

The Title of Your Thesis

Your Name
STUDENT NUMBER: 000000

THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
MASTER OF SCIENCE IN COGNITIVE SCIENCE & ARTIFICIAL INTELLIGENCE OR DATA
SCIENCE & SOCIETY
DEPARTMENT OF COGNITIVE SCIENCE & ARTIFICIAL INTELLIGENCE
SCHOOL OF HUMANITIES AND DIGITAL SCIENCES
TILBURG UNIVERSITY

Thesis committee:

Supervisor
Second Reader

Tilburg University
School of Humanities and Digital Sciences
Department of Cognitive Science & Artificial Intelligence
Tilburg, The Netherlands
January 2000

Preface

Some room for acknowledgements.

The Title of Your Thesis

Your Name

The abstract should typically be a few sentences on what you did, much like described in the thesis guidelines. Usually in the style: "In this work [we propose ... /we investigate ...]. Generally you quickly contrast (one sentence) what has been done before and what you offer. You ideally describe what data you used, if you show improvements or not, and any other important findings. End with some conclusive sentence (e.g. this bodes well for, we concluded that this is a good approach /whatever) – should be some kicker. Basically the whole abstract sells your paper, it should be an elevator pitch of your research.

1. Introduction

It's best to see the introduction as a 'triangle' shape in terms of specificity of the subject. You immediately want to start working on describing both the societal and scientific relevance of the paper (usually, but not always in that order). I challenge you not to start with < Recently, Nowadays, Since ..., etc. >; they make easy starters but are heavily overused. Try to make your story more specific with every paragraph, and be sure to have one thread through all of the sentences / paragraphs; it should be a logical flow of zooming in on the material that will eventually bring you to your research question / contributions. While doing so, towards the end, you should be extremely specific about what you are going to do — you set the reader up with required information for grounding the subject, and understanding what your thought process was towards setting up the research, so that they know what to expect. Any citations along the way should include 'related but different' work, or some very important papers that gave rise to the work that you are doing here. Make sure that you state a hypothesis (or a research question, alternatively), this helps you think about the experiments required to test this hypothesis. Some papers list the contributions their works offers as bullet points as an ending, this is somewhat of a stylistic choice. The standard questions should at least be clear at the end: what are you doing, why are you doing this, what data are you using, what techniques, what will your experiments demonstrate, what can we learn from that, and so on.

2. Related Work

By now you have introduced all concepts you'll be working with. Time to acknowledge all prior work on this topic, which generally always include a brief historic overview. Note that this includes work on tasks that are *similar* but not exactly the same as the work that you are doing. Note that this should be related by topic, not technique (as that could blow up very quickly)! A good related work section draws interesting relations and provides a structured overview of the logical progression of all the task's research. You should be able to identify key methodologies, contrasts in approaches, and work towards specifying the *research gap* you will be addressing with your work. Try to refrain from giving full method section descriptions of all the work that you cite here, typically

a core description per piece of related work is enough. Only the *very* related work deserves a more complete explanation, especially if you will be borrowing some of this methodology.

My personal strategy for forming an extensive bibliography is to start with the most recent papers on a task, and seeing which are the early papers that they cite. Look those up and see if those cite even earlier work. At some point, you'll hit somewhat of a 'beginning' of the field. Try to identify which are the papers with the most impact (say well over 50 citations) and for those with the most impact, look at the 'cited by' sections in either google scholar or semanticscholar. This might help you identify papers that are related, but might be difficult to find within the range of queries you can come up with.

2.1 Citations

Citations in parentheses are declared using the `\cite{}` command, and appear in the text as follows: This technique is widely used (Woods 1970). The command `\citep{}` (cite parenthetical) is a synonym of `\cite{}`. Citations used in the sentence are declared using the `\namecite{}` commands, and appear in the text as follows: Woods (1970) first described this technique. The command `\citet{}` (cite textual) is a synonym of `\namecite{}`.

This style file is designed to be used with the BibTeX style file `compling.bst`. Include the command `\bibliographystyle{compling}` in your source file. Citation commands are based on the `natbib` package; for details on options and further variants of the commands, see the `natbib` documentation. In particular, options exist to add extra text and page numbers. For example, `\cite[cf.][ch. 1]{winograd}` yields: (cf. Winograd 1972, ch. 1).

3. Experimental Setup

If your thesis focuses on an existing data set, for example many Artificial Intelligence or Data Science theses, the usual setup is to divide the Experimental Setup section into two major components: one describing the data that you use, and another outlining your algorithms.

If your thesis includes collecting new data, for example in a Cognitive Science or Artificial Intelligence thesis, you will probably need to include more subsections that provide details on the participants and how they were recruited, the experimental procedure. For more information and examples, see APA Guidelines v.6 and related material.

3.1 Participants

This section, for a thesis that involves collecting new data, should explain information about the demographics of participants and how they were recruited.

3.2 Design & Procedure

This section, for a thesis that involves collecting new data, should explain what the design of the experiment was (independent and dependent variables + possible moderating variables) and exactly how the experiment was conducted.

3.3 Data

This section, for a thesis that involves processing existing data, should provide insight about the data. Try to be as verbose as possible about where your data comes from, how it was collected, what your features are, reporting descriptives, maybe give some actual sample data, what your targets are, what the distribution looks like, etc. Don't try to skip on anything worth reporting! Your reader should have a complete 'look' into your data without actually seeing it. Tables and plots are your friend. If you did pre-processing, or any other steps to get your data in a 'clean' format, please state that, in detail, in this subsection.

3.4 Method / Models

This section, for a thesis that involves complex analysis of new or existing data, should provide details about the analysis methodology. Here you describe the rest of your pipeline (i.e. everything after the data part). This includes what models you used (FIND AND CITE THEIR SOURCES!), which libraries (FIND AND CITE THEIR SOURCES, otherwise link!) what parameters you tuned (and how), evaluation metrics, splits, baselines, everything. Again, be as verbose as possible. Your experiment should be almost fully reproducible by text alone.

4. Results

In addition to what's described in the thesis guidelines, make sure to provide *compact* information. Let's say I ran many many different parameter settings, ideally you only want to show the versions of your models that scored highest overall (so classifier + param settings). Maybe you might want to compare between classifiers, but generally only when that's useful. See Table 1 for reporting different pipeline settings. Note that the import numbers (highest per column) are bold for quick lookup. This is also the general table style you should try to adhere to (nicely APA). For figures, please please at least export to PDF or something that scales (SVGs work poorly in \LaTeX). If you're brave enough you can use `tikz` (there are multiple packages for both R and Python that export e.g. `ggplot` or `matplotlib` to `tikz` code). The advantages are that it behaves well within the thesis environment, and all figures / captions are immediately in the same font as your thesis. Save this for the end — it can be quite time consuming. This is also the part to include post-hoc analysis. Make sure that it's clearly separated from your main results though!

5. Discussion

Results have to be objective and dry. Basically a report. The discussion session is where you can start interpreting those results more freely, get more of your own insight in, and maybe even draw some similarities to related work that you have reviewed. It's also a good time to be critical of your own set-up; what worked, what did not, what can be improved, what are the limitations of the study, etc. Most of all, you should directly talk about your hypothesis or research questions. Were they confirmed / answered? What does this do for the field? You can see the discussion as the inverse of the introduction where you go from your very specific results section, and start zooming out further throughout.

Table 1

Best performing models classifying bots, on Twitter and Facebook respectively. F_1 scores report positive (bot) class. Outline text left (l) and numbers right (r). Don't center align!

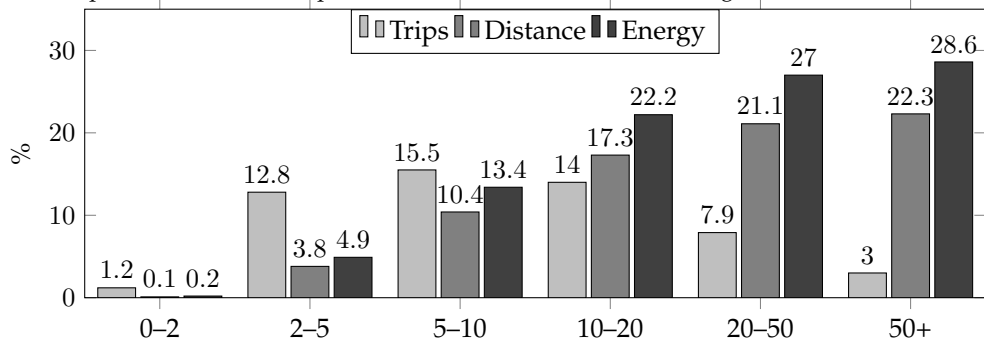
PCA	Models	F_1 score	
		Twitter	Facebook
300	Linear SVM ($C = 0.1$)	0.51	0.91
	Random Forest ($S = 5, F = 5$)	0.71	0.85
	Naive Bayes	0.61	0.73
500	Linear SVM ($C = 0.1$)	0.55	0.84
	Random Forest ($S = 5, F = 5$)	0.76	0.71
	Naive Bayes	0.41	0.64
	Majority	0.50	0.60

6. Conclusion

The conclusion is a short list of what you did, what the most important findings were, and what the most important implications are. Not much more to add than the thesis guidelines here — you can view it as a compression of the last part of your introduction and discussion. Your abstract should be an even shorter version of the conclusion.

Figure 1

Some caption. Note that all captions should be above the Tables / Figures!



References

- Winograd, Terry. 1972. *Understanding Natural Language*. Academic Press, New York.
- Woods, William A. 1970. Transition network grammars for natural language analysis. *Communications of the ACM*, 13(10):591–606.

Appendix A: First item

Here is the appendix. You can put anything you want here, with or without descriptions, and can refer to the sections by labelling them. Appendix 6.

