Learning, Markets, and Exponential Families

Jacob Abernethy

University of Michigan
Department of Computer Science and Engineering

November 14, 2014

A Bird-Eye view of Learning Theory

We want to design algorithms that take data as input and return predictions as output. But there are fundamental limits to our ability to predict and how quickly we can achieve good performance.

Two driving questions

- How well can we learn given very limited data?
- What are the computational challenges of prediction?

An Economic Translation

Thinking in terms of the *economic tradeoffs*, our goal is to determine the equilibrium point among the following:

- The marginal cost of additional data

An Economic Translation

Thinking in terms of the *economic tradeoffs*, our goal is to determine the equilibrium point among the following:

- The marginal cost of additional data
- The marginal value of performance improvement (i.e. better decision making)

An Economic Translation

Thinking in terms of the *economic tradeoffs*, our goal is to determine the equilibrium point among the following:

- The marginal cost of additional data
- The marginal value of performance improvement (i.e. better decision making)
- The marginal cost of computational resources

An Economic Translation

Thinking in terms of the *economic tradeoffs*, our goal is to determine the equilibrium point among the following:

- The marginal cost of additional data
- The marginal value of performance improvement (i.e. better decision making)
- The marginal cost of computational resources
- The marginal value of time

Financialization of ML

In 6 Slides

1. Data Brokerage

In the world of Big Data, buying and selling information is a growing industry.

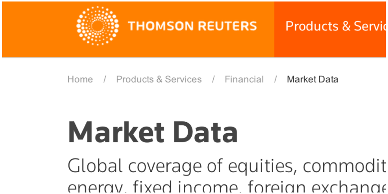

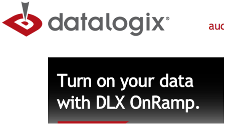

2. Algorithms as a Service

All-purpose ML algorithms are being provided as a web service and sold to developers.

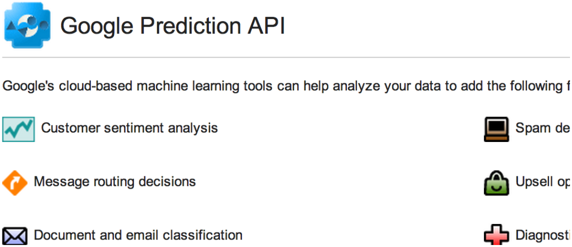

3. Information Markets

Markets built entirely for speculative purposes, where traders can buy/sell securities on elections results to football matches, have flourished in recent years.

betfair

Predictious

4. A Market for Cycles

There is an emerging competitive market where unit of computation are sold like a commodity

Google Cloud Platform

Why Google Products Solutions Customers Developers

amazon web services

	vCPU	ECU	Memory (GiB)	Instance Storage (GB)	Linux/UNIX Usage
General Purpose - Current Generation					
m3.medium	1	3	3.75	1 x 4 SSD	$0.070 per Hour
m3.large	2	6.5	7.5	1 x 32 SSD	$0.140 per Hour
m3.xlarge	4	13	15	2 x 40 SSD	$0.280 per Hour
m3.2xlarge	8	26	30	2 x 80 SSD	$0.560 per Hour

5. A Market for Solutions

Companies are starting to turn towards the *prize-driven competition* to solve big data challenges, rather than hiring in-house data scientists.

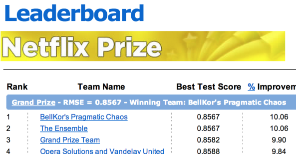

Rank	Team Name	Best Test Score	% Improvem
1	BellKor's Pragmatic Chaos	0.8567	10.06
2	The Ensemble	0.8567	10.06
3	Grand Prize Team	0.8582	9.90
4	Opera Solutions and Vandelay United	0.8588	9.84

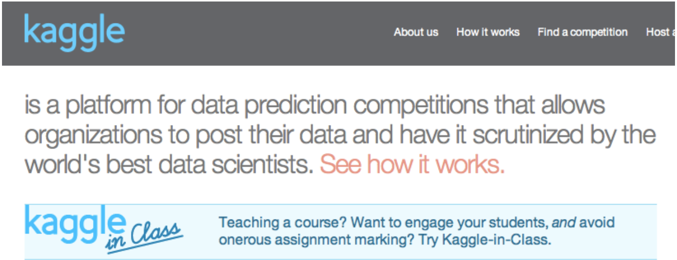

6. Market for Academics

ML Practitioners (including many academics and graduate students) have apparently risen in value in recent years.

DEEPMIND

This Talk

We will discuss some recent results connecting learning-theoretic ideas to finance and economic questions.

- Intro
- Quick review of regret minimization
- Regret in the context of market making
- Exponential family distributions viewed as a prediction market mechanism

The Typical Regret-minimization Framework

We imagine an online game between Nature and Learner. Learner has a (typically convex) *decision set* $\mathcal{X} \subset \mathbb{R}^d$, and Nature has an action set $\mathcal{Z}$, and there is a loss function $\ell : \mathcal{X} \times \mathcal{Z} \to \mathbb{R}$ defined in advance.

The Typical Regret-minimization Framework

We imagine an online game between Nature and Learner. Learner has a (typically convex) *decision set* $\mathcal{X} \subset \mathbb{R}^d$, and Nature has an action set $\mathcal{Z}$, and there is a loss function $\ell : \mathcal{X} \times \mathcal{Z} \to \mathbb{R}$ defined in advance.

Online Convex Optimization

For $t = 1, \ldots, T$:

- Learner chooses $x_t \in \mathcal{X}$
- Nature chooses $z_t \in \mathcal{Z}$
- Learner suffers $\ell(x_t, z_t)$

Learner is concerned with the *regret*:

$$\sum_{t=1}^{T} \ell(x_t, z_t) - \min_{x \in \mathcal{X}} \sum_{t=1}^{T} \ell(x, z_t)$$

The Typical Regret-minimization Framework

We imagine an online game between Nature and Learner. Learner has a (typically convex) *decision set* $\mathcal{X} \subset \mathbb{R}^d$, and Nature has an action set $\mathcal{Z}$, and there is a loss function $\ell : \mathcal{X} \times \mathcal{Z} \to \mathbb{R}$ defined in advance.

Online Convex Optimization

For $t = 1, \ldots, T$:

- Learner chooses $x_t \in \mathcal{X}$
- Nature chooses $z_t \in \mathcal{Z}$
- Learner suffers $\ell(x_t, z_t)$

Learner is concerned with the *regret*:

$$\sum_{t=1}^{T} \ell(x_t, z_t) - \min_{x \in \mathcal{X}} \sum_{t=1}^{T} \ell(x, z_t)$$

This talk we assume ℓ is *linear* in x; WLOG $\ell(x_t, z_t) = x^\top z_t$.

Follow the Regularized Leader

FTRL – Primal Version

1: Input: learning rate $\eta > 0$, regularizer $R : \mathcal{X} \to \mathbb{R}$

2: **for** $t = 1 \ldots T$, $\quad x_t \longleftarrow \arg\min_{x \in \mathcal{X}} R(x) + \eta \sum_{s=1}^{t-1} x^\top l_s.$

Follow the Regularized Leader

FTRL – Primal Version

1: Input: learning rate $\eta > 0$, regularizer $R : \mathcal{X} \to \mathbb{R}$

2: **for** $t = 1 \ldots T$, $\quad x_t \longleftarrow \arg\min_{x \in \mathcal{X}} R(x) + \eta \sum_{s=1}^{t-1} x^\top l_s.$

FTRL – Dual Version

1: **for** $t = 1 \ldots T$, $\quad x_t \longleftarrow \nabla R^* \left(-\eta \sum_{s=1}^{t-1} l_s \right).$

Follow the Regularized Leader

FTRL – Primal Version

1: Input: learning rate $\eta > 0$, regularizer $R : \mathcal{X} \to \mathbb{R}$

2: **for** $t = 1 \ldots T$, $\quad x_t \longleftarrow \arg\min_{x \in \mathcal{X}} R(x) + \eta \sum_{s=1}^{t-1} x^\top l_s.$

FTRL – Dual Version

1: **for** $t = 1 \ldots T$, $\quad x_t \longleftarrow \nabla R^* \left(-\eta \sum_{s=1}^{t-1} l_s \right).$

FTRL is essentially the “only” algorithm we have.
(This COLT: even Follow the *Perturbed* Leader is a special case of FTRL [Abernethy, Lee, Sinha, and Tewari, 2014b]

Regret Bounds on FTRL

Theorem (now classical)

Let $l_1, \ldots, l_T$ be an arbitrary sequence of vectors, and let $L_t := l_1 + \ldots l_t$. Assume $R(x_0) = 0$. Then

$$\begin{aligned} \text{Regret}_T &\leq \frac{R(x^*)}{\eta} + \sum_{t=1}^{T} D_R(x_t, x_{t+1}) \\ &\leq \frac{R(x^*)}{\eta} + \eta \sum_{t=1}^{T} (x_t - x_{t+1})^\top l_t \\ \implies \text{Regret}_T &\leq O\left(\sqrt{\textstyle\sum_{t=1}^{T} \|l_t\|^2}\right) \end{aligned}$$

where $D_R(\cdot, \cdot)$ is the *Bregman divergence* w.r.t. R, and the last line follows from tuning η and assuming some curvature properties of R.

Market Making as Regret Minimization

A lot of the big money in finance is made through *market making*: a market maker (MM) is an agent always willing to buy *and* sell shares/securities at sequentially-set prices.

Market Making as Regret Minimization

A lot of the big money in finance is made through *market making*: a market maker (MM) is an agent always willing to buy *and* sell shares/securities at sequentially-set prices. Assume we have a stock/bond/derivative for sale. For $t = 1, \ldots, T$:

- MM sets bid and ask prices $\overline{p}_t, \underline{p}_t \in \mathbb{R}_+$
- A trader purchases $r_t \in \mathbb{R}$ shares (short sale $\equiv r_t < 0$)
- MM receives $g_t = \$\overline{p}_t r_t$ if $r_t > 0$ or $g_t = \$\underline{p}_t r_t$ if $r_t \leq 0$

Market Making as Regret Minimization

A lot of the big money in finance is made through *market making*: a market maker (MM) is an agent always willing to buy *and* sell shares/securities at sequentially-set prices.
Assume we have a stock/bond/derivative for sale. For $t = 1, \ldots, T$:

- MM sets bid and ask prices $\overline{p}_t, \underline{p}_t \in \mathbb{R}_+$
- A trader purchases $r_t \in \mathbb{R}$ shares (short sale $\equiv r_t < 0$)
- MM receives $g_t = \$\overline{p}_t r_t$ if $r_t > 0$ or $g_t = \$\underline{p}_t r_t$ if $r_t \leq 0$

All shares eventually liquidate at a price of p^*.

$$\text{Loss of MM} = \sum_{t=1}^{T} r_t p^* - \sum_{t=1}^{T} r_t (\overline{p}_t \mathbf{1}[r_t > 0] + \underline{p}_t \mathbf{1}[r_t \leq 0])$$

(More at Abernethy and Kale [2013])

Market Making for Complex Security Spaces

Often we want to sell shares in *multiple related securities* and we want to price these securities jointly.

- Traders can purchase bundles of shares $r \in \mathbb{R}^d$.
- Payout function $\phi : \mathcal{X} \to \mathbb{R}^d$
- In event of x, payout for purchasing bundle r is $r^\top \phi(x)$.

Market Making for Complex Security Spaces

Often we want to sell shares in *multiple related securities* and we want to price these securities jointly.

- Traders can purchase bundles of shares $r \in \mathbb{R}^d$.
- Payout function $\phi : \mathcal{X} \to \mathbb{R}^d$
- In event of x, payout for purchasing bundle r is $r^\top \phi(x)$.

The canonical pricing strategy, which has now been well-studied, is the following:

- Construct a convex $C : \mathbb{R}^d \to \mathbb{R}$ in order that $\{\nabla C\}$ coicides with the rel.int. of $\text{Hull}(\{\phi(x) : x \in \mathcal{X}\})$
- Market maker maintains cumulative outstanding share vector q, announces marginal price vector $\nabla C(q)$
- Trader buying r is charged $C(q + r) - C(q)$

Market Making $\approx$ Online Learning

How to construct C? Choose a "liquidity function" $R : \text{Hull}(\{\phi(x) : x \in \mathcal{X}\}) \to \mathbb{R}$, and let

$$C(q) = \sup_{\mu \in \text{Hull}(\{\phi(x) : x \in \mathcal{X}\})} \mu^\top q - R(\mu)$$

Market Making $\approx$ Online Learning

How to construct C? Choose a "liquidity function" $R : \text{Hull}(\{\phi(x) : x \in \mathcal{X}\}) \to \mathbb{R}$, and let

$$C(q) = \sup_{\mu \in \text{Hull}(\{\phi(x) : x \in \mathcal{X}\})} \mu^\top q - R(\mu)$$

Loss of MM $= C(q_T) - C(\mathbf{0}) - q_T^\top \phi(x)$

Market Making $\approx$ Online Learning

How to construct C? Choose a "liquidity function" $R : \text{Hull}(\{\phi(x) : x \in \mathcal{X}\}) \to \mathbb{R}$, and let

$$C(q) = \sup_{\mu \in \text{Hull}(\{\phi(x) : x \in \mathcal{X}\})} \mu^\top q - R(\mu)$$

Loss of MM $= C(q_T) - C(\mathbf{0}) - q_T^\top \phi(x) \quad = \quad$ Pricing Regret!

Market Making ≈ Online Learning

How to construct C? Choose a "liquidity function" $R : \text{Hull}(\{\phi(x) : x \in \mathcal{X}\}) \to \mathbb{R}$, and let

$$C(q) = \sup_{\mu \in \text{Hull}(\{\phi(x) : x \in \mathcal{X}\})} \mu^\top q - R(\mu)$$

$$\text{Loss of MM} = C(q_T) - C(\mathbf{0}) - q_T^\top \phi(x) \quad = \quad \text{Pricing Regret!}$$

With this connection, we get a set of natural equivalences:

Market Making	Online Learning

Market Making ≈ Online Learning

How to construct C? Choose a "liquidity function" $R : \text{Hull}(\{\phi(x) : x \in \mathcal{X}\}) \to \mathbb{R}$, and let

$$C(q) = \sup_{\mu \in \text{Hull}(\{\phi(x):x\in\mathcal{X}\})} \mu^\top q - R(\mu)$$

Loss of MM $= C(q_T) - C(\mathbf{0}) - q_T^\top \phi(x) \quad = \quad$ Pricing Regret!

With this connection, we get a set of natural equivalences:

Market Making	Online Learning
Market Maker Loss	Learning Regret

Market Making $\approx$ Online Learning

How to construct C? Choose a "liquidity function" $R : \text{Hull}(\{\phi(x) : x \in \mathcal{X}\}) \to \mathbb{R}$, and let

$$C(q) = \sup_{\mu \in \text{Hull}(\{\phi(x):x\in\mathcal{X}\})} \mu^\top q - R(\mu)$$

$$\text{Loss of MM} = C(q_T) - C(\mathbf{0}) - q_T^\top \phi(x) \quad = \quad \text{Pricing Regret!}$$

With this connection, we get a set of natural equivalences:

Market Making	Online Learning
Market Maker Loss	Learning Regret
Seq. Pricing Strat.	FTRL

Market Making ≈ Online Learning

How to construct C? Choose a "liquidity function" $R : \text{Hull}(\{\phi(x) : x \in \mathcal{X}\}) \to \mathbb{R}$, and let

$$C(q) = \sup_{\mu \in \text{Hull}(\{\phi(x) : x \in \mathcal{X}\})} \mu^\top q - R(\mu)$$

Loss of MM $= C(q_T) - C(\mathbf{0}) - q_T^\top \phi(x) \quad = \quad$ Pricing Regret!

With this connection, we get a set of natural equivalences:

Market Making	Online Learning
Market Maker Loss	Learning Regret
Seq. Pricing Strat.	FTRL
Liquidity at price p	$\nabla^2 R(p)$

Please see Chen and Vaughan [2010] and Abernethy, Chen, and Vaughan [2013] for details

Exp. Family Distributions and Prediction Markets

Let's now switch gears and see how exp families relate can be viewed through an entirely probability-free lens.

Exponential Family Distributions

Many dist. families we encounter are *exponential families*. Let $\beta \in \mathbb{R}^d$ be params, $\phi : \mathcal{X} \to \mathbb{R}^d$ some "statistics". The pdf of dist. corresponding to β is

$$\text{For } x \in \mathcal{X} : \qquad P_\beta(x) \propto \exp(\beta^\top \phi(x))$$

Exponential Family Distributions

Many dist. families we encounter are *exponential families*. Let $\beta \in \mathbb{R}^d$ be params, $\phi : \mathcal{X} \to \mathbb{R}^d$ some "statistics". The pdf of dist. corresponding to β is

For $x \in \mathcal{X}$: $$P_\beta(x) = \exp(\beta^\top \phi(x) - \Psi(\beta))$$

Where $$\Psi(\beta) = \log \int_{\mathcal{X}} \exp(\beta^\top \phi(x'))dx'$$

Exponential Family Distributions

Many dist. families we encounter are *exponential families*. Let $\beta \in \mathbb{R}^d$ be params, $\phi : \mathcal{X} \to \mathbb{R}^d$ some "statistics". The pdf of dist. corresponding to β is

For $x \in \mathcal{X}$: $$P_\beta(x) = \exp(\beta^\top \phi(x) - \Psi(\beta))$$

Where $$\Psi(\beta) = \log \int_{\mathcal{X}} \exp(\beta^\top \phi(x')) dx'$$

- $\phi(x)$ is called the "sufficient statistics" of x

Exponential Family Distributions

Many dist. families we encounter are *exponential families*. Let $\beta \in \mathbb{R}^d$ be params, $\phi : \mathcal{X} \to \mathbb{R}^d$ some "statistics". The pdf of dist. corresponding to β is

$$\text{For } x \in \mathcal{X}: \qquad P_\beta(x) = \exp(\beta^\top \phi(x) - \Psi(\beta))$$

$$\text{Where} \qquad \Psi(\beta) = \log \int_{\mathcal{X}} \exp(\beta^\top \phi(x'))dx'$$

- $\phi(x)$ is called the "sufficient statistics" of x
- $\Psi(\beta)$ is called the "log partition function"

Exponential Family Distributions

Many dist. families we encounter are *exponential families*. Let $\beta \in \mathbb{R}^d$ be params, $\phi : \mathcal{X} \to \mathbb{R}^d$ some "statistics". The pdf of dist. corresponding to β is

$$\text{For } x \in \mathcal{X}: \qquad P_\beta(x) = \exp(\beta^\top \phi(x) - \Psi(\beta))$$

$$\text{Where} \qquad \Psi(\beta) = \log \int_{\mathcal{X}} \exp(\beta^\top \phi(x'))dx'$$

- $\phi(x)$ is called the "sufficient statistics" of x
- $\Psi(\beta)$ is called the "log partition function"
- A wonderful fact: $\mathbb{E}_{X \sim P_\beta}[\phi(X)] = \nabla \Psi(\beta)$

Forget That: Exponential Family Market

- Imagine $x \in \mathcal{X}$ is some future uncertain outcome, and a FIRM wants predictions on $\phi(x)$.

Forget That: Exponential Family Market

- Imagine $x \in \mathcal{X}$ is some future uncertain outcome, and a FIRM wants predictions on $\phi(x)$.
- FIRM will create a *prediction market*
- Prices should corresp. to aggregate belief

$$\mathbb{E}_{x \sim \text{crowd belief}}[\phi(x)]$$

Forget That: Exponential Family Market

- Imagine $x \in \mathcal{X}$ is some future uncertain outcome, and a FIRM wants predictions on $\phi(x)$.
- FIRM will create a *prediction market*
- Prices should corresp. to aggregate belief

$$\mathbb{E}_{x \sim \text{crowd belief}}[\phi(x)]$$

- FIRM will sell bundles of shares $\delta \in \mathbb{R}^d$ to trader
- Upon outcome x, reward for purchasing δ:

$$\text{payoff}(\delta | x) = \phi(x)^\top \delta$$

Forget That: Exponential Family Market

- Imagine $x \in \mathcal{X}$ is some future uncertain outcome, and a FIRM wants predictions on $\phi(x)$.
- FIRM will create a *prediction market*
- Prices should corresp. to aggregate belief

$$\mathbb{E}_{x \sim \text{crowd belief}}[\phi(x)]$$

- FIRM will sell bundles of shares $\delta \in \mathbb{R}^d$ to trader
- Upon outcome x, reward for purchasing δ:

$$\text{payoff}(\delta | x) = \phi(x)^\top \delta$$

- Let sum of all outstanding shares be $\Theta := \delta_1 + \ldots + \delta_m$
- The *price* of buying δ:

$$\Psi(\Theta + \delta) - \Psi(\delta)$$

Benefits of the Market Interpretation

- Given that Θ represents market state

$$\text{Marginal prices} = \nabla \Psi(\Theta),$$

which correspond to mean parameters in P_Θ!

Benefits of the Market Interpretation

- Given that Θ represents market state

$$\text{Marginal prices} = \nabla \Psi(\Theta),$$

which correspond to mean parameters in P_Θ!

- If the true distribution over x is Q, then

$$\mathbb{E}\text{TraderProfit}(\delta) = KL(Q; P_\Theta) - KL(Q; P_{\Theta+\delta})$$

Benefits of the Market Interpretation

- Given that Θ represents market state

$$\text{Marginal prices} = \nabla\Psi(\Theta),$$

  which correspond to mean parameters in P_Θ!

- If the true distribution over x is Q, then

$$\mathbb{E}\text{TraderProfit}(\delta) = KL(Q; P_\Theta) - KL(Q; P_{\Theta+\delta})$$

- FIRM has to pay

$$\mathbb{E}\text{FirmCost}(\Theta_{\text{final}}) = KL(Q; P_{\mathbf{0}}) - KL(Q; P_{\Theta_{\text{final}}})$$

(Results in Abernethy, Kutty, Lahaie, and Sami [2014a])

Interpreting Market Behavior

Let us imagine traders in such a market that has a belief on the outcome x distributed according to P_β. Assume trader has *exponential utility* (with risk-aversion param a):

$$\text{Utility}(\$99) = 1 - \exp(-a \cdot 99)$$

Interpreting Market Behavior

Let us imagine traders in such a market that has a belief on the outcome x distributed according to P_β. Assume trader has *exponential utility* (with risk-aversion param a):

$$\text{Utility}(\$99) = 1 - \exp(-a \cdot 99)$$

Proposition: Belief $\equiv$ Investment

In terms of optimal trading behavior

$$\text{Buying } \delta \text{ shares} \iff \text{updating belief } \beta \leftarrow \beta + \delta$$

Interpreting Market Behavior

Let us imagine traders in such a market that has a belief on the outcome x distributed according to P_β. Assume trader has *exponential utility* (with risk-aversion param a):

$$\text{Utility}(\$99) = 1 - \exp(-a \cdot 99)$$

Proposition: Belief $\equiv$ Investment

In terms of optimal trading behavior

$$\text{Buying } \delta \text{ shares} \iff \text{updating belief } \beta \leftarrow \beta + \delta$$

Proposition: Equilibrium $\equiv$ MAP-estimate for Gaussian

Assume we have n traders with belief parameters $\beta_1, \ldots, \beta_n$ with risk aversion parameters $a_1, \ldots, a_n$. If they all trade to maximize expected utility, then *in equilibrium* we have:

$$\text{EquilibriumState } \Theta_{\text{final}} := \frac{\Theta_{\text{init}} + \sum_i \beta_i a_i^{-1}}{1 + \sum_i a_i^{-1}}$$

The Vision

We would have a number of major benefits if we were able to cast a broader class of ML algorithms through the lens of market equilibria.

- Robustness on solution
- *Real* decentralization of learning tasks
- Possible model for distributed computing

THANK YOU

Jacob Abernethy and Satyen Kale. Adaptive market making via online learning. In *(NIPS) Advances in Neural Information Processing Systems*, pages 2058–2066, 2013.

Jacob Abernethy, Yiling Chen, and Jennifer Wortman Vaughan. Efficient market making via convex optimization, and a connection to online learning. *ACM Transactions on Economics and Computation*, 1(2):12, 2013.

Jacob Abernethy, Sindhu Kutty, Sébastien Lahaie, and Rahul Sami. Information aggregation in exponential family markets. In *(EC) Proceedings of the 15th ACM conference on Economics and computation*, pages 395–412. ACM, 2014a.

Jacob Abernethy, Chansoo Lee, Abhinav Sinha, and Ambuj Tewari. Online linear optimization via smoothing. In *(COLT) Conference on Learning Theory*, 2014b.

Y. Chen and J. Wortman Vaughan. A new understanding of prediction markets via no-regret learning. In *Proceedings* *of the 11th ACM Conference on Electronic Commerce*, pages 189–198, 2010.