
Professional Certificate in AI for Automated Skin Lesion Analysis

Evaluation Metrics for Skin Lesion Analysis

Accuracy: A measurement of how often the classifier's predictions match the actual labels. It's calculated as the number of true positives and true negatives divided by the total number of samples. While accuracy is a commonly used metric, it may not always be the best choice, especially in imbalanced datasets.

AUC-ROC (Area Under the Receiver Operating Characteristic Curve): A performance measurement for binary classifiers. ROC is a probability curve plotting the true positive rate against the false positive rate. AUC-ROC measures the entire two-dimensional area underneath the entire ROC curve. A higher AUC-ROC indicates a better classifier.

Confusion Matrix: A table used to describe the performance of a classification model on a set of test data. The number of true positives, true negatives, false positives, and false negatives are presented in the matrix. It helps visualize the performance of an algorithm.

F1 Score: A harmonic mean of precision and recall. F1 score reaches its best value at 1 and its worst score at 0. It is a better metric than accuracy for imbalanced datasets because it takes into account both false positives and false negatives.

False Negative (Type II Error): The model incorrectly predicts the negative class. In skin lesion analysis, a false negative might mean that a malignant skin lesion is classified as benign.

False Positive (Type I Error): The model incorrectly predicts the positive class. In skin lesion analysis, a false positive might mean that a benign skin lesion is classified as malignant.

Hausdorff Distance: A measure of similarity between two sets of points. It calculates the maximum distance between each point in one set to the closest point in the other set. In skin lesion analysis, Hausdorff Distance can be used to measure the similarity between segmented lesions and ground truth.

Intra-class Correlation Coefficient (ICC): A measure of the reliability of measurements made by different raters when the true score is unknown. It can be used to assess the agreement between two dermatologists when manually annotating skin lesions.

Intersection over Union (IoU): A measure of the overlap between two bounding boxes. It is calculated as the area of overlap divided by the area of union. In skin lesion analysis, IoU can be used to evaluate the performance of object detection models.

Jaccard Similarity Index: A statistic used to gauge the similarity between two sample sets. It calculates the size of the intersection divided by the size of the union of the sample sets. The Jaccard Similarity Index is a

special case of Intersection over Union (IoU).

K-fold Cross-Validation: A technique for assessing how well a model will generalize to an independent dataset. The original sample is randomly partitioned into k equal sized subsamples. Of the k subsamples, a single subsample is retained as the validation data for testing the model, and the remaining $k - 1$ subsamples are used as training data. The cross-validation process is then repeated k times, with each of the k subsamples used exactly once as the validation data.

Lesion-level Analysis: An evaluation method in skin lesion analysis that assesses each lesion separately, ignoring multiple lesions in the same image.

Matthews Correlation Coefficient (MCC): A metric for binary classification problems that takes into account true and false positives and negatives. It returns a value between -1 and 1, where 1 is a perfect prediction, 0 is no better than random prediction, and -1 is total disagreement between prediction and observation.

Negative Predictive Value (NPV): The proportion of negatives that are correctly identified. It is calculated as the true negatives divided by the sum of true negatives and false negatives.

Precision: A measurement of the relevancy of the model's predictions. It is calculated as the number of true positives divided by the sum of true positives and false positives.

Positive Predictive Value (PPV): The proportion of positives that are correctly identified. It is calculated as the true positives divided by the sum of true positives and false positives.

Receiver Operating Characteristic (ROC) Curve: A graphical representation of the performance of a binary classifier. The ROC curve plots the true positive rate against the false positive rate.

Recall (Sensitivity): A measurement of how well the model identifies the positive class. It is calculated as the number of true positives divided by the sum of true positives and false negatives.

Specificity: A measurement of how well the model identifies the negative class. It is calculated as the number of true negatives divided by the sum of true negatives and false positives.

Standard Deviation: A measure of the amount of variation or dispersion in a set of values. It is calculated as the square root of the variance.

Statistical Power: The probability that the test will reject the null hypothesis when the null hypothesis is false. A high statistical power means that there is a high probability of detecting an effect if there is an effect to be detected.

Train-test Split: A technique for evaluating the performance of a machine learning model. The original sample is split into two sets: a training set to train the model, and a test set to evaluate the model.

Variance: A measure of how spread out the data points are in a dataset. High variance indicates that the

data points are spread out over a wider range, while low variance indicates that the data points are closer to the mean.

Youden's Index: A summary measure of a ROC curve. It is calculated as sensitivity + specificity - 1. Youden's Index can be used to find the optimal cutoff point for a diagnostic test.