Computational Statistics and Data Analysis 67 (2013) 68–83

Contents lists available at SciVerse ScienceDirect

Computational Statistics and Data Analysis

journal homepage: www.elsevier.com/locate/csda

Bayesian computing with INLA: New features

Thiago G. Martins *, Daniel Simpson, Finn Lindgren, Håvard Rue

Department of Mathematical Sciences, Norwegian University of Science and Technology, N-7491 Trondheim, Norway

ARTICLE INFO

Article history:
Received 25 September 2012
Received in revised form 23 April 2013
Accepted 23 April 2013
Available online 2 May 2013

Keywords:
Approximate Bayesian inference
INLA
Latent Gaussian models

ABSTRACT

The INLA approach for approximate Bayesian inference for latent Gaussian models has been shown to give fast and accurate estimates of posterior marginals and also to be a valuable tool in practice via the R-package R-INLA. New developments in the R-INLA are formalized and it is shown how these features greatly extend the scope of models that can be analyzed by this interface. The current default method in R-INLA to approximate the posterior marginals of the hyperparameters using only a modest number of evaluations of the joint posterior distribution of the hyperparameters, without any need for numerical integration, is discussed.

1. Introduction

The Integrated Nested Laplace Approximation (INLA) is an approach proposed by Rue et al. (2009) to perform approximate fully Bayesian inference on the class of latent Gaussian models (LGMs). INLA makes use of deterministic nested Laplace approximations and, as an algorithm tailored to the class of LGMs, it provides a faster and more accurate alternative to simulation-based MCMC schemes. This is demonstrated in a series of examples ranging from simple to complex models in Rue et al. (2009). Although the theory behind INLA has been well established in Rue et al. (2009), the INLA method continues to be a research area in active development. Designing a tool that allows the user the flexibility to define their own model with a relatively easy to use interface is an important factor for the success of any approximate inference method. The R package INLA, hereafter referred to as R-INLA, provides this interface and allows users to specify and perform inference on complex LGMs.

The breadth of classical Bayesian problems covered under the LGM framework, and therefore handled by INLA, is – when coupled with the user-friendly R-INLA interface – a key element in the success of the INLA methodology. For example, INLA has been shown to work well with generalized linear mixed models (GLMM) (Fong et al., 2010), spatial GLMM (Eidsvik et al., 2009), Bayesian quantile additive mixed models (Yue and Rue, 2011), survival analysis (Martino, Akerkar et al., 2011), stochastic volatility models (Martino, Aas et al., 2011), generalized dynamic linear models (Ruiz-Cárdenas et al., 2011), change point models where data dependency is allowed within segments (Wyse et al., 2011), spatio-temporal disease mapping models (Schrödle and Held, 2011), models to complex spatial point pattern data that account for both local and global spatial behavior (Illian et al., 2011), and so on.

There has also been a considerable increase in the number of users that have found in INLA the possibility to fit models that they were otherwise unable to fit. More interestingly, those users come from areas that are sometimes completely

* Corresponding author. Tel.: +47 46937429.
E-mail addresses: thigm85@gmail.com, thiago.guerrera@math.ntnu.no (T.G. Martins), daniel.simpson@math.ntnu.no (D. Simpson), finn.lindgren@math.ntnu.no (F. Lindgren), Havard.Rue@math.ntnu.no (H. Rue).
URLs: http://www.math.ntnu.no/~guerrera (T.G. Martins), http://www.math.ntnu.no/~daniesi/ (D. Simpson), http://www.math.ntnu.no/~finnkrl/ (F. Lindgren), http://www.math.ntnu.no/~hrue/ (H. Rue).

0167-9473/$ – see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.csda.2013.04.014

unrelated to each other, such as econometrics, ecology, climate research, etc. Some examples are bi-variate meta-analysis of diagnostic studies (Paul et al., 2010), detection of under-reporting of cases in an evaluation of veterinary surveillance data (Schrödle et al., 2011), investigation of geographic determinants of reported human Campylobacter infections in Scotland (Bessell et al., 2010), the analysis of the impact of different social factors on the risk of acquiring infectious diseases in an urban setting (Wilking et al., 2012), analysis of animal space use metrics (Johnson et al., 2011), animal models used in evolutionary biology and animal breeding to identify the genetic part of traits (Holand et al., 2011), analysis of the relation between biodiversity loss and disease transmission across a broad, heterogeneous ecoregion (Haas et al., 2011), identification of areas in Toronto where spatially varying social or environmental factors could be causing higher incidence of lupus than would be expected given the population (Li et al., 2011), and spatio-temporal modeling of particulate matter concentration in the North-Italian region Piemonte (Cameletti et al., 2012). The relative black-box format of INLA allows it to be embedded in external tools for a more integrated data analysis. For example, Beale et al. (2010) mention that INLA has been used by tools embedded in a Geographical Information System (GIS) to evaluate the spatial relationships between health and the environment data. The model selection measures available in INLA are also something very much appreciated in the applied work mentioned so far. Such quantities include marginal likelihood, deviance information criterion (DIC) (Spiegelhalter et al., 2002), and other predictive measures.

Some extensions to the work of Rue et al. (2009) have also been presented in the literature; Hosseini et al. (2011) extends the INLA approach to fit spatial GLMM with skew normal priors for the latent variables instead of the more standard normal priors, Sørbye and Rue (2010) extend the use of INLA to joint inference and present an algorithm to derive analytical simultaneous credible bands for subsets of the latent field based on approximating the joint distribution of the subsets by multivariate Gaussian mixtures, Martins and Rue (2012) extend INLA to fit models where independent components of the latent field can have non-Gaussian distributions, and Cseke and Heskes (2011) discuss variations of the classic Laplace-approximation idea based on alternative Gaussian approximations (see also Rue et al. (2009, pp. 386–387) for a discussion on this issue).

A lot of advances have been made in the area of spatial and spatial–temporal models; Eidsvik et al. (2011) address the issue of approximate Bayesian inference for large spatial datasets by combining the use of prediction process models as a reduced-rank spatial process to diminish the dimensionality of the model and the use of INLA to fit these reduced-rank models. INLA blends well with the work of Lindgren et al. (2011) where an explicit link between Gaussian Fields (GFs) and Gaussian Markov Random Fields (GMRFs) allows the modeling of spatial and spatio-temporal data to be done with continuously indexed GFs while the computations are carried out with GMRFs, using INLA as the inferential algorithm.

The INLA methodology requires some expertise in numerical methods and computer programming to be implemented, since all procedures required to perform INLA need to be carefully implemented to achieve a good speed. This can, at first, be considered a disadvantage when compared with other approximate methods such as (naive) MCMC schemes that are much easier to implement, at least on a case-by-case basis. To overcome this, the `R-INLA` package was developed to provide an easy to use interface to the stand-alone `C` coded *inla program*.[1] To download the package one only needs one line of the `R` code that can be found in the download section of the INLA website (http://www.r-inla.org/). In addition, the website contains several worked out examples, papers and even the complete source code of the project.

In Rue et al. (2009) most of the attention was focused on the computation of the posterior marginals of the elements of the latent field since those are usually the biggest challenge when dealing with LGMs given the high dimension of the latent field usually found in the models of interest. On the other hand, it was mentioned that the posterior marginals of the unknown parameters not in the latent field, hereafter referred to as hyperparameters, are obtained via numerical integration of an interpolant constructed from evaluations of the Laplace approximation of the joint posterior of the hyperparameters already computed in the computation of the posterior marginals of the latent field. However, details of such an interpolant were not given. The first part of the paper will show how to construct this interpolant in a cost-effective way. Besides that, we will describe the algorithm currently in use in the `R-INLA` package that completely bypasses the need for numerical integration, providing accuracy and scalability.

Unfortunately, when an interface is designed, a compromise must be made between simplicity and generality, meaning that in order to build a simple to use interface, some models that could be handled by the INLA method might not be available through that interface, hence not available to the general user. The second part of the paper will formalize some new developments already implemented on the `R-INLA` package and show how these new features greatly extend the scope of models available through that interface. It is important to keep in mind the difference between the models that can be analyzed by the INLA method and the models that can be analyzed through the `R-INLA` package. The latter is contained within the first, which means that not every model that can be handled by the INLA method is available through the `R-INLA` interface. Therefore, this part of the paper will formalize tools that extend the scope of models within `R-INLA` that were already available within the theoretical framework of the INLA method.

Section 2 will present an overview of the latent Gaussian models and of the INLA methodology. Section 3 will address the issue of computing the posterior marginal of the hyperparameters using a novel approach. A number of new features

[1] The dependency on the stand-alone `C` program is the reason why `R-INLA` is not available on `CRAN`.