Final project

Task definition

Approach

Analysis

Task definition

- Topic: anything you're excited about (games, music, recommendation systems, optimizing society, personal productivity); see website for ideas; must be well-motivated; okay if done before
- Scope: not too broad (building an AI that organizes your life), not too narrow (training linear classifier for standard dataset)
- Evaluation metric: quantitative measure of success; speed? accuracy? user studies?
- Dataset: manually create 3-10 examples, scrape data, and/or create simulator

Approach

- Baseline: simple method (logistic regression with SGD); reveals challenges to tackle; provides lower bound on performance
- Oracle: "cheating" method that is unrealistically optimistic; inner annotator agreement; provides upper bound on performance

Example task definition

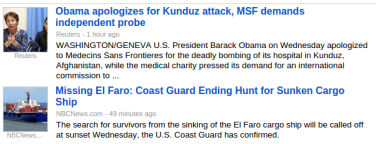

- Task: automatically illustrate a news article with images (motivation: pictures make reading more fun)
- Scope: focus on single paragraphs from Google News
- Evaluation metric: given two paragraphs and their two correct images, figure out which one is which (simplified)
- Dataset: scrape 5000 articles from Google News

Example approach

- Baseline: for each image, run standard object detector; return image with highest overlap with words in paragraph
- Oracle: have multiple humans do it and measure agreement
- Baseline-oracle: run baseline assuming perfect object recognition
- Algorithms: (i) predict an object (or n/a) for each word in the paragraph, perform word similarity with objects detected from images; or (ii) Google image search using words from paragraph for candidate images and perform object similarity

Example analysis

- Question: which is more reliable, word similarity or object similarity?
- Question: which types of articles are easier? concrete events easier than abstract topics
- Question: was the model able to figure out which words are relevant? common nouns are easy, but sometimes spurious correlations get in the way

Machine learning "versus" AI

Wrong question: Can I just do a "machine learning" project, or do I have to do an "AI" project?

Way to think about it:

- Machine learning applied to various models (reflex, state-based, variable-based)
- Strongly encourage to apply to non-reflex models

Examples:

- Predict eye gaze given images (correlations across time)
- Recommend news articles (recommend a set of articles)
- Generate blog posts

always solving collection of related prediction problems

Using same project in other classes

- Encourage sharing code/data across classes
- Ideally, would explore different aspects of classes

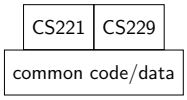

- Cite what you did / turned in for each class
- Expectations are higher if sharing between classes

Project interest form

`https://stanford-cs221.github.io/spring2024/project.html#p-interest`

- We encourage all students who might be interested in a project to fill this form out.
- We will use the form to assign a CA mentor to help you with the project proposal.
- Filling out this form also lets us help you find a group if needed.

due April 19

Get help: come to office hours!

Project proposal

https://stanford-cs221.github.io/spring2024/project.html#p-proposal

- Define input-output behavior; give **concrete example**

- Have **baseline** and **oracle** implemented
- Discuss potential solutions (modeling, inference, learning)

due May 3

Get help: come to office hours!