Neural AI: Machine Learning and Deep Learning

Image classification: Iris dataset

COMP 741/841 Week 4 - Spring 2024

Agenda

- Review: Search
- Machine learning and deep learning
- Getting started on Lab3
- Discussion: Assigned reading discussion

Search Problem

- Description
  - Given a **start state**
  - Reach the **goal state**
  - By moving from state to state based on **actions** taken at each state.
- Representation of the problem's **state space** is **graph**
  - nodes: representing states
  - edges: representing actions to move from one state to another

State Search Tree

- Root: start state
- Node: current state
- Children nodes: neighboring (successor) states from current state
- Path: a plan to reach the current state from start state

Search Example

- Review search algorithms on the course website
- Consider the graph example of a search problem
  - Edges have actual costs
  - Nodes have esimated (heuristic) costs of the path to the goal node
- Examine what different algorithms do by tracing
  - Frontier list
  - Explored nodes list

Neural AI: Machine Learning

- Subfield of AI
- Mathematical basis
  - Statistics (regression, decision trees)
  - Linear algebra (matrix computations)
  - Calculus (gradient descent)
- Idea: AI system **learns from data**
  - Needs large amounts of data (100 data points per feature)
  - Uses complex mathematical/statistical models
  - Finds correlations between known inputs and and known outputs
  - Predicts outputs for unknown inputs
- Requires significant computing power

Neural AI: Deep Learning

- Subfield of machine learning (ML)
- Mathematical basis
  - Same mathematical basis as ML
- Idea: AI system **learns from data using neural networks**
- Uses much larger data sets (thousands data points per feature)
- Requires much more computing power

Source: AWS. 2023. What's the difference between ML and DL. (https://aws.amazon.com/compare/the-difference-between-machine-learning-and-deep-learning/)

Machine Learning

How does ML learn from data?

- Requires ***feature engineering***
  - human intervention to
    - Extract features, label data, assign weights
- Suited for ***structured data***

  - E.g. Predict customer's cancel subscription based on their usage of the service

  - E.g. Recommend movies based on customer's history of movie watching
- Humuan understanding of predictions?
  - Depends on the mathematical model, e.g., decision trees

Deep Learning

How does DL learn from data **using neural networks**?

- Removes or minimizes need for humans
- Suited for ***unstructured data***
  - Challenge: identify complex relationships
  - E.g. Predict user “sentiment” from social media data
  - E.g. Image classification, natural language processing
- Human understanding of predictions?
  - Not possible (yet?)

Source : AWS. 2023. What’s the difference between ML and DL. https://aws.amazon.com/compare/the-difference-between-machine-learning-and-deep-learning/

Lab3: Sequential Neural Networks

- Keras Sequential model
  - Brief tutorial: https://www.dataquest.io/blog/tutorial-introduction-to-keras/
- Log in your SageMaker Studio Lab account, CPU instance
- Get a copy of `lab3` from the GitHub course org by accepting the GitHub classroom invitation
- Clone `lab3` to your SageMaker CPU instance
- Create Markdown cells before each code cell detailing what the code is performing
  - Markdown cells *must* be ***entirely*** in your own words