Conformal changepoint detection in continuous model situations

Ilia Nouretdinov I.R.NOURETDINOV@RHUL.AC.UK
Vladimir Vovk V.VOVK@RHUL.AC.UK
Alex Gammerman A.GAMMERMAN@RHUL.AC.UK
Centre for Reliable Machine Learning, Royal Holloway, University of London, Egham, Surrey, UK

Editor: Lars Carlsson, Zhiyuan Luo, Giovanni Cherubin, and Khuong An Nguyen

Abstract

Conformal prediction provides a way of testing the IID assumption, which is the standard assumption in machine learning. A natural question is whether this way of testing is efficient. A typical situation where the IID assumption is broken is the existence of a changepoint at which the distribution of the data changes. We study the case of a change from one continuous distribution to another with both distributions belonging to standard parametric families. Our conclusion is that the conformal approach to testing the IID assumption is efficient, at least to some degree.

Extended abstract

The conformal approach to testing the IID assumption (Vovk et al., 2005, Section 7.1) consists in constructing a conformal test martingale (CTM), which is a nonnegative process with initial value 1 that is a martingale under any IID distribution. Informally, the value of a CTM is the capital of a gambler betting against the IID assumption, and its large values provide evidence against the IID assumption.

A typical case where the IID assumption is violated is the presence of a changepoint. It has been claimed (Vovk, 2020, 2021) that in the binary case, where the observations are either 0 or 1, the conformal approach is efficient in that the presence of a changepoint can be detected using CTMs. However, the binary case is of a very limited interest in applications. In this note we show, in a simulation study, that CTMs are still efficient, at least to some degree, beyond the binary case.

Our results are presented in Table 1. The total number of independent observations is $N = 10000$ and the changepoint is at $T = 5000$, so that the first T observations $y_1, \dots, y_T$ are IID and the last $N - T$ observation $y_{T+1}, \dots, y_N$ are also IID (but typically with a different distribution). We consider three kinds of changes: a Gaussian distribution $N(\mu, \sigma)$ changes one of its parameters (mean μ or standard deviation σ), an exponential distribution $\mathrm{Exp}(\lambda)$ changes its rate λ, or the "almost uniform" distribution $\mathrm{AU}(c)$ on $[0, 1]$ with the cumulative distribution function $F_1(y) = y^c$ and parameter $c > 0$ changes to its "reflected" version with the cumulative distribution function $F_2(y) = 1 - (1 - y)^c$.

Let d_1 and d_2 be the pre-change and post-change probability density functions, respectively. As our benchmark we will use the likelihood ratio

$$\frac{\Pi_{i=1}^{T} d_1(y_i) \Pi_{i=T+1}^{N} d_2(y_i)}{\Pi_{i=1}^{N} \left(\frac{T}{N} d_1(y_i) + \frac{N-T}{N} d_2(y_i) \right)}$$

pre-change	post-change	benchmark	CTM
$N(0,1)$	$N(0,1)$	0 (0)	−1.6 (1.0)
$N(0,1)$	$N(0.5,1)$	130.8 (10.8)	126.6 (11.3)
$N(0,1)$	$N(0.2,1)$	21.3 (4.4)	18.5 (5.2)
$N(0,1)$	$N(0.1,1)$	5.3 (2.2)	3.0 (3.2)
$N(0,1)$	$N(0,1.5)$	154.3 (7.9)	150.1 (8.3)
$N(0,1)$	$N(0,1.1)$	8.8 (2.3)	6.2 (2.4)
$N(0,1)$	$N(0,0.9)$	12.3 (2.8)	9.7 (3.6)
$N(0,1)$	$N(0,0.7)$	125.8 (8.6)	121.8 (9.5)
Exp(1)	Exp(0.7)	65.2 (7.6)	61.6 (8.1)
Exp(1)	Exp(0.9)	5.6 (2.3)	3.4 (3.2)
AU(0.7)	reflected	196.3 (12.8)	191.5 (13.4)
AU(0.9)	reflected	19.1 (4.2)	16.3 (5.1)

Table 1: Results for Gaussian, exponential, and almost uniform distributions.

of the true distribution to a distribution (the minimax distribution in the sense of Vovk, 2021) in the IID family.

Each CTM is determined by specifying a conformity measure and a betting function. The former will be based on the Neyman–Pearson lemma, while the latter will be in the spirit of the plug-in approach of Fedorova et al. (2012). Namely, the conformity score of the ith observation y_i is defined as $\log d_1(y_i) - \log d_2(y_i)$ (the Neyman-Pearson statistic on the log scale), and the betting function is 1 before the change and is calculated in the following two steps after the change. First, for each time step we calculate an empirical probability density function f for the conformal p-values using 5000 simulations from the same true stochastic mechanism but for different seeds for the pseudorandom number generator. Namely, the value of f between two adjacent simulated conformal p-values is inversely proportional to the distance between them. Second, the probability density function f is forced to be monotonically decreasing by applying the procedure of isotonic regression to it (namely, by applying the `lsqisotonic` function in MATLAB). The resulting function is then used as the betting function. Possibly, the second step would not be needed if 5000 were replaced by a much larger number.

When presenting our results in Table 1, we report the decimal logarithms of the final values of the CTMs, averaged over 50 random seeds. The estimated standard deviations of those logarithms are given in parentheses.

Notice that, when designing our CTMs, we use the full knowledge of the data-generating mechanism. Our goal is to demonstrate that conformal testing does not impose any intrinsic limitations, and that in principle it allows us to compete with the benchmark. Conformal testing is the only method known at this time that provides processes that are test martingales with respect to any IID distribution. Our experiments show that the gap between the performance of conformal testing and that of the benchmark is not excessive, and so the conformal approach to testing the IID assumption is not severely limited in its potential.

References

Valentina Fedorova, Ilia Nouretdinov, Alex Gammerman, and Vladimir Vovk. Plug-in martingales for testing exchangeability on-line. In John Langford and Joelle Pineau, editors, *Proceedings of the Twenty Ninth International Conference on Machine Learning*, pages 1639–1646. Omnipress, 2012.

Vladimir Vovk. Testing randomness online. Technical Report arXiv:1906.09256 [math.PR], arXiv.org e-Print archive, March 2020. To appear in *Statistical Science*.

Vladimir Vovk. Conformal testing in a binary model situation. *Proceedings of Machine Learning Research*, 152:1–20, 2021. To appear in COPA 2021.

Vladimir Vovk, Alex Gammerman, and Glenn Shafer. *Algorithmic Learning in a Random World*. Springer, New York, 2005.