Theoretical Foundations for Deep Learning: Research Project

Instructor: Ankur Moitra

Due Dates: April 9th (proposal), May 19th

The main assignment for the semester is to do an independent reading or research project related to the material we've covered in class. You can work in groups of up to four students or you can work alone. Some guidelines:

(1) For a reading project, you should select 2-3 related papers and give an exposition of their main results. The goal is **not** to merely reproduce what's already in the papers, but to create something that would be independently valuable as a resource. This could involve finding a more self-contained and intuitive proof of some of the key results. Often there are easier ways to see what is going on if you can identify a toy version of the main theorems. It could also involve explaining the connections between the works. Sometimes these connections are not explicitly explained in the papers themselves, or only become clear in retrospect. You could also put the works in context by explaining connections to material we've covered in class, or other classic learning results. As much as possible, your writeup should emphasize what your new contribution is.

(2) You can also do independent research. This is often challenging, and if you go this route I encourage you to think about what might be intermediate results on the way towards what you are aiming for. It is always good to have a range of problems, with at least some of them being approachable enough that you have a good reason to believe you can make some kind of progress in a limited amount of time. Even if you end up not proving what you set out to prove, you can still write up some preliminary ideas along the way or at the very least write up a literature review focused on the open question you worked on. If you choose this type of project, it is important that you come talk to me about it so that I can talk to you about your ideas or point you to other papers that I think might be relevant.

(3) In class, we've focused on provable guarantees for approximation, optimization, generalization and representation learning. Thus I expect that, even if you focus on experiments, there will be at least some theoretical contribution. I will be quite flexible in what exactly this means. It could be in formulating a new question, designing an experiment to test out some of the ideas we've seen in class, or even formulating a toy model that you feel provides new insights.

(4) You are welcome to work with students in the Harvard sister class. The final project reports can overlap, but be aware of the different deadlines. This semester, Harvard and MIT are on rather different schedules. Also the research project should not be the same as any other work you are doing on your own (e.g. if you are a graduate student and already working on a project with your advisor or another team of students that intersects with the course material). I want you to pick something you wouldn't otherwise have done.

(5) **Project Proposal and Meetings**: You should write a one-age project proposal and submit it by April 9th via Gradescope. It should contain an introduction to the reading or research topic that would be understandable to someone else taking the class,. It should also have logistical information about who is in your group and what papers you plan to read, or what research questions you plan to work on. Also I will want to meet with each group sometime in late April or the beginning of May.

Many of you may already have an idea of what you want to work on. But for those of you that don't, I will create a list of suggested topics, which I will continue to update as I come up with more. If you pick one of these topics, you should also explore some of the surrounding literature and look for other papers that study related questions.

- **Generalization Bounds I**
  "Spectrally-normalized Margin Bounds for Neural Networks", Bartlett, Foster, Telgarsky.
- **Generalization Bounds II**
  "Size-Independent Sample Complexity of Neural Networks", Golowich, Rakhlin, Shamir.
- **Neural Tangent Kernels I**
  "On Exact Computation with an Infinitely Wide Neural Net", Arora, Du, Hu, Li, Salakhutdinov, Wang.
- **Neural Tangent Kernels II**
  "Tensor Programs I: Wide Feedforward or Recurrent Neural Networks of Any Architecture are Gaussian Processes", Yang.
- **Depth Separations I**
  "Depth Separations in Neural Networks: What is Actually Being Separated?", Safran, Eldan, Shamir.
- **Depth Separations II**
  "On the Expressive Power of Deep Learning: A Tensor Analysis", Cohen, Sharir, Shashua.

- **Learning Deep Nets I**
  "Learning and Generalization in Overparameterized Neural Networks, Going Beyond Two Layers", Allen-Zhu, Li, Liang

- **Learning Deep Nets II**
  "Learning Neural Networks with Two Nonlinear Layers in Polynomial Time", Goel, Klivans

- **Learning Deep Nets III**
  "Learning Two-layer Neural Networks with Symmetric Inputs", Ge, Kuditipudi, Wang, Li

- **Implicit Regularization**
  "Implicit Regularization in Deep Learning May Not Be Explainable by Norms", Razin, Cohen