Asymptotic uniqueness in long-term prediction

Vladimir Vovk

Department of Computer Science
Royal Holloway, University of London

COPA 2024
Politecnico di Milano
Milano, 9 September, 2024

Not my topic. . .

Vladimir Vovk.
Conformal e-prediction.
arXiv 2001.05989 (August 2024).
Conformal e-prediction (review).

Would have been a nice complement to the previous talk.

This talk

- Not about conformal prediction.
- The "probabilistic" bit of COPA (Conformal and Probabilistic Prediction with Applications).
- About the uniqueness of successful probabilities (Jeffreys's law).

> . . . scientific disagreements tend to disappear. . . when new data accumulate. . . .
>
> Harold Jeffreys, 1938

Plan

1. Jeffreys's law
2. Context

Game-theoretic testing

- Suppose we have Forecaster outputting predictions. How do we test him?
- Idea: we gamble against him.
- My previous work with Shafer and Dawid: only testing one-step-ahead forecasts.
- This talk: forecasts for the infinite future.
- Reality produces observations $y_n \in \mathbf{Y}$ from a finite observation space $\mathbf{Y}$.

Prediction protocol

$\mathcal{K}_0 := 1$

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

Forecaster announces $P_n \in \mathfrak{P}(\mathbf{Y}^\infty)$

IF $n > 1$:

$$\mathcal{K}_{n-1} := \mathcal{K}^{-}_{n-1} + \sum_{x \in \mathbf{Y}^+} f_{n-1}(y_{n-1}x) P_n([x]) - \sum_{x \in \mathbf{Y}^*: |x| > 1} f_{n-1}(x) P_{n-1}([x])$$

Sceptic announces $f_n \in \mathbb{R}^{\mathbf{Y}^+}$ such that

$f_n(x) = 0$ for all but finitely many $x \in \mathbf{Y}^+$

Reality announces $y_n \in \mathbf{Y}$

$\mathcal{K}^{-}_n := \mathcal{K}_{n-1} + f_n(y_n) - \sum_{y \in \mathbf{Y}} f_n(y) P_n([y])$.

Interpretation (1)

- At each step n, Forecaster announces a probability measure P_n for the infinite future $y_n, y_{n+1}, \ldots$.
- The betting interpretation of P_n: for each $x \in \mathbf{Y}^+$, $P_n([x])$ is the price of a ticket (the x-ticket) that pays $1_{\{x \subseteq (y_n, y_{n+1}, \ldots)\}}$.
- Forecaster allows his opponent to buy any real number of such tickets.
- Testing the forecasts is performed by Sceptic: at step n, for each $x \in \mathbf{Y}^+$, Sceptic announces the number $f_n(x)$ of x-tickets that he chooses to buy at this step.
- The numbers $f_n(x)$ are allowed to be different from zero only for finitely many x-tickets, and so the sums $\sum_x$ are uncontroversial.

Interpretation (2)

- At the end of each step Reality announces the actual observation $y_n \in \mathbf{Y}$, and the y-tickets for $y \in \mathbf{Y}$ are cashed in. [One-step-ahead component.]
- The x-tickets for longer x are sold at the next step at the new prices and the gain/loss due to their total cost is recorded also at the next step.
- $\mathcal{K}_n$ Sceptic's capital at time n; large $\mathcal{K}_n$ means that Forecaster's predictions have been discredited ("play money").
- Sceptic is never allowed to go into debt: as soon as $K_n < 0$, the game is stopped and Sceptic's attempt at discrediting Forecaster fails.

Double testing protocol

- Suppose there are two Forecasters; at step n Forecaster I announces $P^{\mathrm{I}}_n \in \mathfrak{P}(\mathbf{Y}^\infty)$ and Forecaster II announces $P^{\mathrm{II}}_n \in \mathfrak{P}(\mathbf{Y}^\infty)$.
- Sceptic chooses f^{I}_n against Forecaster I and f^{II}_n against Forecaster II;
- The resulting capitals are $\mathcal{K}^{\mathrm{I}}_n$ and $\mathcal{K}^{\mathrm{II}}_n$.
- Forecasters might not even know about each other.

Theorem

The total variation distance:

$$\|P - Q\| := 2 \sup_E |P(E) - Q(E)| \in [0, 2].$$

Theorem

Sceptic has a strategy that guarantees the disjunction of

- $\|P_n^{\mathrm{I}} - P_n^{\mathrm{II}}\| \to 0$ *as* $n \to \infty$,
- $\mathcal{K}_n^{\mathrm{I}} \to \infty$ *as* $n \to \infty$,
- $\mathcal{K}_n^{\mathrm{II}} \to \infty$ *as* $n \to \infty$.

Comments

- At least one of the two Forecasters will be discredited if they do not issue almost identical forecasts in the long run.
- On the positive side, the forecasts are for the infinite future.
- On the negative side, the result is merely asymptotic (no inequalities in finite time).

Plan

Idea of the proof

- If the two Forecasters issue very different forecasts, we can mix those forecasts (as scaled up geometric mean) beating the geometric mean of the Forecasters.
- The corresponding strategy for Sceptic will beat at least one of the Forecasters.
- "Beat" in the sense $\limsup = \infty$; can be turned into $\lim = \infty$.

Blackwell–Dubins theorem (two-sided)

- Blackwell and Dubins (who were Bayesians) proved a measure-theoretic version of Jeffreys's law.
- My version (game-theoretic) implies theirs.
- Their result: if two probability measures (expressing beliefs of Forecaster I and Forecaster II) agree about the events of measure 0, their forecasts will converge in total variation a.s.
- Interpretation: both Forecasters believe they will converge. (No testing is involved; it's all about beliefs.)

One-step-ahead prediction

- In our books, Glenn Shafer and I proved Jeffreys's law for one-step-ahead forecasts.
- Their advantage is that they are not merely asymptotic; namely, they can be expressed as inequalities.
- They develop earlier measure-theoretic work by Kakutani and Kabanov/Liptser/Shiryaev.

Radical probabilism

- This talk: I assumed that the Forecasters and Sceptic learn the observations y_n.
- Richard Jeffrey introduced the picture of radical probabilism, where we never learn any observations for sure (but the probability of some observations can become very close to 1).
- The paper in the Proceedings: a version of Jeffreys's law (more general) under radical probabilism.

Bibliography

Glenn Shafer and Vladimir Vovk.
Game-Theoretic Foundations for Probability and Finance.
Hoboken, NJ: Wiley, 2019.
Section 10.7: one-step-ahead Jeffreys law.

Thank you for your attention!