
CS152 Neural Networks Final Project Template

Gabe Hope
Harvey Mudd College
ghope@hmc.edu

Abstract

Your abstract should contain a concise, 1-paragraph description of your project. Should aim to describe your goals, your approach and the gist of your results in 1-2 sentences each.

1 Introduction and related work

The first section should be an introduction (at most 1 page) that describes the problem you are trying to solve with your project. Make sure to provide any background about why this is an interesting or important problem, how it could be used in practice, what approaches have previously been applied to approach this problem (including citing any papers you referred to) and how your work addresses this problem. From this section it should be clear what your goals were and what your project accomplished.

2 Datasets

In this section you should provide a detailed description of the datasets that you used for this project. You should include specifics such as: the features and labels provided by each dataset, the size of each dataset and where the data was obtained (cite any sources!). It is also useful to provide examples of the observations of each dataset as well as high-level statistics, such as the distribution of classes. You should clearly state why each dataset was chosen. If you collected the data yourself, you should describe how the data was collected in detail.

3 Methods

In this section you should describe the approach that you used for your project. You should describe any algorithms that you used (beyond things like gradient descent covered in class), the details of the neural network architectures that you used, and how your model was trained (e.g. what loss function did you use?). Be specific and be sure to provide a typeset equation for any relevant formula that appears. Make sure to describe how you made decisions about your model: did you choose hyperparameters like the learning rate, number of layers etc. using validation data? what ranges of options did you consider?

4 Results and experiments

In this section you should describe the results that you obtained from the methods and datasets described in the previous sections. Start by describing the metrics that you used (i.e. accuracy, MSE) and how you measured them (e.g. how much data did you hold out?). Report any quantitative metrics in a clear, understandable way. Make sure to include any results from competing methods that you were able to find and/or run. For qualitative results make sure to include figures that make it clear how

your model performs (e.g. for style transfer, show restyled images). Discuss any conclusions you can draw from these results: did your model perform as expected? Why/why not? How did it compare to other methods? Are the results conclusive? Remember that this project is judged on your best efforts to understand/implement a model, not on the results. It's fine if your model didn't work! If that's the case, try to understand why it didn't.

5 Conclusion and future work

Briefly summarize any takeaways from this project. How would you recommend others use your results? Describe at least one way this project could be extended in the future (new datasets, a modified approach, extensions to new problems?).

6 Broader impacts

Describe any impacts this project or model might have on society. How might your technique be used? How might it be misused? Does your data contain any harmful biases and could your model end up reflecting or amplifying them? Are there other potential ethical concerns (privacy concerns, copyright concerns, misinformation, etc.)? Describe the resources used to train your model. How long did it take and on what hardware? Did this require a significant amount of energy?

7 Format

You are free to use your preferred format for the technical report, however, I recommend typesetting with Latex using this template format. You can access the template on Overleaf here (<https://www.overleaf.com/read/sgrzqjprrcqv>).

8 Code

Provide a link to a Github repository containing any code you wrote or used for this project. Some guidelines:

- Your code should be sufficient to replicate all of the experiments described in the technical report. Ideally, you should provide a script that could be run to regenerate the reported results. **DO NOT** include your data directly in the Github repo.
- You are not required to write all code from scratch, you may build off of the code of others. In your repository, make it clear which pieces of code you wrote and which are from others. You should attempt to implement the core of the project yourself, particularly if your project focuses on a specific architecture (e.g. vision transformers) or algorithm (e.g. style transfer). If you're unsure about the expectations, please talk to me.
- Make sure your code is organized and readable, with comments where necessary.

9 Presentation

You will also submit a 5-10 minute video presentation of your work. I would recommend a recorded PowerPoint presentation (you could do such a recording through Zoom). You should consider the following guidelines for your presentation:

- The main goal of the presentation should be to communicate the approach that you used. Describe your problem and at a high-level how the algorithm or architecture you explored works. Consider your target audience to be your classmates, rather than an expert in the field. I would encourage you to reuse figures from your technical report and to include visual references to help others understand the techniques. You may use visualizations from other sources provided you give them appropriate credit.
- You should also discuss the results you found and any practical recommendations you encountered that would help another student trying to apply this technique.

- Each member of your team should contribute a roughly equal amount of time to the presentation.

References

Include references to any papers, articles, datasets, etc/ you discuss at the end of your report. You may use your choice of citation style as long as it is consistent. Note that the Reference section does not count towards the page limit.

- [1] Alexander, J.A. & Mozer, M.C. (1995) Template-based algorithms for connectionist rule extraction. In G. Tesauro, D.S. Touretzky and T.K. Leen (eds.), *Advances in Neural Information Processing Systems 7*, pp. 609–616. Cambridge, MA: MIT Press.
- [2] Bower, J.M. & Beeman, D. (1995) *The Book of GENESIS: Exploring Realistic Neural Models with the GEneral NEural SImulation System*. New York: TELOS/Springer-Verlag.
- [3] Hasselmo, M.E., Schnell, E. & Barkai, E. (1995) Dynamics of learning and recall at excitatory recurrent synapses and cholinergic modulation in rat hippocampal region CA3. *Journal of Neuroscience* **15**(7):5249–5262.