

Paper Writing Workshop

Introduction to Deep Learning

A written report describing **original** research results

The purpose: to communicate new scientific findings, clearly.

The paper includes:

- A statement about the problem or the question
- A presentation of evidence
- An assessment of the strengths/weaknesses of the presented evidence
- Conclusions

Begin by conducting a literature search

- Search on Google Scholar and Papers With Code for readings that are relevant to your interests
- Use keywords to search for related works on the topic
- Decide what topic you want to pursue for the project by reading about other experiments
- Generate ideas about how you will contribute to existing research on the own and/or with your teammates, after searching/reading other papers

Standard format

- Title
- Authors
- Abstract
- Introduction
- Literature Review (“Prior Art”)
- Methods (Proposed Solution + baseline)
- Results (Experimental Validation)
- Discussion (Analyses)
- Conclusions and Future Work
- References

Title (and authors)

- Use the fewest possible words to describe the contents of the paper
- Consider the key words used for searching Google Scholar
- The title should capture the gist of the experiments and **findings**
- Never use abbreviations in titles
- Authors and their order
 - The person who writes the most gets their name listed first
 - (Establish paper writing responsibilities early)

Section 0: Abstract

It's a summary of the information in your final paper.

- Less technical
 - States the objective or scope of the experiment
 - Describes the methods
 - Summarizes the results
 - States main conclusions
- Does not have citations
- One paragraph, 250 words or less

(Usually written last, after the paper is finished)

Section 1: Introduction

An introduction presents the nature and scope of the problem you're investigating

- Make a clear problem statement with the data and objective explicitly defined
- State your hypothesis and propose a solution to the problem
- To facilitate understanding, you should have a high-level diagram
 - To better articulate the setup, a separate inference and training diagram may be necessary
- Terminology and basic methods should be defined at a high level

Section 2: “Prior Art” (Literature Review)

In most cases, your problem will not be new -- it will have been thought of before and may have had a different motivation or application.

- What were the earlier solutions to your problem?
- How well did they perform in the context of your objective?
- How is your way of solving the problem different?
 - Did you apply a state-of-the-art model from one domain to a different domain?
 - Did you use a new loss to improve on an existing state-of-the-art result?
 - Did you create a new dataset and methodology that solves a novel problem?
- Important: include key ideas behind the papers cited in the literature review
- *End the literature review by distinguishing yourself from the prior work*

Section 3: Proposed Solution (Methodology)

To help you write this portion, we have divided the writing into three sections:

- **Section 3.a: “High-Level” Concept**
 - What empirical evidence motivates your hypothesis?
 - Why do you think your solution will work?
 - What makes your problem challenging and important to solve?
- **Section 3.b: Technical details**
 - Create a detailed diagram with core components labeled, then describe your methodology.
 - Formalize the mathematical definitions necessary to make your methodology work.
- **Section 3.c: Implementation details**
 - Explain what your algorithms are doing with explicit definitions
 - Outline any approximations and shortcuts you had to take
 - Provide enough detail for another researcher to replicate the study

Section 3: Proposed Solution (Baseline)

You must contextualize your proposed solution with respect to some baseline.

- The baseline is an existing solution to your problem
- It is a competitive solution, but has room for improvement
- The entire goal of your paper is to demonstrate that
 - The baseline you chose is valid
 - That you created a novel and working solution to your problem
 - How your solution performs relative to the baseline
- Justify why comparison of your approach to the baseline is necessary

Section 4: Experimental Validation

This section contains only facts, no analyses. State clearly what you are showing.

- Explain your data in a way such that the reader knows *why* you are using the data and *why* they are appropriate for your work
- Explain the “high-level” understanding of your results
- Explain *why* you got these results, including *what* you varied with the architecture/hyperparameters to show *how* it did/didn't work

Section 5: Analyses

- Analyze your experiment outcomes and relate them to the original hypothesis
 - Answer the research question and hypothesis you presented in the introduction
 - Discuss if your original hypothesis has been validated or invalidated.
 - Point out the strengths **and** limitations of the study
 - Compare and contrast your results with prior works, contextualized with your motivation
- Show the significance of the results you observed in the experiment
 - Present the relationships/generalizations that are shown by the results
 - Discuss the theoretical implications and/or practical applications of the results

Section 6: Conclusions & Future Work

Summarize the evidence

- What did you conclude overall in your analyses?
- Did you solve the problem? If not, what is the way forward?
- What future works or developments are needed to improve?

Some Rules...

If it's hard to read, you'll get a bad grade

- Use the present tense when referring to work that has already been published
 - Example: “Deep learning is...”
- Use the past tense when referring to your study
 - Example: “We conducted...”
- If the method has been previously published, you should cite the reference
 - **Good:** “We used Demucs, which is a state-of-the-art music source separation model, currently capable of separating drums, bass, and vocals from the rest of the accompaniment.. Demucs is based on a U-Net convolutional architecture inspired by Wave-U-Net. The v4 version features Hybrid Transformer Demucs, a hybrid spectrogram/waveform separation model using Transformers. It is based on Hybrid Demucs (also provided in this repo) with the innermost layers are replaced by a cross-domain Transformer Encoder.” (A Defossez, et al., 2019)
[source](#)
 - **Not good:** “We used Demucs from GitHub for source separation of different types of audio.”

References

Use the APA format (like these Star Trek examples, below)

- **Book**
 - Okuda M, Okuda D. *Star Trek Chronology: The History of the Future*. New York: Pocket Books; 1993.
- **Book Article or Chapter**
 - James NE. Two sides of paradise: the Eden myth according to Kirk and Spock. In: Palumbo D, ed. *Spectrum of the Fantastic*. Westport, Conn: Greenwood; 1988:219-223.
- **Journal Article with volume numbers**
 - Wilcox RV. Shifting roles and synthetic women in Star trek: the next generation. *Stud Pop Culture*. 1991;13:53-65.
- **Journal Article without volume numbers**
 - Di Rado A. Trekking through college: classes explore modern society using the world of Star trek. *Los Angeles Times*. March 15, 1995:A3.
- **Journal Article on the Internet**
 - McCoy LH. Respiratory changes in Vulcans during pon farr. *J Extr Med* [serial online]. 1999;47:237-247. Available at: http://infotrac.galegroup.com/itweb/nysl_li_liu. Accessed April 7, 2015.
- **Website**
 - Crawford A. Amy's Website. Amy's Web Page. 2022. Available at: <http://www.some.org/info/fish/DR9/>. Accessed March 8, 2023.

Examples of Good Papers with similar structure

Sarenne Carrol Wallbridge, Catherine Lai and Peter Bell

“Investigating perception of spoken dialogue acceptability through surprisal”

Hansheng Chen, Pichao Wang, Fan Wang, Wei Tian, Lu Xiong, Hao Li

“EPro-PnP: Generalized End-to-End Probabilistic Perspective-n-Points for Monocular Object Pose Estimation”

Shengjia Zhao, Abhishek Sinha, Yutong (Kelly) He, Aidan Perreault, Jiaming Song, Stefano Ermon

“Comparing Distributions by Measuring Differences that Affect Decision Making”

Questions During Grading (Peer Review)

Do the authors make clear their applied or theoretical problem?

Do the authors review and explain enough papers to clearly contextualize their work?

Is the core model / theory / algorithm clear?

Is the process of evaluation clear? Is it clear what the evaluation metrics are and why they are used?

Is the process of training clear such that you could reproduce it given their model / theory / algorithm?

Do the authors do a sufficient job experimenting with different specifications?

Are the results presented with a clear interpretation? Are the implications of these findings clear?

Is it clear what progress they made or what future work needs to be done?

How close are the authors on providing an answer to their problem statement?

