Bayes test for monotonicity

Jean-Bernard Salomond (jean.bernard.salomond@ensae.fr)
CREST & Université Paris Dauphine

Abstract

We propose a Bayesian procedure for testing monotonicity of a regression function. Our test is proved to be consistent and to achieve the optimal separation rate (up to a $\log(n)$ factor). We propose a choice for the prior distribution and study the behaviour of our test for finite samples.

Context and Aim

We consider the model

$$Y_i = f(i/n) + \epsilon_i, \quad \epsilon_i \overset{iid}{\sim} \mathcal{N}(0;\sigma^2), \quad (1)$$

with σ^2 fixed or unknown. We want to test

$$H_0 : f \searrow \text{ vs } H_1 : f \text{ not } \searrow$$

We thus test a non parametric null, versus a non parametric alternative

Bayesian approach to monotonicity testing

We build a prior on f by considering a piecewise constant approximation

$$f_{\omega,k}(.) = \sum_{i=1}^{k} \omega_i \mathbb{1}_{I_i}(.) \qquad (2)$$

and then consider the prior

$$\pi : \begin{cases} k & \sim \pi_k \\ \omega_i & \overset{iid}{\sim} g \\ \sigma & \sim h \end{cases}$$

The Bayes Factor approach fails under H_0 when f has flat parts (see [McVinish and Rousseau, 2011]). We thus consider a modified version of the BF. For a given M_n^k

$$\delta_n^\pi = \mathbb{1}\Big\{\pi\big(H(\omega,k) > M_n^k | Y_n\big) > 1/2\Big\}$$

where

$$H(\omega,k) = \max_{j>i}(\omega_j - \omega_i)$$

it is thus **straightforward to implement**.

We want our test to be consistent against α-Hölderian alternatives

$$\sup_{f\searrow} \mathrm{E}_0^n(\delta_n^\pi) = o(1)$$
$$\sup_{d(f,\searrow)>\rho, f\in H_\alpha(L)} \mathrm{E}_0^n(1-\delta_n^\pi) = o(1) \qquad (3)$$

and to achieve an optimal separation rate ρ_n obtained in [Akakpo et al., 2012]

$$\sup_{f\in\searrow} \mathrm{E}_0^n(\delta_n^\pi) = o(1)$$
$$\sup_{d(f,\searrow)>\rho_n, f\in H_\alpha(L)} \mathrm{E}_0^n(1-\delta_n^\pi) = o(1) \qquad (4)$$

Theorem

Let $M_n^k = M_0\sqrt{k\log(n)/n}$ and let π be a prior on $f_{\omega,k}$ such that $\omega_i \overset{iid}{\sim} g$ and $k \sim \pi_k$, $\sigma \sim h$. Assume that g and h put mass on $\mathbb{R}$ and $\mathbb{R}^{+*}$ respectively and that π_k is such that there exist positive constants C_d and C_u such that

$$e^{C_d k L(k)} \leq \pi_k(k) \leq e^{C_u k L(k)}$$

Where $L(k)$ is either $\log(k)$ or 1. Consider the test

$$H_0 : f \searrow \ versus\ H_1 : f \text{ not } \searrow, f \in H_\alpha(L)$$

let $\rho_n = M\big(n/\log(n)\big)^{-\alpha/(2\alpha+1)}$, then the test δ_n^π is consistent and achieve the separation rate ρ_n.

Examples

- Simulated data with σ known
- Regression functions adapted from the frequentist literature
- We choose for the prior

$$\pi_k := \mathcal{P}(\lambda)$$
$$g := \mathcal{N}(m; s^2)$$

- explicit posterior → easy to sample from
- The hyperparameter λ has a great influence on the results

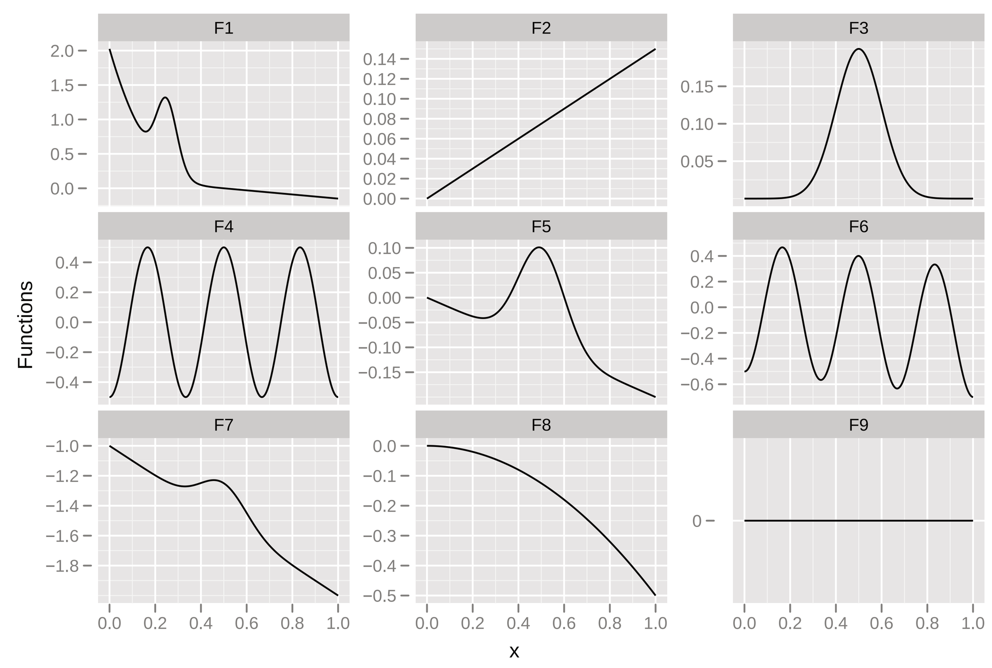

FIGURE: Plot of the regression function

FIGURE: Result of the simulation study

Comparison with frequentist methods

- Frequentist methods → computationally difficult
- We obtain similar results for $\{f_1, \ldots, f_6\}$
- Loss of power for f_7
- Similar Type I error for $\lambda = 3\sigma_y$

Extensions

Our method could easily be adapted to test for other qualitative assumptions

- H_0 : f is positive
- H_0 : f is convex
- H_0 : f is unimodal
- ...

References

[Akakpo et al., 2012] Akakpo, N., Balabdaoui, F., and Durot, C. (2012).
Testing monotonicity via local least concave majorants.

[McVinish and Rousseau, 2011] McVinish, R. and Rousseau, J. (2011).
Bayesian testing of decreasing densities.