Game-theoretic probability and some of its applications: *references and addenda*

Vladimir Vovk
vovk@cs.rhul.ac.uk
http://vovk.net

May 4, 2007

In the talk I state several results without proof (or only with a proof sketch, such as the game-theoretic strong law of large numbers). This note gives references to papers and technical reports where full proofs and more general statements can be found.

Slide 4

A more detailed description of the theory's forecasts is as follows. Given the parents' genotypes F_1F_2 and M_1M_2, the theory's forecast for the child's genotype is the probability measure that assigns a weight of $1/4$ to each of the four genotypes F_1M_1, F_1M_2, F_2M_1, F_2M_2 (with the letters in each of the four genotypes sorted alphabetically, and if some of these genotypes happen to coincide, their weights should be combined).

Slides 5–7

For other examples of open theories, described from the point of view of game-theoretic probability, see, e.g., [16], Section 8.4 (quantum mechanics as formalized by von Neumann, Cox's regression model), and [21] (the "classical theory of probability").

In principle, even open theories can be tested using measure-theoretic probability. These are two examples of how this could be done in the case of the theory of human blood group inheritance:

- To test the theory, we might only record the genotypes of the children born, say, to OO/AB parents. According to the theory, the genotypes of such children with be independent and will take values AO or BO with equal probabilities 1/2; this is just like tossing a fair coin, and one could use standard methods to reject the theory when the long-term frequency of AO is significantly different from 0.5.

- We could record all observations (the parents' and the child's genotypes) and consider the composite null hypothesis consisting of probability measures with the conditional probabilities of the child's genotype given the parents' genotypes agreeing with the theory. Testing this composite null hypothesis using standard methods appears to be the closest approximation to prequential testing in measure-theoretic probability (cf. the strong prequential principle in [9]).

However, prequential testing seems to be the cleanest general procedure.

Another interesting case is where Forecaster is a human, such as a weather forecaster. It can be argued that weather forecasters never follow a forecasting strategy: their methods only work for a narrow range of conditions similar to the usual ones. (And so measure-theoretic probability is not a suitable tool for testing their forecasts.)

For A. P. Dawid's exposition of the prequential principle, see, e.g., [5] (the principle introduced), [6] (where the prequential principle is recast as a metacriterion, M2, for choosing a criterion of success for forecasting systems), [8] (where the prequential principle is complemented by the prequential and partial prequential models) and [9] (where the prequential principle is complemented by strong and super-strong prequential principles). My personal view is that the (weak) prequential principle remains the most fundamental element of the prequential approach.

Slide 8

It is Sceptic's responsibility to ensure that f_n is P_n-integrable.

Slide 9

The measure-theoretic formalization of (1) on slide 9 can be found in, e.g., [16], Corollary 8.1 (this result is derived in [16] from a game-theoretic result; for a derivation not using game-theoretic probability see, e.g., [17], Theorem VII.5.4).

Slide 13

The web site for Iowa Electronic Markets is `http://www.biz.uiowa.edu/iem/`.

For the general definition of game-theoretic probability, see [9], Section 7.2, and [16], Section 7.1; for the case of probability forecasting, see also [21], Section 4. The more general notion of game-theoretic expectation can also be defined ([16], Section 7.1).

Slides 14–17

For the details of this proof (going back to Ville's [18] martingale proof of the law of the iterated logarithm), see [16], Section 3.2.

Slide 15

The "weakly force" in the statement of the lemma means that Sceptic has a strategy that does not risk bankruptcy and ensures that either $\sup_n \mathcal{K}_n = \infty$ or

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

The "a.s." means that the event is almost certain. It is clear that we can set $\mu_n := 0$ without loss of generality.

Slide 18

Kolmogorov's SLLN, generalizing the simple SLLN given in the talk to unbounded observations, is considered in Chapter 4 of [16]. In the forecasting protocol used for this result, Forecaster is asked to announce not only μ_n (his expectation for y_n) but also v_n (his expectation of the squared deviation of y_n from μ_n). The assumption that y_n are bounded is relaxed to

$$\sum_{n=1}^{\infty} \frac{v_n}{n^2} < \infty.$$

An opposite statement (asserting the existence of a strategy for Reality) is also proved.

Kumon, Takemura and Takeuchi have studied game-theoretic versions of the classical SLLN (also due to Kolmogorov) for independent identically distributed observations [14].

The law of the iterated logarithm is treated in Chapter 5 of [16]. The game-theoretic treatment reveals new features of the iterated-logarithm phenomenon; for example, it turns out that the two parts of LIL (validity and sharpness) require different assumptions.

A game-theoretic version of Lindeberg's CLT is proved in Chapter 7 of [16]. The versatility of the game-theoretic framework is demonstrated in the previous chapter (Section 6.3), where an important one-sided CLT is proved (based on Doob's parabolic potential theory [10]).

Continuous-time processes are treated in [16] using non-standard analysis (see, in particular, the treatment of stochastic differential equations in Chapter 14). A game-theoretic treatment of two simple kinds of continuous-time processes (counting and diffusion processes) not using non-standard analysis is presented in [20].

Slide 19

The most influential advocate of the idea that a theory is rejected only when it is replaced by a better theory was perhaps Kuhn [13], but this idea also pervades much of the classical theory of testing statistical hypotheses, with its emphasis not only on the null but also on the alternative hypothesis.

Slides 20–21

Proposition 2 is proved in [30] (see Theorem 2 of that paper, which also contains an opposite statement to Proposition 2).

The first Jeffreys's law was proved by Dawid ([6], Theorem 7.1), although Blackwell and Dubins [2] were among important predecessors. A result similar to Proposition 2 but stated in terms of absolute continuity and singularity of probability measures (analogously to Blackwell and Dubins's result) was obtained by Kabanov, Liptser and Shiryaev [12].

Slide 22

The statement that any continuous strategy for Sceptic can be transformed into a strategy for Forecaster that does not allow Sceptic's capital to grow is true for a wide range of prediction protocols: see [31] (Theorem 1) and [28] (Lemma 3).

Slides 23–24

For details and implications of Takemura's procedure, see [32].

Slides 25–32

The idea of defensive forecasting was applied to calibration and resolution in [25]. The exposition closest to that of this talk is in [26]. In particular, Proposition 4 follows from [26], Theorem 2 in combination with (15) (cf. (37)).

Slide 26

Derivation of a strategy for Sceptic ensuring

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

using defensive forecasting from a law of large numbers can be found in [32], Section 4.1.

Slide 27

A convoluted LLN is proved in [31], the end of Section 5.

Slides 28–29

Good sources for the theory of Besov and equally important Triebel–Lizorkin spaces, $B^s_{p,q}$ and $F^s_{p,q}$ respectively, are [1] and [11]. These two scales of function spaces include many classical function spaces, such as Sobolev spaces and Hölder–Zygmund spaces (see slide 45). In the talk I concentrate on Besov spaces, but most of the results carry over to Triebel–Lizorkin spaces.

Important function classes not covered by these two scales are various classes of analytic functions.

Slide 32

Dawid's example is described in [7].

There is no contradiction with Proposition 4 because the μ_n output by the forecasting algorithm of Proposition 4 in response to Dawid's strategy for Reality will tend to 0 so fast that moderate-norm $F \in B^s_{p,q}$ (namely, F with $\|F\| \ll n^{1/p}$) will not be able to tell on which side of 0 the μ_n is.

For a further discussion of Dawid's example, see [32], Section 4.4.

Slides 33–42

The idea of using defensive forecasting for competitive decision making was introduced in [24] and further developed in [26] and [28]. In particular, Proposition 5 is proved in [26] (see (5), which, however, assumes $p = q$). The 1988 result by Cobos and Edmunds used in the proof is Theorem 3 of [4].

In the proof of Proposition 5, which involves a very simple loss function, probability forecasting (as used on slides 40–42) can be replaced by mean-value forecasting (the protocol of slide 10). The sketch on slides 40–42 is formalized in [29].

Slide 43

The Aggregating Algorithm was introduced in [19]; for a more general description, see [23]. Its optimality, in a natural sense, was proved in [22]. There are many other similar algorithms, some of which more computationally efficient than the AA; for a recent review of the field, see [3].

Slides 44–47

Proposition 6 is proved in [27] (inequality (52)). Edmunds and Triebel's result used in the proof is Theorem 3.5 in [11] (applied to $s_1 := s$, $p_1 := p$, $q_1 := q$, $s_2 := 0$, $p_2 := \infty$ and $q_2 := 1$, in combination with (2.3.3/3)). This is a very general result giving precise estimates of the metric entropy of the unit ball in various Besov spaces considered a subset of another Besov space; we needed a small part of it. Once the metric entropy of the unit ball is known to be not too large, the unit ball can be approximated by a finite set of functions, which can be merged using the AA. Further merging over balls of various sizes gives a decision strategy competitive with the Besov space.

The "regret term" (such as the second line of the formula in Proposition 6) grows as $O(n^{-\alpha})$ for some $\alpha \in (0, 1)$ in the case of Besov spaces. For finite-dimensional benchmark classes (such as the class of linear functions on a subset of Euclidean space) the regret term can be as small as $O(\log n/n)$ ([23], Theorem 1). Even for some infinite-dimensional classes dense in the class of all continuous functions, such as various classes of analytic functions, it can be as small as $O(\log^c n/n)$ for a constant c ([27], Section 5).

The argument of slides 46–47 uses the Sobolev embedding theorem: see, e.g., [11], (3.5/1).

References

[1] Robert A. Adams and John J. F. Fournier. *Sobolev Spaces*, volume 140 of *Pure and Applied Mathematics*. Academic Press, Amsterdam, second edition, 2003.

[2] David Blackwell and Lester Dubins. Merging of opinions with increasing information. *Annals of Mathematical Statistics*, 33:882–886, 1962.

[3] Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, Cambridge, England, 2006.

[4] Fernando Cobos and David E. Edmunds. Clarkson's inequalities, Besov spaces and Triebel–Sobolev spaces. *Zeitschrift für Analysis und ihre Anwendungen*, 7:229–232, 1988.

[5] A. Philip Dawid. Statistical theory: the prequential approach. *Journal of the Royal Statistical Society* A, 147:278–292, 1984.

[6] A. Philip Dawid. Calibration-based empirical probability (with discussion). *Annals of Statistics*, 13:1251–1285, 1985.

[7] A. Philip Dawid. Self-calibrating priors do not exist: Comment. *Journal of the American Statistical Association*, 80:340–341, 1985. This is a contribution to the discussion in [15].

[8] A. Philip Dawid. Fisherian inference in likelihood and prequential frames of reference (with discussion). *Journal of the Royal Statistical Society* B, 53:79–109, 1991.

[9] A. Philip Dawid and Vladimir Vovk. Prequential probability: principles and properties. *Bernoulli*, 5:125–162, 1999.

[10] Joseph L. Doob. *Classical Potential Theory and Its Probabilistic Counterpart*. Springer, New York, 1984.

[11] David E. Edmunds and Hans Triebel. *Function Spaces, Entropy Numbers, Differential Operators*, volume 120 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, England, 1996.

[12] Yury M. Kabanov, Robert Sh. Liptser, and Albert N. Shiryaev. To the question of absolute continuity and singularity of probability measures. *Mathematics of the USSR—Sbornik*, 33:203–221, 1977.

[13] Thomas S. Kuhn. *The Structure of Scientific Revolutions*. University of Chicago Press, Chicago, 1962. Third edition: 1996.

[14] Masayuki Kumon, Akimichi Takemura, and Kei Takeuchi. Game-theoretic versions of strong law of large numbers for unbounded variables. To appear in *Stochastics*, 2007.

[15] David Oakes. Self-calibrating priors do not exist (with discussion). *Journal of the American Statistical Association*, 80:339–342, 1985.

[16] Glenn Shafer and Vladimir Vovk. *Probability and Finance: It's Only a Game!* Wiley, New York, 2001.

[17] Albert N. Shiryaev. *Probability*. Springer, New York, second edition, 1996. Third Russian edition published in 2004.

[18] Jean Ville. *Etude critique de la notion de collectif*. Gauthier-Villars, Paris, 1939.

[19] Vladimir Vovk. Aggregating strategies. In Mark Fulk and John Case, editors, *Proceedings of the Third Annual Workshop on Computational Learning Theory*, pages 371–383, San Mateo, CA, 1990. Morgan Kaufmann.

[20] Vladimir Vovk. Forecasting point and continuous processes: prequential analysis. *Test*, 2:189–217, 1993.

[21] Vladimir Vovk. A logic of probability, with application to the foundations of statistics (with discussion). *Journal of the Royal Statistical Society* B, 55:317–351, 1993.

[22] Vladimir Vovk. A game of prediction with expert advice. *Journal of Computer and System Sciences*, 56:153–173, 1998.

[23] Vladimir Vovk. Competitive on-line statistics. *International Statistical Review*, 69:213–248, 2001.

[24] Vladimir Vovk. Competitive on-line learning with a convex loss function. Technical Report `arXiv:cs.LG/0506041` (version 3), `arXiv.org` e-Print archive, September 2005. Also published in: Sanjay Jain, Hans Ulrich Simon and Etsuji Tomita, editors, *Proceedings of the Sixteenth International Conference on Algorithmic Learning Theory*, volume 3734 of *Lecture Notes in Artificial Intelligence*, pages 444–458, Springer, Berlin, 2005.

[25] Vladimir Vovk. Non-asymptotic calibration and resolution. Technical Report `arXiv:cs.LG/0506004` (version 3), `arXiv.org` e-Print archive, August 2005. Also published in: Sanjay Jain, Hans Ulrich Simon and Etsuji Tomita, editors, *Proceedings of the Sixteenth International Conference on Algorithmic Learning Theory*, volume 3734 of *Lecture Notes in Artificial Intelligence*, pages 429–443, Springer, Berlin, 2005. To appear in *Theoretical Computer Science*.

[26] Vladimir Vovk. Competing with wild prediction rules. Technical Report arXiv:cs.LG/0512059 (version 2), arXiv.org e-Print archive, January 2006. Also published in: Gábor Lugosi and Hans Ulrich Simon, editors, *Proceedings of the Nineteenth Annual Conference on Learning Theory*, volume 4005 of *Lecture Notes in Artificial Intelligence*, pages 559–573, Springer, Berlin, 2006. To appear in *Machine Learning*.

[27] Vladimir Vovk. Metric entropy in competitive on-line prediction. Technical Report arXiv:cs.LG/0609045, arXiv.org e-Print archive, September 2006.

[28] Vladimir Vovk. Predictions as statements and decisions. Technical Report arXiv:cs.LG/0606093, arXiv.org e-Print archive, June 2006. A one-page abstract published in: Gábor Lugosi and Hans Ulrich Simon, editors, *Proceedings of the Nineteenth Annual Conference on Learning Theory*, volume 4005 of *Lecture Notes in Artificial Intelligence*, page 4, Springer, Berlin, 2006.

[29] Vladimir Vovk. Predictions as statements and decisions. Technical Report arXiv:cs.LG/0606093, arXiv.org e-Print archive, June 2006.

[30] Vladimir Vovk. Merging of opinions in game-theoretic probability. Technical Report arXiv:0705.0372 [math.PR], arXiv.org e-Print archive, May 2007.

[31] Vladimir Vovk, Ilia Nouretdinov, Akimichi Takemura, and Glenn Shafer. Defensive forecasting for linear protocols. Technical Report arXiv:cs.LG/0506007 (version 2), arXiv.org e-Print archive, September 2005. Also published in: Sanjay Jain, Hans Ulrich Simon and Etsuji Tomita, editors, *Proceedings of the Sixteenth International Conference on Algorithmic Learning Theory*, volume 3734 of *Lecture Notes in Artificial Intelligence*, pages 459–473, Springer, Berlin, 2005.

[32] Vladimir Vovk, Akimichi Takemura, and Glenn Shafer. Defensive forecasting. Technical Report arXiv:cs.LG/0505083, arXiv.org e-Print archive, May 2005. Also published in the Proceedings of the Tenth International Workshop on Artificial Intelligence and Statistics, January 6–8, 2005, Savannah Hotel, Barbados.