

Final Project – Deep Neural Networks (CS/DS 541, Whitehill, Spring 2021)

1 Introduction & Scope

For the remainder of the course, you will define and conduct a small-scale research project consisting of some machine learning problem that you will tackle using neural networks as the computational framework. The topic is mostly up to you, but you will need to submit a brief proposal (see below) describing what it is you are trying to accomplish. In addition to defining the problem, implementing neural network(s) to try to solve it, and conducting computational experiments, you will also need to write a short (4-5 page, depending on your group size) describing your work.

All of the following kinds of projects are acceptable:

- **New domain:** Tackling a novel (i.e., never before tried) machine learning problem using neural networks.
- **Old domain, new approach:** Devising a novel neural network architecture to try to outperform an existing baseline on a machine learning problem.
- **Fancy network implementation:** Implementing and training a non-trivial neural network architecture yourself **from scratch** (using any standard neural networks software package such as TensorFlow, PyTorch, etc.).

You may **not** do a project on handwritten digit recognition using the MNIST dataset – it’s time for something new.

2 Proposal

By 11:59pm on Monday, 15 March 2021, please submit a concise and precise **3-4 paragraph** description of the problem you want to tackle including:

1. Who are the members of the project team, and how will the work of the project be distributed among the individual team members; **Make sure you discuss how you will be communicating with each other during the collaboration**, e.g., Zoom, Slack, fax, etc.
2. What computational problem are you tackling, and what do you want to accomplish.
3. What are the data you have available for training. What do they consist of (inputs, target labels), how many do you have, how much diversity do they exhibit, etc.
4. How will you measure performance of your neural network. In particular, you should define a **reasonable baseline** that you are trying to beat. If the problem has never before been tackled and it’s unclear whether a neural network could be useful at all, then the baseline might consist of just “guessing” the target labels (e.g., the mean value for regression, or a uniform distribution for classification). If there’s already a lot of prior research, then outperforming some well-known machine learning approach (e.g., ImageNet) would be reasonable.

For all four points above, it is crucial that you **elaborate** and **be detailed**. Do not simply write (for example), “We will explore semantic segmentation of images.” Instead, describe which specific techniques you will use, e.g., “We will consider the modification XYZ (explain) to the standard W-Net approach to semantic segmentation. The motivation is that XYZ can better achieve ABC because of MNO.”

Also, it is fine to use somebody else’s code (e.g., via GitHub) as the starting point for your research. However, you **must** make clear in your proposal that your own contributions will extend this prior work **significantly** – say **how** this will work. Obviously, you will not know all the details up front, but you should give a compelling plan as to what your project might consist of.

I will provide feedback on your project proposal. You **must** obtain approval from me before proceeding with the project.

3 Paper

As described in class, every team is required to produce a project report. See `CS541FinalProject2021.zip` in the Files section of Canvas for more details. It is fine (and recommended) to use the LaTeX template; however, if you prefer another text-editing environment (e.g., Google Docs, Microsoft Word), that is fine too – **just make sure that your paper looks very similar (in formatting, not content) to the template I provided.**

Groups of size 3-4 should produce a 3-page report, whereas groups of size 5 should produce a 4-page report. These page limits do not include references, which can extend up to infinity pages.

4 Presentation

During the last day of class, each team will give a short presentation of their project. Groups of size 3-4 should record a YouTube-based 3-minute presentation, whereas groups of size 5 should produce a 4-minute presentation. The top 10 groups, as rated by peer reviews (more details later), will present live during class (via Zoom), and only they will be eligible to win the Fabulous Prize.

5 Software

In contrast to the previous homework assignments, you are free to use whatever neural network library you wish, e.g., TensorFlow, PyTorch, Keras, etc.

6 Deliverables

Your final project submission will consist of (1) all your code that you used to implement your project and conduct experiments; (2) your final report; and (3) a link to the YouTube video containing your project presentation.

7 Teams

You are required to form a team, ideally of 3-4 people, but up to 5 people is also ok (as noted above).

8 Grading

Projects will be graded primarily on the basis of the presentation and the paper. The maximum score is 100 points, divided as follows:

- **Methodology (25 points):** What procedures did you use to design and optimize both the parameters and hyperparameters of your network?
- **Performance (20 points):** How much more accurate, relative to the difficulty of the problem you defined, did your network get?
- **Thoroughness (20 points):** Obtaining one research result often opens up many more questions. Does your project offer clear “lessons learned”, or does it simply report a few numbers with little insight? Were there obvious research questions that could have easily been explored but weren’t?

- **Writing (20 points):** Is the paper understandable? Are all important details clear? Is it well organized? Does the paper cite relevant prior research and explain the significance of the results relative to the prior state-of-the-art? (Note: for “fancy network implementation” projects, make sure that you relate the network you are implementing to prior architectures and motivate why it is computationally interesting, effective, or promising.)
- **Oral Presentation (15 points):** How well did your presentation motivate the topic; how well did it explain your approach to tackling the problem; and how clearly did it present your results?

To earn a good grade, please make sure that your paper clearly describes your methodology and empirical results (to evaluate performance), and highlights your thoroughness in examining your research problem. Note that, for a course-based project such as this one, **negative results are acceptable**. While desirable, it is **not** necessary to obtain excellent performance in order to earn a good grade. As long as your methodology is sound and you conducted a thorough (within the confines of a 6-week project) exploration of your target domain, you can still earn an A.