Learning Algorithms From Game Theory: Calibrated Prediction

Aaron Roth

University of Pennsylvania

April 25 2024

Overview

- In this class, we've frequently used techniques from machine learning to solve game theory problems: Equilibrium computation, online auctions, dynamic pricing, ...

Overview

- In this class, we've frequently used techniques from machine learning to solve game theory problems: Equilibrium computation, online auctions, dynamic pricing, . . .
- Today: the reverse direction: We can derive ML algorithms from game theoretic arguments (the minimax theorem)

Overview

- In this class, we've frequently used techniques from machine learning to solve game theory problems: Equilibrium computation, online auctions, dynamic pricing, . . .
- Today: the reverse direction: We can derive ML algorithms from game theoretic arguments (the minimax theorem)
- In fact, in a strong sense, learning algorithms like polynomial weights are *equivalent* to the minimax theorem.

Calibration

- You turn on your TV, and the weatherman tells you that there is a 10% chance of rain.

Calibration

- You turn on your TV, and the weatherman tells you that there is a 10% chance of rain.
- What does this mean? Today only happens once; not a repeatable event.

Calibration

- You turn on your TV, and the weatherman tells you that there is a 10% chance of rain.
- What does this mean? Today only happens once; not a repeatable event.
- If it doesn't rain, was he wrong? What if it rains?

Calibration

- You turn on your TV, and the weatherman tells you that there is a 10% chance of rain.
- What does this mean? Today only happens once; not a repeatable event.
- If it doesn't rain, was he wrong? What if it rains?
- Is there any way we can test if the weatherman knows what he is doing?

Model

Lets write down a simple model — the weather prediction game.
In rounds $t = 1$ to T:

1. The prediction player predicts some probability p_t of rain, for $p_t \in \{0, 1/m, 2/m, \ldots, (m-1)/m, 1\}$.
2. The outcome $y_t \in \{0, 1\}$ is revealed: it either rains ($y_t = 1$) or it does not ($y_t = 0$).

Model

Lets write down a simple model — the weather prediction game. In rounds $t = 1$ to T:

1. The prediction player predicts some probability p_t of rain, for $p_t \in \{0, 1/m, 2/m, \ldots, (m-1)/m, 1\}$.
2. The outcome $y_t \in \{0, 1\}$ is revealed: it either rains ($y_t = 1$) or it does not ($y_t = 0$).

- Can we devise a test to determine whether the weatherman knows what he is doing?

Devising a Test

- Suppose every day, a probability p_t^* is revealed to the weatherman, and then it rains with that probability: $\Pr[y_t = 1] = p_t^*$.

Devising a Test

- Suppose every day, a probability p_t^* is revealed to the weatherman, and then it rains with that probability: $\Pr[y_t = 1] = p_t^*$.
- If the weatherman predicts $p_t = p_t^*$ he should pass the test. Call him "the oracular weatherman"

Devising a Test

- Suppose every day, a probability p_t^* is revealed to the weatherman, and then it rains with that probability: $\Pr[y_t = 1] = p_t^*$.
- If the weatherman predicts $p_t = p_t^*$ he should pass the test. Call him "the oracular weatherman"
- It should also be possible to fail the test.

Devising a Test

- Suppose every day, a probability p_t^* is revealed to the weatherman, and then it rains with that probability: $\Pr[y_t = 1] = p_t^*$.
- If the weatherman predicts $p_t = p_t^*$ he should pass the test. Call him "the oracular weatherman"
- It should also be possible to fail the test.
- A first attempt:

Definition (Average Consistency)

A prediction strategy satisfies ϵ average consistency if for every sequence of outcomes, the sequence of predictions it generates $(p_1, y_1, \ldots, p_T, y_T)$ satisfies

$$\mathrm{E}\left[\left|\frac{1}{T}\sum_{t=1}^{T} p_t - \sum_{t=1}^{T} y_T\right|\right] \leq \epsilon$$

We say it satisfies average consistency if $\epsilon \to 0$ as $T \to \infty$.

Devising a Test

- The oracular weatherman passes this test (Remember the Chernoff-Hoeffding bound!)

Devising a Test

- The oracular weatherman passes this test (Remember the Chernoff-Hoeffding bound!)
- But the test seems too easy to pass...

Devising a Test

- The oracular weatherman passes this test (Remember the Chernoff-Hoeffding bound!)
- But the test seems too easy to pass...
- Consider the "yesterday weatherman": "On day 1, predict $p_1 = 0$, and on day t, predict $p_t = y_{t-1}$".

Devising a Test

- The oracular weatherman passes this test (Remember the Chernoff-Hoeffding bound!)
- But the test seems too easy to pass...
- Consider the "yesterday weatherman": "On day 1, predict $p_1 = 0$, and on day t, predict $p_t = y_{t-1}$".
- (Just predicts that whatever happened yesterday happens today)

Devising a Test

- The oracular weatherman passes this test (Remember the Chernoff-Hoeffding bound!)
- But the test seems too easy to pass...
- Consider the "yesterday weatherman": "On day 1, predict $p_1 = 0$, and on day t, predict $p_t = y_{t-1}$".
- (Just predicts that whatever happened yesterday happens today)
- $\left| \frac{1}{T} \sum_{t=1}^{T} p_t - \sum_{t=1}^{T} y_T \right| = y_T / T \leq 1/T$

Devising a Test

- Easy to differentiate the yesterday weatherman from the oracular weatherman. (How?)

Devising a Test

- Easy to differentiate the yesterday weatherman from the oracular weatherman. (How?)
- When the oracular weatherman predicts a 100% chance of rain, it *always* rains. But the yesterday weatherman frequently predicts a 100% chance of rain and is wrong.

Devising a Test

- Easy to differentiate the yesterday weatherman from the oracular weatherman. (How?)
- When the oracular weatherman predicts a 100% chance of rain, it *always* rains. But the yesterday weatherman frequently predicts a 100% chance of rain and is wrong.
- The yesterday weatherman violates *prediction conditioned average consistency*.

Devising a Test

- Easy to differentiate the yesterday weatherman from the oracular weatherman. (How?)
- When the oracular weatherman predicts a 100% chance of rain, it *always* rains. But the yesterday weatherman frequently predicts a 100% chance of rain and is wrong.
- The yesterday weatherman violates *prediction conditioned average consistency*.
- Bucket the weatherman's predictions into 100 buckets. Say p_t is in bucket i ($p_t \in B(i)$) if it is closer to $i/100$ than any other point $j/100$.

Devising a Test

Definition

Given a sequence of predictions and outcomes $(p_1, y_1, \ldots, p_T, y_T)$, let $n_T(i) = |\{t : p_t \in B(i)\}|$ be the number of rounds on which the prediction was in bucket i. The sequence satisfies ϵ-prediction conditioned average consistency for a bucket i if:

$$\left| \frac{\sum_{t:p_t \in B(i)} y_t - p_t}{n_T(i)} \right| \leq \epsilon$$

Devising a Test

Definition

Given a sequence of predictions and outcomes $(p_1, y_1, \ldots, p_T, y_T)$, let $n_T(i) = |\{t : p_t \in B(i)\}|$ be the number of rounds on which the prediction was in bucket i. The sequence satisfies ϵ-prediction conditioned average consistency for a bucket i if:

$$\left| \frac{\sum_{t:p_t \in B(i)} y_t - p_t}{n_T(i)} \right| \leq \epsilon$$

- i.e. conditioned $p_t \approx i/100$ probability of rain, it should rain roughly a $i/100$ fraction of the time.

Devising a Test

Definition

Given a sequence of predictions and outcomes $(p_1, y_1, \ldots, p_T, y_T)$, let $n_T(i) = |\{t : p_t \in B(i)\}|$ be the number of rounds on which the prediction was in bucket i. The sequence satisfies ϵ-prediction conditioned average consistency for a bucket i if:

$$\left| \frac{\sum_{t:p_t \in B(i)} y_t - p_t}{n_T(i)} \right| \leq \epsilon$$

- i.e. conditioned $p_t \approx i/100$ probability of rain, it should rain roughly a $i/100$ fraction of the time.
- Idea for calibration: Forecaster should be correct on average, conditioned on her forecast.

Calibration

- Idea: Ask for conditional consistency for all 100 buckets.

Calibration

- Idea: Ask for conditional consistency for all 100 buckets.
- Problem: Even the oracular weatherman wouldn't satisfy this for buckets that were infrequently used.

Calibration

- Idea: Ask for conditional consistency for all 100 buckets.
- Problem: Even the oracular weatherman wouldn't satisfy this for buckets that were infrequently used.
- But can ask for it on average:

Definition

A prediction strategy satisfies ϵ-average calibration if for all sequences of outcomes, the sequence of predictions it generates $(p_1, y_1, \ldots, p_T, y_T)$ satisfies:

$$\mathrm{E}\left[\sum_{i=1}^{100} \frac{n_T(i)}{T} \cdot \left| \frac{\sum_{t:p_t \in B(i)} y_t - p_t}{n_T(i)} \right| \right] =$$

$$\frac{1}{T}\mathrm{E}\left[\sum_{i=1}^{100} \left| \sum_{t=1}^{T} \mathbb{1}[p_t \in B(i)](y_t - p_t) \right| \right] \leq \epsilon$$

We say it satisfies average calibration if $\epsilon \to 0$ as $T \to \infty$

Calibration

- More convenient to instead work with a "Euclidean" metric of calibration error:

$$L_T = \sum_{i=1}^{100} \left(\sum_{t=1}^{T} \mathbb{1}[p_t \in B(i)](y_t - p_t) \right)^2$$

Calibration

- More convenient to instead work with a "Euclidean" metric of calibration error:

$$L_T = \sum_{i=1}^{100} \left(\sum_{t=1}^{T} \mathbb{1}[p_t \in B(i)](y_t - p_t) \right)^2$$

- Can confirm ("Cauchy-Schwartz inequality") that the calibration error is upper bounded by:

$$\frac{1}{T} \mathrm{E} \left[\sum_{i=1}^{100} \left| \sum_{t=1}^{T} \mathbb{1}[p_t \in B(i)](y_t - p_t) \right| \right] \leq \mathrm{E} \left[\frac{10}{T} \sqrt{L_T} \right] \leq \frac{10}{T} \sqrt{\mathrm{E}[L_T]}$$

Calibration

- More convenient to instead work with a "Euclidean" metric of calibration error:

$$L_T = \sum_{i=1}^{100} \left(\sum_{t=1}^{T} \mathbb{1}[p_t \in B(i)](y_t - p_t) \right)^2$$

- Can confirm ("Cauchy-Schwartz inequality") that the calibration error is upper bounded by:

$$\frac{1}{T}\mathrm{E}\left[\sum_{i=1}^{100} \left| \sum_{t=1}^{T} \mathbb{1}[p_t \in B(i)](y_t - p_t) \right| \right] \leq \mathrm{E}\left[\frac{10}{T}\sqrt{L_T} \right] \leq \frac{10}{T}\sqrt{\mathrm{E}[L_T]}$$

- Our goal: Develop an algorithm to allow a fraudulent weatherman to pass this test *no matter what*.

Deriving The Fraudulent Weatherman's Algorithm

- Suppose our weatherman has made predictions up through day $s - 1$, and is considering what to predict on day s.

Deriving The Fraudulent Weatherman's Algorithm

- Suppose our weatherman has made predictions up through day $s-1$, and is considering what to predict on day s.
- Let $V^i_{s-1} = \sum_{t=1}^{s-1} \mathbb{1}[p_t \in B(i)](y_t - p_t)$

Deriving The Fraudulent Weatherman's Algorithm

- Suppose our weatherman has made predictions up through day $s-1$, and is considering what to predict on day s.
- Let $V_{s-1}^{i} = \sum_{t=1}^{s-1} \mathbb{1}[p_t \in B(i)](y_t - p_t)$
- If he predicts $p_s \in B(i)$ and the outcome is y_s, then the increase in the loss function is:

$$\Delta_s(p_s, y_s) \quad = \quad L_s - L_{s-1}$$

Deriving The Fraudulent Weatherman's Algorithm

- Suppose our weatherman has made predictions up through day $s-1$, and is considering what to predict on day s.
- Let $V_{s-1}^{i} = \sum_{t=1}^{s-1} \mathbb{1}[p_t \in B(i)](y_t - p_t)$
- If he predicts $p_s \in B(i)$ and the outcome is y_s, then the increase in the loss function is:

$$\begin{aligned}
\Delta_s(p_s, y_s) &= L_s - L_{s-1} \\
&= \left(\sum_{t=1}^{s} \mathbb{1}[p_t \in B(i)](y_t - p_t)\right)^2 \\
&\quad - \left(\sum_{t=1}^{s-1} \mathbb{1}[p_t \in B(i)](y_t - p_t)\right)^2
\end{aligned}$$

Deriving The Fraudulent Weatherman's Algorithm

- Suppose our weatherman has made predictions up through day $s-1$, and is considering what to predict on day s.
- Let $V_{s-1}^i = \sum_{t=1}^{s-1} \mathbb{1}[p_t \in B(i)](y_t - p_t)$
- If he predicts $p_s \in B(i)$ and the outcome is y_s, then the increase in the loss function is:

$$\begin{aligned}
\Delta_s(p_s, y_s) &= L_s - L_{s-1} \\
&= \left(\sum_{t=1}^{s} \mathbb{1}[p_t \in B(i)](y_t - p_t)\right)^2 \\
&\quad - \left(\sum_{t=1}^{s-1} \mathbb{1}[p_t \in B(i)](y_t - p_t)\right)^2 \\
&= \left(V_{s-1}^i + (y_s - p_s)\right)^2 - \left(V_{s-1}^i\right)^2
\end{aligned}$$

Deriving The Fraudulent Weatherman's Algorithm

- Suppose our weatherman has made predictions up through day $s-1$, and is considering what to predict on day s.
- Let $V_{s-1}^i = \sum_{t=1}^{s-1} \mathbb{1}[p_t \in B(i)](y_t - p_t)$
- If he predicts $p_s \in B(i)$ and the outcome is y_s, then the increase in the loss function is:

$$\begin{aligned}
\Delta_s(p_s, y_s) &= L_s - L_{s-1} \\
&= \left(\sum_{t=1}^{s} \mathbb{1}[p_t \in B(i)](y_t - p_t) \right)^2 \\
&\quad - \left(\sum_{t=1}^{s-1} \mathbb{1}[p_t \in B(i)](y_t - p_t) \right)^2 \\
&= \left(V_{s-1}^i + (y_s - p_s) \right)^2 - \left(V_{s-1}^i \right)^2 \\
&\leq 2V_{s-1}^i \cdot (y_s - p_s) + 1
\end{aligned}$$

Deriving The Fraudulent Weatherman's Algorithm

$$\Delta_s(p_s, y_s) \leq 2V_{s-1}^i \cdot (y_s - p_s) + 1$$

Deriving The Fraudulent Weatherman's Algorithm

$$\Delta_s(p_s, y_s) \leq 2V^i_{s-1} \cdot (y_s - p_s) + 1$$

- Suppose our predictions guaranteed: $\mathrm{E}[\Delta_s(p_s, y_s)] \leq \frac{2T}{m} + 1$

Deriving The Fraudulent Weatherman's Algorithm

$$\Delta_s(p_s, y_s) \leq 2V^i_{s-1} \cdot (y_s - p_s) + 1$$

- Suppose our predictions guaranteed: $\mathrm{E}[\Delta_s(p_s, y_s)] \leq \frac{2T}{m} + 1$
- Then we would have:

$$\mathrm{E}[L_T] = \sum_{t=1}^{T} \mathrm{E}[\Delta_t(p_t, y_t)] \leq \frac{2T^2}{m} + T = O\left(\frac{T^2}{m} + T\right)$$

Deriving The Fraudulent Weatherman's Algorithm

$$\Delta_s(p_s, y_s) \leq 2V^i_{s-1} \cdot (y_s - p_s) + 1$$

- Suppose our predictions guaranteed: $\mathrm{E}[\Delta_s(p_s, y_s)] \leq \frac{2T}{m} + 1$
- Then we would have:

$$\mathrm{E}[L_T] = \sum_{t=1}^{T} \mathrm{E}[\Delta_t(p_t, y_t)] \leq \frac{2T^2}{m} + T = O\left(\frac{T^2}{m} + T\right)$$

- And our calibration loss would be bounded by:

$$\epsilon \leq \frac{10}{T}\sqrt{\mathrm{E}[L_T]} = O(\frac{1}{\sqrt{m}} + \frac{1}{\sqrt{T}})$$

Deriving The Fraudulent Weatherman's Algorithm

$$\Delta_s(p_s, y_s) \leq 2V^i_{s-1} \cdot (y_s - p_s) + 1$$

- Suppose our predictions guaranteed: $\mathrm{E}[\Delta_s(p_s, y_s)] \leq \frac{2T}{m} + 1$
- Then we would have:

$$\mathrm{E}[L_T] = \sum_{t=1}^{T} \mathrm{E}[\Delta_t(p_t, y_t)] \leq \frac{2T^2}{m} + T = O\left(\frac{T^2}{m} + T\right)$$

- And our calibration loss would be bounded by:

$$\epsilon \leq \frac{10}{T}\sqrt{\mathrm{E}[L_T]} = O(\frac{1}{\sqrt{m}} + \frac{1}{\sqrt{T}})$$

- $O(1/\sqrt{T})$ if we choose $m = T$. This is our goal.

Learning Via Game Theory

- At round s, define a zero-sum game to guide the Learner's strategy.

Learning Via Game Theory

- At round s, define a zero-sum game to guide the Learner's strategy.
- The Learner (the minimization player) has action set $A_1 = \{1/m, 2/m, \ldots, 1\}$.

Learning Via Game Theory

- At round s, define a zero-sum game to guide the Learner's strategy.
- The Learner (the minimization player) has action set $A_1 = \{1/m, 2/m, \ldots, 1\}$.
- The Adversary (the maximization player) has action set $A_2 = \{0, 1\}$.

Learning Via Game Theory

- At round s, define a zero-sum game to guide the Learner's strategy.
- The Learner (the minimization player) has action set $A_1 = \{1/m, 2/m, \ldots, 1\}$.
- The Adversary (the maximization player) has action set $A_2 = \{0, 1\}$.
- The cost function is:

$$C_s(p, y) = 2V_{s-1}^i \cdot (y_s - p_s) + 1$$

Learning Via Game Theory

- At round s, define a zero-sum game to guide the Learner's strategy.
- The Learner (the minimization player) has action set $A_1 = \{1/m, 2/m, \ldots, 1\}$.
- The Adversary (the maximization player) has action set $A_2 = \{0, 1\}$.
- The cost function is:

$$C_s(p, y) = 2V^i_{s-1} \cdot (y_s - p_s) + 1$$

- Recall: $\Delta_s(p_s, y_s) \leq C_s(p_s, y_s)$

Learning Via Game Theory

$$C_s(p, y) = 2V_{s-1}^i \cdot (y_s - p_s) + 1$$

- What is the $\min\max = \max\min$ value of this game?

Learning Via Game Theory

$$C_s(p, y) = 2V^i_{s-1} \cdot (y_s - p_s) + 1$$

- What is the $\min\max = \max\min$ value of this game?
- Prediction is easier if you know the answer already, so lets consider the $\max\min$ value: corresponds to Adversary *committing* to the probability of rain q_s and telling Learner.

Learning Via Game Theory

$$C_s(p, y) = 2V^i_{s-1} \cdot (y_s - p_s) + 1$$

- What is the $\min\max = \max\min$ value of this game?
- Prediction is easier if you know the answer already, so lets consider the $\max\min$ value: corresponds to Adversary *committing* to the probability of rain q_s and telling Learner.
- $\mathrm{E}_{y\sim q}[C_s(p, y)] = 2V^i_{s-1} \cdot (q_s - p_s) + 1$

Learning Via Game Theory

$$C_s(p, y) = 2V^i_{s-1} \cdot (y_s - p_s) + 1$$

- What is the $\min\max = \max\min$ value of this game?
- Prediction is easier if you know the answer already, so lets consider the $\max\min$ value: corresponds to Adversary *committing* to the probability of rain q_s and telling Learner.
- $\mathrm{E}_{y\sim q}[C_s(p, y)] = 2V^i_{s-1} \cdot (q_s - p_s) + 1$
- Learner can best respond choosing $p_s = \arg\min_{p\in A_1} |q_s - p|$: $|q_s - p_s| \leq 1/m$.

Learning Via Game Theory

$$C_s(p, y) = 2V^i_{s-1} \cdot (y_s - p_s) + 1$$

- What is the $\min\max = \max\min$ value of this game?
- Prediction is easier if you know the answer already, so lets consider the $\max\min$ value: corresponds to Adversary *committing* to the probability of rain q_s and telling Learner.
- $\mathrm{E}_{y \sim q}[C_s(p, y)] = 2V^i_{s-1} \cdot (q_s - p_s) + 1$
- Learner can best respond choosing $p_s = \arg\min_{p \in A_1} |q_s - p|$: $|q_s - p_s| \leq 1/m$.
- So:

$$\max_{q \in \Delta A_2} \min_{p \in A_1} \mathrm{E}_{y \sim q}[C_s(p, y)] \leq \frac{2 \max_i V^i_{s-1}}{m} + 1 \leq \frac{2T}{m} + 1$$

Learning Via Game Theory

- So by the minimax theorem:

$$\min_{\hat{p}\in\Delta A_1}\max_{y\in A_2}\mathrm{E}_{p\sim\hat{p}}[C_s(p,y)] \leq \frac{2\max_i V^i_{s-1}}{m}+1 \leq \frac{2T}{m}+1$$

Learning Via Game Theory

- So by the minimax theorem:

$$\min_{\hat{p}\in\Delta A_1}\max_{y\in A_2} \mathrm{E}_{p\sim\hat{p}}[C_s(p,y)] \leq \frac{2\max_i V_{s-1}^i}{m}+1 \leq \frac{2T}{m}+1$$

- Thus: At every round s, Learner has a strategy $\hat{p}_s$ guaranteeing for all weather outcomes y_s:

$$\mathrm{E}_{p_s\sim\hat{p}_s}[\Delta_s(p_s,y_s)] \leq \frac{2\max_i V_{s-1}^i}{m}+1 \leq \frac{2T}{m}+1$$

Learning Via Game Theory

- So by the minimax theorem:

$$\min_{\hat{p}\in\Delta A_1}\max_{y\in A_2}\mathrm{E}_{p\sim\hat{p}}[C_s(p,y)] \leq \frac{2\max_i V^i_{s-1}}{m}+1 \leq \frac{2T}{m}+1$$

- Thus: At every round s, Learner has a strategy $\hat{p}_s$ guaranteeing for all weather outcomes y_s:

$$\mathrm{E}_{p_s\sim\hat{p}_s}[\Delta_s(p_s,y_s)] \leq \frac{2\max_i V^i_{s-1}}{m}+1 \leq \frac{2T}{m}+1$$

- And so we have proven:

Theorem

There exists a prediction strategy that against an arbitrary adversarially chosen sequence of T outcomes satisfies ϵ-average calibration for $\epsilon = O(1/\sqrt{T})$

The Algorithm?

- We need to compute the $\min \max$ strategy for the learner in the zero sum game.

The Algorithm?

- We need to compute the min max strategy for the learner in the zero sum game.
- We know how to do that with efficiently polynomial weights!

The Algorithm?

- We need to compute the $\min\max$ strategy for the learner in the zero sum game.
- We know how to do that with efficiently polynomial weights!
- But maybe there is a more efficient direct solution...

The Algorithm?

$$\Delta_s(p_s, y_s) \leq 2V^i_{s-1} \cdot (y_s - p_s) + 1$$

Needed: A strategy guaranteeing $\mathrm{E}[\Delta_s(p_s, y_s)] \leq \frac{2T}{m} + 1$.

The Algorithm?

$$\Delta_s(p_s, y_s) \leq 2V^i_{s-1} \cdot (y_s - p_s) + 1$$

Needed: A strategy guaranteeing $\mathrm{E}[\Delta_s(p_s, y_s)] \leq \frac{2T}{m} + 1$.

- Case 1: $V^i_{s-1} \geq 0$ for all i: Predict $p_s = 1$. Then:

$$\Delta_s(p_s, y_s) \leq 2V^i_{s-1} \cdot (y_s - 1) + 1 \leq 1$$

The Algorithm?

$$\Delta_s(p_s, y_s) \leq 2V_{s-1}^i \cdot (y_s - p_s) + 1$$

Needed: A strategy guaranteeing $\mathrm{E}[\Delta_s(p_s, y_s)] \leq \frac{2T}{m} + 1$.

- Case 1: $V_{s-1}^i \geq 0$ for all i: Predict $p_s = 1$. Then:

$$\Delta_s(p_s, y_s) \leq 2V_{s-1}^i \cdot (y_s - 1) + 1 \leq 1$$

- Case 2: $V_{s-1}^i \leq 0$ for all i: Predict $p_s = 0$. Then:

$$\Delta_s(p_s, y_s) \leq 2V_{s-1}^i \cdot (y_s - 0) + 1 \leq 1$$

The Algorithm?

$$\Delta_s(p_s, y_s) \leq 2V^i_{s-1} \cdot (y_s - p_s) + 1$$

Needed: A strategy guaranteeing $\mathrm{E}[\Delta_s(p_s, y_s)] \leq \frac{2T}{m} + 1$.

- Case 1: $V^i_{s-1} \geq 0$ for all i: Predict $p_s = 1$. Then:

$$\Delta_s(p_s, y_s) \leq 2V^i_{s-1} \cdot (y_s - 1) + 1 \leq 1$$

- Case 2: $V^i_{s-1} \leq 0$ for all i: Predict $p_s = 0$. Then:

$$\Delta_s(p_s, y_s) \leq 2V^i_{s-1} \cdot (y_s - 0) + 1 \leq 1$$

- Otherwise: There must exist an i such that $V^i_{s-1} \geq 0$ and $V^{i+1}_{s-1} \leq 0$ or vice versa.

The Algorithm?

$$\Delta_s(p_s, y_s) \le 2V^i_{s-1} \cdot (y_s - p_s) + 1$$

Needed: A strategy guaranteeing $\mathrm{E}[\Delta_s(p_s, y_s)] \le \frac{2T}{m} + 1$.

The Algorithm?

$$\Delta_s(p_s, y_s) \leq 2V^i_{s-1} \cdot (y_s - p_s) + 1$$

Needed: A strategy guaranteeing $\mathrm{E}[\Delta_s(p_s, y_s)] \leq \frac{2T}{m} + 1$.

- Case 3: Let $q \in [0, 1]$ be a probability such that $qV^i_{s-1} + (1-q)V^{i+1}_{s-1} = 0$.

The Algorithm?

$$\Delta_s(p_s, y_s) \leq 2V_{s-1}^i \cdot (y_s - p_s) + 1$$

Needed: A strategy guaranteeing $\mathrm{E}[\Delta_s(p_s, y_s)] \leq \frac{2T}{m} + 1$.

- Case 3: Let $q \in [0, 1]$ be a probability such that $qV_{s-1}^i + (1-q)V_{s-1}^{i+1} = 0$.
- Let $p = \arg\max\{p \in B(i)\}$, $p' = \arg\min\{p' \in B(i+1)\}$. Note $p' = p + 1/m$.

The Algorithm?

$$\Delta_s(p_s, y_s) \leq 2V^i_{s-1} \cdot (y_s - p_s) + 1$$

Needed: A strategy guaranteeing $\mathrm{E}[\Delta_s(p_s, y_s)] \leq \frac{2T}{m} + 1$.

- Case 3: Let $q \in [0, 1]$ be a probability such that $qV^i_{s-1} + (1 - q)V^{i+1}_{s-1} = 0$.
- Let $p = \arg\max\{p \in B(i)\}$, $p' = \arg\min\{p' \in B(i + 1)\}$. Note $p' = p + 1/m$.
- Play $p_s = p$ with probability q and $p_s = p'$ w.p. $(1 - q)$

The Algorithm?

$$\Delta_s(p_s, y_s) \le 2V^i_{s-1} \cdot (y_s - p_s) + 1$$

Needed: A strategy guaranteeing $\mathrm{E}[\Delta_s(p_s, y_s)] \le \frac{2T}{m} + 1$.

- Case 3: Let $q \in [0, 1]$ be a probability such that $qV^i_{s-1} + (1-q)V^{i+1}_{s-1} = 0$.
- Let $p = \arg\max\{p \in B(i)\}$, $p' = \arg\min\{p' \in B(i+1)\}$. Note $p' = p + 1/m$.
- Play $p_s = p$ with probability q and $p_s = p'$ w.p. $(1-q)$
- Then:

$$\mathrm{E}[\Delta_s(p_s, y_s)] \le 2qV^i_{s-1} \cdot (y_s - p) + 2(1-q)V^{i+1}_{s-1}(y_s - p - 1/m) + 1$$

$$\le \frac{2|V^{i+1}_{s-1}|}{m} + 1 \le \frac{2T}{m} + 1$$

Reflecting

- Argument was generic to any linear (i.e. based on bounding sums or averages) test aimed at distinguishing the oracular weatherman from a fraud.

Reflecting

- Argument was generic to any linear (i.e. based on bounding sums or averages) test aimed at distinguishing the oracular weatherman from a fraud.
- Because the minimax theorem literally is allowing us to analyze the Learner as if she is the oracular weatherman!

Reflecting

- Argument was generic to any linear (i.e. based on bounding sums or averages) test aimed at distinguishing the oracular weatherman from a fraud.
- Because the minimax theorem literally is allowing us to analyze the Learner as if she is the oracular weatherman!
- What does this mean about what we can learn from empirical tests of probabilistic models?

Thanks!

See you next class — stay healthy!