Game-theoretic probability and some of its applications

Vladimir Vovk

Department of Computer Science
Royal Holloway, University of London
Egham, Surrey

vovk@cs.rhul.ac.uk

Oxford, 4 May 2007

My plan:

- Basics of game-theoretic probability: only strong laws
- Jeffreys's law: successful forecasters must agree with each other
- Defensive forecasting: game-theoretic laws of probability $\mapsto$ forecasting algorithms; its applications

Philosophical introduction

Three ways of modelling reality (or its pieces):

- deterministic;
- probabilistic;
- too complicated to be modelled, except islands of deterministic and/or stochastic regularities.

Example: human blood group inheritance (simplified mechanism)

Possible genotypes:

- AB (leading to the blood group AB),
- AA, AO (both leading to A),
- BB, BO (both leading to B),
- OO (leading to O).

The theory is that the child's genotype is obtained by plucking a letter at random from each parent.

Example: cont.

For example: when the parents are OO and AB, the forecast will assign weight 1/2 to each of AO and BO.

The theory

predicts the child's genotype from the parents' genotype;

does not predict (even probabilistically) who decides to have children with whom.

It is open: some indispensable events are not modelled (this is typical).

Prequential principle (A. P. Dawid, 1984), as applied to the theory of blood group inheritance

When testing the theory, only use the actual forecasts (e.g.: $P\{\text{AO}\} = P\{\text{BO}\} = 1/2$) and observations (e.g.: AO).

May sound trivial but in fact calls for a revision of the foundations of statistics. All standard approaches to statistics are based on measure-theoretic probability, $(\Omega, \mathcal{F}, \mathbb{P})$; for them to work, events such as "X and Y decide to have a child" (given the past) should be also assigned a probability [if measurable].

Once you have a prequential way of testing theories, you can extract predictions from theories that you trust: predict that the sequences of forecasts/observations that would lead to rejecting the theory will not happen.

Prequential testing

Probability forecasting protocol:

$\mathcal{K}_0 := 1$.
FOR $n = 1, 2, \dots$:
 Forecaster announces $P_n \in \mathcal{P}(\mathbf{Y})$.
 Sceptic announces $f_n : \mathbf{Y} \to \mathbb{R}$ such that $\int f_n \mathrm{d}P_n \leq 0$.
 Reality announces $y_n \in \mathbf{Y}$.
 $\mathcal{K}_n := \mathcal{K}_{n-1} + f_n(y_n)$.
END FOR.

$\mathcal{K}_n$: Sceptic's capital (the degree to which he has managed to discredit Forecaster). He loses as soon as $\mathcal{K}_n < 0$.

Strong law of large numbers (SLLN)

Let $y_n \in [-c, c]$. We expect

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} \left(y_i - \int y P_i(\mathrm{d}y) \right) = 0 \tag{1}$$

unless Forecaster is doing a bad job.

Measure-theoretic formalization: (1) happens with probability one.

Game-theoretic formalization: (1) happens unless Sceptic becomes infinitely rich.

To make the result stronger: notice that only $\mu_n := \int y P_n(\mathrm{d}y)$ matters; cut off the rest.

Game-theoretic SLLN for bounded observations

Bounded forecasting protocol:

$\mathcal{K}_0 := 1$.
FOR $n = 1, 2, \ldots$:
  Forecaster announces $\mu_n \in [-c, c]$.
  Sceptic announces $s_n \in \mathbb{R}$.
  Reality announces $y_n \in [-c, c]$.
  $\mathcal{K}_n := \mathcal{K}_{n-1} + s_n(y_n - \mu_n)$.
END FOR.

$\mathcal{K}_n$: Sceptic's capital.

Proposition 1 (game-theoretic SLLN) Sceptic has a strategy which guarantees that

- $\mathcal{K}_n$ is never negative
- either

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} (y_i - \mu_i) = 0$$

(μ_n are unbiased) or

$$\lim_{n\to\infty} \mathcal{K}_n = \infty.$$

The measure-theoretic SLLN follows easily: if Reality is oblivious (does not pay attention to what her opponents do) and uses a randomized strategy (probability measure $\mathbb{P}$ on the sequences of Reality's moves) and Forecaster computes his moves as conditional expectations w.r. to $\mathbb{P}$: $\mathcal{K}_n$ is a non-negative martingale, and so $\mathcal{K}_n \to \infty$ with probability 0.

Game-theoretic SLLN:

- Reality need not be oblivious (or follow a strategy)
- Forecaster need not ignore Sceptic (or follow a strategy)

Caveat: I assumed that Sceptic's strategy was measurable. Empirical fact: for all kinds of limit theorems, Sceptic's strategies people construct are measurable.

Interpretation

Is usually based on the belief that Sceptic will not become very rich. Not always; e.g., the forecasts in Iowa Electronic Markets [the probability that the next American president will be a Democrat was 0.61 yesterday] either allow us to become infinitely rich or are unbiased. Which?

General definition: an event E is almost certain if Sceptic has a strategy that does not risk bankruptcy and makes him infinitely rich if E fails to happen.

Or: Sceptic can force E.

I will almost prove Proposition 1.

Usual tricks:

- we can replace $\mathcal{K}_n \to \infty$ with $\sup_n \mathcal{K}_n = \infty$ [wait until $\mathcal{K}_n$ reaches C and stop playing; combine this for different $C \to \infty$]

- if $E_1, E_2, \ldots$ are almost certain, $\cap E_i$ is also almost certain [combine the corresponding strategies]

Lemma. Suppose $\epsilon > 0$. Then Sceptic can "weakly force"

$$\limsup_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} y_i \leq \epsilon.$$

The same argument, with $-\epsilon$ in place of ϵ:

$$\liminf_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} y_i \geq -\epsilon \quad \text{a.s.}$$

Combine this for all ϵ.

Proof of the lemma. Sceptic always sets $s_n := \epsilon \mathcal{K}_{n-1}$; then

$$\mathcal{K}_n = \prod_{i=1}^{n} (1 + \epsilon y_i) \, .$$

On the paths where $\mathcal{K}_n$ is bounded:

$$\prod_{i=1}^{n} (1 + \epsilon y_i) \le C$$

$$\sum_{i=1}^{n} \ln (1 + \epsilon y_i) \le D;$$

since $\ln(1 + t) \ge t - t^2$ whenever $t \ge -\frac{1}{2}$,

$$\epsilon \sum_{i=1}^{n} y_i - \epsilon^2 \sum_{i=1}^{n} y_i^2 \le D$$

$$\epsilon \sum_{i=1}^{n} y_i - \epsilon^2 n \leq D$$

$$\epsilon \sum_{i=1}^{n} y_i \leq \epsilon^2 n + D$$

$$\frac{1}{n} \sum_{i=1}^{n} y_i \leq \epsilon + \frac{D}{\epsilon n}.$$

Other limit theorems

A book-length review of game-theoretic probability:

Glenn Shafer and Vladimir Vovk, Probability and finance: it's only a game. New York: Wiley, 2001

In particular, it contains:

- Kolmogorov's SLLN (unbounded observations);
- Kolmogorov's LIL;
- Lindeberg's CLT (and a game-theoretic notion of probability);
- various results from mathematical finance.

Problem with prequential [and any other statistical] testing

We reject the theory when Sceptic becomes very rich when gambling against it. How rich? 20? 100?

In practice, the threshold often does not matter: if we have a better new theory, the old theory will be eventually rejected for any threshold (especially if a crucial experiment has been found: a reproducible situation in which the two theories give different probability forecasts). This follows from

Jeffreys's law

Competitive testing protocol:

$\mathcal{K}_0^{\mathrm{I}} := 1$.
$\mathcal{K}_0^{\mathrm{II}} := 1$.
FOR $n = 1, 2, \dots$:
  Forecaster I announces $P_n^{\mathrm{I}} \in \mathcal{P}(\mathbf{Y})$.
  Forecaster II announces $P_n^{\mathrm{II}} \in \mathcal{P}(\mathbf{Y})$.
  Sceptic I announces $f_n^{\mathrm{I}} : \mathbf{Y} \to \mathbb{R}$ such that $\int f_n^{\mathrm{I}} \mathrm{d} P_n^{\mathrm{I}} \le 0$.
  Sceptic II announces $f_n^{\mathrm{II}} : \mathbf{Y} \to \mathbb{R}$ such that $\int f_n^{\mathrm{II}} \mathrm{d} P_n^{\mathrm{II}} \le 0$.
  Reality announces $y_n \in \mathbf{Y}$.
  $\mathcal{K}_n^{\mathrm{I}} := \mathcal{K}_{n-1}^{\mathrm{I}} + f_n^{\mathrm{I}}(y_n)$.
  $\mathcal{K}_n^{\mathrm{II}} := \mathcal{K}_{n-1}^{\mathrm{II}} + f_n^{\mathrm{II}}(y_n)$.
END FOR.

Suppose $\mathbf{Y}$ is finite. The Hellinger distance:

$$D\left(P^{\mathrm{I}} \parallel P^{\mathrm{II}}\right) = \sum_{y \in \mathbf{Y}} \left(\sqrt{P^{\mathrm{I}}\{y\}} - \sqrt{P^{\mathrm{II}}\{y\}}\right)^2 .$$

Proposition 2 (Jeffreys's law) In the competitive testing protocol, the Sceptics have a joint strategy ensuring $\mathcal{K}_n^{\mathrm{I}} \ge 0$, $\mathcal{K}_n^{\mathrm{I}} \ge 0$ and guaranteeing that, for all n,

$$\ln \mathcal{K}_n^{\mathrm{I}} + \ln \mathcal{K}_n^{\mathrm{II}} = \sum_{i=1}^{n} D\left(P_i^{\mathrm{I}} \parallel P_i^{\mathrm{II}}\right) .$$

There is an opposite statement (in this sense the Hellinger distance is the best possible).

Defensive forecasting

Sceptic's strategies constructed for various laws of probability usually are (or can be made) continuous.

For any continuous strategy for Sceptic there exists a strategy for Forecaster that does not allow Sceptic's capital to grow.

Therefore: continuous game-theoretic law of probability $\mapsto$ forecasting algorithm guaranteed to satisfy it.

Modified bounded forecasting protocol:

$\mathcal{K}_0 := 1$.

FOR $n = 1, 2, \ldots$:

Reality announces $x_n \in \mathbf{X}$.

Sceptic announces continuous $S_n : [-c, c] \to \mathbb{R}$.

Forecaster announces $\mu_n \in [-c, c]$.

Reality announces $y_n \in [-c, c]$.

$\mathcal{K}_n := \mathcal{K}_{n-1} + S_n(\mu_n)(y_n - \mu_n)$.

END FOR.

Proposition 3 (Takemura) Forecaster has a strategy that ensures $\mathcal{K}_0 \geq \mathcal{K}_1 \geq \mathcal{K}_2 \cdots$.

Proof

- choose μ_n such that $S_n(\mu_n) = 0$
- if the equation $S_n(\mu) = 0$ has no roots (in which case S_n never changes sign),

$$\mu_n := \begin{cases} c & \text{if } S_n > 0 \\ -c & \text{if } S_n < 0 \end{cases}$$

QED

Can be easily generalized; Intermediate Value Theorem $\mapsto$ numerous fixed point and minimax theorems in topological vector spaces.

General scheme of defensive forecasting

- Decide which law(s) of probability you want Forecaster's moves to satisfy.
- Prove the corresponding game-theoretic result.
- Apply Proposition 3.
- If necessary, streamline the resulting forecasting algorithm.

What does it give in the case of LLN?

In fact, nothing interesting: Forecaster performs his task too well. E.g., he can choose

$$\mu_n := \begin{cases} 0 & \text{if } n = 1 \\ y_{n-1} & \text{otherwise,} \end{cases}$$

ensuring

$$\left| \sum_{i=1}^{n} (y_i - \mu_i) \right| \le c$$

for all n (much better than using the true probabilities).

We need a “convoluted” LLN. Suppose $\Phi : [-c, c] \times \mathbf{X} \to H$ (feature mapping to an inner product space) and

$$\mathbf{c}_\Phi := \sup_{\mu,x} \|\Phi(\mu, x)\| < \infty.$$

The convoluted LLN: for any $\delta \in (0, 1)$,

$$\left\| \frac{1}{n} \sum_{i=1}^{n} (y_i - \mu_i)\Phi(\mu_i, x_i) \right\| \le \frac{c\mathbf{c}_\Phi}{\sqrt{n\delta}}$$

with probability at least $1 - \delta$. An easy modification of the standard statement ($\Phi \equiv 1$: Kolmogorov 1929). True both measure-theoretically (with Φ measurable) and game-theoretically.

Example: calibration and resolution for Besov spaces

Kolmogorov's method can be easily extended to Banach spaces that are not too far from being Hilbert spaces (as measured by "Clarkson's modulus of convexity").

Important function spaces on a set Ω: Besov spaces $B^s_{p,q}(\Omega)$.

- s is responsible for smoothness;
- p for convexity ($p = 2$ is the Hilbert case);
- q is much less important (set, e.g., $q := p$: Slobodetsky spaces).

Ω will always be assumed a bounded "Lipschitz domain" in Euclidean space $\mathbb{R}^m$ (a mild regularity condition).

$s > m/p$ ensures that the elements of $B^s_{p,q}$ are genuine functions (even continuous): always assumed.

Proposition 4 Let $B^s_{p,q}$ be a Besov space on $[-c,c] \times \mathbf{X}$ with $s > m/p$, $p \geq 2$ and $q \in [p', p]$. A forecasting algorithm ensures

$$\left| \frac{1}{n} \sum_{i=1}^{n} (y_i - \mu_i) F(\mu_i, x_i) \right| \leq c C_{\mathbf{X},s,p,q} \|F\|_{B^s_{p,q}} n^{-1/p}$$

for some constant $C_{\mathbf{X},s,p,q}$, all n and all $F \in B^s_{p,q}$.

The algorithm (obtained from Kolmogorov's proof of LLN by defensive forecasting) can be implemented efficiently.

When $F(\mu, x)$ depends only on μ: calibration (take a "soft neighbourhood" of $\mu^* \in [-c, c]$ as F).

When $F(\mu, x)$ depends only on x: resolution (take a "soft neighbourhood" of $x^* \in \mathbf{X}$ as F).

When $F(\mu, x)$ depends on both: joint calibration and resolution (more than calibration and resolution separately).

Bordering on paradox

There is no forecasting algorithm that "works" for every sequence. Dawid's (1985) example:

$$y_n := \begin{cases} c & \text{if } \mu_n < 0 \\ -c & \text{otherwise.} \end{cases}$$

The algorithm producing μ_n is always wrong on this sequence!

Application to decision making

Decision-making protocol:

Loss$_0 := 0$.
FOR $n = 1, 2, \ldots$:
  Reality announces $x_n \in \mathbf{X}$.
  Decision Maker announces $\gamma_n \in \Gamma$.
  Reality announces $y_n \in \mathbf{Y}$.
  Loss$_n :=$ Loss$_{n-1} + \lambda(y_n, \gamma_n)$.
END FOR.

λ: the loss function.

The difference between the two protocols

- In the forecasting protocol, our goal to produce probabilistic statements (in principle, they can turn out to be false).
- In the decision-making protocol, we are merely minimizing our loss.

General scheme of defensive decision making

- Choose a goal that could be achieved if you knew the true probabilities generating the observations.
- Construct a decision strategy provably achieving your goal.
- Isolate a continuous law of probability on which the proof depends.
- Use defensive forecasting to get rid of the true probabilities.

The goal should be:

1. in terms of observables;
2. achievable regardless of what the true probabilities are.

The goal has to be relative.

Competitive decision making

Decision rule $D : \mathbf{X} \to \Gamma$.

We are given a benchmark class of decision rules and our goal is to perform almost as well as the best rules in the class. No assumptions about Reality.

Example: square-loss prediction competitive with Besov spaces

Proposition 5 Let $\mathbf{Y} = [-c, c]$, $p \geq 2$, $q \in [p', p]$ and $s > m/p$. There exist a constant $C_{\mathbf{X},s,p,q} > 0$ and a strategy for Decision Maker producing γ_n that are guaranteed to satisfy

$$\frac{1}{n}\sum_{i=1}^{n}(y_i - \gamma_i)^2 \leq \frac{1}{n}\sum_{i=1}^{n}(y_i - D(x_i))^2 + cC_{\mathbf{X},s,p,q}\left(\|D\|_{B^s_{p,q}} + c\right)n^{-1/p}$$

for all $n = 1, 2, \ldots$ and all $D \in B^s_{p,q}(\mathbf{X})$.

This uses a 1988 result by Cobos and Edmunds on the modulus of convexity of Besov spaces.

No upper bound on $\|D\|$, so we have universal consistency: for any continuous decision rule D,

$$\limsup_{n\to\infty}\left(\frac{1}{n}\sum_{i=1}^{n}(y_i-\gamma_i)^2-\frac{1}{n}\sum_{i=1}^{n}(y_i-D(x_i))^2\right)\le 0.$$

How to prove results such as Proposition 5 (case of probability forecasting)

Fix a choice function $G : \mathcal{P}(\mathbf{Y}) \to \Gamma$:

$$G(P) \in \arg\min_{\gamma \in \Gamma} \lambda(P, \gamma),$$

where

$$\lambda(P, \gamma) := \int_{\mathbf{Y}} \lambda(y, \gamma) P(\mathrm{d}y).$$

Informal statement The decisions $\gamma_n := G(P_n)$ ("ELM principle"), with P_n output by defensive forecasting for suitable laws of probability, satisfy

$$\frac{1}{n}\sum_{i=1}^{n}\lambda(y_i,\gamma_i) \lessapprox \frac{1}{n}\sum_{i=1}^{n}\lambda(y_i,D(x_i))$$

for all n and all regular decision rules D.

Derivation Using LLN:

$$\sum_{i=1}^{n} \lambda(y_i, \gamma_i) = \sum_{i=1}^{n} \lambda(y_i, G(P_i)) \approx \sum_{i=1}^{n} \lambda(P_i, G(P_i))$$
$$\leq \sum_{i=1}^{n} \lambda(P_i, D(x_i)) \approx \sum_{i=1}^{n} \lambda(y_i, D(x_i)).$$

Another approach to competitive decision making

Aggregating Algorithm (1990): mix all decision strategies in the benchmark class (generalization of Bayes mixture).

AA for Besov spaces

Combining the AA with Edmunds and Triebel's (1996) result about the metric entropy of the unit balls in Besov spaces:

Proposition 6 Under conditions of Proposition 5,

$$\frac{1}{n}\sum_{i=1}^{n}(y_i-\mu_i)^2 \le \frac{1}{n}\sum_{i=1}^{n}(y_i-D(x_i))^2 + c^{1+\frac{s}{m+s}} C_{\mathbf{X},s,p,q}\left(\|D\|_{B^s_{p,q}}+c\right)^{\frac{m}{m+s}} n^{-\frac{s}{m+s}}$$

for all $D \in B^s_{p,q}(\mathbf{X})$ from some n on.

p is irrelevant (unlike DF, where s was irrelevant)

Comparison for the Hölder–Zygmund spaces $\mathcal{C}^s(\mathbf{X}) := B^s_{\infty,\infty}(\mathbf{X})$

For $s = k + \alpha$, where k is integer and $\alpha \in (0, 1)$, $\mathcal{C}^s(\mathbf{X})$ consists of the functions whose kth partial derivatives exist and are all Hölder continuous of order α.

Defensive forecasting works better than aggregation at the "rough" end of the scale $\mathcal{C}^s(\mathbf{X})$:

- Suppose $s \in (0, m/2]$. The DF exponent $-1/p$ of n can be taken arbitrarily close to $-s/m$, and we can see that it is then better than the AA exponent of n:

$$\frac{s}{m} > \frac{s}{m+s}.$$

For example, if $m = 1, s \approx 1/2$ (typical trajectories of the Brownian motion are of this type) defensive forecasting gives approximately $n^{-1/2}$ whereas aggregation gives approximately $n^{-1/3}$.

- Suppose $s \in (m/2, m)$. The DF exponent of n can always be taken as $-1/2$, and it is still better than the AA exponent of n:

$$\frac{1}{2} > \frac{s}{m+s}.$$

- Suppose $s \in [m, \infty)$. A weakness of the method of defensive forecasting (in its current state) is that it cannot give regret terms better than $O(n^{-1/2})$. Therefore, the method of aggregation beats defensive forecasting for smooth $\mathcal{C}^s(\mathbf{X})$, $s > m$.

Limitations of defensive decision making

Competitive decision making: its goal implicitly assumes a small decision maker.

Remember a typical guarantee:

$$\frac{1}{n}\sum_{i=1}^{n}\lambda(y_i,\gamma_i)\le\frac{1}{n}\sum_{i=1}^{n}\lambda(y_i,D(x_i))+O\left(n^{-\alpha}\right).$$

Ideal probability forecasts (actual) are not enough in big decision making!

Simple example: $\Gamma = \{0, 1\}$, λ is given by the matrix

Reality \ Decision Maker	0	1
0	1	2
1	2	0

Reality's strategy: $y_n := \gamma_n$. Decision Maker's theory: Reality always chooses $y_n = 0$.

Decision Maker's mistake: he was being greedy (concentrated on exploitation and completely neglected exploration). But:

- he acted optimally given his beliefs,
- his beliefs have been verified by what actually happened.

We have to worry about what would have happened if we had acted in a different way.

Big decision making

Perhaps the Popperian recipe is unavoidable:

1. constantly devise [open stochastic] theories;
2. test them severely; when they fail, go to 1;
3. constantly devise decision strategies; to evaluate their future performance use the surviving theories.

Game-theoretic probability may be useful in steps 2 and 3.

Open problems (random selection)

- Prove formal results about big decision making based on testing open stochastic theories [might be too ambitious].
- Develop defensive forecasting for other laws of probability.
- Are there ways to make use of both smoothness and convexity of benchmark classes? More specifically: prove matching upper and lower bounds for decision making competitive with Besov spaces.