

Machine (incl. Deep) Learning Research @

Predictive Analytics and Computer Vision wing, A3 Lab



Dr. S. Ghosh (senior postdoc) & Dr. H. Ayoobi (Asst Prof.)



Data Exploration, finding outliers and natural 'groups' in data

Machine Learning

Learns to 'map' a relation between a set of variables and a ground truth label

Unsupervised learning

Supervised learning

Clustering

Anomaly detection

Regression

Classification

Genome sequencing, metabolomics, finding 'natural groups' in RCTs, etc

Real-time patient monitoring based alerts, fraud in billing/insurance, etc

Risk score of disease incidence, survival analysis (e.g., Cox regression), healthcare cost prediction, dose-response

Treatment-response prediction, disease diagnosis, screening for disease, risk stratification, etc


EXplainable AI (XAI)

Counterfactual Explanations
What (all) could have changed my outcome? How can I get better?

Transparent models (inherently interpretable)

Pre-model explainability

Post-hoc explainability (model agnostic, might use surrogate models)



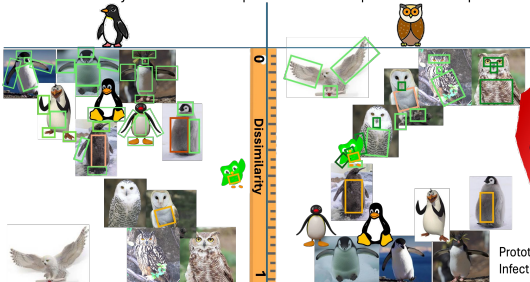
Uncertainty Quantification
How sure is the model, and what all can confuse it and by how much?

Dissimilarity between new samples and learned representative sample

Principal component analysis, Factor analysis, feature engineering, feature selection

Local/Interpretable Model Explanations (LIME), SHapley Additive exPlanations (SHAP), permutation feature importance, counterfactual explanations, partial dependency plots, sensitivity analysis

LVQ, Logistic regression, Decision Trees, etc.



Psst: We have some really sick data, and A3 lab is now part of the ICAI network