Lower Bounds and Impossibility Results

Ian Chen

University of Illinois Urbana-Champaign

July 5, 2026

Adversaries

Information Theoretic Bounds

No Free Lunch

Adversaries

Information Theoretic Bounds

No Free Lunch

Deterministic Adversaries

- A majority tree is perfect ternary tree (3^n leaves), s.t.
  - Nodes are labeled 0 or 1
  - Each internal node has label of majority of children
- Is there a deterministic algorithm that has worst case runtime $o(3^n)$?

Deterministic Adversaries (cont.)

- An adversary takes in an algorithm A and constructors a bad instance.
- It must make no assumptions on A
  - i.e. it does not assume that it computes the value of any internal node
  - it can assume that A is a function of the input it has seen.
- Let (x_k, y_k) be the k-th leaf and its value A looks at
- Adversary will iteratively build a tree so that A must look at every leaf.

Deterministic Adversaries (cont.)

- Let B_i, $i = \{0, 1\}$, so that it sets $y_k = f_i(x_k; (x_1, y_1), \ldots, (x_{k-1}, y_{k-1}))$,

$$f_i(x \mid \bullet) = \begin{cases} i & \text{if } x \text{ is the root} \\ 0 & \text{elif no siblings of } x \text{ has value } 0 \\ 1 & \text{elif no siblings of } x \text{ has value } 1 \\ f_i(\mathit{parent}(x)) & \text{otherwise} \end{cases}$$

- Identical until x_{3^n}, so any correct A must look at every leaf.

Randomized Adversaries

- Randomized algorithms take in an instance of a problem, and a random bitstring
- Adversary can design the instance of problem, and observe what A does for bitstrings
- Let approximate median algorithm of array by:
  1. Sample each element with probability $p = o(1/n)$ independently
  2. Run A on the sample
- Claim: exists instances that are off by a fact of 2 with probability $> 1/3$ (for sufficiently large n)

Randomized Adversaries (cont.)

- Let $x = 1$ and $y = 10$
- Let X and Y be arrays size $2n + 1$
  - X has $n + 1$ occurrences of x, n occurrences of y
  - Y has $n + 1$ occurrences of y, n occurrences of x
- Adversary checks probability that $A(X)$ is good (i.e. 2-approximation):
  - If $\Pr\{A(X) \text{ bad}\} > 1/2$, then return X as bad instance
  - Else, return Y as bad instance
- $\Pr\{A(Y) \text{ bad}\} > (1 - \frac{1}{2})(1 - p) > 1/3$

Adversaries

Information Theoretic Bounds

No Free Lunch

Shannon's Source Coding

- Let $X = \{ x_1, \ldots, x_n \}$ encoded/decoded across a channel
- Then, $\mathrm{Ev}(|\mathrm{Enc}(x)|) > H(X) = -\sum_{i=1}^{n} p_x \log_2 p_x$
- Proof by relaxing to a number of bits, and solving optimization problem

Comparisons and Sorting

- Consider any array A with n distinct elements
- Consider X the uniform distribution of all permutations
- $H(X) = \log_2(n!) = \Omega(n \log n)$
- The results of a comparison algorithm leads to an encoding/decoding algorithm

Adversaries

Information Theoretic Bounds

No Free Lunch

PAC and VC-Dimension

- Fundamental theorem of learning theory: TFAE
  1. PAC learnability
  2. Agnostic PAC learnability
  3. Finite VC-dimension
- Finite VC-dimension $\implies$ agnostic PAC (ϵ-nets)
- Agnostic PAC $\implies$ PAC (trivially)
- PAC $\implies$ Finite VC-dimension
  - Infinite VC-dimension $\implies$ not PAC learnable

PAC and VC-Dimension (cont.)

- H of infinite VC-dimension, A learner, m sample size
- Let D of size $2m$ be shattered by H, F all 2^{2^m} labelings
- Average error over all labelings $f \in F$ is $\geq 1/4$
  - Any point not in the sample has average error $1/2$
- For $f*$ labeling with error $\geq 1/4$, then

$$\Pr_S\left(\epsilon_D(A(S)) \leq \frac{1}{8}\right) \leq \frac{1 - 1/4}{1 - 1/8} = 1 - \frac{1}{7}$$