PSL Young researcher Starting Grant 2024

Detailed research project

PASDOS
Physics-Aware Statistics for Data On Surfaces

Mike PEREIRA

Department of Geosciences and Geoengineering, Mines Paris - PSL University

Abstract

In the context of geostatistical modeling of spatio-temporal data, it is usual to assume that these data can be represented by a Gaussian random field defined over a spatial domain. This approach allows not only to model the uncertainty one might have on the data, but also to perform such tasks as simulations and predictions while accounting for the spatio-temporal correlations observed in the data. This motivates the need to develop models for spatio-temporal Gaussian random fields flexible enough to represent complex patterns of correlations in the data, but simple enough to allow for numerically efficient algorithms for their inference, simulation and prediction.

Building on this idea, the aim of this project is to propose new Gaussian process models to describe spatio-temporal data in the particular case where these lie on a surface. Such models could be used in various domains of applications, for instance to model brain activity across the cortical surface in neuroimaging, or to model measurements across the surface of the Earth in climate and environmental sciences.

We propose to define the Gaussian random fields used to model the data as solutions of some problem-specific stochastic partial differential equations (SPDE) defined on a surface. This "SPDE" approach to Gaussian processes has been extensively used to model spatio-temporal data on Euclidean domains (Lindgren et al., 2022), and extended to model spatial data on surfaces (see eg. Mejia et al. (2020)) and more generally on Riemannian manifolds (Lang and Pereira, 2023; Pereira et al., 2022). Our aim is now to further generalize this approach to the spatio-temporal setting by considering SPDEs which for instance include diffusion and/or advection terms in order to capture the spatio-temporal correlations observed in the data. Our challenges will be on the one hand, to properly define and solve such SPDEs on surfaces, and on the other hand to propose efficient algorithms for the inference of the resulting statistical model, and for simulations and predictions based on the model. To answer the latter question, we plan on building on a link between discretization of SPDEs and graph signals (Pereira, 2019) to propose novel and efficient graph neural network architectures to tackle these tasks.

PSL UNIVERSITÉ PARIS

Contents

Note: *This document contains* ***links to animated figures colored in red*** . *All these animations are collected at: https://mike-pereira.github.io/PSL-SG/.*

1. Introduction

Many aspects of climate and environmental modeling are related to complex physical systems. For instance, forecasting the behavior of such systems is paramount to questions related to risk and hazard assessment. In response, physics-based models able to simulate such behaviors have been developed, and are routinely used in practice (eg. to predict the evolution global sea levels or temperatures in the future). However, due to the complexity of the physical processes involved, these models are often very heavy to run. Generating large numbers of forecasts for sensitivity analysis or uncertainty quantification purposes then becomes infeasible.

A first answer to these challenges comes from Geostatistics which is a branch of statistics that focuses on modeling data correlated in space and time. Within the geostatistical paradigm, the observed variable is modeled as (a realization of) a space-time Gaussian random field (GRF), so that the mere characterization of its mean and covariance functions suffices to fully describe its statistical properties. In particular, these two functions are chosen to mimic the spatiotemporal variability and structure observed in the data Wackernagel (2003).

This probabilistic framework has many advantages. On the one hand, it allows to efficiently perform simulations of random fields with the same spatiotemporal structure as the one observed in the data, and predictions of at unobserved locations. On the other hand, uncertainties can be quantified both on the variable behavior at unobserved locations (using so-called conditional simulations) or on the model parameters (through Bayesian approaches) (Chiles and Delfiner, 2009) in a closed form.

However, when applying GRF-based models to spatiotemporal data arising from climate and environmental applications, two main obstacles arise. On the one hand, the complex nature of the physical processes involved translates into intricate spatiotemporal correlation in the data. For instance, when modeling the dispersion of pollutants in the atmosphere, transport phenomena due to winds and currents will clearly affect the spatiotemporal correlations observed in the data (see eg. Figure 1.1). Moreover, when considering global datasets, the spatial domain on which the data lie is not "flat", but rather a sphere representing our planet (see eg. Rayner et al. (2020)). Then, using the GRF-based approach would require specifying a flexible family of covariance models capable of handling the geometry of the domain and the complex correlation of the data. On the other hand, inferring the parameters of a GRF model is known to be a bottleneck as this operation usually scales as $\mathcal{O}(N^3)$ for N observations when using naive approaches (eg. maximum likelihood).

These problems are not arising only in climate and environmental applications, but more generally in any application that requires handling data related to complex physical processes across surfaces. Hence the main motivation of this project: proposing a framework for spatiotemporal GRFs capable of handling complex patterns of correlations observed in data lying on arbitrary surfaces. In particular, as in our setting the data in question are related to physical processes, and our framework aims at being "physics-aware", meaning that the construction of the GRFs used in practice incorporates some physical knowledge of the data-generating process (eg. transport phenomena). Finally, our goal is to propose numerically efficient algorithms to solve the computational problems linked to the inference of GRF parameters, and to uncertainty quantification.

2. Context and state-of-the art

When it comes to defining and building spatiotemporal GRFs to model data, there exists two main approaches: either through the definition of valid covariance functions (which are then "fitted" on the data), or through dynamical models describing the evolution in space and time of the GRFs. Let us review the principle and limits of both approaches.

2.1. The covariance-based approach to GRFs

Since we consider GRFs, tasks such as sampling from (un)-conditional distributions, predictions (through conditional expectations),and likelihood-based inference can all be performed by solving linear

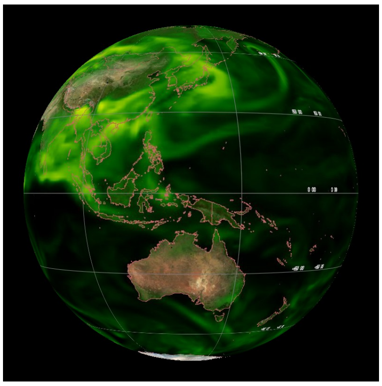

Figure 1.1: Simulation of the presence of sulfate in the Earth atmosphere (Source: NASA Global Modeling and Assimilation Office)

systems or adequately factorizing covariance matrices of the field (Wackernagel, 2003). Hence, the most straightforward (and classical) approach to spatiotemporal geostatistical modeling consists in fitting valid space-time covariance functions on the data, so that these covariance matrices may be built. Consequently, extensive literature on which covariance functions may be used to model spatiotemporal data, even with complex correlation patterns, is available (see Chen et al. (2021); Porcu et al. (2018, 2021) for recent reviews).

Nonetheless, the covariance-based approach has two main drawbacks. First, the matrix factorizations required in sampling, prediction and inference tasks have a complexity that scales as the cube of the number of observations and/or target points, thus making them unfeasible when this last number is large. To circumvent this, simplifying assumptions on the covariance model must be made, such as the separability of space and time dependencies or stationarity. This in turn may result in a lack of realism of the model. Secondly, since they rely on Euclidean or arc-length distances, most of the covariance models available in the literature are restricted to the setting where the spatial domain is either Euclidean or the sphere . Hence, they hardly generalize to other surfaces.

2.2. The dynamic approach to GRFs

As foretold by its name, this approach relies on models of the dynamic evolution of the GRF in time and space. These models take the form of stochastic partial differential equations (SPDE), the solutions of which are GRFs. This "SPDE approach" to GRF modeling has been popularized by Lindgren et al. (2011), and builds on a result from Whittle (1954) which states that isotropic GRFs Z on $\mathbb{R}^d$ ($d \in \mathbb{N}$) with a Matérn covariance function are stationary solutions of the SPDE given by

$$(\kappa^2 - \Delta)^{\alpha/2} Z = \tau \mathcal{W}, \tag{1}$$

where $\kappa > 0$, $\alpha > d/2$, $\tau > 0$, and $(\kappa^2 - \Delta)^{\alpha/2}$ is a pseudo-differential operator (defined as $(\kappa^2 - \Delta)^{\alpha/2}[\cdot] = \mathcal{F}^{-1}[\boldsymbol{w} \mapsto (\kappa^2 + \|\boldsymbol{w}\|^2)^{\alpha/2}\mathcal{F}[\cdot](\boldsymbol{w})]$), and $\mathcal{W}$ is a Gaussian white noise on $\mathbb{R}^d$. Solving numerically this SPDE using stochastic finite elements allows to directly obtain an expression for the precision matrix (i.e. the inverse of the covariance matrix) of a GRF with Matérn covariance. Then, at the price of a minor approximation (called mass lumping), these expressions yield a Gaussian Markov random field representation of the GRF, characterized by a sparse precision matrix (Lindgren et al., 2022; Rue and Held, 2005). This in turn results in significant computational gains since sparse matrix algorithms can be used to deal with the matrix factorizations and linear system solving involved when performing sampling, prediction and inference (Davis, 2006; Lindgren et al., 2011).

The SPDE approach has been extensively used to model spatial data on Euclidean domains (see Lindgren et al. (2022) for a recent review), and extended to model spatial data on surfaces (see eg. Bonito et al. (2024); Coveney et al. (2019); Mejia et al. (2020)) and more generally on Riemannian manifolds (see eg. Herrmann et al. (2020); Lang and Pereira (2023)) by replacing the Laplace operator $-\Delta$ in SPDE (1) by a Laplace–Beltrami operator.

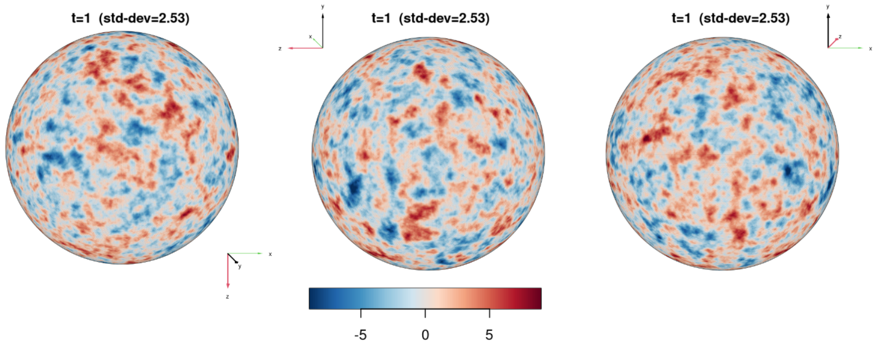

Figure 2.1: Simulation a spatiotemporal diffusion SPDE on the sphere represented from three different viewpoints on the surface (t represents the time step, std-dev the standard deviation of the field value across the surface).

Extensions of the SPDE approach to the spatiotemporal setting have also been proposed. Cameletti et al. (2013) propose an approach where the spatial SPDE is coupled with an AR(1) process in time, thus yielding a separable model. Non-separable models based on a direct generalization of SPDE (1) have been proposed by Bakka et al. (2020) and Rayner et al. (2020), who consider the solutions of diffusion SPDE defined on Euclidean domains and on surfaces by

$$\frac{\partial Z}{\partial t} + (\kappa^2 - \Delta)^{\alpha/2} Z = \tau \mathcal{W}_T \otimes \mathcal{W}_S,$$

where $\mathcal{W}_T$ denotes a temporal white noise and $\mathcal{W}_S$ denotes either a white or colored noise in space (cf. Figure 2.1 for a simulation from this model). But these models result in random fields with even covariance functions, meaning that changing the sign of the spatial or temporal lag at which the covariance is evaluated does not change the value of the covariance. Consequently, these models are incapable of accounting for transport effects such as advection phenomena (which are intrinsically asymmetrical in time). Note however that in the Euclidean setting, extensions of the SPDE approach allowing to deal with asymmetries in the covariance structure have been proposed by Sigrist et al. (2015), Liu et al. (2022), and Clarotto et al. (2022). But, to the best of our knowledge, the generalization of such models to more complex geometries is left open.

2.3. Machine Learning as a solution?

On the one hand, in recent years, numerous Machine Learning (ML) approaches have been proposed to deal with the complexity of spatiotemporal datasets arising from climate and environmental studies (see eg. Donnelly et al. (2024); Scher (2018); Taylor and Feng (2022)). The premise is that once trained, an ML algorithm would be able to instantaneously generate forecasts, as compared to a run from a physics-based model. However two caveats are still the subject of ongoing research. On the one hand, the question of adequately propagating uncertainties to the forecasts produced by the algorithm is still left open. Indeed, most applications of ML algorithms to forecasting tasks rely on deterministic approaches and do not directly yield ways to incorporate uncertainty in the forecasts (Donnelly et al., 2024). On the other hand, the forecasts produced by such algorithms lack interpretability as they often come from black-box algorithms that merely try to reproduce the training data without any additional knowledge of the physical processes behind the data.

To circumvent this last caveat, physics-informed neural networks have been proposed Cuomo et al. (2022). They allow to impose PDE-based constraints on the forecasts produced by a neural network. If such approaches show promising results, they are not particularly suited to handle data on surfaces, and their successful application to real-world (large) datasets and processes is scarce.

On the other hand, ML also provides promising solutions to deal with the computational problems associated with the inference of GRF-based models. Two main approaches can be mentioned. The first aims to train a neural network capable of identifying the parameters of a covariance function, from a realization of a GRF with this covariance (Gerber and Nychka, 2021; Lenzi et al., 2023; Sainsbury-Dale et al., 2023; Wikle and Zammit-Mangion, 2022). The second aims instead to approximate likelihood as a function of parameters and observations, resulting in an easy-to-compute proxy (Walchessen et al., 2023). The neural likelihood surface can then be maximized for a fixed set of observations to obtain an estimator of the model parameters associated with these observations. In both cases, the computational burden of GRF parameter inference is avoided through the use of neural networks, and therefore opens new possibilities for the use of GRF-based models in real-world applications.

3. Motivations and contributions of the project

The PASDOS project aims to propose new GRF-based models for spatiotemporal data on surfaces that combine the SPDE and Machine Learning approaches described in the previous section. More precisely, the idea is to adopt a SPDE-based approach to GRFs, so that the resulting random fields can be by construction "aware" of physical phenomena like advection or diffusion. This would therefore result in a higher interpretability of the forecasts produced by the GRF. Besides, since GRFs are considered, they can used in a Bayesian framework to yield tractable methods for uncertainty quantification. Finally, the computational problems linked to the inference of GRF models, which now amounts to the inference of SPDE parameters, will be tackled using ML approaches, and in particular graph neural networks.

The challenges and contributions of the PASDOS project are summarized in the next bullet points:

- Generalizing the spatiotemporal SPDE approach to GRF modeling in order to jointly account for diffusion and transport phenomena on arbitrary (smooth compact) surfaces, with a thorough quantification of the errors introduced when solving numerically the resulting SPDE.
- Proposing inference algorithms based on graph neural networks for the SPDE parameters, based on a link between stochastic graph signals and discretizations of SPDEs.
- Proposing a Bayesian framework for uncertainity quantification (of the model parameters and of the produced forecasts).
- Generalizing the proposed approach to multivariate modeling and to more complex SPDEs.
- Collecting and sharing these algorithms through their inclusion in an open source C++ library.

Hence, the impact of the PASDOS project would be multiple. Indeed, we hope that our work on inference will help shed new light on graph neural networks, as the idea would ultimately be to see them as the numerical discretization of SPDEs defined on specific Riemannian manifolds. Besides, through freely available code, we hope to empower practitioners with tools to model and forecast complex spatiotemporal data defined across the globe, for climate and environmental applications. More generally, since we plan on working on models suited for data on arbitrary (compact, smooth) surfaces, our framework could spark interest from practitioners from other communities as well. An example would, for instance, be neuroimaging, in which case the fMRI data lie on the cortical surface (ie. the surface of the brain) (Mejia et al., 2020).

In the next section, we detail how we plan to achieve the goals set above.

4. Detailed research plan

4.1. Starting point and preliminary results

4.1.1 On the SPDE approach to GRFs

In their work, Clarotto et al. (2022) propose an extension of the SPDE approach to an advection-diffusion SPDE defined on a Euclidean domain, and derive analytically the expression of the covariance and precision matrices needed to perform sampling, prediction and inference (through likelihood maximization). Hence, the starting point of this project is to generalize their approach to the Riemannian setting by following the framework we proposed in earlier work (Lang and Pereira, 2023) to define spatial random fields on Riemannian manifolds. More precisely, given a compact Riemannian manifold $(\mathcal{M}, g)$ defined from a smooth compact surface $\mathcal{M}$ equipped with a Riemannian metric g, we now consider the SPDE defined by

$$\frac{\partial Z}{\partial t}(t,x) + \frac{1}{c}\left(g_x(\gamma, \nabla_{\mathcal{M}} Z(t,\cdot)) + (\kappa^2 - \Delta_{\mathcal{M}})Z(t,x)\right) = \tau \mathcal{W}_T \otimes \mathcal{W}_S(t,x), \quad t \in \mathbb{R}_+,\, x \in \mathcal{M}, \tag{2}$$

where $\kappa, \tau, c > 0$, $-\Delta_{\mathcal{M}}$ denotes the Laplace–Beltrami operator on $(\mathcal{M}, g)$, $\mathcal{W}_T$ and $\mathcal{W}_S$ are temporal and spatial noises, and γ is a vector field (defining the advection). Since our goal is to work with general surfaces, we further assume that the advection vector field γ is given as the gradient of a scalar and differentiable function $\phi : \mathcal{M} \to \mathbb{R}$, i.e. $\gamma = \nabla_{\mathcal{M}}\phi$, which eliminates (for now) the question of how to properly define and parametrize vector fields on arbitrary surfaces. Note that in some specific cases, more general parametrization are are easily parametrized. For instance, if $\mathcal{M} = \mathbb{S}^2$ is the 2-sphere, we can decompose any tangent vector field γ as

$$\gamma(s) = \nabla\xi(s) + \vec{n}(s) \times \nabla\chi(s) \in T_s\mathbb{S}^2, \tag{3}$$

for some scalar functions $\xi, \chi : \mathbb{S}^2 \to \mathbb{R}$, and with $\vec{n}(s)$ denoting the vector normal to $\mathbb{S}^2$ and pointing outwards Molina and Slevinsky (2018).

We have already formulated a space-time discretization scheme for SPDE (2) by combining the implicit Euler scheme proposed by Clarotto et al. (2022) with the Galerkin–Chebyshev approximation of Lang and Pereira (2023). In particular, the preliminary results we obtained are described in details in a preprint (Pereira, 2023). Similarly to the Euclidean case, the numerical approximations of solutions SPDE (2) are subject to instabilities when the advection term dominates the diffusion term (Mekuria et al., 2016). This phenomenon can be observed in Figure 4.1b, where a GRF on a cortical surface is simulated using the scheme described earlier: as one can notice in the animation, the values taken by the field quickly explode and tend to oscillate. To avoid such instabilities, it is usual to introduce a stabilization term to the SPDE (Codina, 2000). Consequently, we adapted the Streamline Diffusion approach of Clarotto et al. (2022) to the Riemannian setting, which numerically allowed us to get rid of the instabilities in the simulation (see Figure 4.1c).

We end this section by presenting in Figure 4.2 an illustration on the sphere of the complexity of spatiotemporal patterns that can already me modeled with our approach, by adequately choosing the tangent vector field γ. We also present in Figure 4.3 the spatiotemporal evolution of the covariance between the value of the resulting field at time $t = 0$ at three reference points (in blue), and the values of the field elsewhere and at later times. As one can note in the animation, the zone of high-correlation "moves" along the advection direction, as expected in an advection problem.

4.1.2 On the link between Graph neural networks and discretization of SPDEs

As part of my PhD thesis, I worked on a link between the notion of convolution of signals defined on graphs and the discretization of random fields defined on compact surfaces (Pereira, 2019). The idea was to see the finite element discretization of a linear time-independent SPDE on a triangulated surface as a graph signal, for the graph whose vertices and edges are those of the triangulation. A similar link has been exploited by Sidén and Lindsten (2020) when they proposed their so-called Deep Gaussian Markov models, at least for spatial data on a is a flat square. Their approach resulted in deep neural network architecture for lattice data inspired by the discretization of SPDEs. Their approach was later extended to handle spatial and spatiotemporal data on general graphs using Graph neural networks (Lippert et al., 2023; Oskarsson et al., 2022), but without putting forward any link to SPDEs.

Starting from these preliminary results, we propose to organize the next steps of the project in three work packages (WP), which we now detail.

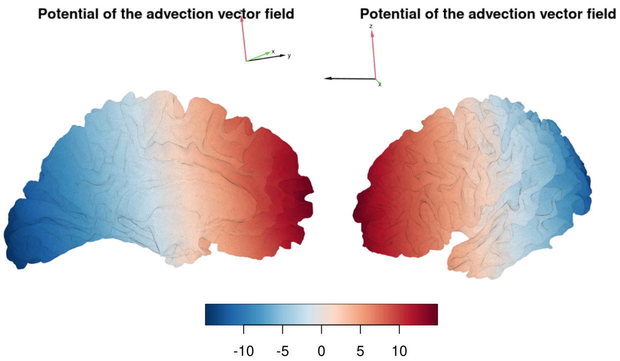

(A) Advection potential function ϕ (whose gradient defines the advection).

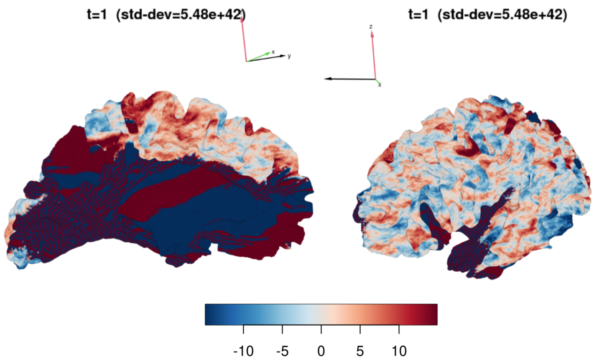

(B) Numerical solution of the SPDE, without stabilization.

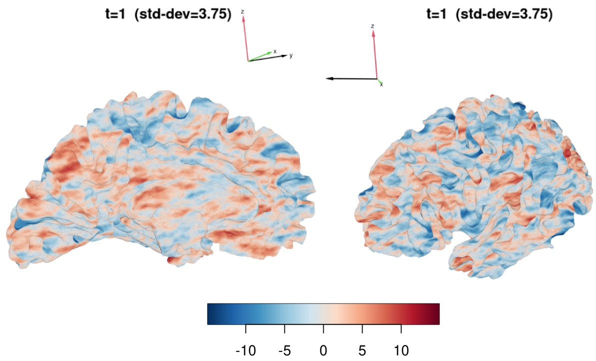

(C) Numerical solution of the SPDE, with stabilization.

Figure 4.1: Simulation of the spatiotemporal advection-diffusion SPDE on a cortical surface, represented from three different viewpoints on the surface (t represents the time step, std-dev the standard deviation of the field values across the surface).

4.2. WP1 : Stochastic and statistical analysis of the model

In this first work package, the emphasis is put on the theoretical analysis of the properties the GRFs obtained by solving SPDE (2).

First off, the errors induced by the spatial and temporal discretization schemes introduced above are assessed. Following the approach we used in earlier work (Lang and Pereira, 2023), mean-square errors and/or errors in covariance between the solution of SPDE (2) and its numerical counterpart are computed. Particular attention is put on the effect of the stabilization term added to avoid numerical instabilities.

Then, further analysis of the covariance of the solutions to SPDE (2) is conducted, so that direct links between the parameters of the SPDE and the properties of the covariance of the field can be deduced. For instance, the separability of the model, the smoothness of the field, the correlation ranges in space and time, and the advection term are linked explicitly to the SPDE parameters.

4.3. WP2 : ML-based inference of SPDE parameters

In this second work package, emphasis is put on the link between the (discretization of) the SPDE-based spatiotemporal GRFs studied in (WP1) and graph neural networks, and how this link can be exploited to propose efficient inference methods for the GRFs.

First, the link between graph signals and discretization of advection-diffusion SPDEs is formalized, thus generalizing previous work made on the spatial case (Pereira, 2019). Then, following an approach similar to the one proposed Sidén and Lindsten (2020) on their seminal work on Deep Gaussian Markov

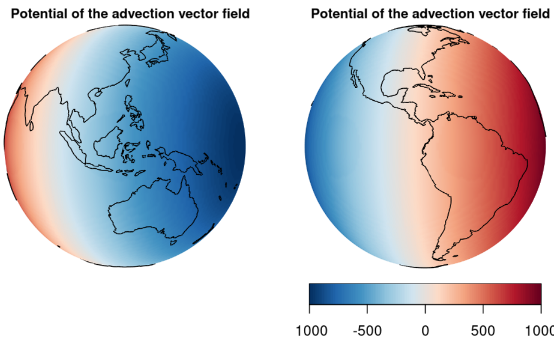

(A) Advection potential function ξ (irrotational component).

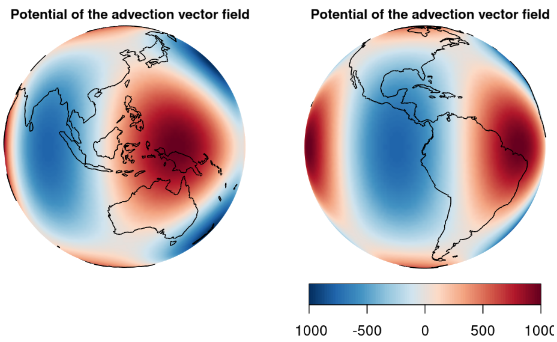

(B) Advection potential function χ (divergence-free component).

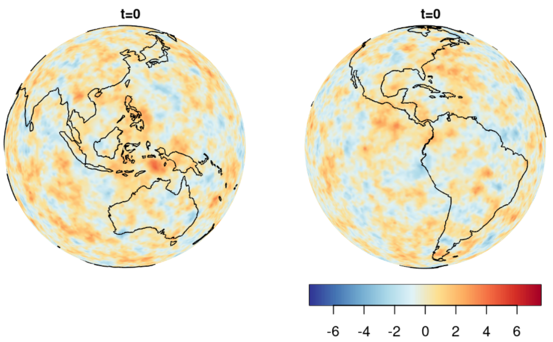

(C) Numerical solution of the SPDE.

Figure 4.2: Simulation of a solution to the advection-diffusion SPDE on a meshed sphere, represented from three different viewpoints on the surface.

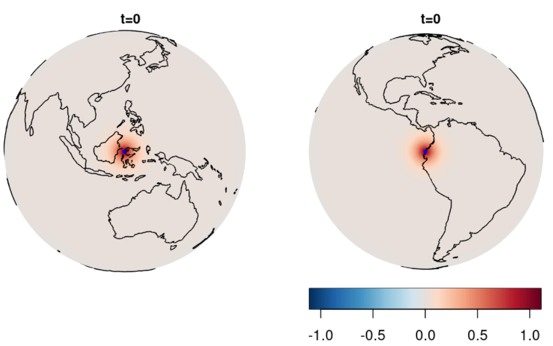

Figure 4.3: spatiotemporal evolution of the covariance between a reference point (in blue) and the rest of the points in the domain, for the spatiotemporal model simulated in Figure 4.2. The color (red to blue) represent the value of the covariance.

random fields, graph neural networks architecture inspired by the discretization of advection-diffusion SPDEs are proposed, and links to the architecture proposed by Oskarsson et al. (2022) are made. In particular, since the graph used in these graph neural networks arises from the triangulation of the surface, the resulting architecture will be particularly tailored to handle data on surfaces.

Finally, the question of the inference of the SPDE-based GRFs is tackled. We start from the approach developed by Sainsbury-Dale et al. (2023) which uses a graph neural network to estimate posterior distributions of model parameters with possibly intractable. The idea is to modify their approach for the estimation of SPDE parameters, while using the graph neural architectures created in the previous step.

All the code and algorithms proposed in this WP will be included in `gstlearn`, the new cross-platform and open source C++ library developed by the Geostatistics team of Mines Paris - PSL, hence making them available to a wide audience (statisticians or otherwise). they will be tested on "real" datasets, for instance the ERA5 global reanalysis dataset Hersbach et al. (2020) and the data from the EUSTACE project Rayner et al. (2020).

4.4. WP3 : Generalizations of the model

The final work package focuses on further generalizations of the approach.

First, the model is modified to allow for more general advection vector fields. For instance, the operator approach to tangent vectors proposed by Azencot et al. (2013) can shed a new light on how to easily define and parametrize vector fields on general surfaces.

Then, SPDE (2) is generalized further. First, hyper-diffusions is considered, in the sense that operator $(\kappa^2 - \Delta_{\mathcal{M}})$ is replaced by a more general function of the Laplacian $f(-\Delta_{\mathcal{M}})$. Then, following the models proposed by Carrizo Vergara et al. (2022) in the Euclidean setting, SPDEs with higher and possibly fractional time derivatives are considered, which in turn allow to better handle the smoothness in time of the resulting GRFs.

Finally, the multivariate case is studied by investigating coupled SPDEs.

5. Project organization and schedule

This project is intended to last three years, with a preliminary timeline shown in Figure 5.1. The funding from the grant would be used to hire a PhD student that would collaborate with me on the project.

Close collaborations with researchers from French and international institutes are also planned. For (WP1), we would work jointly with Annika Lang (Professor), Erik Jansson (PhD student) and Ioanna Motschan-Armen (PhD student) from Chalmers University of Technology and the University of Gothenburg (Sweden) on questions related to the stochastic analysis of the SPDEs and their numerical approximation. For (WP2), Nicolas Desassis (Assistant professor), Thomas Romary (Assistant professor) from Mines Paris - PSL University (France), Lucia Clarotto (Assistant professor) from AgroParisTech - Paris-Saclay University (France) and Denis Allard (Professor) from INRAE (France) would contribute with their experience on sampling and inference algorithms for advection-diffusion SPDEs in the Euclidean setting, as well as on variational methods for Machine Learning. Finally, for (WP3), I would work alongside Ricardo Carrizo-Vergara (Post-doctoral researcher) from the Swiss Ornithological Institute (Switzerland) on the definition and analysis of general spatiotemporal SPDEs on manifolds.

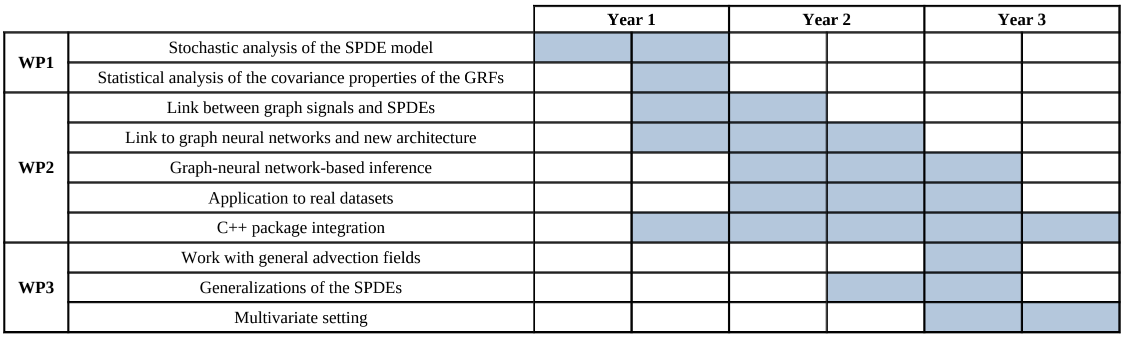

Figure 5.1: Preliminary timeline for the project.

REFERENCES

Azencot, O., Ben-Chen, M., Chazal, F., and Ovsjanikov, M. (2013). An operator approach to tangent vector field processing. In *Computer Graphics Forum*, volume 32, pages 73–82. Wiley Online Library.

Bakka, H., Krainski, E., Bolin, D., Rue, H., and Lindgren, F. (2020). The diffusion-based extension of the Matérn field to space-time.

Bonito, A., Guignard, D., and Lei, W. (2024). Numerical approximation of gaussian random fields on closed surfaces. *Computational Methods in Applied Mathematics*, (0).

Cameletti, M., Lindgren, F., Simpson, D., and Rue, H. (2013). Spatio-temporal modeling of particulate matter concentration through the SPDE approach. *AStA Advances in Statistical Analysis*, 97:109–131.

Carrizo Vergara, R., Allard, D., and Desassis, N. (2022). A general framework for SPDE-based stationary random fields. *Bernoulli*, 28(1):1–32.

Chen, W., Genton, M. G., and Sun, Y. (2021). Space-time covariance structures and models. *Annual Review of Statistics and Its Application*, 8:191–215.

Chiles, J.-P. and Delfiner, P. (2009). *Geostatistics: Modeling Spatial Uncertainty*, volume 497. John Wiley & Sons.

Clarotto, L., Allard, D., Romary, T., and Desassis, N. (2022). The SPDE approach for spatio-temporal datasets with advection and diffusion. *arXiv preprint arXiv:2208.14015*.

Codina, R. (2000). On stabilized finite element methods for linear systems of convection–diffusion-reaction equations. *Computer Methods in Applied Mechanics and Engineering*, 188(1-3):61–82.

Coveney, S., Corrado, C., Roney, C. H., Wilkinson, R. D., Oakley, J. E., Lindgren, F., Williams, S. E., O'Neill, M. D., Niederer, S. A., and Clayton, R. H. (2019). Probabilistic interpolation of uncertain local activation times on human atrial manifolds. *IEEE Transactions on Biomedical Engineering*, 67(1):99–109.

Cuomo, S., Di Cola, V. S., Giampaolo, F., Rozza, G., Raissi, M., and Piccialli, F. (2022). Scientific machine learning through physics–informed neural networks: Where we are and what's next. *Journal of Scientific Computing*, 92(3):88.

Davis, T. A. (2006). *Direct Methods for Sparse Linear Systems*, volume 2. SIAM.

Donnelly, J., Daneshkhah, A., and Abolfathi, S. (2024). Forecasting global climate drivers using gaussian processes and convolutional autoencoders. *Engineering Applications of Artificial Intelligence*, 128:107536.

Gerber, F. and Nychka, D. (2021). Fast covariance parameter estimation of spatial gaussian process models using neural networks. *Stat*, 10(1):e382.

Herrmann, L., Kirchner, K., and Schwab, C. (2020). Multilevel approximation of Gaussian random fields: fast simulation. *Mathematical Models and Methods in Applied Sciences*, 30(01):181–223.

Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., Nicolas, J., Peubey, C., Radu, R., Schepers, D., et al. (2020). The era5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730):1999–2049.

Lang, A. and Pereira, M. (2023). Galerkin–chebyshev approximation of gaussian random fields on compact riemannian manifolds. *BIT Numerical Mathematics*, 63(4):51.

Lenzi, A., Bessac, J., Rudi, J., and Stein, M. L. (2023). Neural networks for parameter estimation in intractable models. *Computational Statistics & Data Analysis*, 185:107762.

Lindgren, F., Bolin, D., and Rue, H. (2022). The SPDE approach for Gaussian and non-Gaussian fields: 10 years and still running. *Spatial Statistics*, 50:100599.

Lindgren, F., Rue, H., and Lindström, J. (2011). An explicit link between Gaussian fields and Gaussian markov random fields: the stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(4):423–498.

Lippert, F., Kranstauber, B., van Loon, E., and Forré, P. (2023). Deep gaussian markov random fields for graph-structured dynamical systems. In Oh, A., Neumann, T., Globerson, A., Saenko, K., Hardt, M., and Levine, S., editors, *Advances in Neural Information Processing Systems*, volume 36, pages 76248–76271. Curran Associates, Inc.

Liu, X., Yeo, K., and Lu, S. (2022). Statistical modeling for spatio-temporal data from stochastic convection-diffusion processes. *Journal of the American Statistical Association*, 117(539):1482–1499.

Mejia, A. F., Yue, Y., Bolin, D., Lindgren, F., and Lindquist, M. A. (2020). A Bayesian general linear modeling approach to cortical surface fMRI data analysis. *Journal of the American Statistical Association*, 115(530):501–520.

Mekuria, G. T., Rao, J. A., et al. (2016). Adaptive finite element method for steady convection-diffusion equation. *American Journal of Computational Mathematics*, 6(03):275.

Molina, J. and Slevinsky, R. M. (2018). A rapid and well-conditioned algorithm for the helmholtz–hodge decomposition of vector fields on the sphere. *arXiv preprint arXiv:1809.04555*.

Oskarsson, J., Sidén, P., and Lindsten, F. (2022). Scalable deep Gaussian Markov random fields for general graphs. In Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., and Sabato, S., editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 17117–17137. PMLR.

Pereira, M. (2019). *Generalized random fields on Riemannian manifolds: theory and practice*. PhD thesis, Université Paris sciences et lettres.

Pereira, M. (2023). A note on spatio-temporal random fields on meshed surfaces defined from advection-diffusion SPDEs. working paper or preprint.

Pereira, M., Desassis, N., and Allard, D. (2022). Geostatistics for large datasets on Riemannian manifolds: A matrix-free approach. *Journal of Data Science*, 20(4):512–532.

Porcu, E., Alegria, A., and Furrer, R. (2018). Modeling temporally evolving and spatially globally dependent data. *International Statistical Review*, 86(2):344–377.

Porcu, E., Furrer, R., and Nychka, D. (2021). 30 years of space–time covariance functions. *Wiley Interdisciplinary Reviews: Computational Statistics*, 13(2):e1512.

Rayner, N. A., Auchmann, R., Bessembinder, J., Brönnimann, S., Brugnara, Y., Capponi, F., Carrea, L., Dodd, E. M., Ghent, D., Good, E., et al. (2020). The EUSTACE project: delivering global, daily information on surface air temperature. *Bulletin of the American Meteorological Society*, 101(11):E1924–E1947.

Rue, H. and Held, L. (2005). *Gaussian Markov Random Fields: Theory and Applications*. CRC press.

Sainsbury-Dale, M., Richards, J., Zammit-Mangion, A., and Huser, R. (2023). Neural bayes estimators for irregular spatial data using graph neural networks. *arXiv preprint arXiv:2310.02600*.

Scher, S. (2018). Toward data-driven weather and climate forecasting: Approximating a simple general circulation model with deep learning. *Geophysical Research Letters*, 45(22):12–616.

Sidén, P. and Lindsten, F. (2020). Deep gaussian markov random fields. In *International conference on machine learning*, pages 8916–8926. PMLR.

Sigrist, F., Künsch, H. R., and Stahel, W. A. (2015). Stochastic partial differential equation based modelling of large space-time data sets. *Journal of the Royal Statistical Society: Series B: Statistical Methodology*, pages 3–33.

Taylor, J. and Feng, M. (2022). A deep learning model for forecasting global monthly mean sea surface temperature anomalies. *Frontiers in Climate*, 4:932932.

Wackernagel, H. (2003). *Multivariate Geostatistics: An Introduction with Applications*. Springer Science & Business Media.

Walchessen, J., Lenzi, A., and Kuusela, M. (2023). Neural likelihood surfaces for spatial processes with computationally intensive or intractable likelihoods. *arXiv preprint arXiv:2305.04634*.

Whittle, P. (1954). On stationary processes in the plane. *Biometrika*, pages 434–449.

Wikle, C. K. and Zammit-Mangion, A. (2022). Statistical deep learning for spatial and spatio-temporal data. *arXiv preprint arXiv:2206.02218*.