The Game-Theoretic Approach to Machine Learning and Adaptation

Nicolò Cesa-Bianchi

Università degli Studi di Milano

Machine Learning

A wide range of applications

- Categorization of documents, speech, images, genes
- Natural language processing
- Robot control
- Search engine quality
- Dynamic allocation of resources

Learning theory

Foundations of machine learning

- Under what conditions can a machine learn from examples?
- How much information (e.g., training examples) is needed to achieve a given predictive performance?
- How many computational resources (time and space)?
- What is the best mathematical framework to study these phenomena?

The statistical learning vision

- The training data are a statistical sample (i.i.d.)
- Relate the empirical error of a predictor to its true error rate
- A finite-sample estimation problem

Vladimir Vapnik

Overfitting

- The best predictor on the data is not guaranteed to have a small error rate if it is chosen from a large set
- Need enough data to guarantee that empirical error is close to true error for each predictor in the set
- This “enough” turns out to depend on a notion of combinatorial dimension of the set of (VC dimension)

The need for a different vision

- The statistical approach is at the basis of the most successful applications of machine learning in the past twenty years
- As the range of machine learning applications widens, new paradigms are needed

Some hard cases for statistical modelling

- Data source is highly nonstationary
- Environment reacts to the learner (e.g., spam)

On a more philosophical level

Is statistics the only language for describing the phenomenon of learning in machines?

Theory of repeated games

James Hannan

David Blackwell

Learning to play a game (1956)

Play a game repeatedly against a possibly suboptimal opponent

Zero-sum 2-person games played more than once

$$\begin{array}{c|cccc|}
 & 1 & 2 & \dots & M \\
\hline
1 & \ell(1,1) & \ell(1,2) & \dots & \\
2 & \ell(2,1) & \ell(2,2) & \dots & \\
\vdots & \vdots & \vdots & \ddots & \\
N & & & & \\
\hline
\end{array}$$

$N \times M$ known loss matrix

- Row player (player) has N actions
- Column player (opponent) has M actions

For each game round $t = 1, 2, \dots$

- Player chooses action i_t and opponent chooses action y_t
- The player suffers loss $\ell(i_t, y_t)$ (= gain of opponent)

Player can learn from opponent's history of past choices $y_1, \dots, y_{t-1}$

Prediction with expert advice

Volodya Vovk

Manfred Warmuth

Opponent's moves $y_1, y_2, \dots$ define a sequential prediction problem with loss function ℓ

1. Play action I_t from $1, \dots, N$
2. Observe next value y_t
3. Incur loss $\ell(I_t, y_t)$

Exponentially weighted forecaster

At time t pick action i with probability proportional to

$$\exp\left(-\eta \, \mathrm{Loss}_{i,t}\right)$$

where $\mathrm{Loss}_{i,t}$ is total loss of action i up to now

Expert's theorem

The average per-round expected loss of the forecaster converges to that of the best action for the observed sequence at rate

$$\sqrt{\frac{\ln N}{T}}$$

where N is number of actions and T is the number of time steps

Note: no dependence on number of opponent's actions

The bandit problem: playing an unknown game

- In order to keep counts $\text{Loss}_{i,t}$ for each action, we need to know the losses $\ell(i, y_t)$ also for the actions i we did not play at round t
- What if we can only observe the loss of the played action I_t?

...

N slot machines

- Dynamic content optimization
- Surprisingly, convergence rate to best action is $\sqrt{\dfrac{N \ln N}{T}}$

Structured actions: adversarial routing

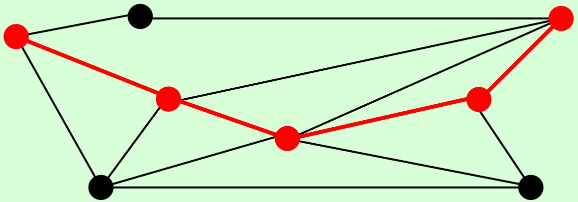

- In certain problems, actions have a combinatorial structure (paths, trees, matchings)
- If loss is linear over the edges, then the bandit convergence rate to best action is

$$\sqrt{\frac{d \ln N}{T}}$$

where d is number of edges and N is the number of actions (typically superpolynomial in d)

Partial monitoring: not observing any loss

Dynamic pricing

1. Post a T-shirt price
2. Observe if next customer buys or not
3. Adjust price

Note: feedback does not reveal the player's loss

Goal: converge to the average return of the best fixed price

Convergence rate to best fixed price is $T^{-1/3}$ rather than $T^{-1/2}$ as in the bandit case

K-person games

- There are K players choosing actions $I_{1,t}, \ldots, I_{K,t}$
- Each player i has its own loss function $\ell_i(I_{1,t}, \ldots, I_{K,t})$
- What happens if all players use exponentially weighted forecasting, or similar algorithms?

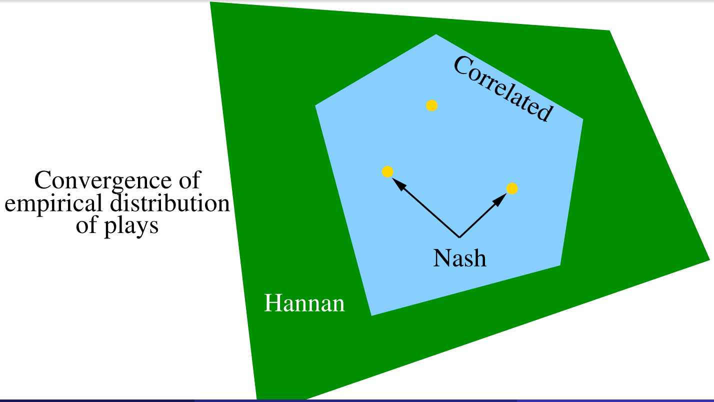

From game theory to machine learning

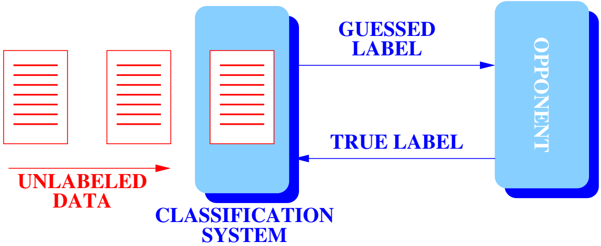

- Now opponent's moves y_t have side information $x_t \in \mathbb{R}^d$ (e.g., text on a document)
- A repeated game between the player choosing a classifier and the opponent choosing an action (x_t, y_t)
- Convergence to performance of best classifier in a given class (e.g., linear classifiers with bounded norm)

Online learning algorithms

- Simple: easy to implement
- Scalable: local optimization vs. global optimization
- Robust: inherit game-theoretic performance guarantees
- Versatile: classification, regression, ranking, structured prediction

Structured Prediction

A combinatorial label space (sequences, trees)

- POS tagging: sentence → sequence of POS tags
- Parsing: sentence → parse tree
- Bilingual alignment: sentence pair → alignment (matching)
- Letter to phoneme: word → phoneme sequence
- Phrase-based translation: source sentence → target sentence

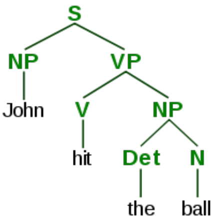

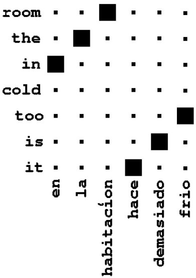

Online learning in general spaces

Some applications

- Reproducing kernel Hilbert spaces: efficiently embed data in high-dim space where linear classifiers can do well

  $\longrightarrow$ Bioinformatics, vision, language

- Linear space of matrices

  $\longrightarrow$ Integrating data sources, learning different tasks at once

- Banach spaces of models

  $\longrightarrow$ Financial data

Tracking linear classifiers

- If data source is not fitted well by any linear model, then comparing to the best linear model f^* is trivial
- We want instead compare to the best sequence $f_1, f_2, \ldots$ of linear models

Adversarial tracking

- Bound on predictive performance reflects the opponent's trade-off between fit of sequence and total shift $\sum_t \|f_t - f_{t-1}\|$
  $\rightarrow$ dynamic overfitting control
- This is achieved by enforcing sparsity of the learner's model (expressed as a linear combination of past $\boldsymbol{x}_t$'s)

Tracking a shifting topic

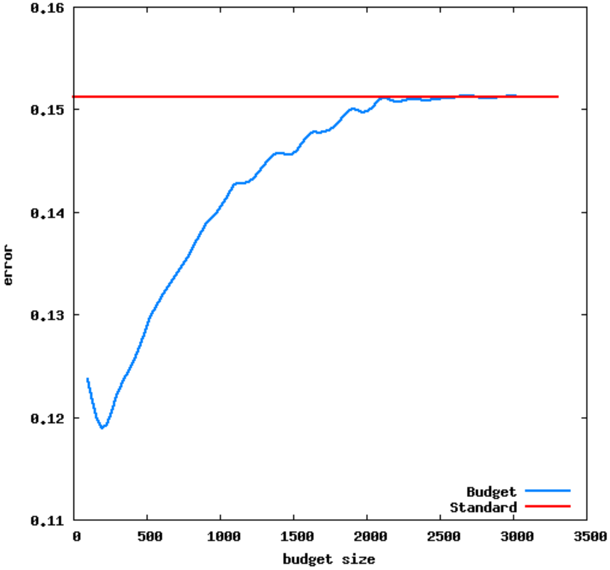

Online active learning

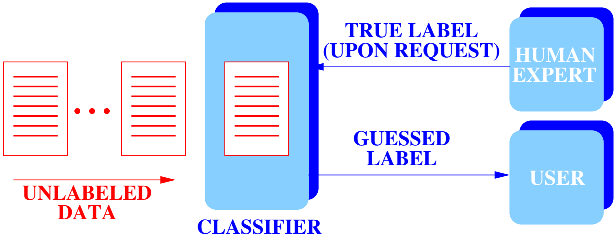

- Observing the data process is cheap
- Observing the label process is expensive
  $\longrightarrow$ need to query the human expert

Question

How much better can we do by subsampling **adaptively** the label process?

A game with the opponent

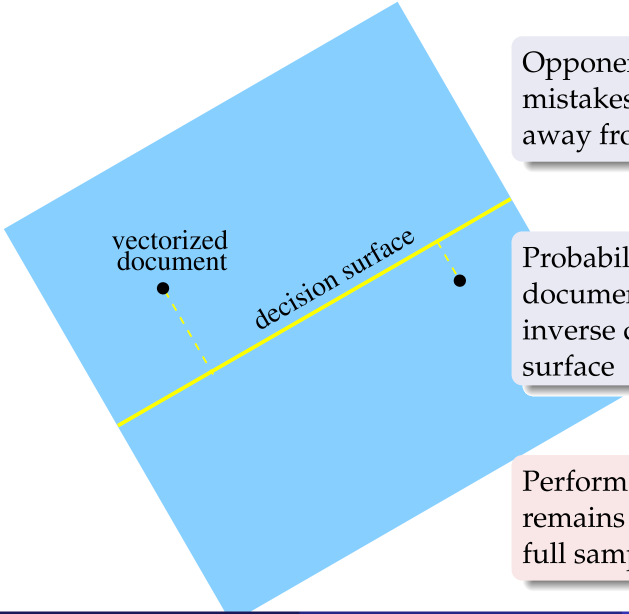

Opponent avoids causing mistakes on documents far away from decision surface

Probability of querying a document proportional to inverse distance to decision surface

Performance guarantee remains unchanged w.r.t. the full sampling case

Experiments on Reuters corpus

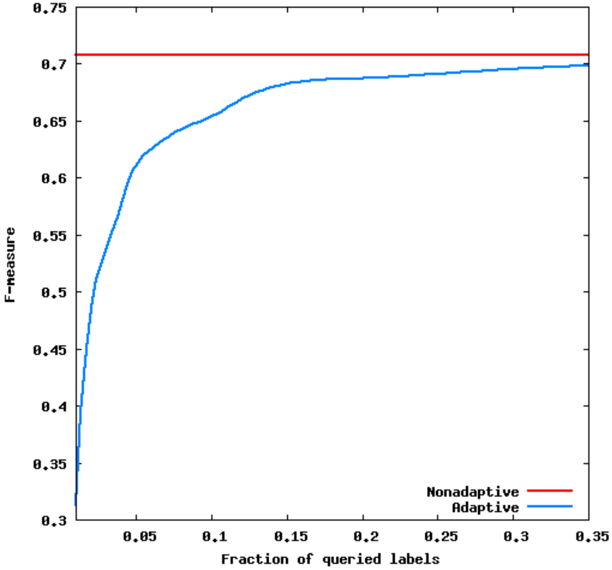

Prediction on graphs

- Web, social networks, biological networks
- Predict labels on nodes (or links)

Game-theoretic framework allows to derive principled algorithms without statistical assumptions

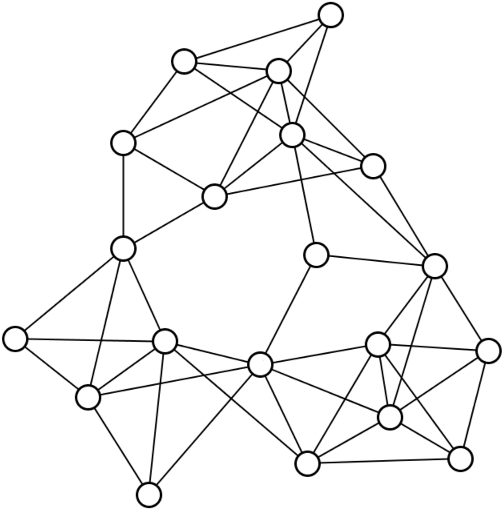

Node prediction

What is the optimal number of mistakes when sequentially predicting the node labels of a given graph?

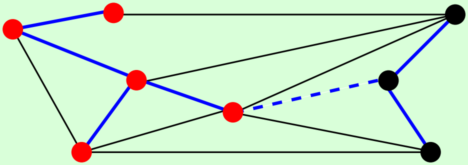

- This number is captured (to within log factors) by the cutsize of the graph's random spanning tree
- This is a density independent regularity measure of the graph labeling and there are efficient predictors that achieve this

Conclusions

- Online game-theoretic analysis provides nonstochastic foundations to machine learning — good for nonstationary, adversarial sources
- Algorithms typically have good scaling properties due to local (rather than global) optimization
- Fruitful exchange of concepts between game theory and machine learning
- Interacting learners
  - Multitask learning: same side information, different objectives
  - Multiview learning: different side information, same objective