Deep Learning — Tiny Case Study 3

The compliance department at XYZ Corp has a system to identify fraudulent customers (customers engaged in fraud). The system is a sequence of SQL statements running against all transactions.

The detected instances are analyzed and manually labeled as being valid or not. Valid instances are passed onto enforcement which may reach out to customers for clarification, close their accounts, or reach out to law enforcement authorities.

Bill, XYZ Corp's new CIO has recently attended a talk by a cloud provider, and is inspired by *Deep Learning* in the cloud platform. A new project is initiated to move the existing system to "Deep Learning".

To start off the project, engineers at XYZ Corp take the labeled data set, split it into training data, and test data, and use the cloud provider's tools to train a deep neural network.

The resulting deep-neural-network model is 99% accurate on the test dataset. Bill declares the project as success, is seen as a visionary, gets a bonus, and is offered a job at the cloud provider.

When the model is used in production, analysts complain that most of the results are erroneous.

Questions:

1. What went wrong? Make a list.
2. How would you make it better (assuming you still wish to use deep-learning).
3. How would you confirm that your approach is better?