



Figure 1. (a) Model training and evaluation scenario: The model's performance is measured using metrics such as accuracy, F1 score, precision, recall, etc on held-out test data. Once the performance is deemed satisfactory, the model is deployed. (b) Deployment scenario: Using gaussian kernel features derived from the trained model's representations, the model is able to assign a score to the incoming test data. (c) The scores assigned to incoming samples are high when the data is very different from the training set distribution (out-of-distribution) and low when it is close to the training-set distribution (in-distribution).