Optimal Strategies from Random Walks

Jacob Abernethy
(UC Berkeley)

Manfred Warmuth
Joel Yellin
(UC Santa Cruz)

The Goal

Wanted:	Optimal **Hedge** Algorithm
Result:	A bit more than expected...

The Game: Gambler Vs Casino

On each day:

I. Casino runs N events with **binary** outcomes

II. Gambler has $1 to bet

III. After bets are placed, Gambler loses all money on lost events

IV. Game stops once all events have **more than K losses**

This is the Hedge Setting

Hedge Algorithm:

$$w_i^{t+1} = w_i^t \cdot \beta^{\ell_i^t}.$$

$$\mathbf{p}^t = \frac{\mathbf{w}^t}{\sum_{i=1}^N w_i^t}.$$

Freund and Schapire in 1997:

$$L_{Hedge} \leq k + \sqrt{2k \ln N} + \ln N$$

Some Notation

$\mathbf{s}$	state vector
$\boldsymbol{\ell}$	loss vector
$V(\mathbf{s})$	value of the game at a state
$\mathbf{w}(\mathbf{s})$	bet distribution at a state

The Adversarial Casino

The Minimax Value

$$V(\mathbf{s}) = \min_{\mathbf{w}\in\Delta_N} \max_{\boldsymbol{\ell}\in\{0,1\}^N} \mathbf{w}\cdot\boldsymbol{\ell} + V(\mathbf{s}+\boldsymbol{\ell})$$

$$V(\langle k+1,\ldots,k+1\rangle) := 0$$

Computable...?

$$\tau(\mathbf{s})$$

Expected number of rounds required until all events are **dead**

$$\hat{p}_i(\mathbf{s})$$

Probability that the i^{th} event is the last to **die**, the i^{th} Survival Probability

Brainteaser:

- Gambler's cumulative expected loss is always the same no matter how he bets, and that is:

$$\frac{\tau(\mathbf{s})}{N}$$

Central Results:

$$V(\mathbf{s}) = \frac{\tau(\mathbf{s})}{N}$$

$$\mathbf{w}^*(\mathbf{s}) = \hat{\mathbf{p}}(\mathbf{s})$$

The Adversarial Casino is no worse than the Random Casino!

The Gambler's optimal betting strategy is to play "Guess The Survivor"!

Interpretations

$$V(\mathbf{s}) = \frac{\tau(\mathbf{s})}{N}$$

$$\mathbf{w}^*(\mathbf{s}) = \hat{\mathbf{p}}(\mathbf{s})$$

- Worst Case Loss:

  Generalized Coupon Collector's Problem

- Regret:

  Expected number of times you inflict loss on an **already dead event/expert**

MAIN LEMMA: $\forall \mathbf{s}, i \quad \tau(\mathbf{s}) - \tau(\mathbf{s} + \mathbf{e}_i) = N \cdot \hat{\mathbf{p}}_i(\mathbf{s})$

PROOF: Let S_t be the random state from $\mathbf{s}$ after t rounds

$$\tau(\mathbf{s}) = \mathbb{E}\left[\sum_{t=0}^{\infty} \mathbf{1}[S_t \text{ not dead}]\right]$$

$$\tau(\mathbf{s}) - \tau(\mathbf{s} + \mathbf{e}_i) = \mathbb{E}\left[\sum_{t=0}^{\infty} \mathbf{1}[S_t \text{ not dead}] - \mathbf{1}[S_t + \mathbf{e}_i \text{ not dead}]\right]$$

$$= \mathbb{E}\left[\sum_{t=0}^{\infty} \mathbf{1}[S_t \text{ is } i\text{th almost-dead state}]\right]$$

$$= \hat{\mathbf{p}}_i(\mathbf{s}) \cdot N$$

Waiting Lemma:

$$\widehat{p}_i(\mathbf{s}) < \widehat{p}_i(\mathbf{s} + \mathbf{e}_j)$$

Which implies Best strategy for the Casino is to inflict ONE loss on each round

Proof By Intuition:

If another event suffers a loss, unchanged event has higher survival probability

Theorem:

$$V(\mathbf{s}) = \frac{\tau(\mathbf{s})}{N}$$

Proof:

Assume we are at state $\mathbf{s}$

Assume Gambler always bets according to $\widehat{p}_i(\mathbf{s})$

If Casino plays ANY unit loss sequence: $\ell = \mathbf{e}_{i_1}, \mathbf{e}_{i_2}, \mathbf{e}_{i_3}, \ldots$

Gambler loses: $\widehat{p}_{i_1}(\mathbf{s}) + \widehat{p}_{i_2}(\mathbf{s} + \mathbf{e}_{i_1}) + \widehat{p}_{i_3}(\mathbf{s} + \mathbf{e}_{i_1} + \mathbf{e}_{i_2}), \ldots$

$$= \frac{\tau(\mathbf{s}) - \tau(\mathbf{s} + \mathbf{e}_{i_1})}{N} + \frac{\tau(\mathbf{s} + \mathbf{e}_{i+1}) - \tau(\mathbf{s} + \mathbf{e}_{i_1} + \mathbf{e}_{i_2})}{N} + \ldots$$

$$= \frac{\tau(\mathbf{s})}{N}$$

Conclusion

Randomness is a Gambler's worst Adversary