Abstract

The focus of this 4-month-long project was to implement advanced machine learning modeling techniques to detect fraud score anomalies in transaction data collected from major banks worldwide. Experimentation was conducted using Support Vector Machines, Local Outlier Factor, DBScan, Isolation Forest, Classifier-Adjusted Density Estimation (CADE), and Autoencoders. The latter three models were selected as models of investigation based on their effectiveness in identifying anomalies in complex datasets.

The models were run on datasets of transaction date, transaction fraud score, and the count of said scores. After preprocessing the data, we engineered temporal features to feed into our model to improve our results. Such features include frequencies of scores relative to the past 3 days, week, 2 weeks, month; and to the same day the previous week, previous 2 weeks, and previous 3 weeks. Feeding these new features into the model resulted in consistent results showing anomalies at shifts within the data.

Our studies showed that the Isolation Forest model was proficient in classifying anomalies, but presented us with the problem of not understanding why certain points were anomalous due to its binary output. We found that Classifier-Adjusted Density Estimation (CADE) model captured more global movements, as it implements discriminative models and has a focus on decision boundaries between dense and sparse regions. Because of this, we found it captured important shifts. However, we found the model not to be practical for monitoring given the large batch data processing jobs that would be required.

We found that autoencoder models identified sufficient anomalies based on the reconstruction error, identifying local shifts, but also can be run on out-of-time data (practical for production). Despite a simple structure, our autoencoder provides us with useful results, and practicality when it comes to deployment of a monitoring solution.

Yusuf Morsi

Data Science Intern

© 2023 FICO. All rights reserved.