


Article

Robust Inference of Dynamic Covariance Using Wishart Processes and Sequential Monte Carlo

Hester Huijsdens *, David Leeftink, Linda Geerligs and Max Hinne 

Donders Institute for Brain, Cognition and Behaviour, Radboud University Nijmegen, Thomas van Aquinostraat 4, 6525 GD Nijmegen, The Netherlands; david.leeftink@ru.nl (D.L.); linda.geerligs@donders.ru.nl (L.G.); max.hinne@donders.ru.nl (M.H.)

* Correspondence: hester.huijsdens@donders.ru.nl

Abstract: Several disciplines, such as econometrics, neuroscience, and computational psychology, study the dynamic interactions between variables over time. A Bayesian nonparametric model known as the Wishart process has been shown to be effective in this situation, but its inference remains highly challenging. In this work, we introduce a Sequential Monte Carlo (SMC) sampler for the Wishart process, and show how it compares to conventional inference approaches, namely MCMC and variational inference. Using simulations, we show that SMC sampling results in the most robust estimates and out-of-sample predictions of dynamic covariance. SMC especially outperforms the alternative approaches when using composite covariance functions with correlated parameters. We further demonstrate the practical applicability of our proposed approach on a dataset of clinical depression ($n = 1$), and show how using an accurate representation of the posterior distribution can be used to test for dynamics in covariance.

Keywords: dynamic covariance; Wishart processes; Bayesian inference; Sequential Monte Carlo; Markov Chain Monte Carlo; variational inference



Citation: Huijsdens, H.; Leeftink, D.; Geerligs, L.; Hinne, M. Robust Inference of Dynamic Covariance Using Wishart Processes and Sequential Monte Carlo. *Entropy* **2024**, *26*, 695. <https://doi.org/10.3390/e26080695>

Academic Editors: Mihaly Mezei and Boris Ryabko

Received: 3 June 2024

Revised: 5 July 2024

Accepted: 13 August 2024

Published: 16 August 2024



Copyright: © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Various domains study the joint behaviour of multiple time series. For example, in the human brain, these time series consist of neuronal activation patterns; in finance, they represent stock indices; and in psychology, they show self-reported measures of mental health. For many research questions in these domains, it is essential to study the covariance structure between different time series. In neuroscience for example, the communication between different brain areas is studied [1,2], which in turn can be used as a marker to diagnose several neurological disorders [3]. Other examples include assessing the risks and returns of stock portfolios by investigating the covariance of different assets [4], and investigating the co-occurrence of symptoms in mental disorders [5–7]. Recently, there has been a shift in focus from hitherto static representations of these interactions, to dynamic covariances, in which the interactions between time series change as a function of an input variable. For example, recent findings in neuroscience suggest that the interactions between brain regions can change over time and that modelling the covariance between brain regions dynamically provides more sensitive biomarkers for cognition [1,2,8]. Similarly, in finance, the dynamic interactions between stock markets are used to study volatility and financial crises [9–12]. Lastly, in psychology, covariance structures between mental health markers are shown to be altered in individuals with neuroticism [13] and major depressive disorder (MDD) [14]. Namely, the covariance between symptoms is stronger in subjects diagnosed with major depressive disorder compared to healthy controls. Importantly, even within a single subject, the covariance structure between symptoms changes near the onset of depressive episodes, providing potential early warning signals [15].

Processes of dynamic covariance can be modelled in several ways. The most prominent approaches include the multivariate generalised autoregressive conditional heteroscedastic

(MGARCH) model commonly used in finance [16–19], and the sliding-window approach that is popular in neuroscience [20,21]. However, both approaches have a number of shortcomings. Importantly, the MGARCH-family of models requires that observations are evenly spaced over the input domain. Although even spacing is often not a problem when the questions concern time series, there are many examples where even spacing is not feasible, for example, when studying cross-sectional age-related differences or how medication dosage affects the covariance of mental health symptoms. Furthermore, the sliding-window approach requires the user to determine a number of parameters, such as window size and stride length, that can greatly affect the dynamic covariance estimates they result in [22,23]. For example, larger window sizes will result in slower observed changes in covariance, while smaller window sizes will result in more noisy measurements of covariance. In an attempt to address these challenges, Wilson and Ghahramani [24] introduced the generalised Wishart process. The Wishart process is a Bayesian nonparametric approach based on Gaussian processes (GPs) [25] that, in contrast to the aforementioned methods, can handle unevenly spaced observations. Furthermore, as a Bayesian approach, it does not provide a point estimate of dynamic covariance, but a distribution over dynamic covariance structures. This in turn allows the model to indicate its estimation uncertainty, which enables the user to perform statistical tests. For example, with a probabilistic estimate of the Wishart process, one can test for the presence of (dynamic) covariance, even when observations are only available for a single subject. This model has been applied in different contexts, for example, for modelling noise covariance in neural populations across trials [26], when studying time-varying functional brain connectivity [27,28], for improving the resolution in diffusion magnetic resonance imaging [29], and in combination with stochastic differential equations [30].

Although the Wishart process overcomes several of the limitations of the other dynamic covariance methods, inference of the model parameters remains challenging. Especially for composite covariance functions, the model is high-dimensional, and several parameters are highly correlated. This makes the posterior distribution potentially multimodal. Wilson and Ghahramani [24] inferred the model parameters using Markov Chain Monte Carlo (MCMC) sampling. Although MCMC samplers are guaranteed to converge to the true distribution, they have difficulties in sampling from high-dimensional distributions efficiently. As a solution, Heaukulani and van der Wilk [31] proposed a variational inference approach based on sparse Gaussian processes [32]. Although this approach is indeed much more scalable than MCMC-based approximations, it is not robust against local minima, and provides no posterior distribution over the hyperparameters of the model.

In this work, we propose a third approximate inference scheme for the Wishart process using Sequential Monte Carlo (SMC) [33,34]. SMC approximations were originally introduced for filtering approaches in state-space models [35], but more recently, they have been gaining popularity as a generic approximate Bayesian inference technique [36,37]. Fundamentally, SMC performs a large number of short MCMC-based inference chains on different initializations of the model parameters, known as particles, in parallel, which are then combined using importance sampling. The parallelisation makes SMC well suited for inference of high-dimensional and multimodal distributions, as it tends not to become stuck in local optima. In addition, the computation that is required within the chains can largely be executed in parallel, which enables the algorithm to benefit from modern parallel compute hardware, such as GPUs. Here, we introduce the SMC inference scheme for the Wishart process and compare it to MCMC [24] and variational inference [31]. In most cases, SMC outperforms these approaches in terms of model fit and predictive performance. Furthermore, although variational inference tends to converge more quickly, SMC provides the full posterior at comparatively little additional running time.

This paper is organised as follows. In Section 2, we describe the Wishart process and SMC sampling. We also briefly recap MCMC sampling and variational inference. In Section 3, we compare the inference methods in different simulation studies that focus on capturing the true covariance process and the latent model parameters, and how accu-

rately capturing these model parameters can be of importance of accurate out-of-sample predictions. In Section 4, we demonstrate how the Wishart process can be used in practice by applying it to a dataset of self-reported depression symptoms [15] and show how the distribution over the covariance can be used to test for dynamics in covariance. Section 5 concludes our comparison and discusses future research directions.

2. Bayesian Inference of Wishart Processes

2.1. Wishart Processes

To understand the generalised Wishart process [24,31,38], we first describe the situation in which we model a constant covariance matrix using the Wishart distribution. Let d be the number of variables (that is, time series), and let $\mathbf{x} = (x_1, \dots, x_n)$ be a vector of n input locations, $x_i \in \mathbb{R}$, and $\mathbf{Y} = (\mathbf{y}_1, \dots, \mathbf{y}_n)^\top$ a matrix of observations with $\mathbf{y}_i \in \mathbb{R}^d$, such that $\mathbf{Y} \in \mathbb{R}^{n \times d}$. We assume that \mathbf{y}_i is drawn from a multivariate normal distribution with a mean of zero (although this can easily be extended to other mean vectors as well), and a covariance Σ :

$$\mathbf{y}_i \sim \mathcal{MVN}_d(\mathbf{0}, \Sigma) \quad , \quad i = 1, \dots, n \quad . \quad (1)$$

To learn the covariance matrix Σ from the observations, we follow a Bayesian approach, which implies we must decide on a prior distribution for the latent variable Σ . A popular choice of prior for covariance matrices is the Wishart distribution, because it is conjugate to the normal distribution, and therefore, the posterior $p(\Sigma \mid \mathbf{x}, \mathbf{Y})$ can be computed analytically (see [39]). The Wishart distribution is parameterised by a scale matrix \mathbf{V} and a scalar degrees of freedom parameter v , and has the following density:

$$p(\Sigma \mid \mathbf{V}, v) = \frac{|\Sigma|^{(v-d-1)/2} \exp(-\text{tr}(\mathbf{V}^{-1}\Sigma)/2)}{2^{vd/2} |\mathbf{V}|^{v/2} \Gamma_d(v/2)} = \mathcal{W}_d(\mathbf{V}, v) \quad (2)$$

where $\text{tr}(\cdot)$ is the trace function and $\Gamma_d(\cdot)$ is the multivariate gamma function. The intuition behind parameters \mathbf{V} and v is as follows. Suppose we have a matrix $\mathbf{F} \in \mathbb{R}^{d \times v}$, of which each column is drawn independently from a multivariate normal distribution with a mean of zero and no covariance between the elements (that is, the covariance matrix is the identity matrix \mathbf{I}), i.e., $\mathbf{f}_l = (f_{1l}, \dots, f_{dl})^\top \sim \mathcal{MVN}_d(\mathbf{0}, \mathbf{I})$; then, the sum over the outer products of the v columns of \mathbf{F} is Wishart-distributed with scale matrix \mathbf{I} and v degrees of freedom. Additionally, we can scale the outer products by the lower Cholesky decomposition \mathbf{L} of scale matrix \mathbf{V} (that is, $\mathbf{V} = \mathbf{L}\mathbf{L}^\top$):

$$\begin{aligned} \mathbf{f}_l &= (f_{1l}, \dots, f_{dl})^\top \sim \mathcal{MVN}_d(\mathbf{0}, \mathbf{I}) \quad , \quad l = 1, \dots, v \\ \Sigma &= \sum_{l=1}^v \mathbf{L}\mathbf{f}_l\mathbf{f}_l^\top \mathbf{L}^\top \sim \mathcal{W}_d(\mathbf{V}, v) \quad , \quad l = 1, \dots, v \quad . \end{aligned} \quad (3)$$

The resulting covariance matrix $\Sigma \in \mathbb{R}^{d \times d}$ is Wishart-distributed with scale matrix \mathbf{V} and degrees of freedom v .

In the Wishart process, the constant covariance matrix Σ is replaced by an input-dependent covariance matrix. This leads to the following definition of the observations, where now Σ is parameterised by x :

$$\mathbf{y}_i \sim \mathcal{MVN}_d(\mathbf{0}, \Sigma(x_i)) \quad , \quad i = 1, \dots, n \quad . \quad (4)$$

Wilson and Ghahramani [24] describe a constructive approach for $\Sigma(x_i)$ that is similar to how the Wishart distribution is made out of normal distributions. This time, the multivariate normal distributed columns of \mathbf{F} in Equation (3) are replaced by i.i.d. GPs evaluated at \mathbf{x} . These GPs have a zero mean function and a kernel function κ_θ , with θ as its hyperparameters:

$$\mathbf{f}_{jl}(x) \sim \text{GP}(\mathbf{0}, \kappa_\theta) \quad , \quad j = 1, \dots, d \quad , \quad l = 1, \dots, v \quad . \quad (5)$$

Under the assumption that $\kappa_\theta(x_i, x_i) = 1$, taking the sum of outer products for every observation x_i results in a Wishart-distributed covariance matrix, that is, $\Sigma(x_i) = \sum_{l=1}^v \mathbf{f}_l(x_i)\mathbf{f}_l(x_i)^\top \sim W_d(\mathbf{I}, v)$, where \mathbf{I} is the identity matrix. Similar to the Wishart distribution, we then construct Σ by scaling this sum of outer products by the scale matrix \mathbf{V} :

$$\Sigma(x_i) = \sum_{l=1}^v \mathbf{L}\mathbf{f}_l(x_i)\mathbf{f}_l(x_i)^\top \mathbf{L}^\top \sim W_d(\mathbf{V}, v) , \quad i = 1, \dots, n , \quad (6)$$

with $\mathbf{V} = \mathbf{L}\mathbf{L}^\top$ as before. Figure 1 provides a visual illustration of the constructive approach to the Wishart process.

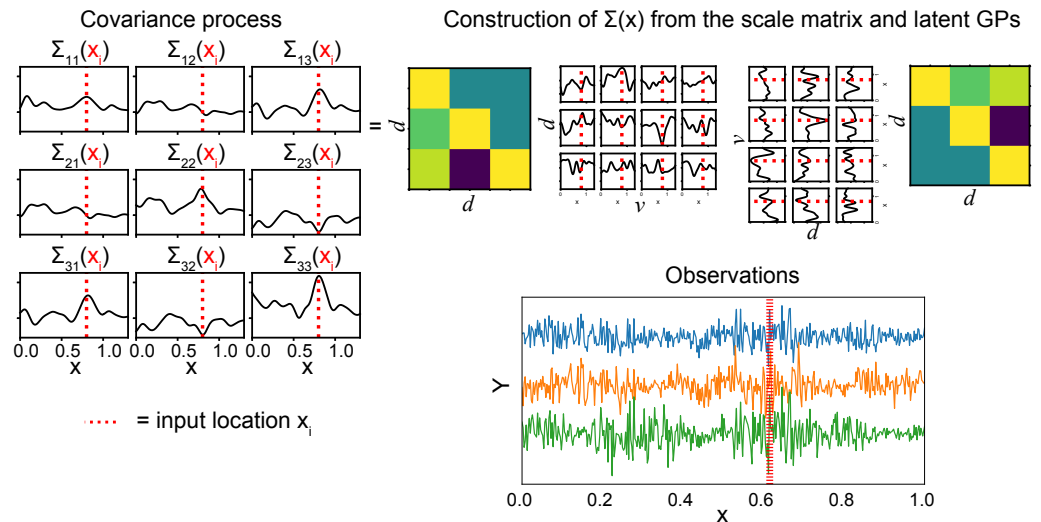


Figure 1. A visualization of the construction of the Wishart process (see Equation (6)). With $d = 3$ variables and $v = 4$ degrees of freedom, the covariance at a single input x_i (represented by the red dotted line) is constructed from the outer product of the GP samples evaluated at this input. The resulting $d \times d$ is subsequently scaled by the $d \times d$ lower Cholesky decomposition of a scale matrix. Below, the time series from which to perform inference are shown.

To complete the Bayesian model, we define the following prior distribution scheme. We set a normal prior on each element of \mathbf{L} independently, and determine the prior of θ based on the covariance function (as will be described in Section 3):

$$\begin{aligned} \theta &\sim p(\theta) \\ \mathbf{f}_{jl} &\sim \text{GP}(\mathbf{0}, \kappa_\theta) \quad \text{and} \quad \mathbf{f}_l = (f_{1l}, \dots, f_{dl})^\top & j = 1, \dots, d , \quad l = 1, \dots, v \\ L_{jo} &\sim \mathcal{N}(0, 1) & j = 1, \dots, d , \quad o = 1, \dots, d \\ \Sigma(x_i) &= \sum_{l=1}^v \mathbf{L}\mathbf{f}_l(x_i)\mathbf{f}_l(x_i)^\top \mathbf{L}^\top & i = 1, \dots, n \\ \mathbf{y}_i &\sim \mathcal{MVN}_d(\mathbf{0}, \Sigma(x_i)) & i = 1, \dots, n . \end{aligned}$$

Finally, if we want to predict observations \mathbf{y}_* at test locations \mathbf{x}_* , we first predict the latent GPs:

$$\mathbf{f}_{jl}^* | \mathbf{f}_{jl} \sim \mathcal{MVN}_{n_*} \left(\mathbf{K}_{*x} \mathbf{K}_{xx}^{-1} \mathbf{f}_{jl}, \mathbf{K}_{**} - \mathbf{K}_{*x} \mathbf{K}_{xx}^{-1} \mathbf{K}_{*x}^\top \right) \quad j = 1, \dots, d , \quad l = 1, \dots, v , \quad (7)$$

where \mathbf{K}_{*x} is formed by evaluating κ_θ at all combinations of test and training inputs, and \mathbf{K}_{**} by evaluating κ_θ at all pairs of test locations. Subsequently, we construct the covariance process using Equation (6) and then sample \mathbf{y}_* using Equation (4).

An important property of the Wishart process is that the covariance function κ_θ can be used to express different qualitative beliefs of the dynamic covariances. For example, when using the Radial Basis Function (RBF) as covariance function, the covariance process becomes autocorrelated and smooth, as covariances corresponding to nearby input locations will be similar. Alternatively, if periodicity is expected in the covariance process, we can model this using a (locally) periodic function for κ_θ . Figure 2 demonstrates a few examples of how different covariance functions and hyperparameters influence the covariance process.

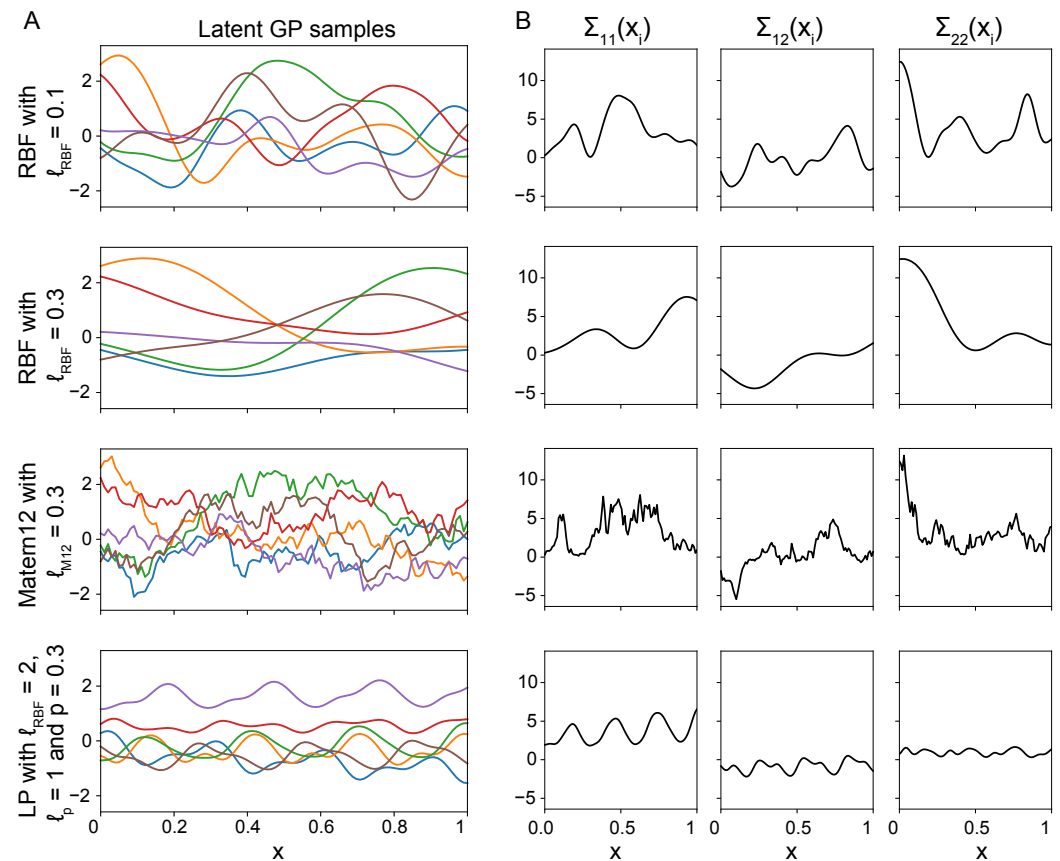


Figure 2. Different GP covariance functions can model different covariance structures. **(A)** GP samples drawn from GP priors with different covariance functions and hyperparameters. The GP samples have an RBF covariance function with $\ell_{\text{RBF}} = 0.1$ or $\ell_{\text{RBF}} = 0.3$, a Matérn 1/2 covariance function with $\ell_{\text{M12}} = 0.3$, or a Locally Periodic covariance function with $\ell_{\text{RBF}} = 2.0$, $\ell_p = 1$ and $p = 0.3$. **(B)** The covariance processes are constructed from the GP samples on the left and the lower Cholesky decomposition of the scale matrix, \mathbf{L} (here set to the identity matrix). The upper triangular elements of $\Sigma(x_i)$ are visualised, showing that the covariance process is a complex combination of GP samples.

2.2. Bayesian Inference

Although the construction of the Wishart process appears to be a straightforward extension of the Wishart distribution, inference of the corresponding posterior distribution $p(\Sigma(x_i) | \mathbf{x}, \mathbf{Y})$ (note that the dependency on \mathbf{x} is sometimes omitted to improve legibility when no confusion is likely to arise) is substantially more involved. Foremost, the likelihood of the Wishart process is not conjugate to the prior, which prohibits exact inference and forces us to opt for approximate methods instead. However, this remains a challenge, as some of the model parameters are highly correlated. Previous studies have sampled from the posterior using Markov Chain Monte Carlo (MCMC) sampling [24], or approximated the posterior using a variational approach [26,31]. Although both approaches showed an

improved performance compared to existing dynamic covariance modelling methods, both methods have trouble inferring high-dimensional and potentially multimodal distributions. Therefore, we introduce a third method to inference of the Wishart process based on Sequential Monte Carlo (SMC) samplers [34]. Before expanding on this new approach, we briefly recap the existing algorithms used to infer the posterior distributions of a Wishart process.

2.2.1. Markov Chain Monte Carlo and Variational Inference

We want to infer the posterior $p(\Sigma(x_i) | \mathbf{x}, \mathbf{Y})$. Since $\Sigma(x_i)$ follows deterministically from $f_{jl}(x_i)$, for all $j \in 1, \dots, d$ and $l \in 1, \dots, v$, this comes down to learning the posterior $p(\mathbf{F}, \mathbf{L}, \theta | \mathbf{x}, \mathbf{Y})$, where \mathbf{F} contains all $d \times v$ independent GP samples. Wilson and Ghahramani [24] use MCMC sampling to infer this posterior distribution. A detailed explanation of their approach can be found in Appendix A, but here, we will briefly describe the sampling algorithm. The MCMC approach uses Gibbs sampling [40] where in each MCMC iteration, the parameters are updated according to their conditional distributions:

$$\begin{aligned} p(\mathbf{F} | \theta, \mathbf{L}, \mathbf{x}, \mathbf{Y}) &\propto p(\mathbf{Y} | \mathbf{F}, \mathbf{L}, \mathbf{x})p(\mathbf{F} | \mathbf{x}, \theta) \\ p(\theta | \mathbf{F}, \mathbf{L}, \mathbf{x}, \mathbf{Y}) &\propto p(\mathbf{F} | \mathbf{x}, \theta)p(\theta) \\ p(\mathbf{L} | \theta, \mathbf{F}, \mathbf{x}, \mathbf{Y}) &\propto p(\mathbf{Y} | \mathbf{F}, \mathbf{L}, \mathbf{x})p(\mathbf{L}) \end{aligned} \quad (8)$$

Here, $p(\mathbf{Y} | \mathbf{F}, \mathbf{L}, \mathbf{x})$ is the multivariate normal likelihood and $p(\theta)$ and $p(\mathbf{L})$ are prior distributions for the covariance function parameters and the scale matrix, respectively. The distribution $p(\mathbf{F} | \mathbf{x}, \theta)$ is a multivariate normal prior on the latent GPs. Although the approach by Wilson and Ghahramani [24] was shown to be effective, the method scales unfavourably to higher dimensions, both in the number of observations n and the number of variables d . To address this issue, Heaukulani and van der Wilk [31] instead propose to approximate the posterior using variational inference, a method that uses optimisation instead of sampling. Additionally, they make use of several techniques commonly found in the GP literature, such as sparse Gaussian processes [41], to make inference more efficient. More details on their method can be found in Appendix B.

2.2.2. A Sequential Monte Carlo Sampler for Wishart Processes

The MCMC and variational inference approaches for approximate inference of the Wishart process both possess a number of drawbacks. First, due to the correlations in the model parameters and potential multimodality in the posterior, depending on the choice of covariance function κ_θ , a standard MCMC approach is inefficient and requires a large number of samples to converge. The variational inference approach by Heaukulani and van der Wilk [31] enables scaling applications up to larger datasets, but in practice, it is prone to becoming stuck in local optima. Furthermore, it does not provide a posterior distribution for the hyperparameters of the model. To overcome these limitations, we here introduce a novel Sequential Monte Carlo (SMC) sampler [34,36] for posterior inference. SMC samplers are efficient at sampling from multimodal distributions because, instead of initialising the parameters at a single location, they initialise a large amount of parameter sets, called *particles*, and iteratively update these particles based on their fit to the observations. Additionally, the updates for the different particles can be performed in parallel. This allows for a substantial speed increase compared to the other approaches, although of course this requires the availability of parallel computation hardware, such as GPUs.

The SMC algorithm starts from an easy-to-sample density such as the prior, and incrementally lets the particles sample from more complex densities, to eventually approach the target density. The outline of the sampler is as follows. First, s sets of parameters (particles) $\{\mathbf{F}^{(i)}, \theta^{(i)}, \mathbf{L}^{(i)}\}_{i=1}^s$ are initialised by drawing them from their prior distributions. Each particle is assigned a weight, which is initially set to $w_0^{(i)} = 1/s$. Next, we iteratively apply a weighting, resampling, and mutation step to adapt the particles based on their fit to the observations:

1. In the weighting step, particles are assigned a weight based on how well each particle fits the data using

$$w_t^{(i)} = \frac{p_t(\mathbf{Y} | \mathbf{F}_t^{(i)}, \theta_t^{(i)}, \mathbf{L}_t^{(i)})}{p_{t-1}(\mathbf{Y} | \mathbf{F}_t^{(i)}, \theta_t^{(i)}, \mathbf{L}_t^{(i)})}, \quad (9)$$

where t represents the current SMC iteration, and i , the particle index. Depending on the implementation, the distribution p_t can change for every SMC iteration t . This distribution will eventually form our approximation of the posterior distribution.

2. Next, the particles are resampled with replacement in proportion to their weights. This means that particles with small weights are discarded, and particles with large weights are duplicated.
3. Lastly, the particles are mutated by performing a number of Gibbs cycles (see Equation (8)) for each particle, using the tempered distribution $p_t(\mathbf{Y} | \mathbf{F}_t^{(i)}, \theta_t^{(i)}, \mathbf{L}_t^{(i)})$ as the likelihood. This avoids the risk of all particles receiving identical parameters after a few iterations.

If we set the tempered distribution to the likelihood, we will mainly explore regions with a high likelihood, because the particles are only weighted based on their fit to the data. This risks the issue known as particle collapse, where all particles consist of the same high-likelihood values. To overcome this, we use an adaptive-tempering variant of SMC [36,42]. Here, the distribution at SMC iteration t is tempered according to

$$p_t(\mathbf{Y} | \mathbf{F}, \theta, \mathbf{L}) = p(\mathbf{Y} | \mathbf{F}, \theta, \mathbf{L})^{\beta_t}, \quad (10)$$

where β_t is the temperature that dampens the influence of the likelihood. β_t is initially set to 0. This means we initially simply sample from the prior, and β is then gradually increased until it reaches a value of 1, at which point we sample from the posterior. The increase in temperature between two successive SMC iterations, $\Delta\beta_t$, is determined via the effective sample size of the weights, s_{eff} . The effective sample size is a measure of particle diversity. The new temperature at every iteration is then determined by finding β_t such that $s_{\text{eff}} = as$, where a is the fraction of particles that we want to be independent [33,43,44].

3. Simulation Studies

We compare MCMC, variational inference, and SMC on two distinct simulations representing different scenarios of dynamic covariances. Before we compare these three methods, we first provide implementation details of all three methods. Next, we describe the data generation procedure for these simulation studies and how we will evaluate each approach.

3.1. Implementation Details

Gibbs MCMC sampling is implemented using the Blackjax Python library [45], which builds on top of the JAX framework [46]. When sampling with MCMC, we sample the covariance function hyperparameters θ and the lower Cholesky decomposition of the scale matrix \mathbf{L} using a Random Walk Metropolis Hastings sampler with a step size of 0.01. We use a thinning of 1000 samples and the number of burn-in steps is determined by the convergence of all model parameters, unless mentioned otherwise in the experiment. Convergence is measured by the Potential Scale Reduction Factor (PSRF) [47], where we interpret a value of less than 1.1 as being converged. We measure the PSRF over four re-runs, known as ‘chains’, of the inference algorithm, each time with a different random initialisation of the model parameters. After convergence across chains, we combine these four chains by randomly taking 250 samples of each chain.

Variational inference is implemented using the GPflow 2 library [48], and the implementation can be found on GitHub (<https://github.com/DavidLeeftink/BANNER.git>, accessed on 10 March 2024). The noise parameter in Equation (A4) was initialised as $\lambda_{jj} = 0.001$ for $j = 1, \dots, d$. Similar to Heaukulani and van der Wilk [31], we approximate the gradients of expectation of the log-likelihood in Equation (A3) using a small number of Monte Carlo estimates. In our results, we have used three Monte Carlo estimates. To optimize Equation (A3), we use the Adam optimiser [49] with an initial learning rate of 0.001. We do not make use of minibatches or inducing points. For variational inference, we can use the PSRF only as a measure for convergence of the covariance process, but not for the latent model parameters, because variational inference only provides point-estimates for these. Therefore, we optimize the ELBO until it has not improved for 10,000 iterations, after which we use a PSRF below 1.1 as a criteria for the convergence of the covariance process. As before, the PSRF is computed over the estimates of four re-runs with a different random initialisation. Moreover, within each re-run, we inspect whether or not the latent model parameters have converged when the ELBO has converged. We then use the run with the highest ELBO for subsequent analyses.

Similar to the MCMC implementation, SMC sampling is also implemented in the Blackjax Python library. Within each SMC cycle, we sample the covariance function hyperparameters θ and the lower Cholesky decomposition of the scale matrix \mathbf{L} using the Gibbs MCMC sampling approach described above. We set the number of particles to 1000, and, unless mentioned otherwise, we base the number of mutation steps on the convergence of all model parameters, as determined by a PSRF below 1.1. This convergence is again measured over four re-runs of the inference algorithm, each time with a different random parameter initialisation. Code for our analyses is available on GitHub (<https://github.com/Hesterhuijsdens/GWP-SMC>).

3.2. Synthetic Data

In order to compare MCMC, variational inference, and SMC, we evaluate each approach on data with a known ground truth covariance process. In our first simulation study, we construct a dynamic covariance process that is drawn from a Wishart process prior with a Radial Basis Function (RBF; also known as the squared-exponential) as covariance function:

$$\kappa_{\text{RBF}}(x, x') = \exp\left(-\frac{(x - x')^2}{2\ell_{\text{RBF}}^2}\right). \quad (11)$$

We set the lengthscale parameter to $\ell_{\text{RBF}} = 0.35$, representing slow dynamics, and the scale matrix \mathbf{V} to the identity matrix. Using $d = 3$ variables and $v = 4$ degrees of freedom, we draw $d \times v$ GP samples and construct the covariance process using Equation (6). We repeat this covariance process generation procedure ten times while keeping the lengthscale parameter ℓ_{RBF} and the scale matrix \mathbf{V} the same. Finally, the ten resulting covariance processes are used to generate ten datasets, by sampling $n = 300$ observations from a multivariate normal distribution with a mean of zero, and the latent covariance process as covariance. An example of a generated ground truth covariance process is shown in Figure 3A.

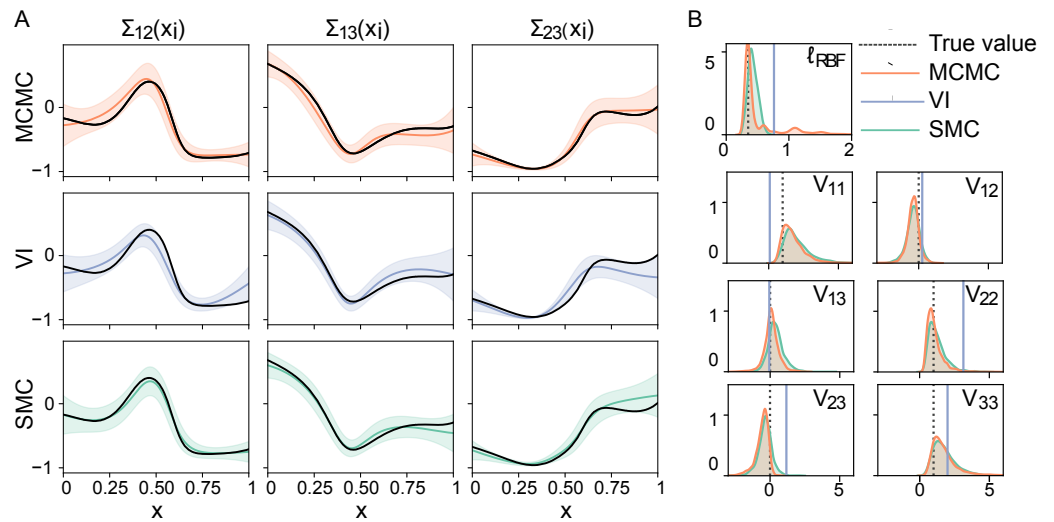


Figure 3. (A) Estimates of a covariance process drawn from a Wishart process prior. The ground truth covariance process is shown in black. (B) Inference results of the lengthscale parameter of the Radial Basis Function covariance function and the scale matrix. The true values ($\ell_{\text{RBF}} = 0.35$ and $\mathbf{V} = \mathbf{I}$) are indicated by the grey dotted lines.

In our second simulation study, we generate a covariance process that follows a rapid state-switching pattern between $d = 3$ variables. This covariance process is generated as follows. The off-diagonal elements of the true latent covariance process alternate every 50 observations between values of 0 and 0.8, and the variances are set to 1. This structure is shown in Figure 4. Again, we generate ten datasets, but this time, we use the same ground truth covariance process for all ten datasets. We draw $n = 600$ observations from a multivariate normal distribution with a mean of zero and the state-switching latent covariance process, of which the first $n_{\text{train}} = 300$ observations are being used for inference, and the remaining $n_{\text{test}} = 300$ observations for out-of-sample prediction. To capture the periodic structure of this covariance process, we use both a Periodic covariance function and a Locally Periodic (LP) covariance function. The Periodic covariance function is defined as

$$\kappa_{\text{Periodic}}(x, x') = \exp\left(-\frac{2 \sin^2(\pi|x - x'|/p)}{\ell_p^2}\right), \tag{12}$$

where p determines the period of the covariance, and ℓ_p determines the fluctuations within each period. The LP covariance function is constructed by multiplying this Periodic function with an RBF covariance function to obtain

$$\kappa_{\text{LP}}(x, x') = \exp\left(-\frac{2 \sin^2(\pi|x - x'|/p)}{\ell_p^2}\right) \exp\left(-\frac{(x - x')^2}{2\ell_{\text{RBF}}^2}\right), \tag{13}$$

where p again determines the period of the covariance, ℓ_p determines the fluctuations within one period, and ℓ_{RBF} allows the repeating covariance to change over time. Both functions should be able to capture the ground truth state switches; however, the LP covariance function allows for more flexibility. We set a log-normal prior on all three parameters.




Figure 4. Covariance process estimates and out-of-sample predictions for MCMC (in orange), VI (in purple) and SMC (in green), using a Periodic covariance function, together with the corresponding distributions over the covariance function parameters. The true covariance process is shown in black, and out-of-sample predictions are shown after the grey dotted line. The Periodic covariance function has two parameters: the period (p) and the lengthscale within each period (ℓ_{RBF}).

3.3. Performance Metrics

We evaluate the different inference approaches based on how well they recover the ground truth. Therefore, we compute the mean squared error (MSE) between the ground truth covariance process and the corresponding mean estimate of the covariance process. This metric is averaged over all d variables and n observations. Additionally, when the model parameters that were used to construct the ground truth covariance process are known (as in our first simulation), we compute the MSE between those parameters and the corresponding mean estimates. Furthermore, since the inference methods provide a distribution over the covariance, we evaluate this full distribution by computing the MSE between samples of the covariance process and the ground truth. This MSE is again averaged over all d variables and n observations. We refer to this metric over the full covariance distribution as $\text{MSE}_{\text{samples}}$. Hence, we use the MSE to evaluate both the accuracy of the mean covariance process estimate and the accuracy of the distribution over the covariance process estimate.

Finally, in the second simulation study, we are interested in making out-of-sample predictions. Here, we evaluate the predictive performances of the three inference methods by means of the log-likelihood (LL) of the observations, given the mean covariance process estimate. This log-likelihood is averaged over the number of observations n . To also validate the predictive performance over the entire predictive posterior distribution, we measure the Kullback–Leibler divergence (KL) between the predictive posterior distribution and the true multivariate normal distribution of the observations.

3.4. Simulation Study 1: Learning the Model Parameters

The aim of our first simulation is to validate the ability of each inference method to capture the ground truth. Therefore, we simulate observations with a covariance process drawn from a Wishart prior (as described in Section 3.2), and use these data to measure the accuracy of MCMC sampling, variational inference, and SMC sampling in inferring the covariance process, the scale matrix, and covariance function hyperparameters. For each approach, we use the RBF covariance function, with a log-normal prior distribution on

the RBF lengthscale parameter. For the current simulation study, MCMC converges after a burn-in of 4 million samples. SMC requires on average 65 SMC adaptation cycles with 2000 mutation steps within each cycle, and VI an average of 21,420 iterations. Notably, the elements of $\Sigma(x)$ converge relatively quickly, while the other model parameters such as the RBF lengthscale (see Appendix D Figure A2) take much longer to reach convergence.

In Figure 3A, we show the estimated dynamic covariance for one exemplar simulation run. The corresponding parameter estimates (here, the RBF lengthscale and the elements of the scale matrix \mathbf{V}) are shown in Figure 3B, together with the ground truth values. MCMC and SMC both infer a distribution over the model parameters, which are visualised using kernel density estimation [50]. Variational inference learns point-estimates of the model parameters, shown as a vertical line. These results are quantified using the performance measures described in Section 3.3, and shown in Table 1.

Table 1. The accuracy of each inference method in capturing the ground truth covariance process (MSE_{Σ}), lengthscale ($\text{MSE}_{\ell_{\text{RBF}}}$), and scale matrix $\text{MSE}_{\mathbf{V}}$ from Simulation study 1. We evaluate the mean covariance process estimate, as well as its full posterior distribution ($\text{MSE}_{\text{samples}}$). The computation time is shown in minutes (for MCMC per chain and for VI per initialisation). The mean and standard deviation over ten datasets are shown, with the best scores shown in bold.

Method	MSE_{Σ}	$\text{MSE}_{\ell_{\text{RBF}}}$	$\text{MSE}_{\mathbf{V}}$	$\text{MSE}_{\text{samples}}$	Runtime (Minutes)
MCMC	0.45 ± 0.36	0.00 ± 0.00	0.28 ± 0.26	0.59 ± 0.51	308.29 ± 7.13
VI	0.44 ± 0.29	0.10 ± 0.05	0.55 ± 0.36	0.51 ± 0.37	54.84 ± 0.82
SMC	0.45 ± 0.35	0.00 ± 0.00	0.31 ± 0.28	0.61 ± 0.52	105.40 ± 2.08

The performance measures, along with the runtime of each inference method in Table 1 indicate that all three approaches are successful in recovering the mean of the ground truth covariance process. When we look at the accuracy of estimating the latent model parameters, we see that these are recovered considerably less well by VI. In particular, MCMC and SMC both outperform VI when we look at the accuracy of inferring the RBF lengthscale and scale matrix. In other words, while variational inference reaches convergence most quickly, this comes at the cost of accurately estimating the latent model parameters, even though all three methods have converged and we took the VI result with the highest ELBO out of four re-runs (see Section 3.1). We investigate the effect of the difference in model parameter estimation in the next simulation study.

3.5. Simulation Study 2: State Switching and Out-of-Sample Prediction

In the previous experiment, we found that all three inference methods are able to accurately estimate the covariance process, but that, unlike MCMC and SMC, variational inference did not recover the latent model parameters well. To explore how this affects the ability to make out-of-sample predictions, we use the covariance process that follows a state-switching pattern (as described in Section 3.2) to make out-of-sample predictions, and use both a Periodic and a Locally Periodic (LP) covariance function to model this covariance process. Both functions should be able to capture the state switches that are present in the ground truth; however, the LP covariance function allows for more flexibility. We set a log-normal prior on all three parameters. In contrast to the RBF covariance function that was used before, which only has a single parameter, these covariance functions have two and three parameters, respectively, and each of these have an important impact on out-of-sample extrapolation. For example, if the period p is estimated poorly, the further away from training data we are, the more out of phase our predictions will be. Therefore, a correct inference of these parameters is crucial, but this is made even more difficult due to multimodality in the posterior distribution.

Recall that we use the first $n_{\text{train}} = 300$ observations for inference, and the remaining $n_{\text{test}} = 300$ observations for out-of-sample prediction. Figure 4 shows an example of an estimate and out-of-sample prediction of the three inference methods using a Periodic covariance function. Upon visual inspection, we observe that all three methods were able

to estimate the periodic ground truth covariance process for the training data, although there are differences in performance. From the MSEs over the mean covariance estimate (MSE_{Σ}) in Table 2, we can see that the estimates using SMC sampling outperform those made by variational inference and MCMC sampling. When looking at the MSE computed over the individual samples (MSE_{samples}), we observe the same pattern. On the training data, SMC sampling is most accurate in estimating the mean covariance process, and the distribution over this covariance process. Moreover, for MCMC and variational inference, we learn that the estimates using an LP covariance function were slightly more accurate than those using a Periodic covariance function.

Table 2. Using either a periodic or Locally Periodic covariance function, we show the accuracy of each inference method in capturing the mean ground truth covariance process and its full distribution. Moreover, for the out-of-sample predictions, we present the average fit to the observations and the fit of the full predictive posterior distribution. Finally, the computation time (for MCMC per chain and for VI per initialisation) is shown in minutes. The mean and standard deviation over ten datasets are shown, with the best scores shown in bold.

Method	Training		Out-of-Sample Prediction				Runtime (Minutes)
	MSE_{Σ}	MSE_{samples}	MSE_{Σ}	MSE_{samples}	LL	KL	
<i>Periodic function</i>							
MCMC	0.08 ± 0.02	0.13 ± 0.01	0.19 ± 0.02	0.34 ± 0.07	−4.10 ± 0.07	0.52 ± 0.07	280.88 ± 4.13
VI	0.06 ± 0.02	0.08 ± 0.02	0.14 ± 0.09	0.22 ± 0.15	−4.01 ± 0.16	0.47 ± 0.20	30.72 ± 7.19
SMC	0.04 ± 0.01	0.06 ± 0.01	0.11 ± 0.08	0.13 ± 0.09	−3.94 ± 0.21	0.48 ± 0.40	152.49 ± 2.11
<i>LP function</i>							
MCMC	0.05 ± 0.01	0.08 ± 0.01	0.12 ± 0.04	0.24 ± 0.08	−3.98 ± 0.13	0.42 ± 0.10	291.20 ± 3.97
VI	0.05 ± 0.01	0.08 ± 0.02	0.10 ± 0.04	0.17 ± 0.08	−3.97 ± 0.15	0.43 ± 0.14	31.62 ± 6.43
SMC	0.04 ± 0.01	0.06 ± 0.01	0.09 ± 0.02	0.19 ± 0.04	−3.93 ± 0.08	0.36 ± 0.06	154.98 ± 2.97

Although the differences in performance between MCMC, variational inference, and SMC on the training data are small, they become more pronounced when looking at the predictive performance. By looking at the out-of-sample predictions in Figure 4, which are shown after the vertical dotted line, we observe that variational inference did not capture the periodicity in the ground truth for this set of observations. A few more estimates can be seen in Appendix C. Over all sets of observations, variational inference captures the periodicity in 5/10 datasets using a Periodic covariance function, and 9/10 datasets using an LP covariance function. This is supported by the model parameter estimates in Figure 4, where it can be seen that the periodicity of approximately 0.33 is not correctly inferred by variational inference. In this case, only the lengthscale parameter has learned the structure of the training data. When we look at the results of Gibbs MCMC, we see that this method too has difficulties in capturing the latent periodicity in the covariance process, especially when using the Periodic covariance function. SMC sampling gives more consistent results. Moreover, the results in Table 2 support the differences we observe in Figure 4, namely that the out-of-sample predictions by SMC more accurately fit the true covariance than MCMC and variational inference. Although the differences in performance between SMC and variational inference seem small when using the LP covariance function, it should be noted that we trained the model four times using different initialisations for variational inference, and selected the estimates with the best ELBO.

Next, we compare the fit of the covariance estimates to the actual observations. We compare this using two metrics: the fit of the mean covariance process estimates is evaluated using the log-likelihood (LL), and the predictive posterior distribution is evaluated using the KL-divergence (KL) between this distribution and the actual distribution over the observations (see Section 3.3). Overall, the LL and KL-divergence results in Table 2 reveal that the differences when looking at the fit to the observations are less pronounced, although the SMC algorithm results in the best fit to the test data overall.

Finally, we can compare the three inference methods in terms of computational efficiency. When using the Periodic covariance function, the covariance samples of MCMC have converged after a burn-in of 5,000,000 burn-in steps per chain, after which we collect the next 1,000,000 samples. We use a thinning of 1000 samples. However, with four chains, this means that we need to perform 24 million Gibbs cycles for all chains in total. This amount of Gibbs cycles remains the same when using the LP covariance function. It should be noted, however, that the results show that MCMC has still not captured the periodic structure accurately after 5000 burn-in steps. SMC does capture this periodicity more accurately, and requires, for both covariance functions, an average of 59 SMC adaptation cycles with 3000 mutation steps per cycle to converge. For 1000 particles, this means that we perform 177 million Gibbs cycles in total. However, unlike for the MCMC, the computations of the SMC mutation steps are parallelised across the particles, and therefore greatly speed up the algorithm. This means that only 177,000 steps are being performed sequentially, which is much less than the 6 million steps of MCMC. Finally, VI converges much faster. Namely, VI requires an average of 46,580 iterations per re-run until convergence of the ELBO, when using the Periodic covariance function. After convergence of the ELBO, the parameters no longer changed. Using four re-runs, this means we optimize the ELBO approximately 186,320 times. When using the LP covariance function, we require 44,700 iterations per re-run and 178,800 iterations in total. However, although variational inference runs faster, we have seen from the results in Table 2 that variational inference has difficulty in capturing the periodicity, and therefore in making out-of-sample predictions.

In short, these results suggest that SMC sampling can reliably capture the periodic structure present in the data by handling the highly correlated covariance function parameters, and therefore, SMC sampling can make more accurate out-of-sample predictions. From the results of Gibbs MCMC and variational inference, we can see that both methods have more difficulty with converging, and this has important consequences for out-of-sample predictions.

4. Empirical Application: Dynamic Correlations in Depression

In this section, we demonstrate how the generalised Wishart process can be used to study the dynamics of psychological processes, and how it enables novel analyses. Recently, there has been a paradigm shift in the study of mental disorders. Instead of defining mental disorders according to the sum score of a set of measurements, they are now increasingly conceptualised as (dynamic) networks of interacting symptoms [6,7,13,14]. For example, recent studies on the onset of depressive episodes in people with Major Depressive Disorder (MDD) have shown that changes in the dynamics between individual symptoms serve as early warning signs of various mental disorders [15,51]. These studies modelled these dynamic correlations using a multilevel vector autoregression method [52]. However, we propose to apply the Wishart process in this context, because this provides us with a distribution over the covariance, which we can use to test for dynamics. Furthermore this approach allows us to work with unevenly spaced data, therefore allowing us to estimate covariance not only as a function of time, but also as a function of some other input variable, such as medication dosage.

4.1. Dataset, Preprocessing, and Model Choices

We use the dataset from Kossakowski et al. [53], which is obtained from a single subject who has been diagnosed with MDD. The subject is a 57 year old male who monitored his mental state over the course of 237 days by filling in a questionnaire of daily life experiences several times a day. Moreover, the subject had been using venlafaxine, an antidepressant, for 8.5 years. Interestingly, during the data collection, the dosage of venlafaxine is gradually reduced to zero in a double-blinded manner according to five experimental phases: baseline (four weeks), before dosage reduction (between zero and six weeks, the exact timings unknown to the subject), during dose reduction (eight weeks), post-assessment (eight weeks), and a follow-up (twelve weeks). These phases are shown in Figure 5A, where we

can also see that the subject became more depressed over the course of the experiment, as measured on a weekly basis by the depression subscale of the Symptom Checklist-Revised (SCL-90-R) [54].

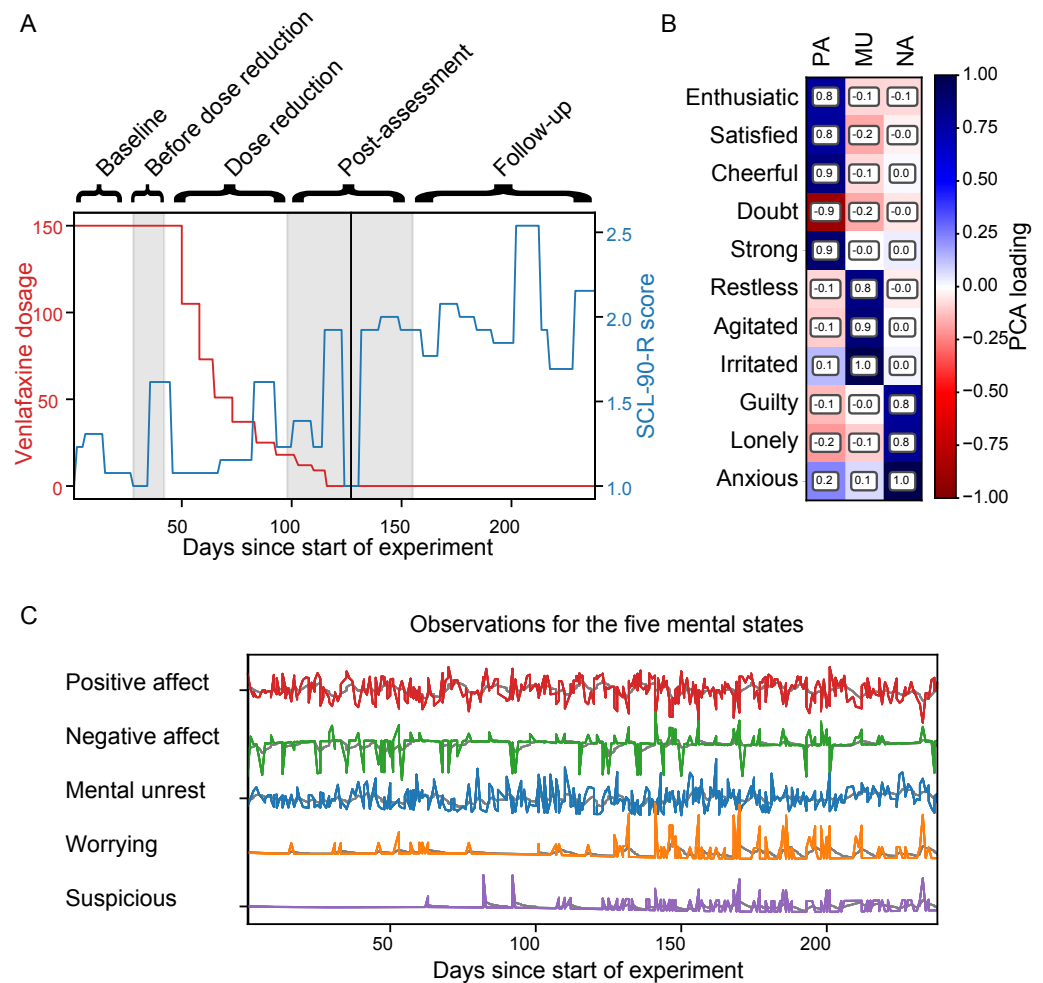


Figure 5. (A) Venlafaxine dosage (in red) and depression measurements (in blue), as measured by the SCL-90-R score over the course of the experiment. The five experimental phases (4 weeks baseline, 0–6 weeks before dose reduction, 8 weeks dose reduction, 8 weeks post-assessment, and 12 weeks follow-up) are shown by the shaded background, and the black vertical line indicates the moment at which the subject relapsed into a depressive episode. (B) The loadings from individual items to the three principal components positive affect (PA), negative affect (NA), and mental unrest (MU). (C) The time series data for the five mental states, together with their moving average (in grey).

Following Wichers et al. [15], we collect the following items from the questionnaires: ‘irritated’, ‘content’, ‘lonely’, ‘anxious’, ‘enthusiastic’, ‘cheerful’, ‘guilty’, ‘indecisive’, ‘strong’, ‘restless’, and ‘agitated’. Subsequently, these symptoms are summarised using principal component analysis together with an oblique rotation [15]. The loadings of each item on these components can be seen in Figure 5B. The components are interpreted as ‘positive affect’, ‘negative affect’, and ‘mental unrest’. Moreover, we use the items ‘worrying’ and ‘suspicious’ as separate variables, again following [15]. Finally, slow non-periodic time trends are removed from the data and, to speed up inference, only every fourth observation is kept. This results in a total of $n = 369$ and $d = 5$ variables (‘positive affect’, ‘negative affect’, ‘mental unrest’, ‘worrying’, and ‘suspicious’), as visualised in Figure 5C.

In order to model both slow and fast changes in covariance between the five mental states, we sum an RBF and a Matérn 1/2 covariance function:

$$\kappa_{\text{RBF} + \text{M12}}(x, x') = \exp\left(-\frac{(x - x')^2}{2\ell_{\text{RBF}}^2}\right) + \exp\left(-\frac{|x - x'|}{2\ell_{\text{M12}}^2}\right). \quad (14)$$

Moreover, to model the slow fluctuations of the level of these symptoms over time regardless of their interactions with other symptoms, we use an exponential moving average (EMA) mean function [55]:

$$\text{EMA}(y_{i+1,j}) = \alpha \left[y_{ij} + (1 - \alpha)y_{i-1,j} + (1 - \alpha)^2 y_{i-2,j} + \dots + (1 - \alpha)^{k-1} y_{i-(k-1),j} \right], \quad (15)$$

where $\alpha = 2/(k + 1)$ influences the smoothness of the mean. We set $k = 10$.

4.2. Hypothesis Test for Dynamic Covariance

With a Bayesian approach to modelling covariance processes, we obtain an estimate of the posterior distribution over the covariance process. The advantage of this is that, once we have inferred this distribution, we can perform a hypothesis test on the covariance process, allowing us to learn what type of covariance is present between a pair of variables. That is, two variables can either be (i) uncorrelated, when $0 \in p(\Sigma_{ij}(x) | \mathbf{x}, \mathbf{Y}), \forall x$; (ii) statically correlated (denoted by 'S'), if $\exists c \neq 0 \in \mathbb{R}$ such that $c \in p(\Sigma_{ij}(x) | \mathbf{x}, \mathbf{Y}), \forall x$; or (iii) dynamically correlated (denoted by 'D'), if $\exists c \notin \mathbb{R}$ such that $c \in p(\Sigma_{ij}(x) | \mathbf{x}, \mathbf{Y}), \forall x$.

To determine whether zero (or c) falls in the distribution $p(\Sigma_{ij}(x) | \mathbf{x}, \mathbf{Y})$, we determine the 95% highest density interval of this distribution and use that instead, since no curve will have a strictly zero posterior probability. Furthermore, covariance estimates that deviate less than 0.005 from zero (or c) are regarded as zero (or c) by using the region of practical equivalence (ROPE) principle [56]. This hypothesis test demonstrates an important advantage of estimating a distribution over the covariance process, instead of only estimating its mean. Namely, without such a distribution over the covariance process, we would not be able to use this approach to test for dynamics in the covariance. This is an important benefit of the probabilistic Wishart process compared to several common approaches, such as non-Bayesian implementations of the sliding-window method and the MGARCH model, which only provide a mean estimate of the covariance process.

4.3. Modelling of Dynamic Correlations between Mental States

We demonstrate how the Wishart process can be used in two different examples in studying MDD, as will be described below. In both experiments, we again sampled or optimised until convergence of all model parameters, as measured by the PSRF (see Section 3.1).

4.3.1. Dynamics between Mental States over Time

In our first experiment, we use the Wishart process to estimate the covariance between each pair of symptoms over the course of the experiment, that is, we use the day number since the onset of the experiment as our input variable. Since we also have weekly SCL-90-R scores available, this allows us to explore whether these covariances change when the subject relapses in a depressive episode, similar to the study by Wichers et al. [15]. We compare our estimates to those made by a DCC-GARCH model (implemented using the R package `rmgarch` [57]; more details on its implementation are provided in Appendix E).

To compare the performances of the Wishart process using MCMC, variational inference and SMC, and the DCC-GARCH model, we use the following 10-fold cross validation scheme. We split the dataset evenly into 10 subsets of training and testing observations. In the first subset, we train each model on the first 36 observations and then predict the next 10 data points. In the next subset, the first 72 observations are used for training and

we again predict the next 10 data points. This pattern is continued for 10-folds. We report the test log-likelihood averaged over all folds.

When modelling the covariance process as a function of time, the convergence of all parameters took 5,000,000 MCMC burn-in steps, with a thinning of 1000, 61 SMC cycles with 5000 mutation steps per cycle, and 21,710 VI optimisation iterations. The estimates of the final fold of the Wishart process, using SMC sampling, and the DCC-GARCH model are shown in Figure 6. Based on visual inspection, the covariance estimates of the methods are qualitatively in agreement. Both the Wishart process and the DCC-GARCH model estimate positive affect to be negatively correlated with all other variables, and negative affect to be positively correlated with all variables except for positive affect. Moreover, previous studies [15,51] have found that the covariance between different mental states tends to become stronger with the increase in depressive score. As we saw in Figure 5A, the subject relapsed into a depressive episode over time, which, together with the different experimental phases of the venlafaxine dosage reduction, is also indicated in Figure 6. The estimates show that, overall, the covariances between the different symptoms increase in strength as the subject becomes more depressed, which is in line with the findings by Wichers et al. [15]. In particular, we observe increased covariance between worrying and negative affect, mental unrest, and suspicion, and a decreased covariance between positive affect and worrying, and positive affect and suspicion. Finally, to determine if this covariance process is dynamic or static, we apply the in Section 4.2 described hypothesis tests for dynamic covariance on the posterior distribution estimated by the Wishart process. These results suggest that all covariance pairs except for those between ‘positive affect’, ‘negative affect’, and ‘mental unrest’ are dynamic. By looking at the covariance function parameters inferred by SMC, also visualised in Figure 6, we observe that these dynamics in covariance are relatively slow, since the distribution over the lengthscales for the Matérn 1/2 function are large compared to the input range.

In Table 3, we show the performance of the Wishart process, inferred with MCMC, variational inference or SMC, and the DCC-GARCH model for this experiment. As described in Section 4.1, we evaluated each method on a 10-fold setup, where within each fold we computed the test log-likelihood over the next 10 test observations. The results show the average and standard deviation over these 10-folds, indicating that the Wishart process with MCMC or SMC and the DCC-GARCH model outperformed variational inference. Even though we again selected the variational inference result with the largest ELBO, the standard deviation of this method is relatively large, indicating that this method is not robust over different initialisations.

Table 3. For both the DCC-GARCH model and the Wishart process using MCMC, variational inference, or SMC, we present the mean and standard deviation of the fit to the test observations by means of the log likelihood.

	DCC-GARCH	Wishart Process		
		MCMC	VI	SMC
LL _{test}	-6.19 ± 2.75	-5.29 ± 2.85	-7.29 ± 4.21	-5.82 ± 3.39

4.3.2. Dynamics between Mental States as a Function of Venlafaxine Dosage

In our second application, we use venlafaxine dosage as an input variable to model the covariance between each pair of symptoms. Unlike time, venlafaxine dosage is an unevenly spaced variable. Therefore, we can only use the Wishart process in this scenario, as the DCC-GARCH model is unable to handle unevenly spaced input data.

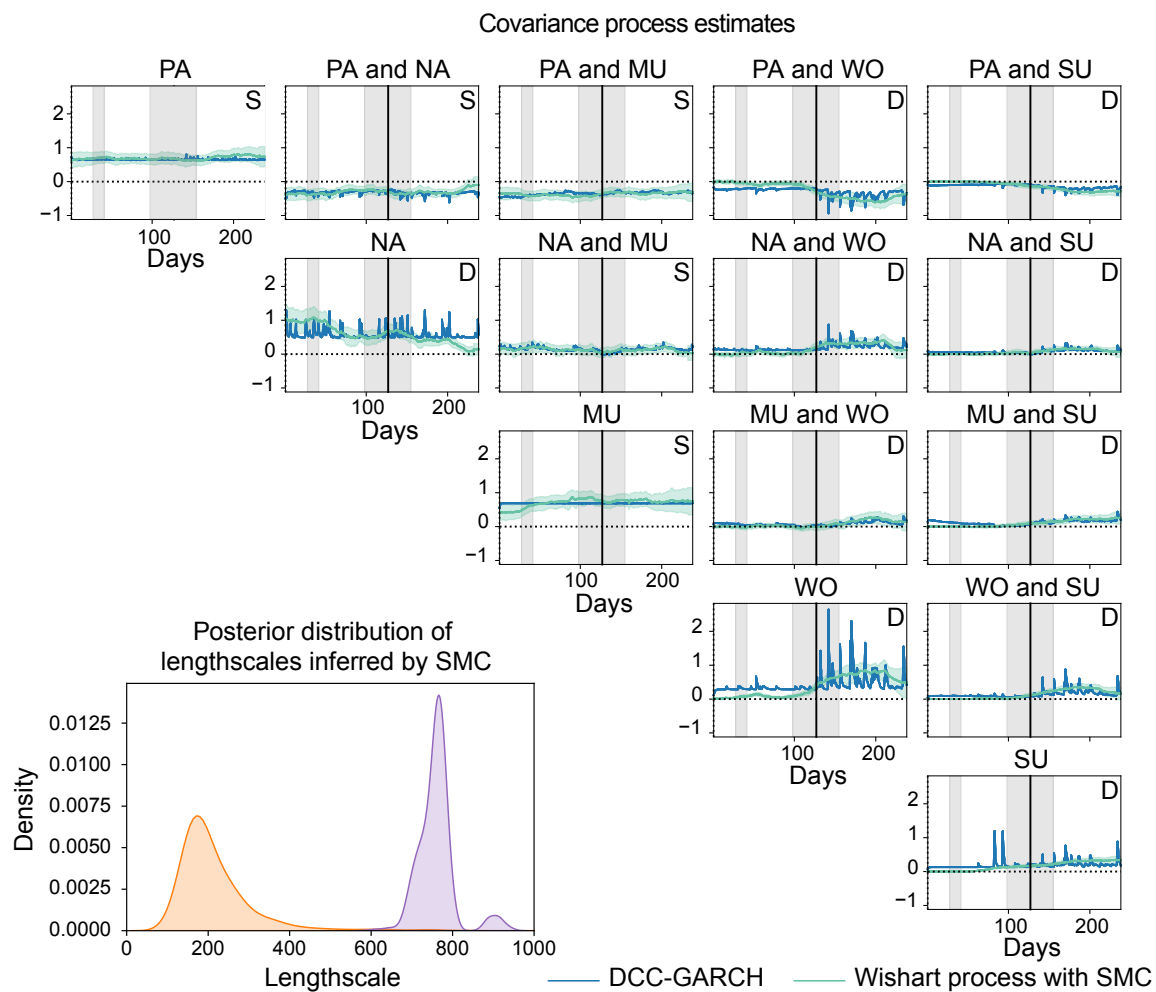


Figure 6. Estimates of the RBF (in orange) and Matérn 1/2 (in purple) lengthscale parameters and the covariances between the five different mental states (PA = positive affect, NA = negative affect, MU = mental unrest, WO = worrying, and SU = suspicious) as a function of the day number. The vertical black lines indicate the day on which the subject relapsed into depression, and the different background shades indicate different phases of the antidepressant dose reduction scheme. We show the estimates and out-of-sample predictions of the final fold. For the Wishart process estimates, we test for dynamic covariance (D) or static covariance (S).

When modelling the covariance process as a function of venlafaxine dosage, the convergence of all model parameters required 84 SMC cycles with 5000 mutation steps per cycle. The estimates by SMC are shown in Figure 7. Similarly to the previous demonstration, only every fourth observation is kept, resulting again in $n = 369$ observations. The Wishart process can handle unevenly spaced input data; therefore, this model allows us to directly model the covariance process based on the dosage reduction scheme instead of indirectly over time. This analysis is not possible for the DCC-GARCH model because this model requires evenly spaced input data. The resulting estimates by the Wishart process with SMC are shown in Figure 7, where the input locations at which observations were measured are indicated by the red dots. As expected, these estimates are in agreement with those of Figure 6, since the dosage was reduced over the course of the experiment. When the dosage is low, the covariances between the different mental states are generally stronger than when the dosage has not been reduced yet.

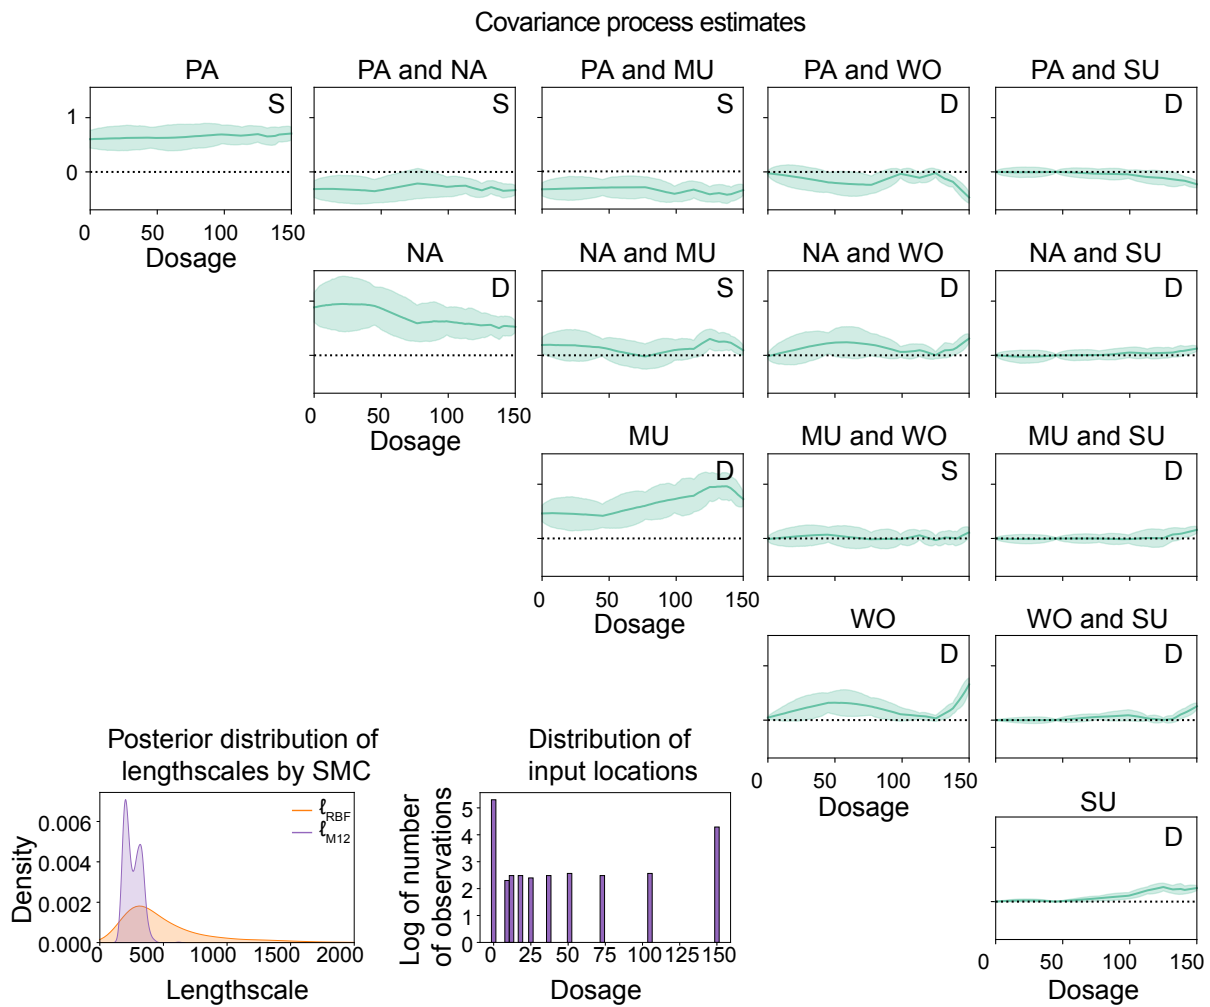


Figure 7. Estimates of the RBF and Matérn 1/2 lengthscales parameters and the covariances between the five different mental states (PA = positive affect; NA = negative affect; MU = mental unrest; WO = worrying, and SU = suspicious) as a function of antidepressant dosage. The bar plot indicates the amount of observations available at each input location, and we tested for dynamic (D) and static (S) covariances.

4.3.3. Differences between Dynamics in Mental State Correlations over Time and Dosage

Finally, we study whether there are any difference in the amount of dynamics estimated as a function of either time or dosage. We evaluate this for every pair of symptoms and every posterior sample by computing the distance between the minimum and maximum estimated covariances. The resulting distributions of these distances are visualised using kernel density estimation in Figure 8. This figure shows that, overall, these covariance process intervals are larger when using time as an input variable than when estimating covariance as a function of antidepressant dosage. Hence, these results imply that the effect of time on dynamic covariance is larger than the effect of dosage on dynamic covariance. This might mean that, apart from antidepressant reduction, there are other factors affecting the dynamics between the five mental states.

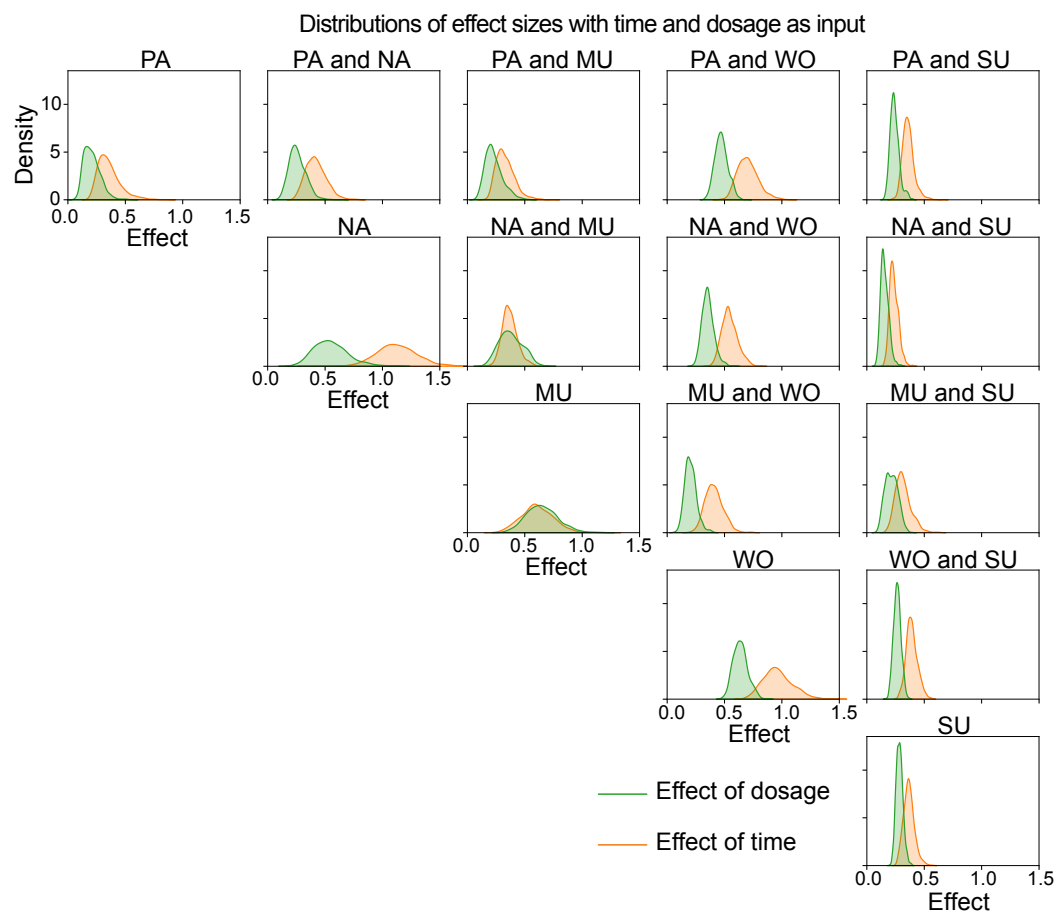


Figure 8. For each combination of mental states (PA = positive affect; NA = negative affect; MU = mental unrest; WO = worrying; and SU = suspicious), and using either time or dosage as predictor variable, the distribution over the difference between the minimum and maximum of the covariance process is shown.

5. Discussion

Across different research fields, there is substantial interest in dynamically modelling the joint behaviour of multiple time series instead of statically. Although the Wishart process is ideally suited for this task, inference of the model parameters is challenging. Wilson and Ghahramani [24] used MCMC to infer the posterior distribution of the Wishart process. Although their study showed an improved performance of modelling dynamic covariance compared to an MGARCH model, this approach is not scalable to larger numbers of observations and variables. Alternatively, Heaukulani and van der Wilk [31] inferred the posterior distribution with variational inference, which scales to larger numbers of observations and variables. However, in our experiments, we found that variational inference did not accurately learn the covariance function parameters, which negatively affected out-of-sample predictions of the covariance process. Moreover, both MCMC and variational inference did not give robust estimates when using composite covariance functions with multiple parameters. A similar problem has been observed in Gaussian process regression, where covariance function hyperparameters can be difficult to identify correctly [58,59]. In an attempt to overcome this limitation, Svensson et al. [60] has demonstrated that Sequential Monte Carlo (SMC) can robustly marginalize over the hyperparameters.

The SMC algorithm approximates the posterior distribution via a large amount of model parameter sets (called particles), which are initialised from their prior, and iteratively updated, weighted, and resampled based on their fit to the observations. Unlike MCMC samplers, which are likely to explore high-density areas of the posterior, the different particles of SMC can cover different modes of a distribution. Therefore, SMC is more

capable of dealing with multimodal distributions than MCMC and variational inference. Additionally, unlike for MCMC, which is inherently sequential, large parts of the SMC algorithm can be performed in parallel, therefore speeding up the inference procedure. Since the Wishart process is constructed from GP elements, we hypothesised that inference of the Wishart process would benefit from SMC as well. Therefore, in this work, we proposed to use SMC for inference of the Wishart process parameters.

We showed that the Wishart process combined with SMC sampling indeed offers a robust approach to modelling dynamic covariance from time series data. In two simulation studies where the true covariance process was known, we found that the SMC covariance process estimates were more robust than those inferred by MCMC and variational inference, since, unlike for SMC, the results of MCMC and variational inference varied across different runs. This became especially pronounced in Simulation study 2 (see Section 3.5), where MCMC and variational inference had difficulties in inferring the covariance function hyperparameters when using composite covariance functions. Moreover, our results demonstrated that this has important consequences when we want to use the model parameters to make out-of-sample predictions.

In general, the Wishart process is a flexible model for estimating dynamic covariance, and our work provides a key contribution in making its estimates reliable. The Wishart process allows us to model dynamic covariance over unevenly spaced data, unlike common dynamic covariance methods such as the sliding-window approach and MGARCH models. Additionally, a Bayesian approach allows for a principled approach to test for dynamics in covariance on a single subject using hypothesis tests. Although it is feasible to obtain a distribution over the covariance using MGARCH as well, such as through a bootstrapping approach, this is not a straightforward approach.

The Wishart process is applicable in various fields, such as finance, where it can be used to study the interactions between stock markets [9–12]; neuroscience, where the interactions between brain regions are studied [1,27–29]; or biological systems, such as the study of social influence within animal groups [61]. In our work, we demonstrated another application in biological systems, namely in psychology, where the Wishart process was used to study the covariances between mental states. With this application, we demonstrated the unique ability of the Wishart process to study covariances over unevenly spaced input variables such as medication dosage. By comparing the change in covariance across different predictors, this allows researchers to infer which predictors are the strongest drivers of changes in covariance. This is not possible with common implementations of existing approaches such as the sliding-window method or MGARCH models and therefore opens up many new possibilities for research questions across many fields. A specific example is the question whether the interactions between brain systems are mainly shaped by the developmental age of children or by other factors such as cognitive development or environmental influences [62].

There are several important aspects that deserve consideration. First of all, the latent GPs in the Wishart process are non-identifiable, since different permutations of latent GPs can result in identical covariance process estimates. However, the non-identifiability of the latent GPs does not affect the covariance processes, and therefore the comparison between different inference methods. Additionally, depending on the GP covariance function, the parameters of the covariance function can be correlated, potentially causing multimodality in the posterior distribution [58,60]. As demonstrated in our simulation studies, SMC was able to sample the covariance function parameters more robustly when these are correlated. Moreover, although we found that SMC was able to infer the posterior more efficiently than MCMC while using a Metropolis sampler, more efficient samplers within the Gibbs cycle might be helpful, such as the Metropolis-adjusted Langevin algorithm [63] or a Hamiltonian Monte Carlo algorithm [64]. Moreover, although the Wishart process is a flexible model, the covariance process estimates depend on the choice of covariance function for the GP. For example, the RBF covariance function will result in relatively smooth covariance process estimates compared to the estimates when using a Matérn 1/2

function. Selecting an appropriate covariance function might be challenging when the expected dynamics in covariance are unknown. In these situations, a potential solution would be to learn the covariance function from the data, for example, by using the approach by Wilson and Adams [65]. Moreover, our approach with SMC currently does not scale to large numbers of observations and variables. This is because the mutation steps within SMC are parallelised, therefore storing large matrices in memory for all particles simultaneously. To improve scalability of the Wishart process with SMC, the current implementation can be augmented in several ways, of which we will provide a few suggestions here. In order to scale to larger numbers of observations, we could make use of inducing points [66], which were already implemented for the Wishart process using variational inference [31]. The challenge of this approach is the trade-off between scalability and the precision of the estimates, as the performance decreases when using fewer inducing points. Another way would be to use a factored variant of the Wishart process, where a mapping from a small number of latent variables to a larger number of observed variables is learned [31,67]. This approach works particularly well when there is a shared underlying dynamic covariance structure among a larger set of variables. Finally, when we work with evenly spaced data, we could make use of the Toeplitz structure to improve scalability [68,69].

Currently, the Wishart process models dynamic covariance, which contains both direct as well as indirect interactions between variables. However, in certain domains, such as neuroscience, it might be more relevant to study direct interactions only, via sparse partial correlations [70,71]. Zero elements in a partial correlation matrix imply that there is no direct interaction between a pair of variables. Therefore, sparse partial correlation matrices can offer researchers valuable insights. A previous study [72] proposed to model sparse partial correlations using a G-Wishart distribution, which is, similarly to the Wishart distribution itself, a distribution over positive-definite symmetric matrices, but it also ensures that there are zero elements in the partial correlation matrices when there are no direct interactions between two variables. Extending the Wishart process to model sparse partial correlation matrices would be an interesting direction for future work.

In conclusion, by inferring the Wishart process using Sequential Monte Carlo sampling, we can robustly estimate and out-of-sample predict dynamic covariance processes. The benefits relative to existing approaches become especially pronounced when composite covariance functions are used, where multiple modes of likely parameter combinations exist. Additionally, we showed how the distribution over the covariance can be used to test for dynamics, and how dynamic covariance can be modelled as a function of an unevenly spaced input. The combination of Wishart processes and SMC can be used to answer research questions related to dynamic covariance in different domains, such as psychology, finance, and neuroscience.

Author Contributions: Conceptualization, H.H., D.L., L.G. and M.H.; Methodology, H.H.; Software, H.H., D.L. and M.H.; Validation, H.H.; Formal analysis, H.H.; Writing—original draft, H.H.; Writing—review & editing, D.L., L.G. and M.H.; Visualization, H.H.; Supervision, L.G. and M.H.; Funding acquisition, M.H. All authors have read and agreed to the published version of the manuscript.

Funding: L.G. was supported by a Vidi grant from the Dutch Research Council (NWO, VI.Vidi.201.150).

Institutional Review Board Statement: Not applicable.

Data Availability Statement: The clinical depression data is publicly available from <https://osf.io/j4fg8/> (accessed on 20 May 2024).

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. Markov Chain Monte Carlo for Wishart Processes

Wilson and Ghahramani [24] infer the posterior distribution of the Wishart process $p(\Sigma(x_i) | \mathbf{x}, \mathbf{Y})$ using MCMC sampling with the following Gibbs cycle [40]:

$$\begin{aligned} p(\mathbf{F} | \theta, \mathbf{L}, \mathbf{x}, \mathbf{Y}) &\propto p(\mathbf{Y} | \mathbf{F}, \mathbf{L}, \mathbf{x})p(\mathbf{F} | \mathbf{x}, \theta) \\ p(\theta | \mathbf{F}, \mathbf{L}, \mathbf{x}, \mathbf{Y}) &\propto p(\mathbf{F} | \mathbf{x}, \theta)p(\theta) \\ p(\mathbf{L} | \theta, \mathbf{F}, \mathbf{x}, \mathbf{Y}) &\propto p(\mathbf{Y} | \mathbf{F}, \mathbf{L}, \mathbf{x})p(\mathbf{L}) \end{aligned} \quad (\text{A1})$$

Hence, the distributions over the latent model parameters are updated in each MCMC iteration by combining the multivariate normal likelihood $p(\mathbf{Y} | \mathbf{F}, \mathbf{L}, \mathbf{x})$ with the prior distributions over the latent GPs $p(\mathbf{F} | \mathbf{x}, \theta)$, covariance function parameters $p(\theta)$, or lower Cholesky decomposition of the scale matrix $p(\mathbf{L})$. The distribution $p(\mathbf{F} | \mathbf{x}, \theta)$ is a multivariate normal prior on the latent GPs. To sample θ and \mathbf{L} , we use a Random Walk Metropolis-Hastings sampler. The latent GPs are sampled using elliptical slice sampling [73], since the GPs are highly correlated. Elliptical slice sampling does not require any input parameters and is efficient because new proposals are always accepted. Moreover, the elliptical slice sampler can benefit from the Kronecker structure of $p(\mathbf{F} | \mathbf{x}, \theta)$ by sampling \mathbf{F} for every block independently, and then combining the resulting samples. In other words, we can independently sample $d \times v$ GPs, and then construct \mathbf{F} . The computational benefit is that, instead of having to invert an $ndv \times ndv$ matrix, we now only have to invert dv matrices that are $n \times n$, which is computationally much faster.

Appendix B. Variational Wishart Processes

Heaukulani and van der Wilk [31] approximate the posterior of the Wishart process $p(\Sigma(x_i) | \mathbf{x}, \mathbf{Y})$ by means of variational inference, a method that uses optimisation instead of sampling. First, a variational distribution $q(\mathbf{F} | \phi)$ is introduced over the dv GP samples \mathbf{F} . The variational parameters ϕ are optimised so that the distribution q is close to the true posterior, with the distance between two distributions being measured by the Kullback–Leibler divergence [74], denoted as $\text{KL}(q \| p)$. Ideally, the form of $q(\mathbf{F} | \phi)$ should capture the shape of the posterior, and is often chosen to follow a multivariate normal distribution with each GP sample having mean \mathbf{m} and covariance matrix \mathbf{S} , such that

$$q(\mathbf{f}_{jl} | \phi) = \mathcal{MVN}_n(\mathbf{m}_{jl}, \mathbf{S}_{jl}) \quad , \quad j = 1, \dots, d \quad , \quad l = 1, \dots, v \quad . \quad (\text{A2})$$

The variational parameters \mathbf{m}_{jl} and \mathbf{S}_{jl} , as well as the Wishart process parameters \mathbf{L} and θ , are optimised by iteratively maximising the evidence lower bound (ELBO) using gradient descent:

$$\text{ELBO} = \sum_{i=1}^n \mathbb{E}_{q(\mathbf{F}(x_i))} [\log p(\mathbf{y}_i | \mathbf{F}(x_i))] - \sum_{j=1}^d \sum_{l=1}^v \text{KL}(q(\mathbf{f}_{jl}) \| p(\mathbf{f}_{jl})) \quad , \quad (\text{A3})$$

where $\mathbf{F}(x_i)$ is the $d \times v$ matrix of latent GPs evaluated at x_i . The first part of the ELBO represents the model fit, computed as the expectation of the log-likelihood, using the variational distribution. The second part of the ELBO pushes the variational distribution closer to the prior distribution via the negative KL-divergence. Maximising the ELBO means that we try to maximise the likelihood, while also keeping the variational distribution close to the prior of \mathbf{F} (via the KL-divergence). After having approximated the variational parameters of the distribution $q(\mathbf{F} | \phi)$, predictions at test locations \mathbf{x}_* can be made by first sampling latent GPs at \mathbf{x}_* from $p(\mathbf{F}_*)$ using Equation (7) and then constructing the covariance process $\Sigma(x_*)$.

Additionally, Heaukulani and van der Wilk [31] show that the parameter estimation is improved by adding an additional noise term in the construction of the covariance matrix

(see Equation (6)). The noise term is a diagonal matrix $\mathbf{\Lambda} \in \mathbb{R}^{d \times d}$. $\mathbf{\Sigma}(x_i)$ is now constructed as follows:

$$\mathbf{\Sigma}(x_i) = \sum_{l=1}^v \mathbf{L} \mathbf{f}_l(x_i) \mathbf{f}_l(x_i)^\top \mathbf{L}^\top + \mathbf{\Lambda}, \quad i = 1, \dots, n. \tag{A4}$$

This regularisation term can also be optimised similar to the other parameters, and improves the approximation of the gradients required in optimisation of the ELBO.

Appendix C. Covariance Process Estimates of the Second Simulation Study

For the second simulation study, we provide the covariance process and covariance function parameter estimates based on only one dataset and covariance function. To illustrate the variety of estimates by all three inference methods, we show three more covariance process estimates in Figure A1, based on different sets of observations, using either a Periodic or Locally Periodic covariance function.

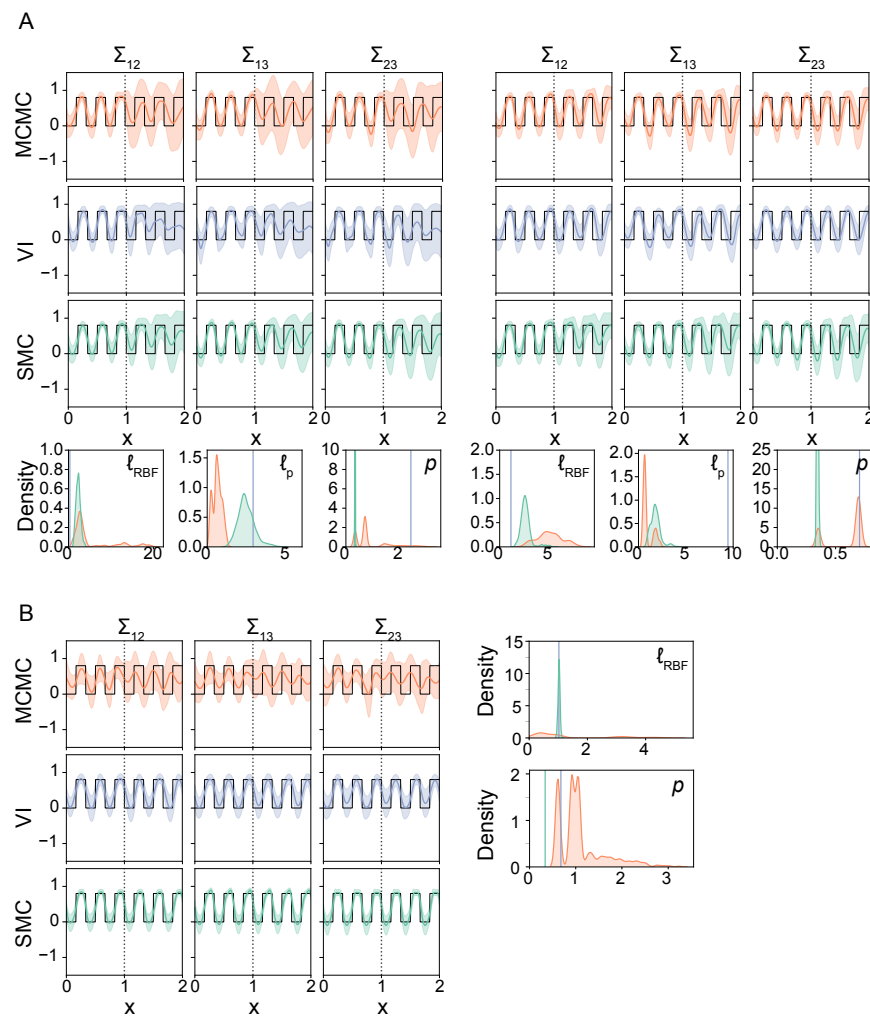


Figure A1. Covariance process estimates and out-of-sample predictions for each inference method, together with the corresponding distributions over the covariance function parameters. The true covariance process is shown in black, and out-of-sample predictions are shown after the grey dotted line. (A) The covariance process is modelled using a Periodic covariance function, which has two parameters: the period (p) and the lengthscale within each period (l_{RBF}). (B) Based on the same observations as in A, the covariance process is modelled using a Locally Periodic covariance function (see Equation (13)), which has three parameters: the period (p), the lengthscale within each period (l_{RBF}) and the lengthscale between periods (l_p).

Appendix D. Convergence of the Covariance Process Samples

We have measured the convergence of all Wishart process parameters by means of the potential scale reduction factor between the posterior distributions resulting from different random initialisations of the model parameters. Here, we found that, although the samples of the covariance process estimates converged relatively fast, convergence of the covariance function parameters and scale matrix required more burn-in steps, optimisation steps, or mutation steps, as shown for the data from Simulation study 1 in Figure A2.




Figure A2. Convergence of the covariances, covariance function hyperparameters, and scale matrix from Section 3.4. The convergence was measured using the Potential Scale Reduction Factor (PSRF), and averaged over the elements (when applicable). Since variational inference learns point estimates of the model parameters, we only show the convergence of the covariance for variational inference. In general, convergence of the covariance requires much less burn-in steps, mutation steps, or iterations than convergence of the scale matrix and covariance function hyperparameters.

Appendix E. Multivariate Generalised Autoregressive Conditional Heteroscedastic Models

Multivariate generalised autoregressive conditional heteroscedastic (multivariate GARCH) models [17,75] are an often used approach in finance [18,19]. Similar to how autoregressive moving average (ARMA) models [76] assume that the observations follow a Gaussian distribution and estimate current observations based on past observations and past residuals (or error terms), univariate GARCH models [16] estimate the variance of a residual as a function of the past variances in residuals and the past residuals itself. Multivariate GARCH models estimate both the variance of a variable itself and the covariance between pairs of variables. There are several well-known versions of multivariate GARCH models.

The Dynamic Conditional Correlation (DCC-GARCH) model by Engle [77] is a variant of the multivariate GARCH model that estimates covariance from a non-linear combination of univariate GARCH models. Namely, the DCC-GARCH model specifies a univariate GARCH model $\mathbf{h}_{jj} \in \mathbb{R}^n$ (for $j = 1, \dots, d$) for every variable:

$$\mathbf{H}_i = \text{diag}\left(h_{jji}^{1/2}, \dots, h_{ddi}^{1/2}\right), \quad i = 1, \dots, n. \tag{A5}$$

The univariate estimates are first transformed as $u_{ij} = y_{ij} / \sqrt{h_{jji}}$ and then used to construct the symmetric positive definite matrix $\mathbf{Q}_i \in \mathbb{R}^{d \times d}$:

$$\mathbf{Q}_i = (1 - \alpha - \beta)\bar{\mathbf{Q}} + \alpha \mathbf{u}_{i-1} \mathbf{u}_{i-1}^T + \beta \mathbf{Q}_{i-1}, \quad i = 1, \dots, n, \tag{A6}$$

where α and β are non-negative scalar parameters such that $\alpha + \beta < 1$, and $\bar{\mathbf{Q}} \in \mathbb{R}^{d \times d}$ is the unconditional variance matrix of \mathbf{u}_i ; it contains the variances of \mathbf{u}_i independent of \mathbf{u}_{i-1} . From matrix \mathbf{Q}_j , we now construct $\mathbf{R} \in \mathbb{R}^{n \times d \times d}$:

$$\mathbf{R}_i = \text{diag}\left(q_{jji}^{1/2}, \dots, q_{ddi}^{1/2}\right) \mathbf{Q}_i \text{diag}\left(q_{jji}^{1/2}, \dots, q_{ddi}^{1/2}\right), \quad i = 1, \dots, n. \tag{A7}$$

Finally, the covariance estimates are constructed by combining \mathbf{H} and \mathbf{R} :

$$\Sigma_i = \mathbf{H}_i \mathbf{R}_i \mathbf{H}_i . \quad (\text{A8})$$

Similarly to the other multivariate GARCH variants, the DCC-GARCH model combines the outer product of the previous observations (via \mathbf{u}_j) and the previous covariances (via \mathbf{Q}_{j-1}) to estimate the covariance at input location j . The DCC-GARCH variant has $(d+1) \times (d+4)/2$ parameters, which is much less than the $(p+q)(d(d+1)/2)^2 + d(d+1)/2$ parameters of the original multivariate GARCH model. This makes the DCC-GARCH variant less likely to overfit.

References

- Lurie, D.J.; Kessler, D.; Bassett, D.S.; Betzel, R.F.; Breakspear, M.; Kheilholz, S.; Kucyi, A.; Liégeois, R.; Lindquist, M.A.; McIntosh, A.R.; et al. Questions and controversies in the study of time-varying functional connectivity in resting fMRI. *Netw. Neurosci.* **2020**, *4*, 30–69. [[CrossRef](#)] [[PubMed](#)]
- Calhoun, V.D.; Miller, R.; Pearlson, G.; Adalı, T. The chronnectome: Time-varying connectivity networks as the next frontier in fMRI data discovery. *Neuron* **2014**, *84*, 262–274. [[CrossRef](#)]
- Fornito, A.; Bullmore, E.T. Connectomics: A new paradigm for understanding brain disease. *Eur. Neuropsychopharmacol.* **2015**, *25*, 733–748. [[CrossRef](#)] [[PubMed](#)]
- Ledoit, O.; Wolf, M. Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *J. Empir. Financ.* **2003**, *10*, 603–621. [[CrossRef](#)]
- Borsboom, D. A network theory of mental disorders. *World Psychiatry* **2017**, *16*, 5–13. [[CrossRef](#)] [[PubMed](#)]
- Cramer, A.O.; Waldorp, L.J.; Van Der Maas, H.L.; Borsboom, D. Comorbidity: A network perspective. *Behav. Brain Sci.* **2010**, *33*, 137–150. [[CrossRef](#)] [[PubMed](#)]
- Schmittmann, V.D.; Cramer, A.O.; Waldorp, L.J.; Epskamp, S.; Kievit, R.A.; Borsboom, D. Deconstructing the construct: A network perspective on psychological phenomena. *New Ideas Psychol.* **2013**, *31*, 43–53. [[CrossRef](#)]
- Liégeois, R.; Li, J.; Kong, R.; Orban, C.; Van De Ville, D.; Ge, T.; Sabuncu, M.R.; Yeo, B.T. Resting brain dynamics at different timescales capture distinct aspects of human behavior. *Nat. Commun.* **2019**, *10*, 2317. [[CrossRef](#)] [[PubMed](#)]
- Chen, M.; Li, N.; Zheng, L.; Huang, D.; Wu, B. Dynamic correlation of market connectivity, risk spillover and abnormal volatility in stock price. *Phys. A Stat. Mech. Its Appl.* **2022**, *587*, 126506. [[CrossRef](#)]
- Mollah, S.; Quoreshi, A.S.; Zafirov, G. Equity market contagion during global financial and Eurozone crises: Evidence from a dynamic correlation analysis. *J. Int. Financ. Mark. Inst. Money* **2016**, *41*, 151–167. [[CrossRef](#)]
- Chiang, T.C.; Jeon, B.N.; Li, H. Dynamic correlation analysis of financial contagion: Evidence from Asian markets. *J. Int. Money Financ.* **2007**, *26*, 1206–1228. [[CrossRef](#)]
- Karanasos, M.; Paraskevopoulos, A.G.; Menla Ali, F.; Karoglou, M.; Yfanti, S. Modelling stock volatilities during financial crises: A time varying coefficient approach. *J. Empir. Financ.* **2014**, *29*, 113–128. [[CrossRef](#)]
- Bringmann, L.F.; Pe, M.L.; Vissers, N.; Ceulemans, E.; Borsboom, D.; Vanpaemel, W.; Tuerlinckx, F.; Kuppens, P. Assessing temporal emotion dynamics using networks. *Assessment* **2016**, *23*, 425–435. [[CrossRef](#)]
- Pe, M.L.; Kircanski, K.; Thompson, R.J.; Bringmann, L.F.; Tuerlinckx, F.; Mestdagh, M.; Mata, J.; Jaeggi, S.M.; Buschkuhl, M.; Jonides, J.; et al. Emotion-network density in major depressive disorder. *Clin. Psychol. Sci.* **2015**, *3*, 292–300. [[CrossRef](#)] [[PubMed](#)]
- Wichers, M.; Groot, P.C.; Psychosystems, E.; Group, E. Critical slowing down as a personalized early warning signal for depression. *Psychother. Psychosom.* **2016**, *85*, 114–116. [[CrossRef](#)] [[PubMed](#)]
- Bollerslev, T. Generalized autoregressive conditional heteroskedasticity. *J. Econom.* **1986**, *31*, 307–327. [[CrossRef](#)]
- Bauwens, L.; Laurent, S.; Rombouts, J.V. Multivariate GARCH models: A survey. *J. Appl. Econom.* **2006**, *21*, 79–109. [[CrossRef](#)]
- Brownlees, C.T.; Engle, R.F.; Kelly, B.T. A practical guide to volatility forecasting through calm and storm. *J. Risk* **2011**, *14*, 3–22. [[CrossRef](#)]
- Hansen, P.R.; Lunde, A. A forecast comparison of volatility models: Does anything beat a GARCH (1, 1)? *J. Appl. Econom.* **2005**, *20*, 873–889. [[CrossRef](#)]
- Sakoğlu, Ü.; Pearlson, G.D.; Kiehl, K.A.; Wang, Y.M.; Michael, A.M.; Calhoun, V.D. A method for evaluating dynamic functional network connectivity and task-modulation: Application to schizophrenia. *Magn. Reson. Mater. Physics, Biol. Med.* **2010**, *23*, 351–366. [[CrossRef](#)] [[PubMed](#)]
- Allen, E.A.; Damaraju, E.; Plis, S.M.; Erhardt, E.B.; Eichele, T.; Calhoun, V.D. Tracking whole-brain connectivity dynamics in the resting state. *Cereb. Cortex* **2014**, *24*, 663–676. [[CrossRef](#)] [[PubMed](#)]
- Shakil, S.; Lee, C.H.; Keilholz, S.D. Evaluation of sliding window correlation performance for characterizing dynamic functional connectivity and brain states. *NeuroImage* **2016**, *133*, 111–128. [[CrossRef](#)] [[PubMed](#)]
- Mokhtari, F.; Akhlaghi, M.I.; Simpson, S.L.; Wu, G.; Laurienti, P.J. Sliding window correlation analysis: Modulating window shape for dynamic brain connectivity in resting state. *NeuroImage* **2019**, *189*, 655–666. [[CrossRef](#)]

24. Wilson, A.G.; Ghahramani, Z. Generalised Wishart processes. In Proceedings of the 27th Conference on Uncertainty in Artificial Intelligence, Barcelona, Spain, 14–17 July 2011; pp. 736–744.
25. Rasmussen, C.E.; Williams, C.K.I. *Gaussian Processes for Machine Learning*; The MIT Press: Cambridge, MA, USA, 2005.
26. Nejatbakhsh, A.; Garon, I.; Williams, A.H. Estimating noise correlations across continuous conditions with Wishart processes. In Proceedings of the Thirty-seventh Conference on Neural Information Processing Systems, New Orleans, LA, USA, 10–16 December 2023.
27. Kampman, O.P.; Ziminski, J.; Afyouni, S.; van der Wilk, M.; Kourtzi, Z. Time-varying functional connectivity as Wishart processes. *Imaging Neurosci.* **2024**, *2*, 1–28. [[CrossRef](#)]
28. Meng, R.; Yang, F.; Kim, W.H. Dynamic covariance estimation via predictive Wishart process with an application on brain connectivity estimation. *Comput. Stat. Data Anal.* **2023**, *185*, 107763. [[CrossRef](#)]
29. Cardona, H.D.V.; Álvarez, M.A.; Orozco, Á.A. Generalized Wishart processes for interpolation over diffusion tensor fields. In Proceedings of the Advances in Visual Computing: 11th International Symposium, ISVC 2015, Las Vegas, NV, USA, 14–16 December 2015; Springer: Berlin/Heidelberg, Germany, 2015; pp. 499–508.
30. Jørgensen, M.; Deisenroth, M.; Salimbeni, H. Stochastic differential equations with variational wishart diffusions. In Proceedings of the International Conference on Machine Learning, Virtual Event, 13–18 July 2020; pp. 4974–4983.
31. Heaukulani, C.; van der Wilk, M. Scalable Bayesian dynamic covariance modeling with variational Wishart and inverse Wishart processes. *arXiv* **2019**, arXiv:1906.09360.
32. Bauer, M.; van der Wilk, M.; Rasmussen, C.E. Understanding probabilistic sparse Gaussian process approximations. In Proceedings of the 30th International Conference on Neural Information Processing Systems, Red Hook, NY, USA, 5–10 December 2016; p. 1533–1541.
33. Chopin, N.; Papaspiliopoulos, O. *An introduction to Sequential Monte Carlo*, 1st ed.; Springer Series in Statistics; Springer: Cham, Switzerland, 2020.
34. Del Moral, P.; Doucet, A.; Jasra, A. Sequential Monte Carlo samplers. *J. R. Stat. Soc. Ser. Stat. Methodol.* **2006**, *68*, 411–436. [[CrossRef](#)]
35. Kantas, N.; Doucet, A.; Singh, S.; Maciejowski, J. An overview of Sequential Monte Carlo methods for parameter estimation in general state-space models. *IFAC Proc. Vol.* **2009**, *42*, 774–785. [[CrossRef](#)]
36. Speich, M.; Dormann, C.F.; Hartig, F. Sequential Monte-Carlo algorithms for Bayesian model calibration—A review and method comparison. *Ecol. Model.* **2021**, *455*, 109608. [[CrossRef](#)]
37. Wills, A.G.; Schön, T.B. Sequential Monte Carlo: A unified review. *Annu. Rev. Control. Robot. Auton. Syst.* **2023**, *6*, 159–182. [[CrossRef](#)]
38. Bru, M.F. Wishart processes. *J. Theor. Probab.* **1991**, *4*, 725–751. [[CrossRef](#)]
39. Zhang, Z. A note on Wishart and inverse Wishart priors for covariance matrix. *J. Behav. Data Sci.* **2021**, *1*, 119–126. [[CrossRef](#)]
40. Geman, S.; Geman, D. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans. Pattern Anal. Mach. Intell.* **1984**, *PAMI-6*, 721–741. [[CrossRef](#)]
41. Hensman, J.; Fusi, N.; Lawrence, N.D. Gaussian processes for big data. In Proceedings of the 29th Conference on Uncertainty in Artificial Intelligence, Bellevue, WA, USA, 11–13 July 2013; pp. 282–290.
42. Jasra, A.; Stephens, D.A.; Doucet, A.; Tsagaris, T. Inference for Lévy-driven stochastic volatility models via adaptive sequential Monte Carlo. *Scand. J. Stat.* **2011**, *38*, 1–22. [[CrossRef](#)]
43. Agapiou, S.; Papaspiliopoulos, O.; Sanz-Alonso, D.; Stuart, A.M. Importance sampling: Intrinsic dimension and computational cost. *Stat. Sci.* **2017**, *32*, 405–431. [[CrossRef](#)]
44. Herbst, E.; Schorfheide, F. Sequential Monte Carlo sampling for DSGE models. *J. Appl. Econom.* **2014**, *29*, 1073–1098. [[CrossRef](#)]
45. Cabezas, A.; Corenflos, A.; Lao, J.; Louf, R. BlackJAX: Composable Bayesian inference in JAX. 2024.
46. Bradbury, J.; Frostig, R.; Hawkins, P.; Johnson, M.J.; Leary, C.; Maclaurin, D.; Necula, G.; Paszke, A.; VanderPlas, J.; Wanderman-Milne, S.; et al. JAX: Composable Transformations of Python+NumPy Programs. 2018. Available online: <https://github.com/google/jax> (accessed on 10 March 2024).
47. Gelman, A.; Rubin, D.B. Inference from iterative simulation using multiple sequences. *Stat. Sci.* **1992**, *7*, 457–472. [[CrossRef](#)]
48. Matthews, A.G.d.G.; van der Wilk, M.; Nickson, T.; Fujii, K.; Boukouvalas, A.; León-Villagrà, P.; Ghahramani, Z.; Hensman, J. GPflow: A Gaussian process library using TensorFlow. *J. Mach. Learn. Res.* **2017**, *18*, 1–6.
49. Kingma, D.; Ba, J. Adam: A Method for Stochastic Optimization. In Proceedings of the International Conference on Learning Representations (ICLR), San Diego, CA, USA, 7–9 May 2015.
50. Scott, D.W. *Multivariate Density Estimation: Theory, Practice, and Visualization*; John Wiley & Sons: Hoboken, NJ, USA, 2015.
51. Cabrieto, J.; Adolf, J.; Tuerlinckx, F.; Kuppens, P.; Ceulemans, E. An objective, comprehensive and flexible statistical framework for detecting early warning signs of mental health problems. *Psychother. Psychosom.* **2019**, *88*, 184–186. [[CrossRef](#)]
52. Bringmann, L.F.; Vissers, N.; Wichers, M.; Geschwind, N.; Kuppens, P.; Peeters, F.; Borsboom, D.; Tuerlinckx, F. A network approach to psychopathology: New insights into clinical longitudinal data. *PLoS ONE* **2013**, *8*, e60188. [[CrossRef](#)] [[PubMed](#)]
53. Kossakowski, J.J.; Groot, P.C.; Haslbeck, J.M.; Borsboom, D.; Wichers, M. Data from ‘critical slowing down as a personalized early warning signal for depression’. *J. Open Psychol. Data* **2017**, *5*, 1. [[CrossRef](#)]
54. Derogatis, L.R.; Rickels, K.; Rock, A.F. The SCL-90 and the MMPI: A step in the validation of a new self-report scale. *Br. J. Psychiatry* **1976**, *128*, 280–289. [[CrossRef](#)] [[PubMed](#)]

55. Benton, G.; Maddox, W.; Wilson, A.G. Volatility based kernels and moving average means for accurate forecasting with Gaussian processes. In Proceedings of the International Conference on Machine Learning, Baltimore, MD, USA, 17–23 July 2022; pp. 1798–1816.
56. Kruschke, J.K. Rejecting or accepting parameter values in Bayesian estimation. *Adv. Methods Pract. Psychol. Sci.* **2018**, *1*, 270–280. [[CrossRef](#)]
57. Galanos, A. *Rmgarch: Multivariate GARCH Models*, R package version 1.3-6; 2019. Available online: <https://cran.r-project.org/web/packages/rmgarch/> (accessed on 10 March 2024).
58. Yao, Y.; Vehtari, A.; Gelman, A. Stacking for non-mixing Bayesian computations: The curse and blessing of multimodal posteriors. *J. Mach. Learn. Res.* **2022**, *23*, 1–45.
59. Lalchand, V.; Rasmussen, C.E. Approximate inference for fully Bayesian Gaussian process regression. In Proceedings of the Symposium on Advances in Approximate Bayesian Inference, Vancouver, BC, Canada, 8 December 2019; pp. 1–12.
60. Svensson, A.; Dahlin, J.; Schön, T.B. Marginalizing Gaussian process hyperparameters using sequential Monte Carlo. In Proceedings of the 2015 IEEE 6th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP), Cancun, Mexico, 13–16 December 2015; pp. 477–480.
61. Sridhar, V.H.; Davidson, J.D.; Twomey, C.R.; Sosna, M.M.; Nagy, M.; Couzin, I.D. Inferring social influence in animal groups across multiple timescales. *Philos. Trans. R. Soc. B* **2023**, *378*, 20220062. [[CrossRef](#)] [[PubMed](#)]
62. Gilmore, J.H.; Knickmeyer, R.C.; Gao, W. Imaging structural and functional brain development in early childhood. *Nat. Rev. Neurosci.* **2018**, *19*, 123–137. [[CrossRef](#)] [[PubMed](#)]
63. Xifara, T.; Sherlock, C.; Livingstone, S.; Byrne, S.; Girolami, M. Langevin diffusions and the Metropolis-adjusted Langevin algorithm. *Stat. Probab. Lett.* **2014**, *91*, 14–19. [[CrossRef](#)]
64. Neal, R.M. MCMC using Hamiltonian dynamics. *Handb. Markov Chain. Monte Carlo* **2010**, *54*, 113–162.
65. Wilson, A.; Adams, R. Gaussian process kernels for pattern discovery and extrapolation. In Proceedings of the International Conference on Machine Learning, Atlanta, GA, USA, 16–21 June 2013; pp. 1067–1075.
66. Rossi, S.; Heinonen, M.; Bonilla, E.; Shen, Z.; Filippone, M. Sparse Gaussian processes revisited: Bayesian approaches to inducing-variable approximations. In Proceedings of the International Conference on Artificial Intelligence and Statistics, Virtual, 13–15 April 2021; pp. 1837–1845.
67. Rowe, D.B. *Multivariate Bayesian Statistics: Models for Source Separation and Signal Unmixing*; Chapman and Hall/CRC: Boca Raton, FL, USA, 2002.
68. Wilson, A.; Nickisch, H. Kernel interpolation for scalable structured Gaussian processes (KISS-GP). In Proceedings of the International Conference on Machine Learning, Lille, France, 6–11 July 2015; pp. 1775–1784.
69. Cunningham, J.P.; Shenoy, K.V.; Sahani, M. Fast Gaussian process methods for point process intensity estimation. In Proceedings of the 25th International Conference on Machine Learning, Helsinki, Finland, 5–9 July 2008; pp. 192–199.
70. Wang, Y.; Kang, J.; Kemmer, P.B.; Guo, Y. An efficient and reliable statistical method for estimating functional connectivity in large scale brain networks using partial correlation. *Front. Neurosci.* **2016**, *10*, 179959. [[CrossRef](#)] [[PubMed](#)]
71. Smith, S.M.; Miller, K.L.; Salimi-Khorshidi, G.; Webster, M.; Beckmann, C.F.; Nichols, T.E.; Ramsey, J.D.; Woolrich, M.W. Network modelling methods for FMRI. *Neuroimage* **2011**, *54*, 875–891. [[CrossRef](#)] [[PubMed](#)]
72. Hinne, M.; Ambrogioni, L.; Janssen, R.J.; Heskes, T.; van Gerven, M.A. Structurally-informed Bayesian functional connectivity analysis. *NeuroImage* **2014**, *86*, 294–305. [[CrossRef](#)] [[PubMed](#)]
73. Murray, I.; Adams, R.; MacKay, D. Elliptical slice sampling. In Proceedings of the 13th International Conference on Artificial Intelligence and Statistics, JMLR Workshop and Conference Proceedings, Chia Laguna Resort, Sardinia, Italy, 13–15 May 2010; pp. 541–548.
74. Kullback, S.; Leibler, R. On information and sufficiency. *Ann. Math. Stat.* **1951**, *22*, 79–86. [[CrossRef](#)]
75. Bollerslev, T.; Engle, R.F.; Wooldridge, J.M. A capital asset pricing model with time-varying covariances. *J. Political Econ.* **1988**, *96*, 116–131. [[CrossRef](#)]
76. Box, G.E.; Jenkins, G.M.; Reinsel, G.C.; Ljung, G.M. *Time Series Analysis: Forecasting and Control*; John Wiley & Sons: Hoboken, NJ, USA, 2015.
77. Engle, R. Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *J. Bus. Econ. Stat.* **2002**, *20*, 339–350. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.