

On my honor as a University of Colorado at Boulder student, I have neither given nor received unauthorized assistance on this work.

*Note regarding unauthorized assistance: I would prefer that you learn as much about Machine Learning as possible throughout your project. Therefore, I encourage you to talk to whoever is willing to talk to you about your project. Feel free to involve collaborators somewhat deeply, but in order to learn, you will of course need to do a significant amount of work yourself. I expect all written assignments to be your own work and reflect **your** understanding of that work. That said, your paper should reflect feedback from your class peer group as well as from me.*

Write this paper in the style you would write it for some relevant conference (even if you don't plan to submit it). **The paper should be approximately 4500 words** – an absolute minimum of 3000 and maximum of 7000 words (not including the sections starting at *Acknowledgements*). If and only if you are planning to submit it to a journal, you can request to write a longer paper – up to the journal's target number of words for an article. All projects are unique – you do not have to include everything in these guidelines, but a good paper will discuss most of these issues.

Title

Author List, with affiliations and email addresses

(Include everyone who has made an *author-worthy* contribution to your project.)

Abstract

~100-150 words

Succinctly state: the problem, the key thing(s) the reader will learn, the key results.

Introduction

What is the problem you are addressing?

Who cares?

Why do they care?

What is lacking from previous attempts to solve the problem?

What do you know about previous attempts to solve it? List key publications that you are aware of.

What have previous results been?

Background/Preliminaries, Definitions, Data, Examples, Key issues/problems, Problem framework

Include any preliminary information that is necessary to understand the problem, but that doesn't belong in a fluent, compelling introduction. E.g.:

Definitions

Describe the data/corpus at a conceptual level: what does it represent, where did it come from, who created it, how was it created and annotated, what are the interesting aspects of the data including class distributions, show some examples of the data... If you created the corpus, provide inter-tagger/annotator/rater agreement (ITA/IRA) and Kappa values if portions were double annotated. Also describe the annotation process in enough detail that readers will be confident that your results (presented later) will be meaningful and generalizable.

Either here or in the intro, consider including a framework/model for solving the problem. Describe the key aspects of most solutions, cast in a way that enables you to refer back to them throughout the paper, comparing and contrasting (in the later sections of the paper) your solution to what is currently being done.

Approach / Methods

You might include an initial subsection or paragraph presenting your hypotheses.

In the first paragraph of the main subsection describing your approach, summarize the approach for training the classifier and how you used the training and test data...

What is your approach to solving the problem? Why did you hypothesize this was better than prior approaches. Discuss the relation to prior relevant work?

What attributes/features did you use?

What specifically did your algorithm be learning?

What algorithms did you utilize, enhance, evaluate, etc?

Experiments

(Papers rarely break experiment descriptions into the following three subsections, but they generally should.)

Design

Describe how the experiment was conducted in enough detail for the reader to be confident that your design was valid, that your results are sound/believable, and that the system should generalize to the target application.

There should be a fair amount of emphasis on how training data and test data were generated, what their relationship is, why one should consider them independent, etc. This could be done earlier, with the key points summarized here.

Results (tables, figures and 1-3 paragraphs)

Present the evaluation metrics for each experimental condition, the baselines, and results from the most closely related prior work in tables, graphs, or whatever style

best represents the hypothesized and actual key differences between your work and that of other researchers. If some of the differences are statistically significant, mark them or include columns stating the associated p values. In the prose, impassively state the key findings, including differences in the results between approaches.

Discussion

This is where you are not so impassive. Expound on why yours is the better mousetrap. Hypothesize or prove why your results are better (or worse). Discuss why the findings and approach are interesting. Emphasize what the world should learn from your experiment(s).

Sometimes authors have a separate *Related Work* section, but I think it is generally a good idea to integrate the discussion of prior work throughout the paper (the intro, the approach, always in the results, and here in the discussion).

Very briefly, describe the future work that your findings suggest should be pursued.

Conclusion (2-3 paragraphs)

Start with a strong value-of-the-work statement. Summarize the approach (optimally in just one sentence), the results, and the discussion, especially the key findings a reader should take away from the paper (the key contributions).

Very briefly, describe the future work that your findings suggest should be pursued. Start with a strong value-of-the-work statement.

Acknowledgments

Acknowledge everyone who has given you help and any relevant grants...

References ...

{Future Work}

What are the key things you learned from the project? Is that what you wanted to learn? If not, how are you going to alter your next experiment?

Detail and justify your next experiment.

If you are working toward a dissertation, discuss how you will use this work and these findings in your dissertation.

{Response to Prior Review}

In the final paper, respond to each significant comment in the prior reviews (how you addressed it, or what the reviewer misinterpreted, etc).