DSC 190/291 · Assignment 3

UCSD · Spring 2026

Released: Friday, April 17 · Due: Friday, April 24, 11:59 PM

AI policy. AI assistance is allowed and encouraged in this course. You may use AI to learn the material, explore proof structure, test examples, debug code or formalizations, and improve exposition. However, you are responsible for checking correctness and for standing behind every proof step, derivation, formalization, experiment, and explanation you submit. Use AI as a collaborator, not as an oracle: do not submit anything you cannot explain and verify. The AI usage report is a required component of the assignment.

Submission. Submit a single PDF on Gradescope containing your write-up, figures, and discussion. Also place any supporting artifacts for the assignment in your course repository under the appropriate assignment directory. This may include code, Lean files, notebooks, scripts, data, or other materials needed to inspect or reproduce your work. Your submission should make it clear how the repository artifacts relate to the write-up.

Part A: A Mini-Course on Concentration Inequalities (65 points)

In the Week 3 lecture we proved the following i.i.d. growth-function ERM guarantee. Let $\hat{h} \in \operatorname{argmin}_{h \in \mathcal{H}} L_n(h)$. Then with high probability over $S \sim \mathcal{D}^n$,

$$L_{\mathcal{D}}\left(\hat{h}\right) \leq \inf_{h \in \mathcal{H}} L_{\mathcal{D}}(h) + O\left(\sqrt{\frac{\log \Gamma_{\mathcal{H}}(2n) + \log\left(\frac{1}{\delta}\right)}{n}}\right).$$

The learning-theoretic structure of the proof was covered in lecture, but several concentration inequalities were used without proof. The inequalities to learn in this mini-course are:

- Hoeffding's inequality,
- Hoeffding's inequality for sampling without replacement,
- McDiarmid's inequality,
- Bernstein's inequality (an additional result worth knowing).

Goal. Create a short mini-course that teaches you these four concentration inequalities well enough to complete the Week 3 ERM proof end-to-end. The mini-course should be a single coherent document, written so that someone unfamiliar with the material could learn from it. Organize it in roughly the following order:

1. **The concept of concentration.** Explain what concentration inequalities are as a class of statements: what kind of random quantity they describe, what it means for such a quantity to concentrate, and why we should expect this phenomenon for sums and well-behaved functions of many independent random variables. Use at least one concrete example to ground the discussion. What is the intuition behind concentration inequalities?
2. **The common technique.** Present the moment-generating-function method that underlies the proofs of Hoeffding and Bernstein, including the Chernoff bound as a general template.

Explain why bounding the MGF translates into a tail bound and what one needs to control about the MGF to obtain a useful bound. McDiarmid's inequality uses the bounded-differences / martingale variant of the same idea; introduce that variant here or in the proofs section, wherever it reads more naturally.

3. **Statements, assumptions, and consequences.** For each of the four inequalities, give a precise statement, state the assumptions under which it holds, and explain what kind of random quantity it controls and how strong the resulting tail bound is. Compare the four concentration inequalities. In particular, make clear what additional structure Bernstein exploits beyond Hoeffding, and what changes when sampling is without replacement.
4. **Proofs.** Derive each of the four inequalities. Use the MGF method developed in the previous section wherever it applies, and the bounded-differences variant for McDiarmid. Is there a common template that can unify the proofs?
5. **Connection to the Week 3 ERM proof.** Combine the previous results to give a full proof of the i.i.d. ERM guarantee stated above.

Reference. Provide detailed references for all results you used, and make sure everything can be checked against these references.

Format. Submit one coherent write-up. Tables, diagrams, and proof-dependency maps are welcome where they help.

Part B: The No-Free-Lunch Theorem and the Fundamental Theorem (20 points)

The Week 3 lecture proved the growth-function ERM guarantee displayed in Part A and, via Sauer-Shelah, the VC corollary

$$L_{\mathcal{D}}\left(\hat{h}\right) \le \inf_{h \in \mathcal{H}} L_{\mathcal{D}}(h) + O\left(\sqrt{\frac{d \log(2en/d) + \log(1/\delta)}{n}}\right),$$

where $d = \mathrm{VCdim}(\mathcal{H})$. This establishes one direction of the Fundamental Theorem of PAC learning: if $\mathrm{VCdim}(\mathcal{H}) < \infty$, then $\mathcal{H}$ is PAC learnable. The sample complexity this yields matches the Fundamental Theorem's tight bound up to a logarithmic factor.

The converse direction, that $\mathrm{VCdim}(\mathcal{H}) = \infty$ implies $\mathcal{H}$ is **not** PAC learnable, was stated in lecture but not proved. The standard route to this lower bound is the No-Free-Lunch theorem.

Theorem (No-Free-Lunch)

Let A be any learning algorithm for binary classification with the 0-1 loss over a domain $\mathcal{X}$, and let n be any integer with $n < |\mathcal{X}|/2$. Then there exists a distribution $\mathcal{D}$ over $\mathcal{X} \times \{0,1\}$ such that (i) some function $f^* : \mathcal{X} \to \{0,1\}$ satisfies $L_{\mathcal{D}}(f^*) = 0$, and (ii) with probability at least $1/7$ over $S \sim \mathcal{D}^n$, the learner's output satisfies $L_{\mathcal{D}}(A(S)) \ge 1/8$.

In this part you will prove NFL, use it to close the lower-bound direction, and demonstrate the construction concretely on objects you choose yourself. Also, you will compare the Week 1 and Week 3 versions of NFL. Complete the following tasks:

1. **Proof.** Prove the NFL theorem above. Before writing the proof, study the statement carefully and identify which objects are universally quantified and which are chosen by the adversary. The standard proof uses an averaging argument over the labelings of a $2n$-point subset of $\mathcal{X}$, followed by the probabilistic method to extract a single bad labeling.
2. **Application: closing the Fundamental Theorem.** State the corollary that NFL implies the lower-bound direction of the Fundamental Theorem for PAC learning. Then make the application concrete: choose a hypothesis class with infinite VC dimension, prove that its VC dimension is indeed infinite, and apply the corollary explicitly to that class.
3. **Worked construction.** The NFL proof guarantees the existence of a bad distribution but does not exhibit one. Provide a concrete example of a learning rule and a domain, and explicitly construct a bad distribution and target labeling for which the rule fails with at least constant probability. Verify the failure by calculation.
4. **Comparison to the Week 1 No-Free-Lunch theorem.** In the first lecture we proved a different version of NFL: for any finite domain $\mathcal{X}$ and any deterministic learner A, there exist a function $f : \mathcal{X} \to \{0, 1\}$ and an enumeration $x_1, \ldots, x_n$ of $\mathcal{X}$ ($n = |\mathcal{X}|$) such that A makes n mistakes on the sequence $(x_t, f(x_t))_{t=1}^{n}$. Compare the Week 1 and Week 3 versions. At minimum, identify the **learning setting** each one lives in, the **kind of adversary** each one allows (in particular, whether the adversary is adaptive or oblivious), and the **form of the conclusion** (deterministic vs probabilistic, number of mistakes vs population loss). Explain what each version is telling us and why both are called "No-Free-Lunch" despite being quantitatively very different statements.

Part C: AI Usage Report (15 points)

Write a short report describing how you used AI in this assignment. Do not just list tools; explain what role AI played in your work and how you checked the result. Address:

1. Describe the parts of the assignment for which you used AI. For example: exploring examples, proposing conjectures, checking algebra, debugging code or formalizations, or improving exposition.
2. Describe concrete AI suggestions you accepted and explain why.
3. Describe concrete AI suggestions you rejected or substantially modified, and explain what was wrong, incomplete, or unhelpful about them.
4. Describe how you verified the correctness of what you submitted. Be specific about the relevant kind of work in this assignment: proof, derivation, code, experiment, or exposition.

AI workflow. Also describe concrete updates to your AI workflow that resulted from this assignment. This may include changes to `CLAUDE.md`, `AGENTS.md`, prompts, checklists, scripts, or skills. **Explain the 5 most recent changes you made to your AI workflow and why**.

If you did not use AI for some part of the assignment, say so explicitly.