes in the relationship between attributes and changes between old and new sections of time series data.
- Integrate the implemented modules to Tupl's machine learning toolkit.

