

CS288 HW2: Machine Translation

Nicholas Tomlin and Dan Klein

Due: 23 February 2022, 5:00PM PST

Overview

This homework will be focused on machine translation. Due to issues with GPU allocation in the previous homework, we're now moving our notebooks from Google Colaboratory to Kaggle. Once you've verified your account with a phone number, Kaggle should provide 30 hours of GPU allocation per user, per week. If you have a Colab Pro or Colab Pro+ account or otherwise prefer using a resource other than Kaggle, you are welcome to download the `.ipynb` file and run it elsewhere. This homework consists of two Kaggle notebooks and a written report.

Background Reading

Please checkout the following resources before beginning this assignment:

- Statistical machine translation: https://cal-cs288.github.io/sp20/slides/cs288_sp20_05_statistical_translation_4up.pdf
- Neural machine translation: <https://web.stanford.edu/~jurafsky/slp3/10.pdf>

Assignment

This homework consists of two Kaggle notebooks, as well as a written report. Below, you'll find a handful of questions. Please answer these questions in \LaTeX and save them to a file named `report.pdf`. Submit this file and all outputs from the notebooks into a single `.zip` file for submission to Gradescope.

1. Alignment with IBM Model 1: <https://www.kaggle.com/nickatomlin/cs288-hw2a-public>

(a) Complete the notebook and save the following files:

- `hw2a.ipynb`
- `example_alignments.pkl`
- `multi30k_alignments.pkl`

(b) Report: find at least one sentence from the Multi30K dataset where the IBM alignment model performs reasonably well, and find another one where it fails catastrophically, and include alignment visualizations for both examples in your report. You may want to consult a German-English dictionary for this part of the problem. Provide a brief explanation for why the alignment model did poorly on the failure case.

2. Neural Machine Translation: <https://www.kaggle.com/nickatomlin/cs288-hw2b-public>

(a) Complete the notebook and save the following files:

- `hw2b.ipynb`
- `predictions.json`

- (b) Report: generate attention visualizations for at least four sentences and describe general trends you observe in your report. What similarities or dissimilarities do you see between the attention maps and the alignment visualizations from Part A?

Debugging and Piazza

Please post questions about this homework in the corresponding homework threads on Piazza. Please limit your use of private questions unless absolutely necessary, as we do not have the staffing capacity to help with all requests for code debugging. If you believe you have found a legitimate bug in the solution, however, please contact us immediately.

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. For this project, you should be able to see your test set accuracies and a confirmation that all required files were found, but you will not be able to see your score until later. Most assignments will be graded primarily on your test set accuracies and written report.

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):

- `predictions.json`
- `example_alignments.pkl`
- `multi30k_alignments.pkl`
- `hw2a.ipynb`
- `hw2b.ipynb`
- `report.pdf`