

TTIC 31190: Natural Language Processing Project Report

Kevin Gimpel

Due: 7:00 pm, Thursday, June 7, 2018

(For graduating students: Due: 7:00 pm, Thursday, May 31, 2018)

Submission: email to `kgimpel@ttic.edu`

Submission Instructions

Email your project report to `kgimpel@ttic.edu` by the deadline. If you are not graduating this quarter, the deadline is 7:00 pm on Thursday, June 7, 2018. If you are graduating this quarter, the deadline is 7:00 pm on Thursday, May 31st.

Formatting Instructions and Requirements

Use LaTeX or Microsoft Word templates/style files prepared for recent NLP conferences. There are links to recent templates on the course webpage.

Your project report can have at most 4 pages of content, though any number of pages of references is permitted. This is the length of a short-paper submission for an NLP conference. Researchers are able to effectively compress many months of work into just 4 pages without losing anything important. This is a useful skill to develop!

Project Report Sections

Below are suggestions of sections that you may want to include in your project report, along with a brief description of each section. Unlike your project proposal, you should include more high-level introductory and motivational material, as well as a more in-depth description of the related work. Also, of course, your final report will contain results and analysis.

Abstract

Provide a brief, high-level overview of the paper, including the problem, methods, and key results.

Introduction

The introduction should describe and motivate the problem you are trying to solve. It can also preview the methods you will be using and mention the key empirical results and/or analysis. For a 4-page paper, a reasonable estimate is for the introduction to be approximately half a page.

Related Work

Include a brief discussion of the most relevant related work. This should probably also be less than half a page.

Methods

Describe the methods you will use or develop in order to solve your problem. Assume that your readers have taken this course, so you do not need to define everything in complete detail. But if you are modifying or extending some existing method or technique, then you will probably want to give formal definitions of what you are modifying (so that you can more clearly explain the novelty). If most of the novelty is in dataset/task creation, then you will want to discuss that here. This section might be anywhere from 0.5 to 1.5 pages in length, depending on the complexity and novelty of your methods.

Experimental Setup

Describe your experimental setup. This description should include enough detail so that someone in the class could replicate what you did (perhaps by additionally asking you a few questions). This may include details like the following:

- datasets used, including how you partitioned the data into training, development, etc. sets; possibly include some dataset statistics
- evaluation metrics used
- baseline results you are comparing to, whether they are results taken directly from other papers or baseline methods that you implemented
- any toolkits used to implement the methods/baselines
- how you tuned hyperparameters of the methods (e.g., number of training epochs, feature engineering, etc.), and the final hyperparameter values / settings you used

Results and Analysis

Present your experimental results (e.g., fill in the cells in the table from your project proposal, if applicable). Also, analyze the results. What are the findings of your empirical investigation? If you see a particular trend in the results, can you analyze the data or model to explain the trend?

Conclusion and Future Work

Summarize the key results / contributions of the project and discuss the natural follow-up work that other researchers (or you) might pursue in the future.

References

Include cited references.