Lecture 1, Jan 8, 2024

Machine Learning

- *Supervised learning*: Mapping an input to an output, based on labelled data/ground truth; regression (continuous values) or classification (categorical/discrete labels)
  - Requires labelled data
  - To select the appropriate model we need to make assumptions about the problem; this is known as *inductive bias* or *learning bias*
    * "No free lunch" theorem says that we need to make assumptions to learn
    * Without proper assumptions, all models tend to perform equally if averaged over all possible tasks
  - We also need to quantify the model's performance; this is done through a *loss function*
  - More complex models have greater capacity to learn, but are more prone to overfitting – the model can learn the peculiarities about the specific training data and fail to generalize
    * To combat this, we partition the dataset into training and testing subsets; we evaluate model performance by running the model on the testing dataset
    * However with too much tuning we can also effectively overfit to the testing set
    * In practice we use 3 subsets: training (to teach the model), validation (to tune hyperparameters), and testing/holdout (to evaluate performance sparingly)
    * Ideally the final holdout set should be used only once (but this is often not possible)
- *Unsupervised learning*: learns patterns from observations without requiring ground truths
  - Includes self-supervised and semi-supervised learning
- *Reinforcement learning*: sparse rewards from the environment; actions affect the environment (e.g. training the model to play a game)