ROYAL STATISTICAL SOCIETY
DATA | EVIDENCE | DECISIONS

Journal of the Royal Statistical Society
Statistics in Society
Series A

J. R. Statist. Soc. A (2019)
182, *Part* 2, *pp.* 379–387

Visualizing spatiotemporal models with virtual reality: from fully immersive environments to applications in stereoscopic view

Stefano Castruccio

University of Notre Dame, USA

and Marc G. Genton and Ying Sun

King Abdullah University of Science and Technology, Thuwal, Saudi Arabia

[*Read before* The Royal Statistical Society *at a meeting on 'Data visualization' at the Society's 2018 annual conference in Cardiff on Wednesday, September 5th, 2018, the President,* Professor D. J. Spiegelhalter, *in the Chair*]

Summary. Recent advances in computing hardware and software present an unprecedented opportunity for statisticians who work with data indexed in space and time to visualize, explore and assess the structure of the data and to improve resulting statistical models. We present results of a 3-year collaboration with a team of visualization experts on the use of stereoscopic view and virtual reality (VR) to visualize spatiotemporal data with animations on non-trivial manifolds. We first present our experience with fully immersive VR with motion tracking devices that enable users to explore global three-dimensional time–temperature fields on a spherical shell interactively. We then introduce a suite of applications with VR mode, freely available for smartphones, to port a visualization experience to any interested people. We also discuss recent work with head-mounted devices such as a VR headset with motion tracking sensors.

Keywords: Data visualization; Global models; Spatiotemporal statistics; Virtual reality

1. Introduction

Data that are indexed in space and time are prevalent in many fields of science, and space–time statistics focuses on understanding their dependence structure (and changes thereof), and quantifying the associated uncertainty. Although the goal of each scientific analysis is topic dependent, the usual outcome of a space–time model is an interpolated map, as well as a measure of its uncertainty.

The study of theoretical properties of optimal interpolation, or kriging, culminating with Stein's seminal book (Stein, 1999), advocates the use of models with structured local behaviour at the expense of global behaviour, which for kriging is (asymptotically) less relevant for most commonly used models. This and many other recent contributions have prompted the development of graphical tools for the study of local behaviour. The most popular of these tools are the covariogram and the variogram, which are routinely used to discern local dependence under a stationarity assumption (Cressie, 1993). An equivalent but more visually straightforward tool to study local behaviour is the estimated spectra or periodogram. Indeed, a periodogram enables

Address for correspondence: Stefano Castruccio, Department of Applied and Computational Mathematics and Statistics, University of Notre Dame, 153 Hurley Hall, South Bend, Notre Dame, IN 46556, USA.
E-mail: scastruc@nd.edu

 0964–1998/19/182379

us to assess the smoothness of the process by observing and quantifying how and to what extent high frequencies are decaying (Stein, 1999).

Although statistical models emphasizing local behaviour provide optimal maps, many emerging applications require the assessment of global patterns and their relative changes with respect to various quantities of interest. For example, in the context of climate models, quantifying relative changes in temperature fields, or other meteorological variables of physical interest, with respect to different input is necessary to assess climate sensitivity. Chang *et al.* (2014, 2016) used principal component analysis (PCA) (which is known as empirical orthogonal functions in the climate community) to produce maps that allowed the main modes of variability to change across the input space. The interest lies in developing a statistical model whose fitted values visually resemble the data. A related and very active area of research is on climate model compression (Guinness and Hammerling, 2018; Castruccio and Guinness, 2017), for which the goal is to reproduce maps that are as similar as possible to the original output, to the extent that practitioners cannot distinguish between them, both visually and in their scientific investigations (Baker *et al.*, 2014).

Several works have focused on the development of effective visualization methods that provide information about point value and uncertainty in a single image. Because it is challenging for our visual memory to assimilate multiple sources of information at the same time, an appropriate choice of visual tool can facilitate the comprehension of complex and simultaneous information indexed in space (Few, 2009), and there is increasing interest in embedding uncertainty in maps (Bonneau *et al.*, 2014). Whereas a uniformly accepted solution is yet to be agreed (MacEachren *et al.*, 2005), several solutions have been devised, especially in the area of census data (Lucchesi and Wikle, 2017), from choropleth maps (Tufte, 1986; MacEachren *et al.*, 2005; Retchless and Brewer, 2015) to map pixellation (Ewans, 1997; MacEachren *et al.*, 2005) and glyph rotation (Wittenbrink *et al.*, 1996; MacEachren *et al.*, 2005).

Very recent work has focused on the use of animated visuals to improve perception of uncertainty and to perform diagnostics on a spatiotemporal model. Genton *et al.* (2015) presented the first work from a collaboration with the Visualization Core Laboratory at King Abdullah University of Science and Technology (KAUST) on diagnostics and visualization of statistical models. Genton *et al.* (2015) discussed limitations of current approaches with static images and the value of movies in enhancing the perception and quantification of uncertainty. Although the work presented an initial discussion on dynamic visualization, they did not provide any guidance on quantification of the visual agreement between images, and they only hinted at the use of visualization on non-trivial geometries.

As part of this on-going collaboration, this work aims at presenting a new framework to assess the global behaviour of a statistical model, focusing on static and animated visuals, as well as virtual reality (VR) and leveraging on diagnostic metrics from image and video processing. This work presents a fully immersive VR environment that is available at the KAUST, and introduces a set of applications (apps) available for iPhone® and Android®, all developed in VR mode and stereoscopic view, for diagnostics of several statistical models, and we briefly describe our experience in presenting these at conferences.

The remainder of our paper is organized as follows. In Section 2, we discuss the use of static and dynamic visuals and their limitations in the perception of three-dimensionality. In Section 3, we introduce the use of VR for diagnostics of non-trivial manifolds such as spherical shells with a fully immersive facility, as well as with apps with VR view. Section 4 includes a discussion on perspectives, opportunities and challenges in bringing advanced visualization media to the statistics community.

2. Diagnostics for assessing global behaviour of spatiotemporal models

In this section, we focus on temperature output from a global climate model; the specifics of the climate model are in the on-line supplementary material. We first consider static diagnostics in Section 2.1, i.e. we look at how to compare the global behaviour of models in static images and at what indices are appropriate for quantifying similarities. In Section 2.2, we introduce dynamic diagnostics through movies and discuss quantification of similarity in this context. In the supplementary material, we present some standard diagnostics of local behaviour (e.g. variograms and periodograms) that also highlight non-stationary structures.

2.1. Static diagnostics for spatiotemporal models

Goodness of fit for a statistical model can be intuitively assessed by visualizing and comparing an original field and the fitted values of the model. Fig. 1 shows an example of a field with 10 and 30 PCA components. It is apparent that an increased number of components considerably improves the resemblance of the reconstructed image to the original data set.

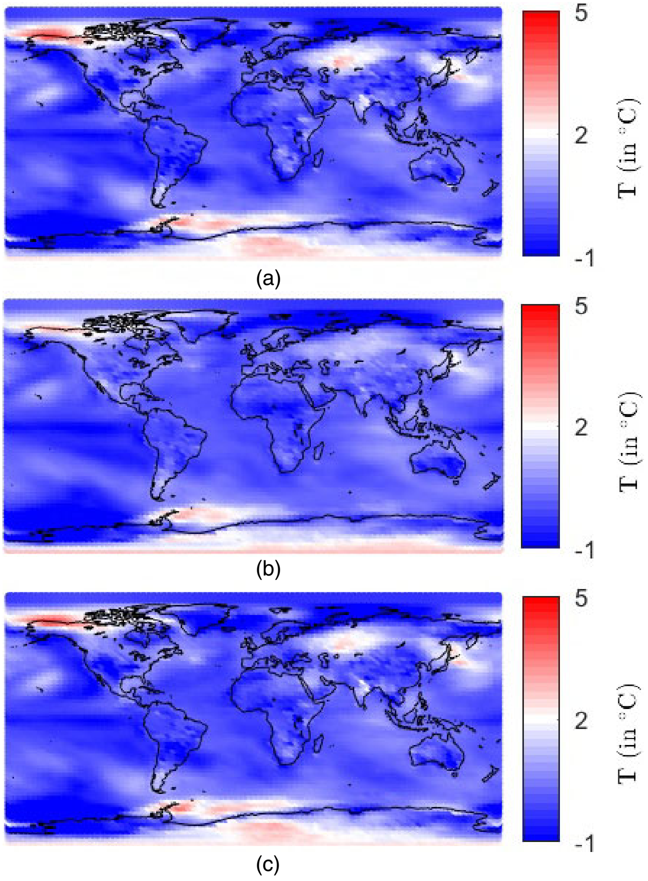

Fig. 1. Analysis of the 2017 annual temperature anomaly (from 2006) with different low rank approaches yielding different perceptions of visual similarity: (a) the original data set (see the on-line supplementary material for details); (b) PCA with 10 components (structural similarity index SSIM = 0.54 (Wang *et al.*, 2004)); (c) PCA with 30 components (SSIM = 0.83)

The visual similarity between the original and reconstructed fields can be quantified through several indices, including for example the root-mean-squared error, the correlation coefficient (Baker *et al.*, 2014) or the peak signal-to-noise ratio. However, a large body of literature on signal processing and visual perception (Girod, 1993; Wang and Bovik, 2009) has argued against these metrics as they do not correspond to the natural visual perception of an image. A more appropriate measure for two aligned image signals $\mathbf{x}$ and $\mathbf{y}$ is the structural similarity index SSIM (Wang *et al.*, 2004):

$$\mathrm{SSIM}(\mathbf{x},\mathbf{y})=\frac{(2\mu_{\mathbf{x}}\mu_{\mathbf{y}}+C_1)(2\sigma_{\mathbf{xy}}+C_2)}{(\mu_{\mathbf{x}}^2+\mu_{\mathbf{y}}^2+C_1)(\sigma_{\mathbf{x}}^2+\sigma_{\mathbf{y}}^2+C_2)}, \tag{1}$$

where $\mu_{\mathbf{x}}$ and $\mu_{\mathbf{y}}$ are the sample means across $\mathbf{x}$ and $\mathbf{y}$ respectively, $\sigma_{\mathbf{x}}^2$ and $\sigma_{\mathbf{y}}^2$ are the sample variances across $\mathbf{x}$ and $\mathbf{y}$ respectively, $\sigma_{\mathbf{xy}}$ is the sample covariance and $C_1=0.01^2$ and $C_2=0.03^2$. SSIM aims at measuring the perceived degradation of quality of the reconstructed field in terms of loss of luminance, contrast and structure. In Fig. 1, the improvement of the model with a higher number of PCA components is also highlighted by a noticeable improvement in SSIM.

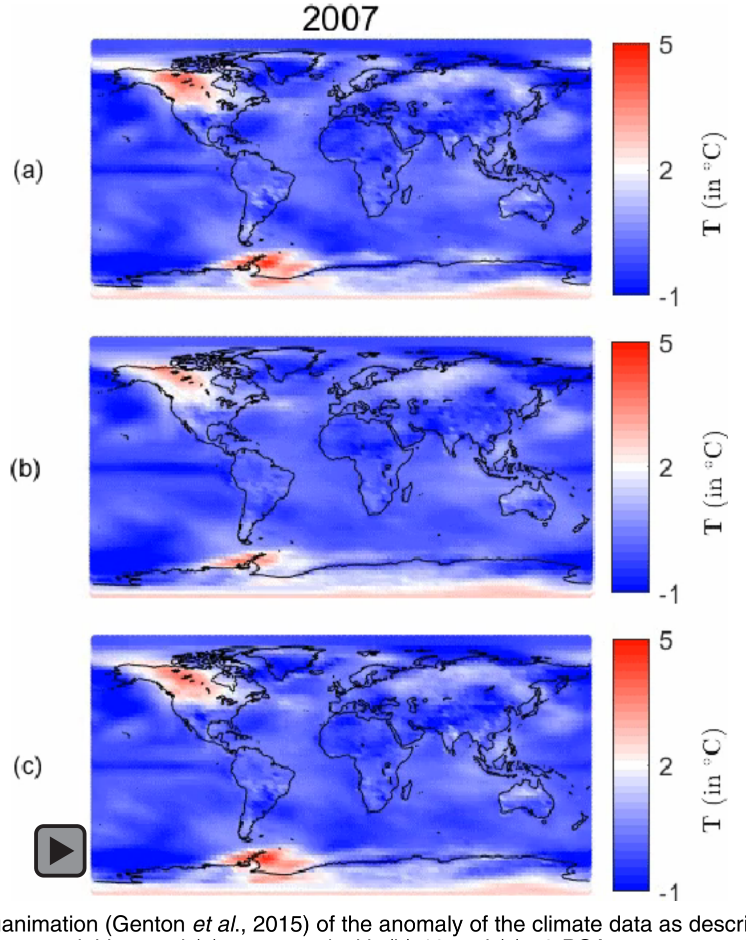

Fig. 2. Visuanimation (Genton *et al.*, 2015) of the anomaly of the climate data as described in the on-line supplementary material in panel (a), compared with (b) 10 and (c) 50 PCA components from year 2006 to 2100 (SSIM as indicated in equation (1) is 0.57 (0.12) for 10 PCA components, whereas it is 0.82 (0.06) for 50 PCA components)

2.2. Dynamic diagnostics for spatiotemporal models

When the data are indexed in space and time, a suitable comparison needs to account for the dynamic nature of the data set, and hence to compare movies instead of figures. Genton *et al.* (2015) introduced the notion of embedding movies in a manuscript ('*visuanimation*'). In the visuanimation in Fig. 2 we compare

(a) the temperature anomaly for all years between 2006 and 2100
(b) with a low rank model with 10 PCA components
(c) and with 50 PCA components.

Although a single image could be efficiently reproduced with only a small basis set (30 bases yield a structurally similar image, as shown in Fig. 1), a dynamic low rank description requires more structure to achieve visually comparable results across multiple times. In terms of quantifying the visual similarities of movies, there is no uniform agreement in the literature on indices that account for structural perception in space and time simultaneously. We therefore compute SSIM for every time point and then communicate the mean and standard deviation in the caption of Fig. 2, which is considerably higher (and less variable) with a higher number of PCA components, as expected.

Dynamic visualization marks an improvement in the assessment of spatiotemporal models. Whereas an ordinary flat screen delivers satisfactory visuals for one- or two-dimensional Euclidean spaces, its use on a spherical domain, or more generally a manifold, is hampered by the lack of three-dimensional perception. It is possible to rotate the view dynamically (see movie 5 in Genton *et al.* (2015)), but this solution mitigates only the intrinsic limitation of the screen, i.e. the lack of a stereoscopic view. A better solution would allow diagnostics on complex geometries with advanced visualization software and hardware.

3. Beyond the Euclidean space: diagnostics for three-dimensional spatiotemporal models on spherical shells in virtual reality environments

In this section, we propose two solutions for three-dimensional visualization, both devised as part of a three-year collaboration with the KAUST Visualization Core Laboratory. In Section 3.1, we discuss an experience with a VR environment that allows fully immersive and interactive visualization; in Section 3.2, we introduce the use of apps with stereoscopic view to allow for portable VR visualization.

3.1. Fully immersive virtual reality: the CORNEA environment

An initial method to compare global properties is to display both the original data set and a conditional simulation from the statistical model within a fully immersive VR environment. CORNEA is a three-dimensional environment enclosed by six walls (DeFanti *et al.*, 2011) and a 200-million pixel stereoscopic display, which enables a complete VR experience. The headset is equipped with a wireless motion capture device to allow each user to explore the data within the environment interactively. In a recent workshop, a model for three-dimensional temperature was fitted and simulations from a statistical model (or *emulator*; full model details are given in Castruccio and Genton (2016, 2018)) were compared with the original climate model. An interactive display was shown. Each participant could explore multiple features of the data and how visually similar the climate model was compared with the statistical simulation. Video footage of the workshop can be found at `vimeo.com/109146573`.

The VR environment allows the observer to explore surfaces whose geometry is considerably more sophisticated than the spherical shell of the planet. An additional showcase of Mount Pinatubo's 1991 eruption (the second largest volcanic eruption of the 20th century) allowed us to visualize the extent of the eruption and its diffusion in the equatorial belt of the planet (see `vimeo.com/103123075`). Current work is focusing on producing a statistical model for the eruption that would allow a visually similar outcome to assess sensitivity with respect to meaningful physical parameters.

3.2. Portable virtual reality: smartphone applications

An immersive VR environment such as CORNEA at the KAUST provides a clear step forward in terms of the visualization of statistical output. However, it is obviously limited to the physical

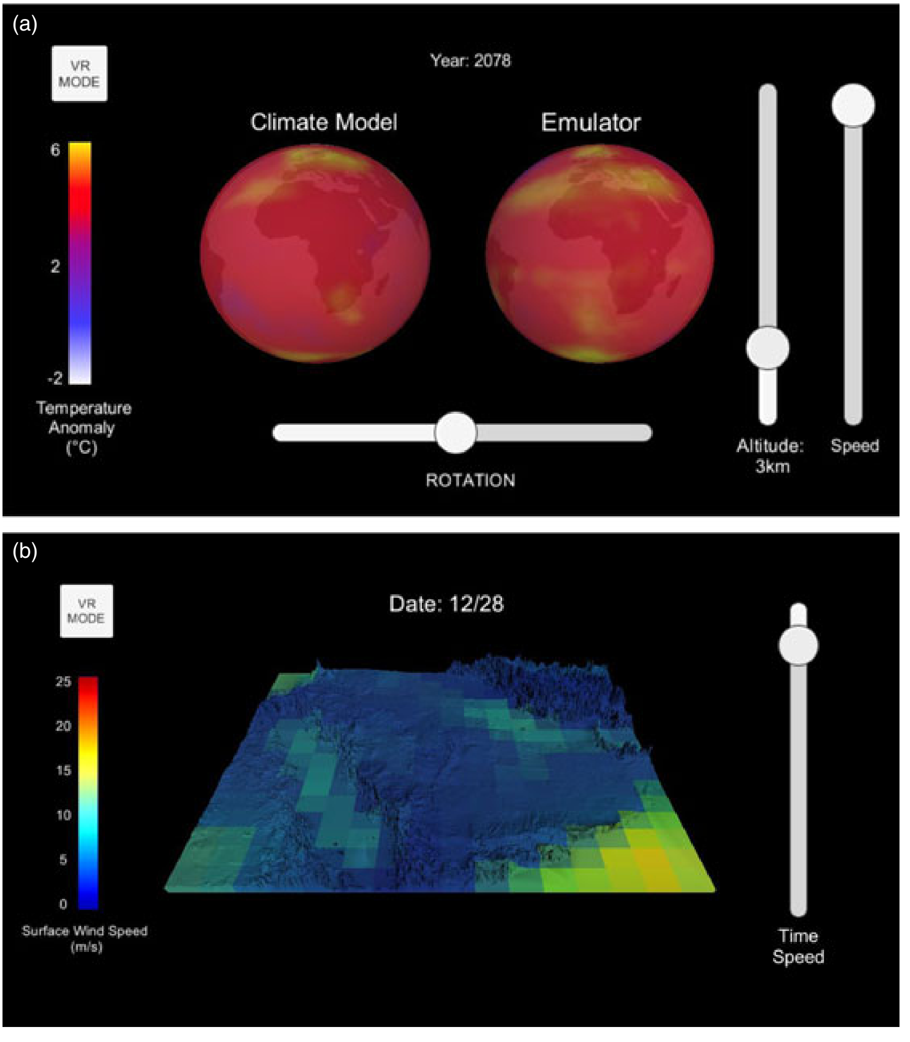

Fig. 3. Snapshot of two apps developed for visualizing and comparing climate model output and output from a statistical model (emulator): (a) the 'statistical emulator' app reproduces annual three-dimensional temperature from 1850 to 2100, whereas (b) 'Saudi surface wind' focuses on daily winds over the Arabian peninsula from 1920 to 2100

space where it is installed, and it lacks the flexibility to port the experience to a wider audience. The development of apps for smartphones is a next step in terms of interactive visualization. As part of the collaborative research project with the KAUST visualization team, several apps that are freely available for both iPhones (in the app store) and Android (on Google Play) were created.

The app 'statistical emulator' can be found and downloaded by searching from these keywords on the search engine for apps. It focuses on three-dimensional global temperature and allows the observer to compare the original statistical model interactively with the emulator on a smartphone. It also provides a VR mode in stereoscopic display with VR headset (currently on the market at prices below $10). When not in VR mode, the data can be explored on the phone by controlling the Earth's rotation on the bottom bar, whereas the altitude of the temperature fields and the time speed can be controlled with the two sidebars (see Fig. 3(a) for a snapshot). A similar app, called 'global surface wind' displays the annual winds between 1920 and 2100 across the globe according to a climate model and another statistical model (full details on the model are provided by Jeong *et al*. (2018)). This app uses a similar visualization scheme to that in the temperature case, but it lacks the three-dimensional component as the model focuses on surface wind speed. An updated version with directional wind fields will be available in the near future.

Although these apps focus on the use of three dimensionality to improve visualization of global data, the VR view also allows realistic perception of altitude with the stereoscopic view. This information can then be used to guide the analysis. 'Saudi surface wind' is a third app that was developed to visualize daily wind in the Arabian peninsula from 1920 to 2100 (see Fig. 3(b) for a snapshot). In this reduced domain, the Earth's curvature is not shown, and the three dimensionality is used to improve perception of the mountain ranges and to cross-reference them with the wind speed. It is apparent how mountain ranges experience higher wind speeds than does flat land throughout the annual cycles. The app currently focuses on visualization of the original climate data. An initial emulator for the region has also been developed (see the details in Tagle *et al*. (2017)) and will be added to the updated version in the near future.

We have presented these apps at several conferences, including two topic sessions on visualizing data at the American Statistical Association's Joint Statistical Meetings in 2016 and 2017. Simple instructions provided a few minutes before the end of the presentation allowed the audience to download the app. We provided inexpensive VR glasses so that audience members could test the VR experience first hand. The experimental sessions with apps were overall received with enthusiasm from the audience. Without an explicit labelling of model output and emulator, there was no perceived difference in the two data sets. There was also an overall consensus on the need to improve interactivity by allowing the user to provide input on the phone when in VR mode.

4. Conclusion and discussion

Traditional applications in spatiotemporal statistics focus on kriging and in models with a structured local dependence. New areas of enquiry are emerging that require analysis of global behaviour. Performing diagnostics on space–time-dependent output with movies, which is a task that allows the user to perceive uncertainty as a time varying visual (Genton *et al*., 2015), avoids the need to embed point estimates and their associated uncertainty in the same static picture.

We present two solutions to visualize data sets with non-trivial geometries, both using stereoscopic view and VR environments. CORNEA allows the user to explore the data in a fully immersive environment and to interact with them by means of a motion tracking device inside the environment. The data can be explored by moving in the physical space. An app with VR mode is considerably more accessible and can and has been used to showcase results without any geographical constraint. Such flexibility, however, comes at the expense of the loss of interactivity with the data, the absence of a motion tracking device (the accelerometer in a smartphone does not capture motion in the physical space) and lower image quality and refreshing frequency.

Although this project focused on visualization on spherical surfaces, spherical shells and regions with mountain ranges, we have also developed apps for visualizing three-dimensional brain data to assess which areas are more active when a subject is performing a given task (Castruccio *et al.*, 2018); see the 'interactive brain activation for fMRI data' in the app store.

Current research is focusing on addressing the main remarks of the audience during the conferences and introducing interactivity on smartphone apps via a wireless control pad, thereby allowing the user to change views of the data set while also choosing the variable of interest. For example, wind power density, which is a useful quantity for energy assessment, is expressed through the cube of the wind speed. A compromise solution between CORNEA and a VR app is a head-mounted device with motion tracking sensors that enables spatial interaction with the data without relying on a six-sided wall environment. We are currently exploring this application in the context of the occurrence of tropical cyclones in the Arabic peninsula. The project is at its preliminary stages and the main limitations arise from the lack of established software and technical expertise to develop new case-studies for this platform. A more sophisticated solution is to embed the VR environment in the real world, and to develop an augmented reality platform to enhance the interaction and the experience in the visualization. Such a development is, however, to be explored once the technology is available to the general public and, more specifically, to statisticians.

Acknowledgements

This publication is based on work supported by the KAUST Office of Sponsored Research under award OSR-2015-CRG4-2640. This work was conducted using resources and services at the Visualization Core Laboratory at KAUST: we thank Dr Madhusudhanan Srinivasan and his team for their support.

References

Baker, A. H., Xu, H., Dennis, J. M., Levy, M. N., Nychka, D., Mickelson, S. A., Edwards, J., Vertenstein, M. and Wegener, A. (2014) A methodology for evaluating the impact of data compression on climate simulation data. In *Proc. 23rd Int. Symp. High-performance Parallel and Distributed Computing*, pp. 203–214. New York: Association for Computing Machinery.

Bonneau, G., Hege, H., Johnson, C., Oliveira, M., Potter, K., Rheingans, P. and Schultz, T. (2014) Overview and state-of-the-art of uncertainty visualization. In *Scientific Visualization*, pp. 3–27. New York: Springer.

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.

Castruccio, S. and Genton, M. G. (2018) Principles for statistical inference on big spatio-temporal data from climate models. *Statist. Probab. Lett.*, **136**, 92–96.

Castruccio, S. and Guinness, J. (2017) An evolutionary spectrum approach to incorporate large-scale geographical descriptors on global processes. *Appl. Statist.*, **66**, 329–344.

Castruccio, S., Ombao, H. and Genton, M. G. (2018) A scalable multi-resolution spatio-temporal model for brain activation and connectivity in fMRI data. *Biometrics*, **74**, 823–833.

Chang, W., Haran, M., Applegate, P. and Pollard, D. (2016) Calibrating an ice sheet model using high-dimensional binary spatial data. *J. Am. Statist. Ass.*, **111**, 57–72.

Chang, W., Haran, M., Olson, R. and Keller, K. (2014) Fast dimension-reduced climate model calibration and the effect of data aggregation. *Ann. Appl. Statist.*, **8**, 649–673.

Cressie, N. (1993) *Statistics for Spatial Data*. New York: Wiley Interscience.

DeFanti, T. A., Acevedo, D., Ainsworth, R. A., Brown, M. D., Cutchin, S., Dawe, G., Doerr, K.-U., Johnson, A., Knox, C., Kooima, R., Kuester, F., Leigh, J., Long, L., Otto, P., Petrovic, V., Ponto, K., Prudhomme, A.,

Rao, R., Renambot, L., Sandin, D. J., Schulze, J. P., Smarr, L., Srinivasan, M., Weber, P. and Wickham, G. (2011) The future of the CAVE. *Cent. Eur. J. Engng*, **1**, 16–37.
Ewans, B. (1997) Dynamic display of spatial data-reliability: does it benefit the map user? *Comput. Geosci.*, **23**, 409–422.
Few, S. (2009) *Now You See It: Simple Visualization Techniques for Quantitative Analysis*. Oaklands: Analytics Press.
Genton, M. G., Castruccio, S., Crippa, P., Dutta, S., Huser, R., Sun, Y. and Vettori, S. (2015) Visuanimation in statistics. *Stat*, **4**, 81–96.
Girod, B. (1993) Whats wrong with mean-squared error? In *Digital Images and Human Vision* (ed. A. B. Watson), pp. 207–220. Cambridge: MIT Press.
Guinness, J. and Hammerling, D. (2018) Compression and conditional emulation of climate model output. *J. Am. Statist. Ass.*, **113**, 56–67.
Jeong, J., Castruccio, S., Crippa, P. and Genton, M. G. (2018) Reducing storage of global wind ensembles with stochastic generators. *Ann. Appl. Statist.*, **12**, 490–509.
Lucchesi, L. and Wikle, C. (2017) Visualizing uncertainty in areal data with bivariate choropleth maps, map pixelation, and glyph rotation. *Stat*, **6**, 292–302.
MacEachren, A., Robinson, A., Hopper, S., Gardner, S., Murray, R., Gahegan, M. and Hetzler, E. (2005) Visualizing geospatial information uncertainty: what we know and what we need to know. *Cartg. Geog. Inform. Sci.*, **32**, 139–160.
Retchless, D. and Brewer, C. (2015) Guidance for representing uncertainty on global temperature change maps. *Int. J. Climtol.*, **36**, 1143–1159.
Stein, M. (1999) *Statistics for Spatial Data: Some Theory of Kriging*. New York: Springer.
Tagle, F., Castruccio, S., Crippa, P. and Genton, M. G. (2017) Assessing potential wind energy resources in Saudi Arabia with a skew-t distribution. *Preprint arXiv.org/abs/1703.04312*.
Tufte, E. (1986) *The Visual Display of Quantitative Information*. Cheshire: Graphics Press.
Wang, Z. and Bovik, A. C. (2009) Mean squared error: love it or leave it?—a new look at fidelity measures. *IEEE Signl Process. Mag.*, **26**, 98–117.
Wang, Z., Bovik, A., Sheikh, H. and Simoncelli, E. (2004) Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Im. Process.*, **13**, 600–612.
Wittenbrink, C., Pang, A. and Lodha, S. (1996) Glyphs for visualizing uncertainty in vector fields. *IEEE Trans. Visualizn Comput. Graph.*, **2**, 266–279.

Supporting information
Additional ‘supporting information’ may be found in the on-line version of this article.