

NIASRA

NATIONAL INSTITUTE FOR APPLIED
STATISTICS RESEARCH AUSTRALIA



National Institute for Applied Statistics Research Australia

University of Wollongong, Australia

Working Paper

05-20

Discussion on “A High-resolution Bilevel Skew- t stochastic
Generator for Assessing Saudi Arabia's Wind Energy Resources”

Andrew Zammit-Mangion

*Copyright © 2020 by the National Institute for Applied Statistics Research Australia, UOW.
Work in progress, no part of this paper may be reproduced without permission from the Institute.*

National Institute for Applied Statistics Research Australia, University of Wollongong, Wollongong
NSW 2522, Australia Phone +61 2 4221 5076, Fax +61 2 4221 4998.

Email: karink@uow.edu.au

Discussion on “A high-resolution bilevel skew- t stochastic generator for assessing Saudi Arabia’s wind energy resources”

Andrew Zammit-Mangion

*National Institute for Applied Statistics Research Australia,
University of Wollongong, Australia*

Abstract

Statistical spatio-temporal environmental data analysis is rarely straightforward, with one having to face challenges relating to big data, non-Gaussianity, non-stationarity, multiple scales of behaviour, deterministic (numerical) model output, and more. One often has to rely heavily on good statistical parallel computing skills and sound knowledge of the application domain. The work of Tagle et al. (2020) overcomes all of these challenges, and is an excellent example of the tangible contributions spatio-temporal modelling and distributional theory can make to the environmental sciences at the policy level. In this discussion piece I focus on a few high-level concepts in the paper of Tagle et al. (2020) that are relevant to related application domains. I also provide some technical suggestions that could be used to facilitate inference.

Introduction

This discussion paper is broadly divided into two parts. In the first part I discuss some aspects of the work by Tagle et al. (2020) at a conceptual level; this part includes a brief discussion on the implications of treating numerical model output as ‘data,’ certain features of the spatio-temporal model used, and the authors’ handling of non-Gaussianity. In the second part I propose some alternative computational strategies that could be used to facilitate inference.

The implications of using numerical model output as ‘data’

Tagle et al. (2020) use high-resolution data from the Weather Research and Forecasting (WRF) model to construct a surrogate statistical model that can, in effect, replace WRF over the region and time horizon of interest. Such a strategy of replacing the deterministic model with a statistical one serves several purposes. One can use the statistical model to elicit a prior distribution when doing data assimilation (e.g., Zammit-Mangion et al., 2015, 2016), for statistical compression (e.g., Castruccio and Genton, 2016), or, as in Tagle et al. (2020), to play the role of a stochastic weather generator that can be simulated from. In such cases one may execute a ‘Turing test’ to ensure that prior predictive realisations from the fitted model cannot be distinguished from a numerical model output by a field expert (see Majumder et al., 2013, for an example in the context of linear models). The maps in this paper certainly seem convincing, but is granularity from the imposed regionalisation apparent at the daily scale? If yes, this may be problematic if one is interested in analysing the output from the stochastic wind generator at a fine temporal resolution.

A common issue when developing a product from a numerical model is that the said numerical model acts as a single point of failure. WRF has a good track record of being fairly accurate but, like any numerical model, is imperfect (e.g., Carvalho et al., 2014). In global carbon

dioxide flux inversion, for example, the lack of consensus between research groups worldwide is largely due to the disparity between atmospheric transport models (Schuh et al., 2019; McNorton et al., 2020). A related issue is that output from atmospheric transport models are generally not accompanied with a measure of uncertainty, which is often substituted with the notion of variability. One should, however, differentiate between the (prior or posterior) uncertainty on the physical process, and the system variability. When working with deterministic models, uncertainty quantification necessitates the consideration of the numerical model parameters (plausible configurations), uncertainties in the boundary conditions (in Tagle et al.’s case, from ECMWF), and model-data discrepancy (if real data is available). In flux inversion, a compromise is generally made by imposing and inflating uncertainties associated with the model outputs until inferences across model groups tend to be in agreement. Which leads me to the question: How would one account for uncertainty associated with the numerical model in this framework?

Another point worth mentioning is that the WRF output used is from a six-year period in the past. Recent studies show that climate change is causing wind speed patterns to change over relatively short time horizons (on the order of a few decades) that span the planning, construction, and production phases of wind energy plants (Zeng et al., 2019). Could this be an issue, and is it possible to somehow adapt the framework to reliably generate wind demand patterns over the next few decades?

The spatio-temporal model

Wind speed, when treated as a stochastic process, is likely to be non-separable in space and time, and the stencil neighbourhood scheme used in Tagle et al. (2020) does allow for non-separability. The model also allows for spatially varying dynamics, which is important since dynamics do tend to vary spatially. Figure S.3 appears to indicate that temporally varying dynamics do not need to be modelled, however this is very uncommon in environmental applications of this nature, especially at the daily scale and over large time horizons. Catering for time varying dynamics is sometimes difficult, but possible. For example, Zammit-Mangion and Wikle (2020) considered spatially varying parameters in a filtering context using convolution neural networks combined with the integro-difference equation. There, sea-surface temperature was modelled daily, and the dynamics were clearly changing on a daily scale. I suspect that modelling temporally varying dynamics would be important if Tagle et al. were to upscale their model to hourly time intervals to yield prior predictive realisations that are more realistic.

In addition, it is somewhat unusual for models of this kind to not incorporate covariates. Elevation, or terrain gradient, seem to be obvious candidates here. They may be used as fixed effects in $\eta(\mathbf{s})$, or in the covariance function (Schmidt et al., 2011). Alternatively, they may be used to parameterise other parameters appearing in the model, such as ν_r . One can expect that certain properties like tail behaviour are dependent on the local topography. Do Tagle et al. (2020) see any scope for adding covariates to their model, and in what way?

Modelling non-Gaussianity and the game of compromises

The use of a non-Gaussian process, and particularly the skew- t distribution, is elegant. However, moving from a latent Gaussian specification comes at a cost. The first casualty here appears to be the spatial continuity when modelling the residuals. The second casualty is the relative computational ease of making inference; a rather involved Metropolis-within-Gibbs-within-expectation-maximisation scheme is used, which I believe is feasible only because of the relatively severe modelling assumptions that have been adopted. While inference does not appear to have been an issue here, single-site Gibbs updating in conjunction with Metropolis-

Hastings is notorious for poor mixing. Further, good mixing and convergence must be difficult to assess/guarantee automatically for each expectation-maximisation (EM) iteration. The approach taken is pragmatic, but in other applications one might need to take a different stance when weighing the use of a pure non-Gaussian model against that of a latent Gaussian model.

Figure 2(a) in Tagle et al. (2020) depicts only mild departures from Gaussianity, so a latent Gaussian model might indeed be appropriate here. Cseke et al. (2016) fitted a spatio-temporal log-Gaussian Cox process model where the spatial dimension was approximately 10,000, and where the data spanned a few hundred time points. The neighbouring interactions were also modelled using a stencil-like approach, with a spike and slab prior on the interaction terms. A laptop was used for fitting, which suggests to me that the scales Tagle et al. are considering can be tackled within a large computing environment. This would allow one to do end-to-end estimation in one stage, and avoid the problem of regionalisation. However this comes at the cost of using a latent Gaussian approach. Tagle et al. provide some discussion on this compromise in the Introduction, but is there compelling evidence in this application that a latent Gaussian approach would have been inadequate?

One upside of regionalisation is that it leads to straightforward modelling of non-stationarity, which in this case extends to orders higher than the conventional first and second orders. With latent Gaussian models it is possible that this regionalisation could be avoided while still accommodating non-stationarity. One possible way forward is to consider a combined parallel model and parallel computation strategy such as that employed by Zammit-Mangion and Rougier (2020), wherein finer scales are associated with increasing degrees of non-stationarity.

Some thoughts on the inference procedure

As noted above, assessing convergence of the Gibbs sampler for each EM iteration is not a trivial problem. For example, the conditional distribution of each $U_{r,t}$ is likely to be bimodal since its prior distribution is symmetric around zero and only its absolute value appears in the likelihood. I suspect that, since a Metropolis–Hastings algorithm was used, only one mode was explored by the Gibbs sampler. This may not be a problem in practice, but one that should be borne in mind.

Since each $U_{r,t}$ and $Z_{r,t}$ is associated with a few hundred data points (since one has 84,000 spatial locations and 200 regions at each time point), it is likely that the posterior distributions over these variables (ignoring any bimodality) are quite tight around the posterior mean, and Gaussian-like. If this is indeed the case, one can possibly justify using Laplace approximations in the EM algorithm instead of a Gibbs sampler, as in Sengupta and Cressie (2013). This would lead to a considerable simplification in the inferential algorithm with potentially little loss of accuracy. If one is wedded to having a non-Gaussian treatment of the conditional distributions over $\{U_{r,t}\}$ and $\{Z_{r,t}\}$, then another alternative is to resort to an approximate inference scheme such as stochastic variational Bayes (SVB), which can be useful in a highly non-Gaussian setting, as I show next.

Consider the simple model

$$Y_i = \theta + \epsilon_i; \quad i = 1, \dots, m,$$

where $\epsilon_i \sim SN(0, 1, \lambda)$ (i.e., ϵ_i is skew-normally distributed with location zero, scale one, and slant λ), and $\theta \sim N(0, 1)$, and assume that one wishes to make inference on θ from the data $\mathbf{Y} \equiv (Y_1, \dots, Y_m)'$. In variational Bayes, one seeks to maximise a lower bound \mathcal{L} on $p(\mathbf{Y})$, given by

$$\mathcal{L} = \mathbb{E}_{q(\theta)} \left[\log \frac{p(\mathbf{Y} | \theta)p(\theta)}{q(\theta)} \right], \quad (1)$$

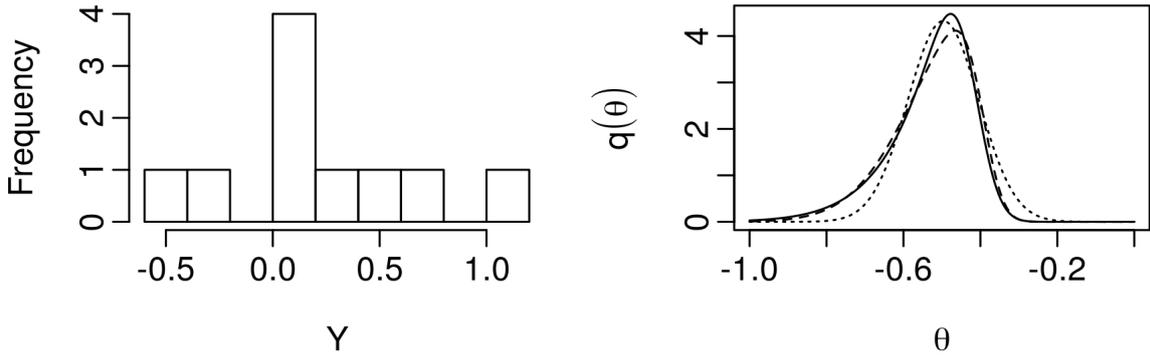


Figure 1: Left panel: Histogram of the simulated data. Right panel: The true posterior distribution $p(\theta \mid \mathbf{Y})$ (solid line), the variational posterior distribution when $q(\theta)$ is assumed to be skew-normal (dashed line), and the variational posterior distribution when $q(\theta)$ is assumed to be Gaussian (dotted line).

where $q(\theta)$ is a variational posterior distribution (in this case an approximation to the posterior distribution). If m is small (less than 30, say), the true conditional distribution of θ , when conditioned on the data \mathbf{Y} , is likely to be highly non-Gaussian. It is therefore reasonable to let $q(\theta)$ be a skew-normal distribution,

$$q(\theta) = SN(\xi, \omega^2, \alpha),$$

where ξ is the unknown location parameter, ω is the unknown scale parameter, and α is the unknown slant parameter of the variational posterior distribution. There have been attempts in the past to use skew-normal distributions in a classical VB framework (see, for example, Ormerod, 2011).

In SVB one approximates (1) via the Monte Carlo approximation

$$\mathcal{L} \approx \frac{1}{N} \sum_{j=1}^N \log \frac{p(\mathbf{Y} \mid \theta^{(j)})p(\theta^{(j)})}{q(\theta^{(j)})},$$

where $\theta^{(j)} \sim SN(\xi, \omega^2, \alpha)$. To obtain gradients with respect to ξ, ω , and α , these parameters need to be explicitly represented in the expression for \mathcal{L} . I therefore now draw on a connection between the distributions used by Tagle et al. (2020) and the reparametrisation trick of Kingma and Welling (2014) by expressing the Monte Carlo samples via the parameterisation of Azzalini and Capitanio (2013, Chapter 2),

$$\theta^{(j)} = \xi + \omega\gamma^{(j)}, \quad j = 1, \dots, N,$$

where, for $\alpha \equiv \delta/\sqrt{1-\delta^2}$, $U_0^{(j)} \sim N(0, 1)$, and $U_1^{(j)} \sim N(0, 1)$,

$$\gamma^{(j)} = \delta|U_1^{(j)}| + \sqrt{1-\delta^2}U_0^{(j)} \sim SN(0, 1, \alpha).$$

The variational distribution over θ is then found by solving the optimisation problem

$$(\hat{\xi}, \hat{\omega}, \hat{\alpha}) = \arg \max_{\xi, \omega, \alpha} \frac{1}{N} \sum_{i=1}^N \frac{p(\mathbf{Y} \mid \theta^{(i)})p(\theta^{(i)})}{q(\theta^{(i)})}, \quad (2)$$

where an expression for $q(\theta)$ can be found in Azzalini and Capitanio (2013, Chapter 2). By virtue of the reparameterisation, the optimisation problem (2) can now be solved quickly and efficiently using parallel computing and automatic differentiation libraries.

For illustration, let $m = 10$, $\lambda = 20$ (assumed known), and $\theta = -0.5$. Figure 1 in this discussion piece, left panel, shows a histogram of the simulated data, while the right panel shows the true posterior distribution over θ (solid line), the approximate posterior distribution when $q(\theta)$ is assumed to be skew-normal (dashed line), and the case where $q(\theta)$ is assumed to be Gaussian (dotted line). SVB was run with $N = 500$ on a graphics processing unit and required only a few seconds to converge. Clearly, the approximated skew-normal variational posterior is reasonably representative of the true posterior distribution. I note that this reparameterisation trick for skew-normal variational posterior distributions has been used recently in the context of copula models by Smith et al. (2020).

Some extensions would be needed to use SVB instead of a Gibbs sampler in Tagle et al. (2020)'s work since, for example, each $Z_{r,t}$ needs to be positive. Zammit-Mangion et al. (2019) used trans-Gaussian variational posterior distributions to achieve a positivity constraint in the context of deep spatial models, while one may consider other functional forms for the approximate posterior distribution, such as the gamma distribution. An interesting avenue of future work is the development of skew- t variational posterior distributions, which would be suitable for approximating posterior distributions with heavy tails. Here, one could proceed by optimising the lower bound for a selection of plausible integer-valued degrees of freedom; the optimal degrees of freedom would then be that which yields the largest variational lower bound.

Acknowledgements

AZ-M is supported by the Australian Research Council (ARC) Discovery Early Career Research Award, DE180100203, and the ARC Discovery Project, DP190100180.

References

- Azzalini, A. and Capitanio, A. (2013). *The Skew-Normal and Related Families*. Cambridge University Press, Cambridge, UK.
- Carvalho, D., Rocha, A., Gómez-Gesteira, M., and Santos, C. S. (2014). Comparison of re-analyzed, analyzed, satellite-retrieved and NWP modelled winds with buoy data along the Iberian Peninsula coast. *Remote Sensing of Environment*, 152:480–492.
- Castruccio, S. and Genton, M. G. (2016). Compressing an ensemble with statistical models: An algorithm for global 3D spatio-temporal temperature. *Technometrics*, 58:319–328.
- Cseke, B., Zammit-Mangion, A., Heskes, T., and Sanguinetti, G. (2016). Sparse approximate inference for spatio-temporal point process models. *Journal of the American Statistical Association*, 111:1746–1763.
- Kingma, D. P. and Welling, M. (2014). Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*.
- Majumder, M., Hofmann, H., and Cook, D. (2013). Validation of visual statistical inference, applied to linear models. *Journal of the American Statistical Association*, 108:942–956.
- McNorton, J., Bousserez, N., Agustí-Panareda, A., Balsamo, G., Choulga, M., Dawson, A., Engelen, R., Kipling, Z., and Lang, S. (2020). Representing model uncertainty for global

- atmospheric CO₂ flux inversions using ECMWF-IFS-46R1. *Geoscientific Model Development*, 13:2297–2313.
- Ormerod, J. (2011). Skew-normal variational approximations for Bayesian inference. Unpublished. URL: <http://talus.maths.usyd.edu.au/u/jormerod/JT0papers/snva.pdf>.
- Schmidt, A. M., Guttorp, P., and O’Hagan, A. (2011). Considering covariates in the covariance structure of spatial processes. *Environmetrics*, 22:487–500.
- Schuh, A. E., Jacobson, A. R., Basu, S., Weir, B., Baker, D., Bowman, K., Chevallier, F., Crowell, S., Davis, K. J., Deng, F., et al. (2019). Quantifying the impact of atmospheric transport uncertainty on CO₂ surface flux estimates. *Global Biogeochemical Cycles*, 33:484–500.
- Sengupta, A. and Cressie, N. (2013). Hierarchical statistical modeling of big spatial datasets using the exponential family of distributions. *Spatial Statistics*, 4:14–44.
- Smith, M. S., Loaiza-Maya, R., and Nott, D. J. (2020). High-dimensional copula variational approximation through transformation. *Journal of Computational and Graphical Statistics*. in press.
- Tagle, F., Genton, M. G., Yip, A., Mostamandi, S., Stenchikov, G., and Castruccio, S. (2020). A high-resolution bilevel skew-*t* stochastic generator for assessing Saudi Arabia’s wind energy resources. *Environmetrics*, e2628.
- Zammit-Mangion, A., Cressie, N., and Ganesan, A. L. (2016). Non-Gaussian bivariate modelling with application to atmospheric trace-gas inversion. *Spatial Statistics*, 18:194–220.
- Zammit-Mangion, A., Ng, T. L. J., Vu, Q., and Filippone, M. (2019). Deep compositional spatial models. *arXiv preprint arXiv:1906.02840*.
- Zammit-Mangion, A. and Rougier, J. (2020). Multi-scale process modelling and distributed computation for spatial data. *Statistics and Computing*. in press.
- Zammit-Mangion, A., Rougier, J., Schön, N., Lindgren, F., and Bamber, J. (2015). Multivariate spatio-temporal modelling for assessing Antarctica’s present-day contribution to sea-level rise. *Environmetrics*, 26:159–177.
- Zammit-Mangion, A. and Wikle, C. K. (2020). Deep integro-difference equation models for spatio-temporal forecasting. *Spatial Statistics*, 37:100408.
- Zeng, Z., Ziegler, A. D., Searchinger, T., Yang, L., Chen, A., Ju, K., Piao, S., Li, L. Z., Ciais, P., Chen, D., Liu, J., Azorin-Molina, C., Chappell, A., Medvigy, D., and Wood, E. F. (2019). A reversal in global terrestrial stilling and its implications for wind energy production. *Nature Climate Change*, 9:979–985.