Card Guessing with Feedback

Sam Spiro, Rutgers University.

Feedback Models

Feedback Models

Start with a deck of mn cards where there are n card types each appearing with multiplicity m.

Feedback Models

Start with a deck of mn cards where there are n card types each appearing with multiplicity m. For example, $n = 13$ and $m = 4$ corresponds to a usual deck of playing cards.

Feedback Models

Let $\mathcal{C}^{+}_{m,n}$ and $\mathcal{C}^{-}_{m,n}$ be the maximum and minimum expected scores that the player can get in the complete feedback model.

Feedback Models

Let $\mathcal{C}^{+}_{m,n}$ and $\mathcal{C}^{-}_{m,n}$ be the maximum and minimum expected scores that the player can get in the complete feedback model.

Theorem (Diaconis-Graham, 1981)

For n fixed,

$$\mathcal{C}^{\pm}_{m,n} = m \pm c_n\sqrt{m} + o_n(\sqrt{m}).$$

Feedback Models

Feedback Models

In the "partial feedback model", the Guesser guesses the next card and is only told whether their guess was correct or not.

Feedback Models

In the "partial feedback model", the Guesser guesses the next card and is only told whether their guess was correct or not. If $\mathcal{P}^{+}_{m,n}$ is the maximum expected score in this model,

$$m \leq \mathcal{P}^{+}_{m,n}$$

Feedback Models

In the "partial feedback model", the Guesser guesses the next card and is only told whether their guess was correct or not. If $\mathcal{P}^+_{m,n}$ is the maximum expected score in this model,

$$m \leq \mathcal{P}^+_{m,n} \leq \mathcal{C}^+_{m,n}$$

Feedback Models

In the "partial feedback model", the Guesser guesses the next card and is only told whether their guess was correct or not. If $\mathcal{P}^+_{m,n}$ is the maximum expected score in this model,

$$m \leq \mathcal{P}^+_{m,n} \leq \mathcal{C}^+_{m,n} = m + o_n(m)$$

Feedback Models

In the "partial feedback model", the Guesser guesses the next card and is only told whether their guess was correct or not. If $\mathcal{P}^+_{m,n}$ is the maximum expected score in this model,

$$m \leq \mathcal{P}^+_{m,n} \leq \mathcal{C}^+_{m,n} = m + o_n(m),$$

What happens when n is large?

Feedback Models

Theorem (Diaconis-Graham-He-S., 2020)

For m fixed,

$$\mathcal{C}^{+}_{m,n} \sim H_m \log(n),$$
$$\mathcal{C}^{-}_{m,n} = \Theta(n^{-1/m}),$$

where H_m is the mth harmonic number.

Feedback Models

Theorem (Diaconis-Graham-He-S., 2020)

For m fixed,

$$\mathcal{C}^{+}_{m,n} \sim H_m \log(n),$$
$$\mathcal{C}^{-}_{m,n} = \Theta(n^{-1/m}),$$

where H_m is the mth harmonic number.

With this we have the trivial bounds

$$m \leq \mathcal{P}^{+}_{m,n} \leq \mathcal{C}^{+}_{m,n} = O_m(\log n).$$

Feedback Models

$$m \leq \mathcal{P}^{+}_{m,n} \leq O_m(\log n).$$

Feedback Models

$$m \leq \mathcal{P}^{+}_{m,n} \leq O_m(\log n).$$

Theorem (Diaconis-Graham-He-S., 2020+)

There exist $c, C > 0$ such that if n is sufficiently large in terms of m, we have

$$m + c\sqrt{m} \leq \mathcal{P}^{+}_{m,n}$$

Feedback Models

$$m \leq \mathcal{P}^{+}_{m,n} \leq O_m(\log n).$$

Theorem (Diaconis-Graham-He-S., 2020+)

There exist $c, C > 0$ such that if n is sufficiently large in terms of m, we have

$$m + c\sqrt{m} \leq \mathcal{P}^{+}_{m,n} \leq m + Cm^{3/4} \log m.$$

A Partial Proof

Lemma

Assume that we have played in the partial feedback model for $t-1$ rounds such that we have guessed card type i a total of g_i times

A Partial Proof

Lemma

Assume that we have played in the partial feedback model for $t-1$ rounds such that we have guessed card type i a total of g_i times, and let S be the total number of points scored.

A Partial Proof

Lemma

Assume that we have played in the partial feedback model for $t-1$ rounds such that we have guessed card type i a total of g_i times, and let S be the total number of points scored. Given this, we have

$$\Pr[\pi_t = i] \leq \frac{m}{mn - g_i - S}.$$

A Partial Proof

Lemma

Assume that we have played in the partial feedback model for $t - 1$ rounds such that we have guessed card type i a total of g_i times, and let S be the total number of points scored. Given this, we have

$$\Pr[\pi_t = i] \leq \frac{m}{mn - g_i - S}.$$

That is, our upper bound is strongest when g_i and S is small.

A Partial Proof

Lemma

Assume that we have played in the partial feedback model for $t - 1$ rounds such that we have guessed card type i a total of g_i times, and let S be the total number of points scored. Given this, we have

$$\Pr[\pi_t = i] \leq \frac{m}{mn - g_i - S}.$$

That is, our upper bound is strongest when g_i and S is small. These conditions are necessary

A Partial Proof

Lemma

Assume that we have played in the partial feedback model for $t-1$ rounds such that we have guessed card type i a total of g_i times, and let S be the total number of points scored. Given this, we have

$$\Pr[\pi_t = i] \leq \frac{m}{mn - g_i - S}.$$

That is, our upper bound is strongest when g_i and S is small. These conditions are necessary: if i has been guessed incorrectly $g_i = mn - m$ times, then we know the card must be an i.

A Partial Proof

Corollary

$$\mathcal{P}^{+}_{m,n} \leq 3m + o(m).$$

A Partial Proof

Corollary

$$\mathcal{P}^{+}_{m,n} \leq 3m + o(m).$$

For all i and t, we have

$$\Pr[\pi_t = i] \leq \frac{m}{mn - g_i - S}$$

A Partial Proof

Corollary

$$\mathcal{P}^{+}_{m,n} \leq 3m + o(m).$$

For all i and t, we have

$$\Pr[\pi_t = i] \leq \frac{m}{mn - g_i - S} \approx \frac{m}{mn - g_i},$$

A Partial Proof

Corollary

$$\mathcal{P}^{+}_{m,n} \leq 3m + o(m).$$

For all i and t, we have

$$\Pr[\pi_t = i] \leq \frac{m}{mn - g_i - S} \approx \frac{m}{mn - g_i},$$

At most one i is guessed more than $mn/2$ times.

A Partial Proof

Corollary

$$\mathcal{P}^{+}_{m,n} \leq 3m + o(m).$$

For all i and t, we have

$$\Pr[\pi_t = i] \leq \frac{m}{mn - g_i - S} \approx \frac{m}{mn - g_i},$$

At most one i is guessed more than $mn/2$ times. Every other j has $\Pr[\pi_t = j] \leq \frac{2}{n}$ for all t.

A Partial Proof

Corollary

$$\mathcal{P}^{+}_{m,n} \leq 3m + o(m).$$

For all i and t, we have

$$\Pr[\pi_t = i] \leq \frac{m}{mn - g_i - S} \approx \frac{m}{mn - g_i},$$

At most one i is guessed more than $mn/2$ times. Every other j has $\Pr[\pi_t = j] \leq \frac{2}{n}$ for all t. Thus in expectation at most $mn \cdot (2/n) = 2m$ cards are guessed correctly from this part

A Partial Proof

Corollary

$$\mathcal{P}^{+}_{m,n} \leq 3m + o(m).$$

For all i and t, we have

$$\Pr[\pi_t = i] \leq \frac{m}{mn - g_i - S} \approx \frac{m}{mn - g_i},$$

At most one i is guessed more than $mn/2$ times. Every other j has $\Pr[\pi_t = j] \leq \frac{2}{n}$ for all t. Thus in expectation at most $mn \cdot (2/n) = 2m$ cards are guessed correctly from this part, and in total at most $3m$ are guessed correctly in expectation. □

Open Problems

Theorem (Diaconis-Graham-He-S., 2020)

$$m + \Omega(\sqrt{m}) \leq \mathcal{P}^{+}_{m,n} \leq m + O(m^{3/4} \log m).$$

Open Problems

Theorem (Diaconis-Graham-He-S., 2020)

$$m + \Omega(\sqrt{m}) \le \mathcal{P}^{+}_{m,n} \le m + O(m^{3/4} \log m).$$

Conjecture (Diaconis-Graham-He-S., 2020)

$$\mathcal{P}^{+}_{m,n} = m + m^{1/2+o(1)}.$$

Open Problems

Recall that $\mathcal{P}^{-}_{m,n}$ is the minimum number of points one can get in expectation with partial feedback.

Open Problems

Recall that $\mathcal{P}^{-}_{m,n}$ is the minimum number of points one can get in expectation with partial feedback. We trivially have

$$m \geq \mathcal{P}^{-}_{m,n} \geq \mathcal{C}^{-}_{m,n} = \Theta(n^{-1/m}).$$

Open Problems

Recall that $\mathcal{P}^{-}_{m,n}$ is the minimum number of points one can get in expectation with partial feedback. We trivially have

$$m \geq \mathcal{P}^{-}_{m,n} \geq \mathcal{C}^{-}_{m,n} = \Theta(n^{-1/m}).$$

Theorem (Diaconis-Graham-S., 2020)

$$\mathcal{P}^{-}_{m,n} \geq 1 - e^{-m} - o_m(1) \geq \frac{1}{2}.$$

Open Problems

Theorem (Diaconis-Graham-S., 2020)

$$\mathcal{P}^{-}_{m,n} \geq 1 - e^{-m} - o_m(1) \geq \frac{1}{2}.$$

Open Problems

Theorem (Diaconis-Graham-S., 2020)

$$\mathcal{P}^{-}_{m,n} \geq 1 - e^{-m} - o_m(1) \geq \frac{1}{2}.$$

Conjecture (Diaconis-Graham-S., 2020)

If n is sufficiently large in terms of m, then

$$\mathcal{P}^{-}_{m,n} \sim m.$$

Practical Strategies

Practical Strategies

Theorem (Diaconis-Graham-He-S., 2020)

There exists a simple strategy showing

$$\mathcal{P}^{+}_{m,n} \geq m + \Omega(\sqrt{m}).$$

Practical Strategies

Theorem (Diaconis-Graham-He-S., 2020)

There exists a simple strategy showing

$$\mathcal{P}^{+}_{m,n} \geq m + \Omega(\sqrt{m}).$$

Guess 1 a total of $mn/2$ times, then do one of two things:

Practical Strategies

Theorem (Diaconis-Graham-He-S., 2020)

There exists a simple strategy showing

$$\mathcal{P}^{+}_{m,n} \geq m + \Omega(\sqrt{m}).$$

Guess 1 a total of $mn/2$ times, then do one of two things:

(1) If you made less than $m/2 + \sqrt{m}$ correct guesses, guess 1 the rest of the game.

Practical Strategies

Theorem (Diaconis-Graham-He-S., 2020)

There exists a simple strategy showing

$$\mathcal{P}^{+}_{m,n} \geq m + \Omega(\sqrt{m}).$$

Guess 1 a total of $mn/2$ times, then do one of two things:

(1) If you made less than $m/2 + \sqrt{m}$ correct guesses, guess 1 the rest of the game.

(2) Else guess 2 the rest of the game.

□

Practical Strategies

A very simple strategy is the *safe strategy*, which guesses 1 until m correct guesses are made, then 2 until m correct guesses are made, and so on.

Practical Strategies

A very simple strategy is the *safe strategy*, which guesses 1 until m correct guesses are made, then 2 until m correct guesses are made, and so on. Elementary arguments give that the score for this strategy is

$$m + 1 - \frac{1}{m+1} + O(e^{-\beta m})$$

for some $\beta > 0$.

Practical Strategies

A very simple strategy is the *safe strategy*, which guesses 1 until m correct guesses are made, then 2 until m correct guesses are made, and so on. Elementary arguments give that the score for this strategy is

$$m + 1 - \frac{1}{m+1} + O(e^{-\beta m})$$

for some $\beta > 0$.

Another simple strategy is the *shifting strategy*, which guesses 1 until a correct guess is made, then 2 until a correct guess is made, and so on.

Practical Strategies

If π is a word where each letter in $\{1, 2, \ldots, n\}$ exactly m times, we define $L(\pi)$ to be the largest integer p so that π contains a subsequence of the form $123 \cdots p$.

Practical Strategies

If π is a word where each letter in $\{1, 2, \dots, n\}$ exactly m times, we define $L(\pi)$ to be the largest integer p so that π contains a subsequence of the form $123 \cdots p$. For example,

$$\pi = 2345124351 \implies L(\pi) = 3.$$

Practical Strategies

If π is a word where each letter in $\{1, 2, \ldots, n\}$ exactly m times, we define $L(\pi)$ to be the largest integer p so that π contains a subsequence of the form $123 \cdots p$. For example,

$$\pi = 2345124351 \implies L(\pi) = 3.$$

Note that $L(\pi)$ is (essentially) the score one gets using the shifting strategy if the deck is shuffled according to π.

Practical Strategies

If π is a word where each letter in $\{1, 2, \ldots, n\}$ exactly m times, we define $L(\pi)$ to be the largest integer p so that π contains a subsequence of the form $123 \cdots p$. For example,

$$\pi = 2345124351 \implies L(\pi) = 3.$$

Note that $L(\pi)$ is (essentially) the score one gets using the shifting strategy if the deck is shuffled according to π.

Corollary

If n is sufficiently large in terms of m, then

$$\mathcal{L}_{m,n} := \mathbb{E}[L(\pi)] \leq m + O(m^{3/4} \log m).$$

Practical Strategies

Conjecture (Diaconis-Graham-He-S., 2020)

If n is sufficiently large in terms of m, then

$$\mathcal{L}_{m,n} \sim m.$$

Practical Strategies

Conjecture (Diaconis-Graham-He-S., 2020)

If n is sufficiently large in terms of m, then

$$\mathcal{L}_{m,n} \sim m.$$

Theorem (Clifton-Deb-Huang-S.-Yoo, 2021)

We have

$$\left| \lim_{n\to\infty} \mathcal{L}_{m,n} - \left(m + 1 - \frac{1}{m+2} \right) \right| \leq O(e^{-\beta m})$$

for some $\beta > 0$.

Practical Strategies

More precisely: if $\alpha_1, \ldots, \alpha_m$ are the zeroes of $\sum_{k=0}^{m} \frac{x^k}{k!}$, then

$$\lim_{n\to\infty} \mathcal{L}_{m,n} = -1 - \sum \alpha_i^{-1} e^{-\alpha_i}.$$

Practical Strategies

More precisely: if $\alpha_1, \ldots, \alpha_m$ are the zeroes of $\sum_{k=0}^{m} \frac{x^k}{k!}$, then

$$\lim_{n \to \infty} \mathcal{L}_{m,n} = -1 - \sum \alpha_i^{-1} e^{-\alpha_i}.$$

This implies $\mathcal{L}_{1,n} \to e - 1$

Practical Strategies

More precisely: if $\alpha_1, \ldots, \alpha_m$ are the zeroes of $\sum_{k=0}^{m} \frac{x^k}{k!}$, then

$$\lim_{n\to\infty} \mathcal{L}_{m,n} = -1 - \sum \alpha_i^{-1} e^{-\alpha_i}.$$

This implies $\mathcal{L}_{1,n} \to e - 1$ and that

$$\mathcal{L}_{2,n} \to e(\cos(1) + \sin(1)) - 1.$$

Card Guessing

Card Guessing

The probability of drawing four aces in a row with a deck shuffled uniformly at random is $1/270725$.

Adversarial Card Guessing

The probability of drawing four aces in a row with a deck shuffled uniformly at random is $1/270725$.

Adversarial Card Guessing

The probability of drawing four aces in a row with a deck shuffled uniformly at random is $1/270725$.

More precisely, we are now considering a two player game played by Shuffler and Guesser.

Adversarial Card Guessing

The probability of drawing four aces in a row with a deck shuffled uniformly at random is $1/270725$.

More precisely, we are now considering a two player game played by Shuffler and Guesser. Let $\mathcal{C}_{m,n}(G, S)$ be the expected number of points Guesser scores when the two players follow strategies G, S.

Adversarial Card Guessing

$$\Theta_m(n^{-1/m}) \leq \mathcal{C}_{m,n}(G, \mathrm{Uniform}) \leq H_m \log n + o_m(\log n).$$

Adversarial Card Guessing

$$\Theta_m(n^{-1/m}) \le \mathcal{C}_{m,n}(G, \mathrm{Uniform}) \le H_m \log n + o_m(\log n).$$

Theorem (S., 2021)

There exists a strategy S' *for Shuffler so that*

$$\mathcal{C}_{m,n}(\mathsf{G}, \mathsf{S}') \le \log n + o_m(\log n),$$

and this bound is best possible.

Adversarial Card Guessing

$$\Theta_m(n^{-1/m}) \le \mathcal{C}_{m,n}(G, \mathrm{Uniform}) \le H_m \log n + o_m(\log n).$$

Theorem (S., 2021)

There exists a strategy S' *for Shuffler so that*

$$\mathcal{C}_{m,n}(\mathsf{G}, \mathsf{S}') \le \log n + o_m(\log n),$$

and this bound is best possible.

This theorem is a first for me, since normally I prove a result, then makes jokes about it during my talk.

Adversarial Card Guessing

Theorem

There exists a strategy S' *for Shuffler so that* $\mathcal{C}_{m,n}(\mathsf{G}, \mathsf{S}') \leq \log n + o_m(\log n)$.

Adversarial Card Guessing

Theorem

There exists a strategy S' *for Shuffler so that* $\mathcal{C}_{m,n}(\mathsf{G}, \mathsf{S}') \leq \log n + o_m(\log n)$.

A strategy that gives this is the "greedy strategy", which is such that if there are r types of cards remaining in the deck, then Shuffler draws each of these card types with probability r^{-1} (regardless of how many copies are left in the deck of each type).

Adversarial Card Guessing

Theorem

There exists a strategy S' *for Shuffler so that* $\mathcal{C}_{m,n}(\mathsf{G}, \mathsf{S}') \leq \log n + o_m(\log n)$.

A strategy that gives this is the "greedy strategy", which is such that if there are r types of cards remaining in the deck, then Shuffler draws each of these card types with probability r^{-1} (regardless of how many copies are left in the deck of each type). E.g. if the deck has a hundred 1's and one 2, we draw a 1 or 2 with probability $\frac{1}{2}$.

Adversarial Card Guessing

Theorem

There exists a strategy S' *for Shuffler so that* $\mathcal{C}_{m,n}(\mathsf{G},\mathsf{S}') \leq \log n + o_m(\log n)$.

A strategy that gives this is the “greedy strategy”, which is such that if there are r types of cards remaining in the deck, then Shuffler draws each of these card types with probability r^{-1} (regardless of how many copies are left in the deck of each type). E.g. if the deck has a hundred 1's and one 2, we draw a 1 or 2 with probability $\frac{1}{2}$. This gives the desired bound due to a variant of the coupon collector problem. $\square$

Adversarial Card Guessing

Theorem (S., 2021)

The greedy strategy is the unique strategy that minimizes the number of correct guesses if Guesser tries to maximize their score.

Adversarial Card Guessing

Theorem (S., 2021)

The greedy strategy is the unique strategy that minimizes the number of correct guesses if Guesser tries to maximize their score.

Interestingly, the greedy strategy is also the "unique" strategy which maximizes the number of correct guesses if Guesser tries to minimize their score.

Semi-restricted Games

There is a classical game called "Matching Pennies" where two players simultaneously choose one of n numbers, and if the two match player A gets a point and otherwise player B gets a point.

Semi-restricted Games

There is a classical game called “Matching Pennies” where two players simultaneously choose one of n numbers, and if the two match player A gets a point and otherwise player B gets a point.

The “semi-restricted” version of this game has mn rounds of Matching Pennies is played where player B must use each number exactly m times.

Semi-restricted Games

More generally, one can consider "semi-restricted" versions of any zero sum game.

Semi-restricted Games

More generally, one can consider "semi-restricted" versions of any zero sum game.

Theorem (S.-Surya-Zeng, 2022)

In semi-restricted Rock, Paper, Scissors the "greedy strategy" is the unique optimal strategy for the restricted player.

Semi-restricted Games

Theorem (S.-Surya-Zeng, 2022)

““Almost every”” semi-restricted game fails to have an optimal strategy which is greedy.

Semi-restricted Games

Theorem (S.-Surya-Zeng, 2022)

““Almost every”” semi-restricted game fails to have an optimal strategy which is greedy.

$$\begin{tikzcd}[column sep=small] & 1 \arrow[dr] & \\ 3 \arrow[ur] & & 2 \arrow[ll] \end{tikzcd}$$

Semi-restricted Games

Given a digraph D, we define its skew adjacency matrix A by $A_{u,v} = +1$ if $u \to v$, $A_{u,v} = -1$ if $v \to u$, and $A_{u,v} = 0$ otherwise.

$$\begin{tikzcd} & 1 \arrow[dr] & \\ 3 \arrow[ur] & & 2 \arrow[ll] \end{tikzcd} \qquad \begin{bmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{bmatrix}$$

Semi-restricted Games

Given a digraph D, we define its skew adjacency matrix A by $A_{u,v} = +1$ if $u \to v$, $A_{u,v} = -1$ if $v \to u$, and $A_{u,v} = 0$ otherwise.

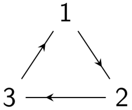

$$\begin{bmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{bmatrix}$$

Question

Which digraphs D are such that their skew-adjacency matrix A satisfies $Null(A) = span(\vec{1})$?