ScienceDirect

Statistics & Probability Letters

Volume 146, March 2019, Pages 206-212

Quantification of the impact of priors in Bayesian statistics via Stein's Method

Fatemeh Ghaderinezhad, Christophe Ley

Show more

Share Cite

https://doi.org/10.1016/j.spl.2018.11.012

Get rights and content

Abstract

We compare two distinct non-uniform choices of prior distributions by quantifying the Wasserstein distance between the respective resulting posterior distributions at any fixed sample size by means of Stein's Method. We illustrate this measure of the prior impact on the normal, Binomial and Poisson models.

Introduction

A key question in Bayesian analysis is the effect of the prior on the posterior, and how we can measure this effect. Will the posterior distributions derived with distinct priors become very similar if more and more data are gathered? It has been proved formally inDiaconis and Freedman (1986a) andDiaconis and Freedman (1986b) that, under certain regularity conditions, the impact of the prior is waning as the sample size increases. From a practical viewpoint it is more important to know what happens at finite sample size n . Recently,Ley et al. (2017a) have provided a partial answer to this question by investigating the Wasserstein distance between the posterior distribution based on a given prior of interest and the no-prior data-only based posterior, which allows detecting at fixed sample size n the effect of the prior of interest. This distance being mostly impossible to calculate exactly, they have provided sharp upper and lower bounds. This work is, to the best of the authors' knowledge, the first to quantify at any sample size the prior effect. However, it strongly relies on the assumption that one prior is the flat uniform prior (or data-only prior), and hence it does not allow a direct comparison between two priors that are both non-uniform. Our aim in the present paper is to extend the methodology ofLey et al. (2017a) to incorporate such general settings with special focus on priors that lead to posteriors which have nested supports.

The paper is organized as follows. In Section2 we define our concept of measure of the distance between two priors, explain the practical relevance of our investigations and state and prove our main theorem which is related to the famous Stein's Method. Then in Section3 we illustrate the strength of our new measure of the difference between two priors by considering three examples of well-known distributions, namely the normal, Binomial and Poisson. In each case, we compare the effects of distinct priors on the posterior distribution of the parameter we are interested in. The choices of our priors are motivated by research papers that have discussed various choices of viable priors for a certain parameter.

Section snippets

Quantification of the effects of two distinct priors

We quantify the different effects of two distinct priors by measuring the distributional distance between the posteriors resulting from these two priors. As inLey et al. (2017a), we opt here for the Wasserstein-1 distance defined as $d_{\mathscr{W}}(P_1, P_2) = \sup_{h \in \mathscr{H}} |\mathrm{E}[h(X_1)] - \mathrm{E}[h(X_2)]|$for X_1 and X_2 random variables with respective distribution functions P_1 and P_2, and where $\mathscr{H}$ stands for the class of Lipschitz-1 functions. In general it is very hard to obtain an exact expression for the Wasserstein distance.…

Comparison of various priors for various distributions

We now illustrate the strength of Theorem1 by comparing popular choices of priors for parameters of three famous distributions, namely the normal, the Binomial and the Poisson models. We will work out in detail the case of the normal distribution. As we shall see, the bounds that we get allow us to conclude that, in all these cases, the difference between the resulting priors vanishes asymptotically independently of the observations, but that at finite sample size n the observations do play a…

Acknowledgments

We would like to thank the Associate Editor as well as two anonymous reviewers for helpful comments that led to an improvement of the present paper. This research is supported by a BOF Starting Grant of of Ghent University .…

Recommended articles

References (9)

DiaconisP. *et al.*

On the consistency of Bayes estimates (with discussion and rejoinder by the authors)

Ann. Statist. (1986)

DiaconisP. *et al.*

On inconsistent Bayes estimates of location

Ann. Statist. (1986)

GelmanA. *et al.*

Bayesian Data Analysis

(2004)

GhaderinezhadF.
New insights into the impact of the choice of the prior for the success parameter of binomial distributions
J. Math. Stat. Oper. Res. (2019)
There are more references available in the full text version of this article.

Cited by (5)

The Wasserstein Impact Measure (WIM): A practical tool for quantifying prior impact in Bayesian statistics
2022, Computational Statistics and Data Analysis

> *Citation Excerpt :*
> …Since it is mostly impossible to calculate this distance explicitly, the authors have provided sharp lower and upper bounds on the Wasserstein distance and their approach relies on a variant of the famous Stein Method. In order to compare any two priors directly, Ghaderinezhad and Ley (2019) recently extended their approach to any two priors for one-dimensional parameters, provided that the posteriors are nested and have finite first moments; see also Ghaderinezhad and Ley (2020). For practical purposes, the power of the Wasserstein distance idea has not been exploited so far.…

Show abstract

Stein's Method Meets Computational Statistics: A Review of Some Recent Developments
2023, Statistical Science

Bayesian Performance Analysis for Algorithm Ranking Comparison
2022, IEEE Transactions on Evolutionary Computation

Stein's Method Meets Computational Statistics: A Review of Some Recent Developments
2021, arXiv

Evaluating sensitivity to the stick-breaking prior in Bayesian nonparametrics
2018, arXiv

View full text