

CS288 HW4: Semantics

Nicholas Tomlin and Dan Klein

Due: 15 April 2022, 11:59PM PST

Overview

This homework will be focused on semantic parsing. This homework is substantially shorter than previous assignments, with much more stencil code. However, the final section on weakly supervised semantic parsing can take a while to train (20-30 minutes), so please start early. The autograder for this homework has been substantially modified to fix a handful of bugs and allow more opportunities for partial credit.

Background Reading

Feel free to consult with the following resources before beginning this assignment:

- Semantic Parsing Tutorial: <https://yoavartzi.com/tutorial/>
- Eisenstein (Ch. 12): <https://cseweb.ucsd.edu/~nnakashole/teaching/eisenstein-nov18.pdf>

Assignment

Notebook: <https://www.kaggle.com/nickatomlin/cs288-hw4-public>

- Complete the the notebook and save the generated files
- Write an error analysis of the semantic parsing model (details in notebook)

Submission to Gradescope

Please submit the assignment to: <https://www.gradescope.com/courses/361823/> (code: 4PBP57)

When you upload your submission to the Gradescope assignment, you should get immediate feedback that confirms your submission was processed correctly. Be sure to check this, as an incorrectly formatted submission could cause the autograder to fail. In particular, please make sure that you're zipping your files and not the folder containing them. Note that Gradescope will allow you to submit multiple times before the deadline, and we will use the latest submission for grading. Make sure you have the following files (with correct names and extensions):

- dev_true_lf_actions.txt
- test_true_lf_actions.txt
- supervised_dev_lfs.txt
- supervised_test_lfs.txt
- latent_dev_lfs.txt
- latent_test_lfs.txt
- hw4.ipynb
- report.pdf