

# CS 771—Natural Language Processing Techniques

## Project Report

Winter 2010

Department of Computer Science and Engineering

Wright State University

---

**Due:** *in class, Wednesday, March 17*

**Worth:** 40% of final grade

**Instructor:** Shaojun Wang, Joshi 387, 775-5140, shaojun.wang@wright.edu

---

Your final project report should cover the following issues and will be graded as shown:

### Background

- **Problem** (4%): As clearly as possible, describe the general NLP problem you are working on, and explain the specific refinements and special cases that you addressed.
- **Literature survey** (6%): Briefly survey the existing work that has been done on your problem, as well as the existing work that has been done on the approach(es) you are considering. You should cite about 10-15 relevant references. Your survey need not be exhaustive, but you should try to cover the most important prior work if you can.

### Methods

- **Approach** (4%): As clearly as possible, describe the approach(es) that you applied to the NLP problem, and clearly specify the final NLP system(s) that were implemented.
- **Rationale** (6%): Explain why you implemented the NLP systems that you did. Specifically, explain the simplifications you made along the way, and why you made these particular choices (and not others). (For example, you might detail other approaches that could be applied to the NLP problem and explain why these were not pursued.)

### Plan

The most important (and most difficult) part of any research project is figuring out how to evaluate the results. You should state concrete goals for your project. In particular, you should identify at least three concrete hypotheses that you think would be interesting to test with your implemented NLP system, and outline a specific plan for experiments that you would need to conduct in order to test each of these hypotheses. Testing these hypotheses should be stated as the concrete goals for your project.

**Note:** You should *not* phrase your project goals as “our technique will achieve great performance at task X”. Rather, you should phrase your goals as tests of specific hypotheses whose outcome would be interesting no matter what happened (so that your project would succeed however the tests turned out). For example, rather than simply claiming “our technique will work great”, it is better to formulate and test a hypothesis that would teach all of us something useful about the technique and problem, like: technique *blah* is robust to certain kinds of noise or perturbations in the problem (or perhaps not), or technique *blah* scales up feasibly in the size of problem *X* (or not), or technique *blah* is inherently limited in the ultimate performance capabilities it can achieve on problem *X* because of *blah blah* about the problem (or not), *etc.* This requires a little more thought than perhaps you are used to in planning a project.

- **Hypotheses** (8%): Clearly state the main questions that you investigated in this project. *These should be identified before you run the experiments!* Ideally, these should be interesting questions whose answer is not obvious beforehand, but the answer would be interesting no matter how the experiments turned out.
- **Experimental design** (10%): Once you have settled on some good questions, it is important to figure out how to answer them. For each main question, describe a series of experiments that is designed to answer the question. Explain the difficulties faced in designing these tests and explain why your experiments will overcome these difficulties to yield a definitive answer.

## Experiments

- **Results** (4%): As clearly as possible, describe the results you obtained with your tests. Use plots, graphs, and tables, if necessary. Make sure it is easy to understand what happened.
- **Critical evaluation** (10%): In the end it is important to critically evaluate the results of the experiments. Did the experimental results answer the questions? If so, then what were the answers? If not, then why did the results fail to yield definitive answers? Is it then possible to formulate a new experimental strategy?

Note that it is important to demonstrate critical thought in this part of your assessment. You will be graded more on the strength of your reasoning rather than how the tests actually turned out. For example, it is perfectly acceptable if you did not achieve definitive answers to your questions, so long as you can recognize this, explain why it happened, and suggest additional tests that might yield more definitive results.

## Conclusion

- **Lessons learned** (8%): What specific things did you learn from doing this project? Did you learn anything about the NLP problem itself, the approaches you tried, or the

experiments you conducted? If there are any good ideas you came up with in the end but did not have time to pursue, this would be a good place to mention them.

A top-notch project is something that with a bit of fine-tuning could be submitted confidently to a conference such as ICML, ACL, EMNLP/SIGDAT, NIPS. The minimum project is something – a survey, introduction, implementation – sufficiently well done that interested people can use it and learn from it for at least the next few years.