Approximating Solutions of Stochastic Differential Equations with Gaussian Markov Random Fields

Geir-Arne Fuglstad
TMA4500 - Industrial Mathematics, Specialization Project
Fall 2010

December 19, 2010

Preface

This report is a part of the subject ”TMA4500 - Industrial Mathematics, Specialization Project” at the Norwegian University of Science and Technology. The subject is a 15 credit subject and the grade is determined only by this final report.

The subject has given the author insight into a more independent type of work than in previous courses at the university. The responsibility for steady progress and completion of the project within the deadline falls mainly on the student. Further the project has allowed work into a new and exciting problem field.

The author wishes to thank the supervisors Håvard Rue and Finn Lindgren for their help and valuable experience.

Abstract

A solution of a stochastic differential equation on a bounded domain can be approximated by its joint distribution on a finite discretization of the domain. By creating a Gaussian Markov random field which approximates this joint distribution, properties of Gaussian Markov random fields which allow fast simulation and inference can be used.

The construction of a Gaussian Markov random field which approximates the joint distribution has previously been done for non-stationary stochastic differential equations. This report creates one such approximation for a non-stationary stochastic differential equation.

First an example of a an earlier approximation for second-order random walk is summarized and an example of how such a model can be applied as part of a larger model, in smoothing of time series, is demonstrated. Then a Gaussian Markov random field approximation is constructed for the stochastic heat equation

$$\frac{\partial}{\partial t}u(x,t) = \frac{\partial^2}{\partial x^2}u(x,t) + \sigma\mathcal{W}(x,t),$$

where $\sigma > 0$ is a constant and $\mathcal{W}(x,t)$ is space-time white noise, on a rectangular domain under the Neumann boundary conditions of zero derivative with respect to x at the spatial boundaries.

The approximation uses a finite volume method with an implicit Euler method in time and shows good results for two different type of initial conditions and merits further investigation.

Contents

Chapter 1

Introduction

Previous work has been done by [6] to create a Gaussian Markov random field which approximates the solution of a second-order random walk on an irregular grid. In addition [5] has demonstrated how a Matérn random field on a sphere can be approximated by a Gaussian Markov random field. Both these examples are approximations of the solution of a stationary stochastic differential equation on some domain.

The problem considered in this report is to construct such a Gaussian Markov random field approximation to a non-stationary stochastic differential equation.

The use of Gaussian Markov random fields for the approximations means that conditional independence between variables gives sparse precision matrices. This sparsity enables the use of numerical tools for sparse matrices to do more efficient simulation and inference. Chapter 2 gives a brief introduction to Gaussian Markov random fields and shows how the sparsity of the precision matrix can be used to improve the efficiency of simulation and calculation of marginal variances. In addition it gives an example of how a permutation of the variables may greatly increase the sparsity of the Cholesky factor of the precision matrix.

Chapter 3 gives a summary of how the approximation in [6] was constructed for the second-order random walk on an irregular grid. Then an example of how this approximation can be applied to smoothing of time series is demonstrated.

Lastly Chapter 4 considers further development of these types of approximations to non-stationary stochastic partial differential equations, by constructing one such approximation for a form of the stochastic heat equation. This is accomplished by the use of a finite volume method and the result is tested with a fixed initial condition and with second-order random walk as initial condition.

The report ends with some conclusions in Chapter 5.

Chapter 2

Gaussian Markov random fields

2.1 Definition

Gaussian Markov random fields are a type of Gaussian random fields. A Gaussian random field can be defined as follows.

Definition 2.1.1 (Gaussian random field (GRF))**.** Let $\mathcal{D} \subset \mathbb{R}^n$, then the random field $\{x(t) : t \in \mathcal{D}\}$ is said to be a Gaussian random field if for all $m \in \mathbb{N}$ for all choices of points $t_1, \ldots, t_m \in \mathcal{D}$, $(x(t_1), \ldots, x(t_m))$ has a multivariate Gaussian distribution.

For the rest of this section the interest will lie in GRFs on some finite set of points, which without loss of generality can be numbered from 1 to n. Call this set $\mathcal{V} = \{1, \ldots, n\}$ and consider the GRF $\{x(t) : t \in \mathcal{V}\}$. As $\mathcal{V}$ is finite, one can consider the joint distribution of all $x(t_i)$ for $i \in \mathcal{V}$, which by definition is a multivariate Gaussian distribution. Assume the covariance matrix of the distribution, Σ, is symmetric positive definite and let $Q = \Sigma^{-1}$. This matrix is called the *precision matrix* of the distribution. The following theorem from [7] characterizes the relationship between conditional independence of $x(i)$ and $x(j)$ for $i, j \in \mathcal{V}$ and element $Q_{i,j}$ of the precision matrix.

Theorem 2.1.1 (From [7])**.** *Let x have a multivariate Gaussian distribution with mean μ and precision matrix $Q > 0$. Then for $i \neq j$, $x(i)|\{x(k) : k \in \mathcal{V} - \{i, j\}\}$ is independent of $x(j)|\{x(k) : k \in \mathcal{V} - \{i, j\}\}$ if and only if $Q_{i,j} = 0$.*

From a GRF on $\mathcal{V} = \{1, \ldots, n\}$ an undirected graph $\mathcal{G}$ can be constructed with $\mathcal{V}$ as the set of vertexes and the set of edges, $\mathcal{E}$, defined by $\{i, j\} \in \mathcal{E}$ if and only if $x(i)$ and $x(j)$ are not conditionally independent. Since all graphs will be undirected, the term graph will be used to mean an undirected graph from this point on. A formal definition for a Gaussian Markov random field can be made by reversing this type of construction and first specifying the conditional independence structure, i.e. the graph.

Definition 2.1.2 (Gaussian Markov random field (GMRF) [7])**.** A random vector $x \in \mathbb{R}^n$ is called a Gaussian Markov random field with respect to a labeled graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with

mean μ and precision matrix $Q > 0$, if and only if its density is of the form

$$\pi(x) = \frac{1}{(2\pi)^{n/2}} |Q|^{1/2} \exp\left(-\frac{1}{2}(x-\mu)^{\mathrm{T}} Q (x-\mu)\right), \tag{2.1}$$

and for all $i \neq j$, $Q_{i,j} \neq 0$ if and only if $\{i, j\} \in \mathcal{E}$.

From this definition one can see that a general GRF (on a labeled graph) is a GMRF if one choses the edges in $\mathcal{E}$ according to the non-zero elements of Q. In the rest of the report the underlying graph will be suppressed in the notation for GMRFs unless it is explicitly needed.

Definition 2.1.2 only allows $Q > 0$, but also GMRFs with positive semi-definite precision matrices will be allowed. The lack of positive definiteness means that there are linear combinations of the elements of the random vector x which has unbounded variance. This means that the distribution is not proper, but it can be used as a prior distribution if the posterior distribution is proper and one can draw samples from the distribution conditioned on that the linear combinations with unbounded variance are zero. More specifically all linear combinations $a^{\mathrm{T}}x$ such that $a \in \mathrm{Null}(Q)$ must be zero. Let $e_1, \ldots, e_k$ be an orthonormal basis for the null space, then $x|Ax = 0$, with $A = [e_1, \ldots, e_k]^{\mathrm{T}}$, is a proper distribution and can be sampled.

2.2 Simulation

Simulations from a GMRF can be done by first Cholesky factorizing the precision matrix Q and then solving a linear system as the following theorem shows.

Theorem 2.2.1. *A simulation from the GMRF, x, with mean μ and precision matrix $Q > 0$ can be made by first drawing $z \sim N_n(0, I_n)$ and then solving*

$$L^T(x - \mu) = z,$$

where L corresponds to a Cholesky decomposition $Q = LL^T$.

Proof. Let $z \sim N_n(0, I_n)$ and $L^{\mathrm{T}}(x - \mu) = z$, where L is the Cholesky factor of Q. L is invertible because Q is invertible, therefore

$$\begin{aligned} \mathrm{Cov}(L^{\mathrm{T}}x) &= \mathrm{Cov}(z) \\ L^{\mathrm{T}}\mathrm{Cov}(x)L &= I_n \\ \mathrm{Cov}(x) &= (LL^{\mathrm{T}})^{-1} = Q^{-1}. \end{aligned}$$

Further x is Gaussian as it is a linear combination of Gaussian iid variables, its covariance is correct from above and its expectation is μ since $\mathrm{E}(z) = 0$. □

2.3 Sparse precision matrix

Let x be a GMRF with precision matrix Q and mean zero. Simulation by Theorem 2.2.1 requires the computation of a Cholesky decomposition and the solution of a system of linear equations of the form $L^{\mathrm{T}}x = z$, where L is lower triangular. This means that the efficiency of the simulation is determined by the efficiency of the Cholesky decomposition and of the solution of the linear system.

For a general (positive definite) precision matrix, the lower triangular matrix L of the Cholesky decomposition can be calculated recursively, column-wise by the formula

$$L_{i,j} = \begin{cases} \sqrt{Q_{i,i} - \sum_{k=1}^{i-1} L_{j,k}^2} & i = j \\ \frac{1}{L_{j,j}} \left(Q_{i,j} - \sum_{k=1}^{j-1} L_{i,k} L_{j,k} \right) & i > j \end{cases}. \tag{2.2}$$

This recursive formula requires $\mathcal{O}(n^3)$ floating point operations and the solution of the equation $L^{\mathrm{T}}x = z$ requires $\mathcal{O}(n^2)$ floating point operations.

However, if Q is sparse the Cholesky factor will often be sparse. Thus if one could know a priori which elements of L that were zero, it would not be necessary to calculate those elements and it would not be necessary to include them in the summands in the recursive formula. Consider for example the case that one knew that the Cholesky factor would have a lower bandwidth of 1, then only $2n - 1$ terms needs to be calculated and all the sums in the recursive formula would have a maximum of one term. Thus the Cholesky decomposition could be calculated with $\mathcal{O}(n)$ floating point operations.

This improvement in asymptotic computational cost motivates the search for general conditions for whether an element $L_{i,j}$ is zero or not. To continue this train of thought it is useful to consider the underlying graph of the GMRF, denote this by $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ is the nodes and $\mathcal{E}$ is the edges. It would be preferable to be able to determine the non-zero elements directly from the graph, that is only based on the knowledge of which $Q_{i,j}$ that are non-zero. This is not possible, but it is possible to find sufficient conditions based on the graph.

The proof of the sufficient condition requires the use of a different type of Markov property than considered previously. Previously in Theorem 2.1.1 the pairwise Markov property was considered. GMRFs also have another Markov property that is called the global Markov property. The two different properties are illustrated in Figure 2.1.

Lemma 2.3.1 (Global Markov property)**.** *Let x be a zero mean GMRF with precision matrix $Q > 0$ and let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be the associated labeled graph. Let $A, B, C \subset \mathcal{V}$ be non-empty pairwise disjoint sets. If C separates A and B, that is all paths from an element in A to an element in B must go through C, then $\{x(i) : i \in A\}$ and $\{x(i) : i \in B\}$ are independent conditioned on $\{x(i) : i \in C\}$.*

Proof. Let Y_1 be the set of all elements of $\mathcal{V}$ than can be reached from A without entering C and Y_2 the set of all elements of $\mathcal{V}$ that can be reached from B without entering C.

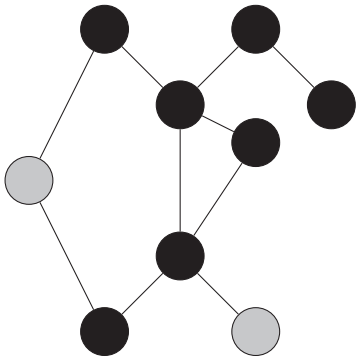

(a) The pairwise Markov property. There is no edge between the two grey nodes, therefore they are independent conditioned on the black nodes.

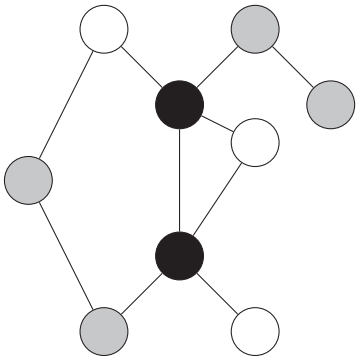

(b) The global Markov property. The black nodes separate the two collections of grey nodes, that is there is no path between the two collections which does not cross the black nodes, therefore the two collections are independent conditioned on the black nodes.

Figure 2.1: The two sub-figures illustrates the pairwise Markov property and the global Markov property, respectively.

These sets are disjoint because C separates A and B. Let Y_3 be all elements not in Y_1 or Y_2, this set includes C.

Write the distribution of x as

$$\pi(x) \propto \exp\left(-\frac{1}{2}x^{\mathrm{T}}Qx\right).$$

Write out $x^{\mathrm{T}}Qx = \sum_{i,j} x_i x_j Q_{i,j}$ and use that $Q_{i,j} = 0$ when i and j are in different sets among Y_1, Y_2 or $Y_3 - C$. This is because the sets are separated by C. The quadratic form can be written

$$x^{\mathrm{T}}Qx = \sum_{i,j\in Y_1} x_i x_j Q_{i,j} + 2\sum_{i\in Y_1, j\in C} x_i x_j Q_{i,j} + \sum_{i,j\in Y_2} x_i x_j Q_{i,j} + 2\sum_{i\in Y_2, j\in C} x_i x_j Q_{i,j} + K, \tag{2.3}$$

where K only depends on $x(i)$ for $i \in Y_3$.

In Equation (2.3) the first two terms only involve elements in Y_1 and C, the third and fourth terms only involve elements in Y_2 and C and K only involves elements in Y_3. Thus the joint distribution can be factored into one factor only involving Y_1 (and C), one factor only involving Y_2 (and C) and one factor only involving Y_3.

This means that the joint distribution for $\{x(i) : i \in Y_1\}$ and the joint distribution for $\{x(i) : i \in Y_2\}$ are independent conditioned on $\{x(i) : i \in C\}$. This implies that $\{x(i) : i \in A\}$ and $\{x(i) : i \in B\}$ are independent conditioned on $\{x(i) : i \in C\}$.

□

By using this lemma a sufficient condition for $L_{i,j} = 0$ based only on the underlying graph of the GMRF can be proved.

Theorem 2.3.2. *Let x be a GMRF with mean zero and precision matrix $Q > 0$ corresponding to the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and let $F(i,j) = (i+1, \ldots, j-1, j+1, \ldots, n)$ for $i < j$. If $F(i,j)$ separates i and j, that is if all paths between i and j must pass through $F(i,j)$, then $L_{j,i} = 0$.*

Proof. Let L be a lower triangular matrix such that $Q = LL^T$ and let $x_{i:n} = (x_i, x_{i+1}, \ldots, x_n)$ and $z_{i:n} = (z_i, \ldots, z_n)$, where z has a multivariate standard Gaussian distribution.

By Theorem 2.2.1 $L^{\mathrm{T}}x \sim z$. This implies that $x_{i:n}$ can be found by backwards substitution from $z_{i:n}$, which means that the distribution of $x_{i:n}$ is determined only by $z_{i:n}$ and that

$$\pi(x_{i:n}) = \pi(z_{i:n}) \propto \exp\left(-\frac{1}{2}(L^{\mathrm{T}}x)^{\mathrm{T}}_{i:n}(L^{\mathrm{T}}x)_{i:n}\right).$$

But if $y = L^{\mathrm{T}}x$, then $y_k = \sum_{l=k}^{n} L_{l,k}x_l$. The only element of $y_{i:n}$ with both an x_i and an x_j term is y_i, but the only cross-term in y_i^2 is $2L_{i,i}L_{j,i}x_ix_j$. This means that if $Q^{i:n}$ is the precision matrix of $x_{i:n}$, then $Q^{i:n}_{j-i+1,1} = L_{j,i}L_{i,i}$.

The global Markov property of x from Lemma 2.3.1 means that since $F(i,j)$ separates x_i and x_j, x_i and x_j are independent conditioned on $x_{F(i,j)}$. By Theorem 2.1.1 this implies $L_{j,i}L_{i,i} = Q^{i:n}_{j-i+1,1} = 0$.

Further $Q > 0$ implies $L > 0$ which means $L_{l,l} > 0$ for $l = 1, \ldots, n$. Thus $L_{j,i} = 0$. □

In the proof of the previous theorem one can see that the sufficient and necessary condition for $L_{j,i} = 0$ is x_i independent of x_j conditioned on $x_{F(i,j)}$, but this is in general inconvenient to determine when x_i and x_j are not separated by $x_{F(i,j)}$. Therefore it will not be given its own theorem. By using Theorem 2.3.2 the relationship in Equation (2.2) can be used to create a recursive formula which does not calculate things which are known to be zero.

Corollary. *Let x be a GMRF with precision matrix $Q > 0$, let $F(i,j) = \{i+1, \ldots, j-1, j+1, \ldots, n\}$ for $i < j$, let $\mathcal{I}(i) = \{k : k > i,\ k$ and i not separated by $F(i,k)\ \}$ and let $\mathcal{K}(i) = \{k : k < i,\ k$ and i not separated by $F(k,i)\ \}$. Then the non-zero elements of the Cholesky factor, L, can be calculated by Algorithm 1.*

Proof. Use Theorem 2.3.2 to remove everything unnecessary from the relation in Equation (2.2). □

A simple example illustrates how powerful Theorem 2.3.2 can be.

Example 2.3.1 (First order autoregressive process)**.** Consider the time series $x_t = 0.5x_{t-1} + \epsilon_t$ for $t = 1, \ldots, n$, where x_0 is a constant and ϵ_t is iid Gaussian white noise. Then $F(i,j)$ separates i and j if and only if $j - i > 1$. By Theorem 2.3.2 $L_{j,i}$ is zero for $j - i > 1$ and L has a lower bandwidth of 1. Thus the Cholesky factorization only requires $\mathcal{O}(n)$ floating point operations. Further the low bandwidth can be taken

Data: matrix $Q > 0$
Result: Cholesky factor L
for $i \leftarrow 1$ **to** n **do**
  $L_{i,i} \leftarrow \sqrt{Q_{i,i} - \sum_{k \in \mathcal{K}(i)} L_{i,k}^2}$;
  for *increasing* j *in* $\mathcal{I}(i)$ **do**
    $L_{j,i} \leftarrow \frac{1}{L_{i,i}} \left(Q_{j,i} - \sum_{k \in \mathcal{K}(i) \cap \mathcal{K}(j)} L_{i,k} L_{j,k} \right)$;
  end
end

Algorithm 1: Calculate the Cholesky factor of an SPD sparse matrix.

advantage of in Theorem 2.2.1, because $L^{\mathrm{T}}(x - \mu) = z$ can be solved for x by $\mathcal{O}(n)$ floating point operations. This means that a simulation can be made with $\mathcal{O}(n)$ floating point operations.

2.4 Marginal variances

Based on a general $n \times n$ precision matrix $Q > 0$ it is costly to compute the marginal variances, because this requires the inversion of the matrix Q and inversion generally requires $\mathcal{O}(n^3)$ floating point operations. It is, however, possible to take advantage of the possibly sparse structure of Q and only compute a partial inverse, in the sense that not all elements of the inverse matrix are computed.

The following relationship can be found in [3, chapter 12].

Lemma 2.4.1. *Let Q be a symmetric positive definite matrix and let $\Sigma = Q^{-1}$. Let L be the Cholesky factor of Q, then*

$$\Sigma_{i,j} = \frac{\delta_{i,j}}{L_{i,i}^2} - \frac{1}{L_{i,i}} \sum_{k=i+1}^{n} L_{k,i} \Sigma_{k,j}, \qquad j \geq i, \tag{2.4}$$

where $\delta_{i,j} = 1$ for $i = j$ and 0 otherwise.

Proof. Let $Q = LL^{\mathrm{T}} > 0$, then

$$L^{\mathrm{T}} \Sigma = L^{\mathrm{T}} (LL^{\mathrm{T}})^{-1} = L^{-1}.$$

Further L^{-1} is lower triangular and $(L^{-1})_{i,i} = 1/L_{i,i}$. Thus

$$\sum_{k=i}^{n} L_{k,i} \Sigma_{k,j} = \frac{1}{L_{i,i}} \delta_{i,j}, \qquad j \geq i,$$

which is equivalent to Equation (2.4). □

[3, chapter 12] shows that this relationship can be used to compute the elements of the covariance matrix recursively and that it is not always necessary to calculate all the terms to find the marginal variances. In fact the reference shows that one only needs to calculate the diagonal of Σ and the $\Sigma_{i,j}$ which corresponds to $i < j$ for which $L_{j,i}$ was calculated in Algorithm 1. The details can be found in the reference, but the argumentation will be summarized in a theorem.

Theorem 2.4.2. *Let x be a GMRF with precision matrix $Q > 0$, let $F(i,j) = \{i+1, \dots, j-1, j+1, \dots, n\}$ for $i < j$ and let*

$$\mathcal{I}(i) = \{k > i : i \text{ and } k \text{ not separated by } F(i,k)\},$$

for $i = 1, \dots, n$.

Then $\Sigma_{i,j}$ for $\{i,j\} \in \{\{i,j\} : i \in \mathcal{I}(j) \cup \{j\}, j = 1, \dots, n\}$ can be computed by Algorithm 2.

Data: matrix $Q > 0$ and Cholesky factor, L, of Q
Result: partial inverse, Σ, of Q
for $i \leftarrow n$ **to** 1 **do**
  for *decreasing j in $\mathcal{I}(i) \cup \{i\}$* **do**
    $\Sigma_{i,j} \leftarrow \frac{\delta_{i,j}}{L_{i,i}^2} - \frac{1}{L_{i,i}} \sum_{k \in \mathcal{I}(i)} L_{k,i}\Sigma_{k,j}$;
    $\Sigma_{j,i} \leftarrow \Sigma_{i,j}$;
  end
end

Algorithm 2: Calculate the partial inverse of an SPD sparse matrix.

Proof. For $j > i$ Theorem 2.3.2 gives $L_{j,i} = 0$ if $F(i,j)$ separates i and j. This is true iff $j \in \mathcal{I}(i)$. Therefore the sum in Lemma 2.4.1 only needs to be taken over $\mathcal{I}(i)$, that is

$$\Sigma_{i,j} = \frac{\delta_{i,j}}{L_{i,i}^2} - \frac{1}{L_{i,i}} \sum_{k \in \mathcal{I}(i)} L_{k,i}\Sigma_{k,j},$$

for $j \geq i$.

Fix $i \in \{1, 2, \dots, n\}$ and $j \in \mathcal{I}(i) \cup \{i\}$ then $\Sigma_{i,j}$ can be computed iff $\Sigma_{k,j}$ is known for all $k \in \mathcal{I}(i)$. That means it is sufficient that $\Sigma_{k,l}$ is known for all $k, l \in \mathcal{I}(i) \cup \{i\}$, where not both k and l are equal to i.

Let $k \in \mathcal{I}(i) \cup \{i\}$. Then either $k = i$ or $k \in \mathcal{I}(i)$, which means that k and i are not separated by $F(i,k)$, which means there is a path from i to k using only the nodes $\{1, \dots, i\}$.

Let $k, l \in \mathcal{I}(i) \cup \{i\}$, with not both k and l equal to i, and assume without loss of generality that $k \leq l$. If $k = l \neq i$, then $\Sigma_{k,k}$ is calculated in iteration $n - k + 1$ of the first for-loop.

If $k < l$, then $k = i$ or there is a path from k to i using only nodes in $\{1, \dots, i\}$ and there is a path from i to l using only nodes in $\{1, \dots, i\}$. Thus there is a path from k to

l using only the nodes $\{1, \ldots, i\}$ and k and l are not separated by $F(k, l)$. This means that $\Sigma_{k,l}$ is computed in iteration $n - k + 1$ of the first for-loop, since $l \in \mathcal{I}(k)$.

Therefore all necessary elements of Σ are known at each iteration and the algorithm calculates all $\Sigma_{i,j}$ and $\Sigma_{j,i}$ for $i \in \{1, 2, \ldots, n\}$ and $j \in \mathcal{I}(i) \cup \{i\}$.

□

Remark (Non-zero pattern of the Cholesky factor). Note that Theorem 2.4.2 does not state that the only elements of Σ that has to be calculated is the diagonal elements of Σ and the elements $\Sigma_{i,j}$ with $i < j$ that corresponds to $L_{j,i} \neq 0$. The theorem states that only the $\Sigma_{i,j}$ which corresponds to an $L_{j,i}$ which is known to be 0 since it satisfies the sufficient condition in Theorem 2.3.2 does not need to be calculated.

Theorem 2.3.2 only specifies a sufficient condition and $L_{j,i}$ for $i < j$ can be 0 even though i and j are not separated by $F(i, j)$. If $L_{j,i} = 0$ by mistake is assumed to imply that i and j are separated by $F(i, j)$ Algorithm 2 may fail because an element of Σ which is needed has not been calculated.

Let x be a GMRF which corresponds to a graph with four vertexes and edges between any two vertexes. Assume for example that the precision matrix, Q, is the positive definite matrix given by

$$Q = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 5 & 1 & 3 \\ 1 & 1 & 5 & 5 \\ 1 & 3 & 5 & 31 \end{bmatrix}.$$

This matrix is dense, but the Cholesky factor

$$L = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 2 & 0 & 0 \\ 1 & 0 & 2 & 0 \\ 1 & 1 & 2 & 5 \end{bmatrix}$$

has a non-zero element below the diagonal even though there is no $i < j$ such that i and j are separated by $F(i, j)$.

Assume that Algorithm 2 is used to calculate the partial inverse and that one from L falsely concludes that 2 and 3 are separated by $F(2, 3)$. Then when calculating $\Sigma_{1,3}$, $\Sigma_{k,3}$ is needed for $k \in \mathcal{I}(1)$. But $\mathcal{I}(1) = \{2, 3, 4\}$ since $L_{2,1}$, $L_{3,1}$ and $L_{4,1}$ are non-zero. $\Sigma_{4,3}$ is found from the symmetry $\Sigma_{4,3} = \Sigma_{3,4}$, but $\Sigma_{2,3}$ has never been calculated since it was assumed that 2 and 3 were separated by $F(2, 3)$. Thus it is impossible to calculate $\Sigma_{1,3}$ and the recursive formula stops.

This shows that the set of $\Sigma_{i,j}$ terms that needs to be calculated should be determined from Q and not directly from L. This is an important point when using a third-party library for the calculation of the Cholesky factor as it may remove elements which are zero or close to zero. Thus either a library which creates a symbolic factorization, that is determines which elements of L that must be zero only from the non-zero pattern of Q, or a library which leaves explicit zeros in the the Cholesky factor is needed.

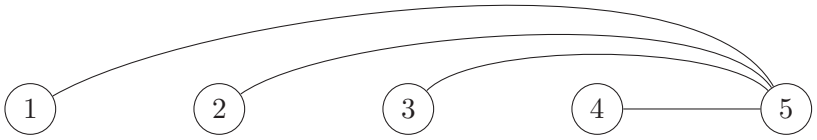

Figure 2.2: An example of a GMRF with a sparse Cholesky factor.

2.5 Permutations

The number of floating point operations and the memory needed for the Cholesky decomposition of a sparse matrix $Q > 0$ will depend on the pattern of the non-zero elements of Q. Thus it may be beneficial to do a pre-permutation of the form PQP^{T}, where P is a permutation matrix. This corresponds to reordering the vector x to Px in definition 2.1.2 and gives a new GMRF.

Example 2.5.1 (Simple graph)**.** The reordering of the elements of a GMRF or equivalently the underlying graph can have dramatic consequences on the sparsity of the Cholesky factor of the precision matrix. Figure 2.2 and and Figure 2.3 shows two different orderings of the same GMRF.

The Cholesky factor from Figure 2.2 has the following sparsity structure

$$L = \begin{bmatrix} * & 0 & 0 & 0 & 0 & 0 \\ 0 & * & 0 & 0 & 0 & 0 \\ 0 & 0 & * & 0 & 0 & 0 \\ 0 & 0 & 0 & * & 0 & 0 \\ 0 & 0 & 0 & 0 & * & 0 \\ * & * & * & * & * & * \end{bmatrix},$$

where * denotes a possibly non-zero elements.

The Cholesky factor from Figure 2.3 has the following sparsity structure

$$L = \begin{bmatrix} * & 0 & 0 & 0 & 0 & 0 \\ * & * & 0 & 0 & 0 & 0 \\ * & * & * & 0 & 0 & 0 \\ * & * & * & * & 0 & 0 \\ * & * & * & * & * & 0 \\ * & * & * & * & * & * \end{bmatrix},$$

where * denotes a possibly non-zero elements.

This shows that the permutation of a GMRF can drastically change the sparsity structure of the GMRF.

The example shows that there is a need for algorithms to minimize the number of non-zero elements of the Cholesky factor of the precision matrix. There exists different

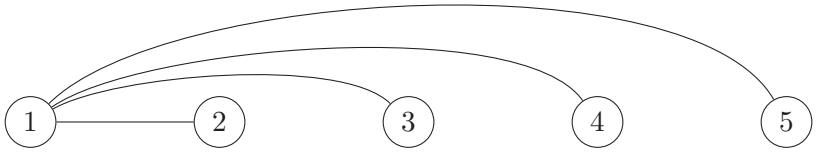

Figure 2.3: An example of a GMRF with a dense Cholesky factor.

strategies and algorithms for finding orderings of the variables of the GMRF which try to minimize the number of non-zero entries in the Cholesky factor. The efficient algorithms do not guarantee a minimum number of non-zeros elements. [7, chapter 2] gives an introduction to some of these algorithms.

2.6 Implementation

For this project the author used the CHOLMOD library [1] for Cholesky decomposition and solution of equations of the type $Ax = b$ and $L^{\mathrm{T}}x = b$ and the EIGEN library [4] for general sparse matrix operations. CHOLMOD is a C library and EIGEN is a template library for C++.

Code for simulation and calculation of marginal variances for GMRFs with sparse precision matrices was written in C++, and some code was written in MATLAB for quicker testing of new ideas. In the C++ implementation the automatic choice of the permutation of the matrix in CHOLMOD for the Cholesky decomposition was used. This code was used to generate the results in later chapters.

Chapter 3

Second-order random walk

3.1 Stochastic differential equations

In this report both stochastic ordinary differential equations and stochastic partial differential equations will be considered. A general stochastic differential equation will be abbreviated SDE and the abbreviation SPDE will be used for stochastic partial differential equation. An example of a SPDE is a form of the stochastic heat equation which will be considered in the next chapter

$$\frac{\partial u}{\partial t}(x,t) = \frac{\partial^2 u}{\partial x^2}(x,t) + \mathcal{W}(x,t), \tag{3.1}$$

where $\mathcal{W}(x,t)$ is space-time white noise. For the purpose of this report the white noise will always be considered to be Gaussian. This equation differs from the corresponding partial differential equation by the white noise term on the right hand side.

The objective will not be to try to solve equations of this type exactly, but rather create a time-space discretization and consider the joint distribution of the solution at the discretization points. The joint distribution will not necessarily be constructed to be exactly correct, but rather an approximation which satisfies some Markov conditions, which gives a GMRF with a sparse precision matrix. First an example of an SDE with only one independent variable will be given in this chapter and then an example of an SPDE in time and one-dimensional space will be considered in the next chapter.

3.2 Second-order random walk as an SDE

The one-dimensional second-order random walk will be used as an example of a stochastic processes indexed only by one variable. This process will serve as an example of a stochastic process with time dependency, but no spacial dependency.

A second-order random walk will be considered to be a (continuous) solution of the SDE

$$\Delta u(t) = \sigma_n \dot{W}(t), \tag{3.2}$$

where Δ is the Laplacian, $\mathrm{d}^2/\mathrm{d}t^2$, $\sigma_n > 0$ is a constant and $W(t)$ is a standard Wiener process. This equation is invariant to the addition of a constant or linear trend to the solution and does not specify a unique solution. However, as will be seen later this will not pose a problem when it is used only as a prior distribution.

The construction of a GMRF approximation for this equation on an irregular grid is already considered in [6] for $\sigma_n = 1$. The steps are nearly identical for $\sigma_n \neq 1$, therefore only a overview which shows the important steps will be given here.

3.2.1 Approximation of the solution

The solution of the SDE given in Equation (3.2) will be approximated on an interval $[a, b]$. Select n points $a = s_1 < \ldots < s_n = b$, and let ϕ_i be the function which is linear between each node and which is 1 at s_i and 0 at all other s_j for $j \neq i$. Further for $i = 1, \ldots, n-1$ define $d_i = s_{i+1} - s_i$.

A stochastic weak solution in $H_n = \mathrm{span}(\phi_1, \ldots, \phi_n) \subset H = L^2[a, b]$ is then found by requiring $u \in H_n$ and

$$\langle \Delta u, \phi_i \rangle_{L^2[a,b]} \stackrel{d}{=} \left\langle \sigma_n \dot{W}, \phi_i \right\rangle_{L^2[a,b]} \qquad i = 1, \ldots, n. \tag{3.3}$$

Let $u(t) = \sum_{k=1}^{n} \phi_k(t) w_k \in H_n$, where w_k are Gaussian random weights. Inserting this into Equation (3.3) gives n equations which can be represented in matrix form as

$$Hw \stackrel{d}{=} C^{1/2} z, \tag{3.4}$$

where H and C are matrices soon to be defined and z is a standard Gaussian multivariate random variable. H consists of all inner products

$$H_{i,j} = \langle \Delta \phi_j, \phi_i \rangle_{L^2[a,b]} \tag{3.5}$$

and C consists of all inner products

$$C_{i,j} = \sigma_n^2 \langle \phi_j, \phi_i \rangle_{L^2[a,b]} \tag{3.6}$$

These matrices can be found in [6] and are given by

$$H_{i,j} = \begin{cases} \frac{1}{d_{i-1}} & j = i-1, 1 < i < n \\ -(\frac{1}{d_{i-1}} + \frac{1}{d_i}) & j = i, 1 < i < n \\ \frac{1}{d_i} & j = i+1, 1 < i < n \end{cases},$$

with zero in all other positions and

$$C_{i,j} = \begin{cases} \frac{d_{i-1}}{6} & j = i-1 \\ \frac{d_i}{6} & j = i+1 \\ \frac{d_{i-1}+d_i}{3} & j = i, 1 < i < n \end{cases},$$

with $C_{1,1} = d_1/3$, $C_{n,n} = d_{n-1}/3$ and zero in all other positions.

It can be shown that $\text{rank}(H) = n - 2$ and that the eigenvectors corresponding to eigenvalues of 0 are $e_1 = (1, \ldots, 1)$ and $e_2 = (s_1, s_2, \ldots, s_n)$. This means that Equation (3.4) does not specify any conditions on $e_1^{\mathrm{T}} w$ and $e_2^{\mathrm{T}} w$. The equation can be rewritten as

$$z \stackrel{d}{=} C^{-1/2} H w, \tag{3.7}$$

where $z \sim N_n(0, I_n)$. This equation cannot have any solutions, because $C^{-1/2}H$ is singular. But it can "nearly" be solved, in the sense that

$$\left\langle C^{-1/2} H w, a \right\rangle \stackrel{d}{=} \langle z, a \rangle \qquad \forall a \in \text{Col}(C^{-1/2}H). \tag{3.8}$$

In fact, as previously mentioned, Equation (3.7) gives no information about what should happen to $a^{\mathrm{T}} w$ for $a \in \text{Null}(H)$, but by choosing $Q = H^{\mathrm{T}} C^{-1} H$ these linear combinations are given unbounded variance and Equation (3.8) holds.

This Q matrix will generally be dense, but as in [6] this matrix will be replaced by $\tilde{Q} = \sigma_n^{-2} H^{\mathrm{T}} \tilde{C}^{-1} H$, where $\tilde{C}$ is a diagonal matrix with $\tilde{C}_{i,i} = < \phi_i, 1 >_{L^2[a,b]}$. This makes $\tilde{Q}$ a sparse band matrix.

3.3 Smoothing of time series by second-order random walk

This section will provide an example of one application of the second-order random walk, namely smoothing of time series. The smoothing will attempt to reduce random fluctuations in the observations and extract the important features of the time series. This will be done by postulating a statistical model for the time series and considering the posterior distribution conditioned on the observed values.

3.3.1 Model

Let $y(t)$ be a real-valued time series. The first part of the model will dictate how the values of $y(t)$ are generated. $y(t)$ will be considered to be the sum of two independent parts, a random fluctuation $\epsilon(t)$, which is Gaussian white noise with constant variance σ_ϵ^2, and a more "regular" part $x(t)$. That is $y(t) = x(t) + \epsilon(t)$. "Regular" is here used as a non-technical term, only to mean that it has no discontinuities and sudden changes in values.

Let $\mathcal{D} = \{s_1 < s_2 < \ldots < s_n\}$ be a discretization of some interval $[a, b]$ and assume that $y(t)$ is observed on $k \leq n$ of these points. For convenience denote the vector of $x(s_i)$ for $i = 1, \ldots, n$ by x and the vector of $y(s_j)$ for the observed k points by y. Then the conditional distribution is $y|x \sim N_k(Ax, \sigma_\epsilon^2 I_k)$, where A is the $k \times n$ matrix selecting the k observed values from x. That is A has exactly one 1 in each row, and each correspond to a location for which $y(t)$ was observed.

The second part of the model is the choice of prior distribution for x. For this purpose a second-order random walk process will be used. The prior distribution of x will be the approximation by GMRF found previously.

From these assumptions one can calculate the posterior distribution of x. From the model one finds

$$\begin{aligned}
\pi(x|y) &\propto \pi(x)\pi(y|x) \\
\pi(x|y) &\propto \exp\left(-\tfrac{1}{2}\left(x^{\mathrm{T}}Qx + (y-Ax)^{\mathrm{T}}\sigma_\epsilon^{-2}I_n(y-Ax)\right)\right) \\
\pi(x|y) &\propto \exp\left(-\tfrac{1}{2}\left(x^{\mathrm{T}}(Q+\sigma_\epsilon^{-2}A^{\mathrm{T}}A)x - 2x^{\mathrm{T}}A^{\mathrm{T}}\sigma_\epsilon^{-2}y\right)\right) \\
\pi(x|y) &\propto \exp\left(-\tfrac{1}{2}\left(x^{\mathrm{T}}\tilde{Q}x - 2x^{\mathrm{T}}\tilde{Q}\tilde{Q}^{-1}A^{\mathrm{T}}\sigma_\epsilon^{-2}y\right)\right),
\end{aligned}$$

where $\tilde{Q} = Q + \sigma_\epsilon^{-2}A^{\mathrm{T}}A$. Comparing this to the expansion of $(x-\mu)^{\mathrm{T}}\tilde{Q}(x-\mu)$ one sees that $\mu = \sigma_\epsilon^{-2}\tilde{Q}^{-1}A^{\mathrm{T}}y$.

The final part of the model will be to specify the scale factor for Q. That is the σ_n from Equation (3.2). As was shown in the previous section the σ_n enters the equation for Q as a σ_n^{-2} factor. Since no investigation of the effect of changing σ_n was done in the previous sections, different values of σ_n will be compared.

3.3.2 Effect of scale parameter

To measure the effect of changing the scale parameter, σ_n, a data set was generated in the following way. Let the interval of interest be $[0, 10]$ and draw 30 points, $D = \{t_i\}_{i=1}^{30}$, uniformly in this interval to create a discretization. Then calculate $y(t)' = \sin(4\pi t/10)$ for $t \in D$ and create one realization of the time series $y(t) = y(t)' + \epsilon(t)$ for $t \in D$, where $\epsilon(t)$ is iid Gaussian white noise with variance 0.01.

Figure 3.1 shows the posterior mean for different values of the scale parameter σ_n. The true underlying regular part of the time series is a sine wave with amplitude 1 and 2 periods on the interval $[0, 10]$. From Figure 3.1(a) one can see that for $\sigma_n = 0.01$ the posterior mean does not follow the true value well. It appears that only the apparent downward trend shown by the observations is detected. The σ_n is too low compared with the variations of the underlying regular part of the time series. Figure 3.1(c) shows that $\sigma_n = 100$ gives a posterior mean with a very "jagged" behavior, which appears to pass very close to all observations, except observations at very close time points. This seems to indicate that the value of σ_n is so large that the type of variations observed in the observations are expected. Lastly Figure 3.1(b) shows the posterior mean when $\sigma_n = 1$. For this value the posterior mean does not follow the data points exactly, but is not extremely different from the data points either.

This shows that it is important to chose a value for σ_n that is appropriate based on knowledge of the variation of the regular part of the time series. This example indicates that increasing σ_n allows larger variations in the posterior mean, whereas decreasing σ_n decreases the allowed variation.

3.3.3 Effect of model variance

The previous section considered the effect of varying σ_n. The other parameter which can be varied is σ_ϵ. This corresponds to the variance assumed for the white noise in the model. This section will use the same data set as generated in the previous section.

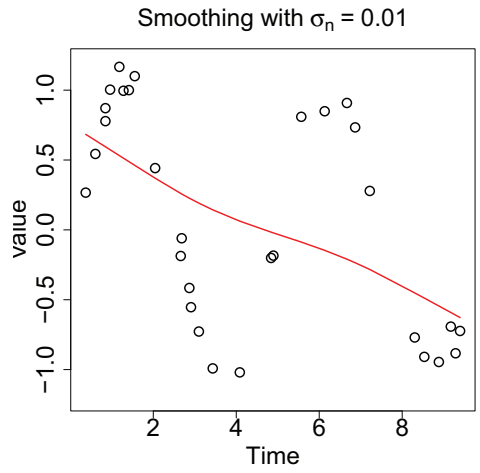

(a) Scaling factor 0.01.

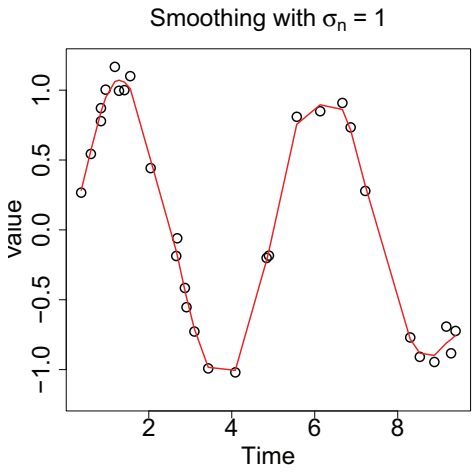

(b) Scaling factor 1.

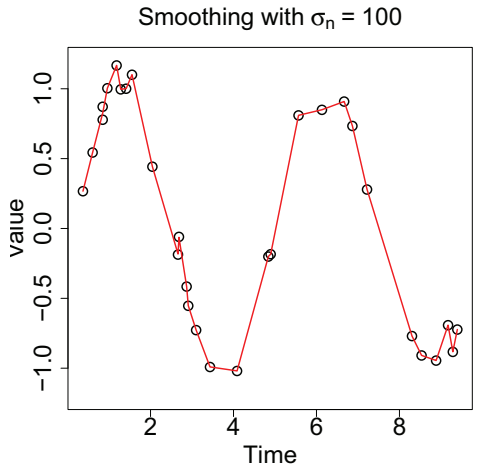

(c) Scaling factor 100.

Figure 3.1: Smoothing of sine-wave data with noise, with different scaling of the precision matrix for the underlying second-order random walk. The black circles are the observed values and the red line is the posterior mean. σ_ϵ was chosen equal to the true value, 0.10.

Figure 3.2 shows the posterior mean and marginal standard deviations for $\sigma_\epsilon = 0.02$, $\sigma_\epsilon = 0.10$ and $\sigma_\epsilon = 0.50$. The Figure shows that as the σ_ϵ parameter is increased the posterior mean is allowed to pass farther and farther away from the observations. However, in Figure 3.2(c) one can see that for $\sigma_\epsilon = 0.5$, the posterior mean lies below the observations for many consecutive observations and similarly above the observations for many consecutive observations. This does not fit well with the model assumption of white noise. Figure 3.2(a) shows that for $\sigma_\epsilon = 0.02$ the posterior mean is forced to pass so close to the observations that the posterior mean becomes "jagged". For $\sigma_\epsilon = 0.10$, which corresponds to setting the model variance equal to the true variance, the behavior appears better. The posterior mean is not overly "jagged" and does not appear to break the assumption of white noise.

These results indicate that it is important to use a model variance which corresponds to the true random variation in the data set. Because a too low value gives a "jagged" behavior, but a too large value gives a posterior distribution which appears to break the assumption of white noise.

3.3.4 Example 1: Missing data

As a first example consider again the data in Section 3.3.2, but with different subsets of data removed. Figure 3.3 shows the results after removing observations $\{12, \ldots, 21\}$ and after removing observations $\{16, \ldots, 20\}$. In Figure 3.3(b) the posterior mean seems to follow the missing observations well, but in Figure 3.3(a) the posterior mean starts to increase much earlier than the missing observations. Further one can see that removing observations seems to increase the marginal variance at the points at which there are no observations, but the variances do not appear to change much for most of the points with observations.

3.3.5 Example 2: Square wave with white noise

In this section the following two data sets will be considered. The first was generated by using a regular grid on $[0, 10]$ with 30 points and evaluating the square wave $x(t) = -\text{sgn}(\sin(4\pi t/10))$ at each point. Then one realization of the time series $y_1(t) = x(t)+\epsilon(t)$ was drawn on the grid, where $\epsilon(t)$ was iid Gaussian white noise with constant variance 0.01. Then the procedure was repeated with 20000 points in the regular grid.

Figure 3.4 shows the posterior mean based on each of these two data sets. Figure 3.4(a) shows that 30 points give a quite good approximation considering only 30 points were used, but as Figure 3.4(b) shows, the posterior mean is not able follow the behavior of the sudden jump even for 20000 observations. However, this is not unexpected as the prior is second-order random walk. A realization of a second-order random walk will almost surely not be a square-wave, as such the underlying model with a second-order random walk prior is not appropriate for these data sets. This shows that this type of smoothing can not be expected to capture the exact behavior at discontinuities.

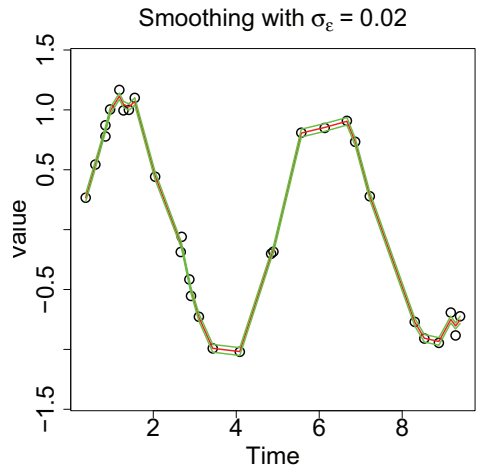

(a) Model standard deviation 0.02.

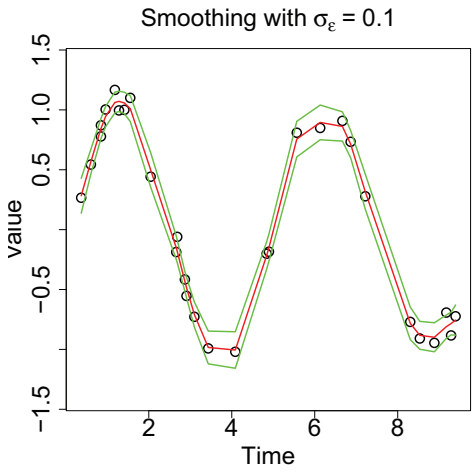

(b) Model standard deviation 0.10.

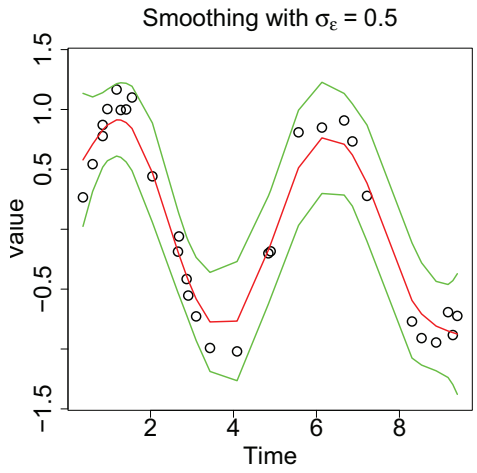

(c) Model standard deviation 0.50.

Figure 3.2: Smoothing of sine-wave data with noise, with use of different values for the model parameter σ_ϵ. The black circles are the observed values, the red line is the posterior mean and the green lines are ± 1.64 times the marginal standard deviations. σ_n was set to 1.

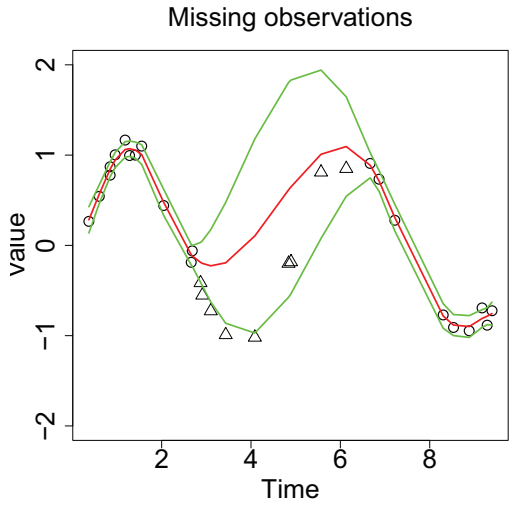

(a) Missing 9 observations.

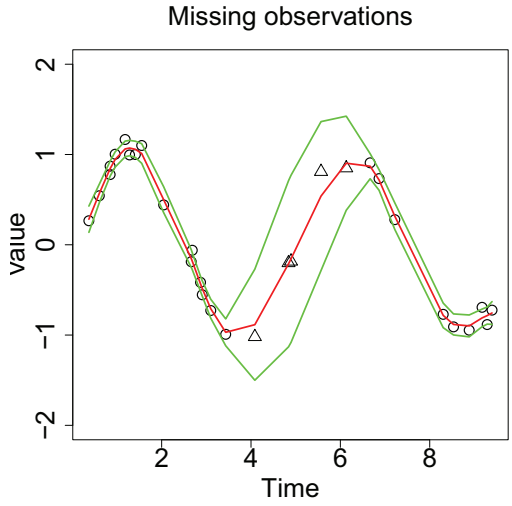

(b) Missing 5 observations.

Figure 3.3: Smoothing of a data set with a sine-wave with white noise. The observations used are marked with black circles, the observations not used are marked with black triangles, the red line indicates the posterior mean and the green lines are ± 1.64 times the marginal standard deviation. $\sigma_\epsilon = 0.1$ and $\sigma_n = 1$.

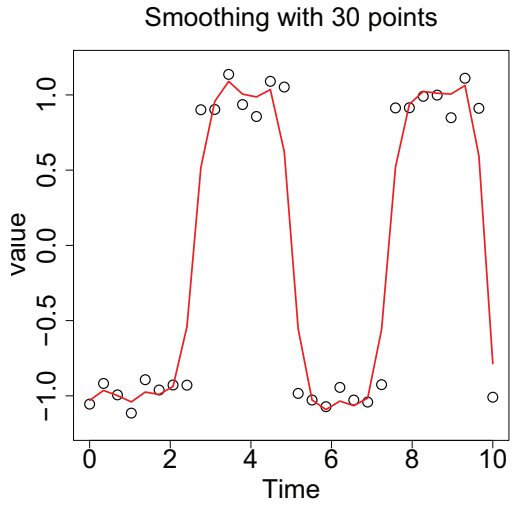

(a) 30 observations.

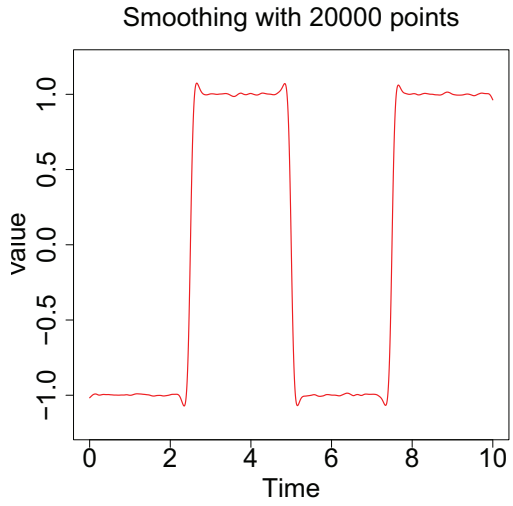

(b) 20000 observations.

Figure 3.4: Smoothing of two data sets of square-wave data with white noise. The first with 30 points, where both the black circles indicating the observations and the red line indicating the posterior mean are plotted, and the second with 20000 points, where only the posterior mean is plotted. $\sigma_\epsilon = 0.1$ and $\sigma_n = 1$.

3.3.6 Conclusion

The model considered in these sections requires the assumption that a time series $y(t)$ can be decomposed into a regular part $x(t)$ and a irregular white noise part $\epsilon(t)$ such that $y(t) = x(t) + \epsilon(t)$. By using second-order random walk as a prior for $x(t)$ certain assumptions are made about the type of behavior likely to be exhibited by $x(t)$. For example lack of discontinuity points and so on. As demonstrated in Section 3.3.5 increasing the number of points at which the time series is observed did not make the smoothed time series follow the exact behavior at the discontinuities.

Further as demonstrated in Section 3.3.2 and Section 3.3.3 the choice of scale parameter and model variance is important for the result of the smoothing. Misspecification can give a smoothed time series with very "jagged" behavior or time series which are very unlikely under the assumption that the observations were generated with white noise.

It is also possible to use this method to make predictions on the behavior at points with missing observations. This was done in Section 3.3.4 for a sine-wave with white noise.

Chapter 4

Stochastic heat equation

4.1 Finite volume methods

Finite volume methods are useful for creating discretizations of conservative laws of the form

$$\frac{\partial}{\partial t}q(x,t) + \nabla F(x,t) = f(x,t), \tag{4.1}$$

where ∇ is the spatial divergence operator. This equation relates the change of the size q in time to the spatial divergence of the flux $F(x,t)$ and the sink-/source-term $f(x,t)$.

The idea in the finite volume methods is to for a fixed t partition the spatial domain into cells, such that the cells "respect" the conservative law, in some sense. The main tool in these methods is the use of the divergence theorem

$$\int_{E_k} \nabla F \,\mathrm{d}V = \oint_{\partial E_k} F \cdot \mathbf{n} \,\mathrm{d}\sigma, \tag{4.2}$$

where $\mathbf{n}$ is the outer normal vector of the surface ∂E_k relative to E_k.

For non-stochastic differential equations a quite thorough treament of finite volume methods can be found in [2]. In this report a finite volume method will be applied to the following form of the stochastic heat equation

$$\frac{\partial}{\partial t}u(x,t) - \nabla \cdot \nabla u(x,t) = \sigma \mathcal{W}(x,t), \tag{4.3}$$

where $\nabla = \frac{\partial}{\partial x}$, $\sigma > 0$ is a constant and $\mathcal{W}(x,t)$ is space-time white noise.

4.2 Stochastic partial differential equation

The following sections will consider the SPDE given in Equation (4.3) on $(x,t) \in [0,L] \times [0,T]$, where $L > 0$ and $T > 0$ are constants. That is

$$\frac{\partial}{\partial t}u(x,t) - \nabla \cdot \nabla u(x,t) = \sigma \mathcal{W}(x,t), \qquad (x,t) \in [0,L] \times [0,T], \tag{4.4}$$

where $\nabla = \frac{\partial}{\partial x}$, $\sigma > 0$ is a constant and $\mathcal{W}(x,t)$ is space-time white noise.

4.2.1 Boundary conditions

Equation (4.4) does not have a unique solution, an arbitrary solution of

$$\frac{\partial}{\partial t}u(x,t) - \frac{\partial^2}{\partial x^2}u(x,t) = 0,$$

can be added to the solution. This means for example that an arbitrary constant can be added to the solution.

For the finite volume scheme which follows in the next section there are two things which must be known if the scheme is going to be used for sequential simulation in time. Some starting condition, $u(x,0)$ for $x \in [0,L]$, and an estimate of spatial flux at the boundaries, $\frac{\partial}{\partial x}u(0,t)$ and $\frac{\partial}{\partial x}u(L,t)$ for $t \in [0,T]$. But in the same manner as for the second-order random walk the interest does not lay in direct simulation, but rather using the solution as a prior distribution. This means that it is not necessary that the distribution is proper.

The problem with the initial condition can be solved by using some Gaussian initial condition for the starting process, $u(x,0) \sim u_0(x)$ for $x \in [0,L]$. That is for each discretization $0 = x_1 < x_2 < \cdots < x_n = L$, $(u(x_1,0),\ldots,u(x_n,0))$ has a Gaussian distribution with precision matrix Q_n, where the size of Q_n is $n \times n$ and the entries only depend on the locations of the points. Q_n does not need to have rank n.

It is hard to make any definite claims about the effects of different boundary conditions, therefore the next section will present the finite volume method on Equation (4.4) with the Neumann boundary conditions

$$\begin{aligned}\frac{\partial}{\partial x}u(0,t) &= 0 \qquad t \in [0,T]\\ \frac{\partial}{\partial x}u(L,t) &= 0 \qquad t \in [0,T],\end{aligned}$$

and an arbitrary Gaussian initial distribution with mean zero.

4.2.2 Approximation by a finite volume method

Consider Equation (4.4) with Neumann boundary conditions $\frac{\partial}{\partial x}u(0,t) = 0$ and $\frac{\partial}{\partial x}u(L,t) = 0$ for $t \in [0,T]$. In addition to $u(x,0) \sim u_0(x)$ for $x \in [0,L]$, where $u_0(x)$ is a Gaussian random field with mean zero.

Begin by choosing a spatial grid $0 = s_1 < \cdots < s_N = L$. Define $d_i = s_{i+1} - s_i$ for $i = 1,\ldots,N-1$. Continue by dividing $[0,L]$ into cells, by letting $E_1 = [s_1, s_1 + d_1/2]$, $E_N = [s_N - d_{N-1}/2, s_N]$ and $E_i = [s_i - d_{i-1}/2, s_i + d_i/2]$ for $i = 2,\ldots,N-1$. Observe that the cells fulfill the properties, $\cup_{i=1}^{n} E_i = [0,L]$ and $E_i \cap E_j = \emptyset$ for $i \neq j$. Further let V_i denote the length of interval E_i for $i = 1,\ldots,N$. For ease of notation let $s_{i+1/2} = s_i + d_i/2$ for $i = 1,\ldots,N-1$, $s_{1/2} = s_1$ and $s_{N+1/2} = s_N$. Further let $0 = t_0 < t_1 < \cdots < t_M = T$, be a regular grid in time and let $k = t_1 - t_0$.

Start by integrating Equation (4.4) over one time step, t_m to t_{m+1}, and one cell, E_n, this gives

$$\int_{E_n}\int_{t_m}^{t_{m+1}} \frac{\partial}{\partial t} u \,\mathrm{d}t\,\mathrm{d}x - \int_{t_m}^{t_{m+1}}\int_{E_n} \nabla\cdot\nabla u \,\mathrm{d}x\,\mathrm{d}t = \sigma\int_{t_m}^{t_{m+1}}\int_{E_n} \mathcal{W}\,\mathrm{d}x\,\mathrm{d}t.$$

By using the divergence theorem given in Equation (4.2) this can be written as

$$\begin{aligned}\int_{E_n}(u(x,t_{m+1})-u(x,t_m))\,\mathrm{d}x - \int_{t_m}^{t_{m+1}}\left(\frac{\partial}{\partial x}u(s_{n+1/2},t)-\frac{\partial}{\partial x}u(s_{n-1/2},t)\right)\mathrm{d}t = \\ = \sqrt{kV_n}\sigma z_n^{m+1},\end{aligned} \tag{4.5}$$

where z_n^{m+1} is a standard Gaussian random variable. No approximations have been made yet, but to continue from this point it is necessary to approximate both the terms on the left hand side of the equation.

First consider the first integral in Equation (4.5). For simplicity assume $u(x,t_m) = U_n^m$ for all $x \in E_n$, that is for a fixed t_m, the solution u is constant on each cell. This means in some sense that the value of $u(s_n,t_m)$ is approximated by the mean of the cell. This gives

$$\int_{E_n}(u(x,t_{m+1})-u(x,t_m))\,\mathrm{d}x = V_n(U_n^{m+1}-U_n^m).$$

Secondly consider the flux terms $\frac{\partial}{\partial x}u(s_{n+1/2},t)$ at the boundaries of each cell. The following simple first order approximation will be used for the the flux at the internal boundaries, i.e. between two cells,

$$\frac{\partial}{\partial x}u(s_{n+1/2},t) = \frac{u(s_{n+1},t)-u(s_n,t)}{d_n}, \qquad n=1,\ldots,N-1.$$

Lastly the second integral in Equation (4.5) must be considered. The integral cannot be solved analytically, therefore a time discretization scheme must be used. Here the implicit Euler method will be used, this gives

$$\begin{aligned}\int_{t_m}^{t_{m+1}}\left(\frac{\partial}{\partial x}u(s_{n+1/2},t)-\frac{\partial}{\partial x}u(s_{n-1/2},t)\right)\mathrm{d}t = \\ = k\left(\frac{\partial}{\partial x}u(s_{n+1/2},t_{m+1})-\frac{\partial}{\partial x}u(s_{n-1/2},t_{m+1})\right).\end{aligned}$$

Inserting all these approximations into Equation (4.5) and considering an internal cell, that is $n \in \{2,\ldots,N-1\}$, gives

$$V_n(U_n^{m+1}-U_n^m) - k\left(\frac{U_{n+1}^{m+1}-U_n^{m+1}}{d_n} - \frac{U_n^{m+1}-U_{n-1}^{m+1}}{d_{n-1}}\right) = \sqrt{kV_n}\sigma z_n^{m+1}. \tag{4.6}$$

For the exterior cells the Neumann boundary conditions

$$\frac{\partial}{\partial x}u(s_{1/2},t) = 0$$

and

$$\frac{\partial}{\partial x}u(s_{N+1/2},t)=0,$$

will be used, that is zero flux at the boundaries. Using this the discretization at the boundary cells can be written as

$$V_1(U_1^{m+1}-U_1^m)-k\frac{U_2^{m+1}-U_1^{m+1}}{d_1}=\sqrt{kV_1}\sigma z_1^{m+1}$$

and

$$V_n(U_n^{m+1}-U_n^m)+k\frac{U_n^{m+1}-U_{n-1}^{m+1}}{d_n}=\sqrt{kV_N}\sigma z_N^{m+1}.$$

The scheme can be written in matrix form as

$$D(U^{m+1}-U^m)-kHU^{m+1}=\sqrt{k}D^{1/2}\sigma Z_{m+1}, \tag{4.7}$$

where $D=\mathrm{diag}(V_1,\ldots,V_N)$, $Z_{m+1}\sim N_N(0,I_N)$, $U^m=(U_1^m,\ldots,U_N^m)$ and H is a matrix with

$$H_{i,j}=\begin{cases}-\frac{1}{d_1} & i=j=1\\ -(\frac{1}{d_{i-1}}+\frac{1}{d_i}) & 1<i=j<N\\ -\frac{1}{d_{n-1}} & i=j=N\\ \frac{1}{d_i} & j=i+1\\ \frac{1}{d_j} & j=i-1\\ 0 & \text{otherwise}\end{cases}.$$

Equation (4.7) can be rewritten as

$$Z_{m+1}=\frac{1}{\sqrt{k}\sigma}D^{-1/2}\left((D-kH)U^{m+1}-DU^m\right). \tag{4.8}$$

In this form one can consider the joint distribution of $Z_1,\ldots,Z_M$, which are independent, and find the joint distribution of $U^0,\ldots,U^M$. By assumptions U_0 has some Gaussian distribution with precision matrix $A\geq 0$ and the joint distribution of $U^0,\ldots,U^M$ can be found by

$$\pi(U^0,\ldots,U^M)\propto\pi(U^0)\pi(Z_1,\ldots,Z_M|U_0)\propto$$
$$\propto\exp\left(-\frac{1}{2}\left((U^0)^{\mathrm{T}}AU^0+Z_1^{\mathrm{T}}Z_1+\ldots+Z_M^{\mathrm{T}}Z_M\right)\right).$$

From this one can deduce that the precision matrix, Q, of $U=\left((U^0)^{\mathrm{T}},\ldots,(U^M)^{\mathrm{T}}\right)$ is a $N(M+1)\times N(M+1)$ block tridiagonal matrix with blocks of size $N\times N$. Let $B=D-kH$, then Q can be written as

$$Q=\frac{1}{\sigma^2k}\begin{bmatrix}D+A\sigma^2k & -B & 0 & \cdots & 0\\ -B^{\mathrm{T}} & D+B^{\mathrm{T}}D^{-1}B & -B & & 0\\ 0 & \ddots & \ddots & \ddots & \vdots\\ \vdots & & -B^{\mathrm{T}} & D+B^{\mathrm{T}}D^{-1}B & -B\\ 0 & \cdots & 0 & -B^{\mathrm{T}} & B^{\mathrm{T}}D^{-1}B\end{bmatrix}. \tag{4.9}$$

4.3 Properties of the approximation

4.3.1 The precision matrix

Let Q be the precision matrix in Equation (4.9), with N nodes in the spatial grid and $M+1$ nodes in the time grid. The following lemma will help find the rank and null space of Q.

Lemma 4.3.1. *Let $u : \mathbb{R}^N \to \mathbb{R}^{N(M+1)}$ be given by*

$$u(a) = \left(a^T, (B^{-1}Da)^T, \ldots, ((B^{-1}D)^M a)^T\right),$$

where B and D are as in Section 4.2.2.

Then u is a linear injection, and $\{u(a_1), u(a_2), \ldots, u(a_l)\}$ is a set of linear independent vectors iff $\{a_1, a_2 \ldots, a_l\}$ is a set of linearly independent vectors.

Proof. The linearity follows from the fact that

$$u(a) = Ma,$$

where M is the block diagonal matrix given by $M = \operatorname{diag}(I_N, B^{-1}D, \ldots, (B^{-1}D)^M)$.

Observe that $u(0) = 0$ and that if $u(a)$ is divided into $M+1$ blocks of size N, then the first block of $u(a)$ is equal to a. This means that $u(a) = 0$ iff $a = 0$, i.e. u is injective.

The linear independence follows from the one-to-one property and the linearity of u, since

$$c_1u(a_1) + c_2u(a_2) + \ldots + c_lu(a_l) = u(c_1a_1 + c_2a_2 + \ldots + c_la_l) = 0$$

iff

$$c_1a_1 + c_2a_2 + \ldots + c_la_l = 0.$$

This implies that $\{u(a_1), u(a_2), \ldots, u(a_l)\}$ is a set of linearly independent vectors iff $\{a_1, a_2, \ldots, a_l\}$ is a set of linearly independent vectors □

The function u defined in Lemma 4.3.1 corresponds to the evolution of the mean using the discretization scheme from the previous section. Let $\mu_k = \mathrm{E}(U^k)$, and take the expected value on both sides of Equation (4.7), then $u(\mu_0) = (\mu_0^{\mathrm{T}}, \mu_1^{\mathrm{T}}, \ldots, \mu_M^{\mathrm{T}})$. The following proposition shows that this function is closely connected to the eigenvectors of Q.

Proposition 4.3.2. *Let Q be the precision matrix found in Section 4.2.2, with size $N(M+1) \times N(M+1)$ and $N \times N$ size blocks, and let $u : \mathbb{R}^N \to \mathbb{R}^{N(M+1)}$ be the function given in Lemma 4.3.1.*

Then $\operatorname{Null}(Q) = \{u(a) : Aa = 0, a \in \mathbb{R}^N\}$ and the nullity of Q is equal to the nullity of A.

Proof. Let $x \in \mathbb{R}^{N(M+1)}$ and divide x into $M+1$ blocks of size N.

$Qx = 0$ iff each block of Qx is zero. The last block is 0 iff

$$-B^{\mathrm{T}}x_{M-1} + B^{\mathrm{T}}D^{-1}Bx_M = 0,$$

which is true iff $x_M = B^{-1}Dx_{M-1}$. This is true since B is strictly diagonally dominant for all time step sizes and therefore always invertible. Assume $x_l = D^{-1}Bx_{l-1}$ for a fixed $l \in \{M, \dots, 2\}$. Then block l of Qx is 0 iff

$$\begin{aligned} 0 = -B^{\mathrm{T}}x_{l-2} + (D + B^{\mathrm{T}}D^{-1}B)x_{l-1} - Bx_l = \\ = -B^{\mathrm{T}}x_{l-2} + (D + B^{\mathrm{T}}D^{-1}B)x_{l-1} - BB^{-1}Dx_{l-1}, \end{aligned}$$

which is true iff $x_{l-1} = B^{-1}Dx_{l-2}$. Thus by mathematical induction the last M blocks of Qx is equal to 0 iff $x_l = B^{-1}Dx_{l-1}$ for $l = M, \dots, 1$.

This means that any x such that $Qx = 0$ must be on the form $x = u(a)$ for some $a \in \mathbb{R}^N$. But if this is true $Qx = 0$ iff the first block of Qx is equal to zero, that is

$$(D + Ak\sigma^2)a - BB^{-1}Da = Aa = 0.$$

Therefore $Qx = 0$ iff $x = u(a)$ for an $a \in \mathbb{R}^N$ such that $Aa = 0$. Thus the null space of Q is given by

$$\mathrm{Null}(Q) = \{u(a) : Aa = 0, a \in \mathbb{R}^N\} \subset \mathbb{R}^{N(M+1)}.$$

From Lemma 4.3.1 the function u is a injective linear transformation. This means that the dimension of $\{u(a) : Aa = 0, a \in \mathbb{R}^N\}$ is equal to the dimension of $\{a : Aa = 0, a \in \mathbb{R}^N\}$ and the nullity of Q and A are the same. □

The proposition shows that $\mathrm{rank}(Q) = NM + \mathrm{rank}(A)$. Let $j = (1, \dots, 1) \in \mathbb{R}^N$, then $D^{-1}Bj = j$, since j is in the null space of H, which means that $j = B^{-1}Dj$. Thus it follows that $(1, \dots, 1) \in \mathbb{R}^{N(M+1)}$ is in the null space of Q iff j is in the null space of A.

4.3.2 Total energy

For the purpose of this section call $\int_0^L u(x,t)\,\mathrm{d}x$ the total energy or total amount of u at time t. This section will show the relationship between the total energy according to Equation (4.4) and the discretization scheme.

From Equation (4.4) one has

$$\int_0^L \int_{t_0}^{t_1} \frac{\partial}{\partial s} u(x,s)\,\mathrm{d}s\,\mathrm{d}x = \int_0^L \int_{t_0}^{t_1} \left(\frac{\partial^2}{\partial x^2} u(x,s) + \sigma \mathcal{W}(x,s) \right) \mathrm{d}s\,\mathrm{d}x,$$

for $0 \le t_0 \le t_1 \le T$. From the divergence theorem in Equation (4.2) this can be written as

$$\int_0^L (u(x,t_1) - u(x,t_0))\,\mathrm{d}x = \int_{t_0}^{t_1} \left(\frac{\partial}{\partial x} u(L,s) - \frac{\partial}{\partial x} u(0,s) \right) \mathrm{d}s + \sigma\sqrt{L}(W(t_1) - W(t_0)),$$

where $W(t)$ is a standard Wiener process. Using the Neumann boundary conditions this becomes

$$\int_0^L u(x,t_1)\,\mathrm{d}x = \int_0^L u(x,t_0)\,\mathrm{d}x + \sigma\sqrt{L}(W(t_1) - W(t_0)).$$

This means that

$$\int_0^L u(x,t)\,\mathrm{d}x = \sigma\sqrt{L}W(t) + \int_0^L u(x,0)\,\mathrm{d}x, \qquad 0 \le t \le T.$$

This gives

$$\int_0^L u(x,t)\,\mathrm{d}x - \int_0^L u(x,0)\,\mathrm{d}x = \sigma\sqrt{L}W(t). \tag{4.10}$$

Let $0 = t_0 < t_1 < \cdots < t_M = T$ be a regular grid in time and $0 = s_1 < \cdots < s_N = L$ a grid in space. Let $k = t_1 - t_0$ and let $B = D - kH$ as in Section 4.2.2, then B is always invertible. Let U_n^m be the approximate solution of $u(s_n, t_m)$ and denote $U^m = (U_1^m, \ldots, U_N^m)$. Let $\tilde{u}(x, t_m)$ be equal to U_n^m for $x \in E_n$, that is constant on each cell for each time. Then for the approximate solution

$$\int_0^L \tilde{u}(x,t_m)\,\mathrm{d}x = j^{\mathrm{T}} D U^m,$$

where $j = (1, \ldots, 1) \in \mathbb{R}^N$.

First observe that

$$j^{\mathrm{T}} B D^{-1} = j^{\mathrm{T}}(I_N - kHD^{-1}) = j^{\mathrm{T}}$$

since $j^{\mathrm{T}} H = 0$. This means that

$$j^{\mathrm{T}} = j^{\mathrm{T}} D B^{-1},$$

which in turn gives

$$j^{\mathrm{T}} D = j^{\mathrm{T}} D B^{-1} D. \tag{4.11}$$

Using Equation (4.7) one has

$$U^{m+1} = B^{-1} D(U^m + \sigma\sqrt{k} D^{-1/2} Z_{m+1}),$$

which gives

$$j^{\mathrm{T}} D U^{m+1} = j^{\mathrm{T}} D B^{-1} D(U^m + \sigma\sqrt{k} D^{-1/2} Z_{m+1}) = j^{\mathrm{T}} D(U^m + \sigma\sqrt{k} D^{-1/2} Z_{m+1}). \tag{4.12}$$

Define $R_m = j^{\mathrm{T}} D U^m$ and observe that

$$\mathrm{Var}(\sigma\sqrt{k} j^{\mathrm{T}} D^{1/2} Z_{m+1}) = \sigma^2 k j^{\mathrm{T}} D j = \sigma^2 k L.$$

This means that

$$R_{m+1} = R_m + \sigma\sqrt{kL} z_{m+1}, \qquad m = 0, \ldots, M-1,$$

where $z_1, \ldots, z_M$ are iid standard Gaussian variables. Thus it can be written in the form

$$\int_0^L \tilde{u}(x, t_m)\,\mathrm{d}x - \int_0^L \tilde{u}(x, 0)\,\mathrm{d}x = \sigma\sqrt{L}\tilde{W}(t_m), \tag{4.13}$$

where $\tilde{W}(t)$ is a standard Wiener process.

Observe further that in Equation (4.10)

$$W(t) = \int_0^t \int_0^L \mathcal{W}(x, t)\,\mathrm{d}x\,\mathrm{d}t$$

and that from Equation (4.5) and Equation (4.12)

$$\begin{aligned}
\tilde{W}(t_m) &= \sum_{i=1}^{m} \sqrt{k} j^{\mathrm{T}} D^{1/2} Z_i \\
&= \sum_{i=1}^{m} \sqrt{k} \sum_{j=1}^{N} \sqrt{V_j} z_j^i \\
&= \sum_{i=1}^{m} \sum_{j=1}^{N} \sqrt{k}\sqrt{V_j} \frac{1}{\sqrt{V_j k}} \int_{t_{i-1}}^{t_i} \int_{E_j} \mathcal{W}(x, t)\,\mathrm{d}x\,\mathrm{d}t \\
&= \sum_{i=1}^{m} \sum_{j=1}^{N} \int_{t_{i-1}}^{t_i} \int_{E_j} \mathcal{W}(x, t)\,\mathrm{d}x\,\mathrm{d}t \\
&= \int_0^{t_m} \int_0^L \mathcal{W}(x, t)\,\mathrm{d}x\,\mathrm{d}t = W(t_m)
\end{aligned}$$

This means that Equation (4.13) gives the discrete equivalent of Equation (4.10). Let the approximate solution be constant on each cell and use the initial condition $u(x, 0) \overset{d}{=} u_0(x)$ and $\tilde{u}(x, 0) \sim \tilde{u}_0(x)$ for $x \in [0, L]$, then

$$\int_0^L u(x, t_m)\,\mathrm{d}x - \int_0^L \tilde{u}(x, t_m)\,\mathrm{d}x = \int_0^L (u(x, 0) - \tilde{u}(x, 0))\,\mathrm{d}x, \qquad m = 0, \ldots, M,$$

If the initial conditions are chosen as U_n^0 equal to the mean of $u_0(x)$ in cell E_n, then

$$\int_0^L u(x, t_m)\,\mathrm{d}x \overset{d}{=} \int_0^L \tilde{u}(x, t_m)\,\mathrm{d}x, \qquad m = 0, \ldots, M.$$

This means that for this choice of distribution for U^0 the discretization scheme gives the correct distribution for the total energy of u at the time steps $t_0, \ldots, t_M$. Further if

$$\mathrm{E}\left[\left|\int_0^L u(x, t_0)\,\mathrm{d}x - \int_0^L \tilde{u}(x, t_0)\right|^2\right] = C \neq 0,$$

then the mean square error between the total energy of the true solution and the discretization scheme is the same for all time steps independent of the number of time steps used.

4.3.3 Initial condition

As shown in Proposition 4.3.2 the precision matrix has the same nullity as A. This means that a diffuse prior for U^0 gives a nullity of N and it can be hard to make the distribution proper when conditioning on variables. Therefore this is not a good strategy unless one condition on U^0, with hard or soft constraints.

Another choice is to use a discretization of some known stochastic process $u_0(x)$. For example the second order random walk studied in Chapter 3. This choice gives the total precision matrix a nullity of 2. In a strict sense this is not allowed unless one conditions the second-order random walk on zero derivative for $x = 0$ and $x = L$, but this will be ignored.

4.4 Examples

4.4.1 Second-order random walk as initial condition

Consider the stochastic heat equation on $(x, t) \in [0, 5] \times [0, 10]$ with the Neumann boundary conditions $\frac{\partial}{\partial x}u(0, t) = \frac{\partial}{\partial x}u(5, t) = 0$ for $t \in [0, 10]$ with $\sigma = 1$. Under the initial condition $u(x, 0)$ is a second-order random walk with $\sigma_n = 1$.

Let $0 = x_1 < \cdots < x_{100} = 5$ be a regular spatial grid and $0 = t_0 \cdots < t_{399} = 10$ be a regular time grid. Let $U_n^m = u(x_n, t_m)$ and use the precision matrix found in Section 4.2.2 with $\sigma = 1$ and A equal to the precision matrix found in the section for second-order random walk with $\sigma_n = 1$. This gives 40000 variables.

By Proposition 4.3.2 this gives a precision matrix with nullity 2, so it is necessary to condition on some variables, either with soft or with hard constraints. To simplify simulation assume the values were measured with Gaussian iid errors with standard deviation 0.00001

$$\begin{aligned} u(x_{10}, t_0) &= U_{10}^0 = 10, \\ u(x_{90}, t_1) &= U_{90}^1 = 8, \\ u(x_{80}, t_{50}) &= U_{80}^{50} = 8, \\ u(x_{40}, t_{200}) &= U_{40}^{200} = 7, \\ u(x_{20}, t_{350}) &= U_{20}^{350} = 9, \\ u(x_{60}, t_{390}) &= U_{60}^{390} = 4. \end{aligned}$$

Figure 4.1(a) shows the marginal variances of the conditional distribution. As expected the variances at the observed points are close to zero, but it appears that the variance increases more slowly in the spacial direction than in the temporal direction near each of the observed point.

Figure 4.1(b) shows the mean of the conditional distribution. It appears that the differences in the mean in space for a given time equalizes in a couple of time units for this distribution between observations. This is connected with the ratio between the constant in front of the $\frac{\partial}{\partial t}u(x, t)$ term and the constant in front of the $\frac{\partial^2}{\partial x^2}u(x, t)$ term in

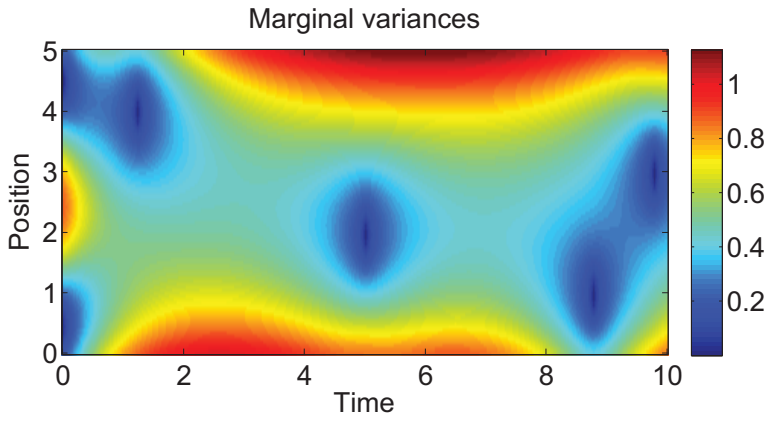

(a) Marginal variances.

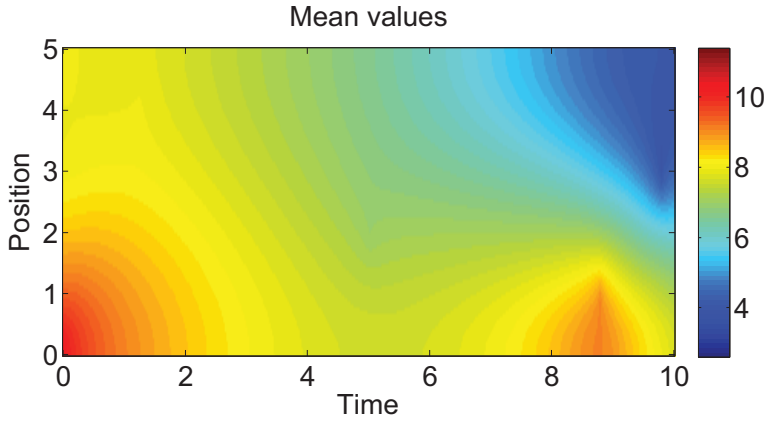

(b) Mean values.

Figure 4.1: Mean values and marginal variances as a function of time and position for the conditional distribution with $\sigma = 1$ in the stochastic heat equation and $\sigma_n = 1$ in the second-order random walk initial condition.

Equation (4.4). In this distribution the ratio is one since a factor of one is used for both terms, but decreasing the factor in front of the $\frac{\partial^2}{\partial x^2}u(x,t)$ term relative to the factor in front of the $\frac{\partial}{\partial t}u(x,t)$ term would decrease the spatial flux, relative to the time scale, and increase the time it would take for the mean to stabilize.

Figure 4.2(a) and Figure 4.2(b) shows two different realizations from the conditional distribution as functions of time and space. The figures show that the realizations from the conditional distribution do not appear to look like ”smoothly” changing functions of time and space. The realizations look noisy.

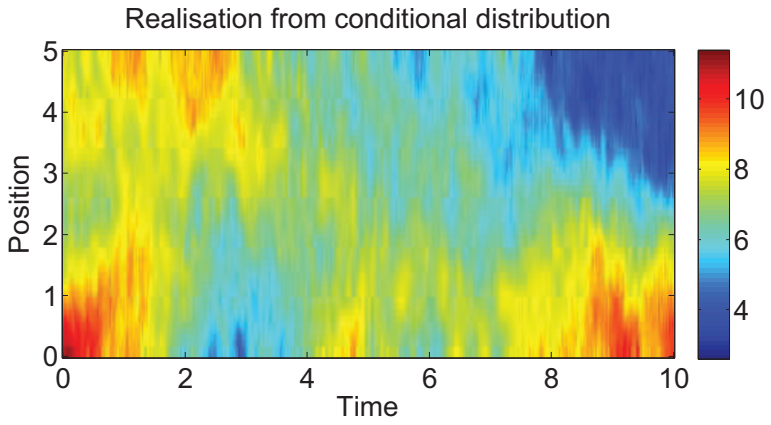

(a) A realization from the conditional distribution.

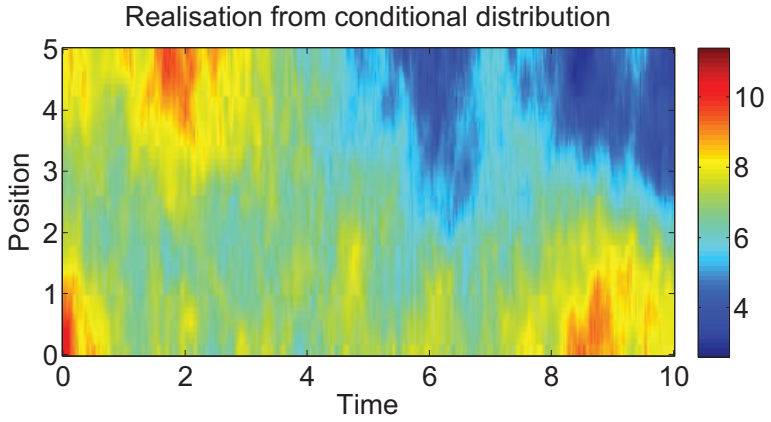

(b) A realization from the conditional distribution.

Figure 4.2: Two realizations as a function of time and position from the conditional distribution with $\sigma = 1$ in the stochastic heat equation and $\sigma_n = 1$ in the second-order random walk initial condition.

4.4.2 Fixed initial condition

A 100×400 grid was constructed in the same way as in Section 4.4.1, with $U_n^m = u(x_n, t_n)$ and the initial condition

$$u_0(x) = \begin{cases} 10 & 0 \leq x < \frac{5}{2} \\ -10 & \frac{5}{2} \leq x \leq 5 \end{cases}. \tag{4.14}$$

Since the initial condition specifies exact values for U^0, the specification of A is irrelevant. This is because only the conditional distribution $\pi(U^1, \ldots, U^M | U^0)$ is needed. To make it easy to use the model in the same fashion as in the previous example $A = 0$ will be used and the distribution will be conditioned on observed $U_n^0 + \epsilon_n$ for $n = 1, \ldots, N$, where ϵ_n are iid Gaussian with mean 0 and standard deviation 0.00001. That is the values are observed nearly exact.

Figure 4.3(a) shows the marginal variances of the conditional distribution. The variances start at zero at time zero and then increases with time. The variances appear to increase most quickly at the edges and more slowly in the middle of the spatial domain.

Figure 4.3(b) shows the the mean of the conditional distribution. Since an exact initial condition was used, this corresponds to the solution of the non-stochastic differential equation

$$\frac{\partial}{\partial t} u(x,t) = \frac{\partial^2}{\partial x^2} u(x,t),$$

with the Neumann boundary condition previously specified. This equation is a conservative law and many properties of this equation is known. It corresponds to a diffusion where no "mass" can leave the boundaries. The solution will stabilize to a constant level in space. The equation only redistributes the "mass" to a constant level in space.

Figure 4.4(a) and Figure 4.4(b) shows two different realizations from the conditional distribution as functions of time and space. These realizations illustrate that as the time increases the mean of the variables tends to zero, and that there is a balance between random fluctuations and equalization of nearby values.

4.5 Conclusion

Section 4.2.2 shows that the finite volume method can be used to construct a GMRF approximation to the discretized solution of one form of the stochastic heat equation.

Section 4.3.2 shows that the discretization scheme created by the finite volume method gives the correct distribution of the total "energy" of the solution at each time, that is the integral of the solution over the spatial domain, when it is used to sequentially simulate the solution from a fixed initial condition.

The discretization scheme appears to give good results for a fixed initial state, as demonstrated in Section 4.4.2, and for one special case where the initial condition is second-order random walk, as demonstrated in Section 4.4.1. The example with second-order random walk also shows that the conditional distribution has reasonable properties when one conditions on values at some locations in time and space. The variances are lower around the observed values and the mean passes through the observed values.

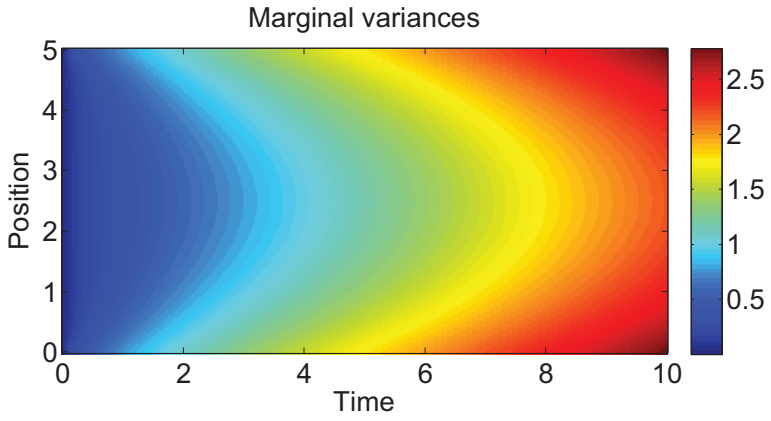

(a) Marginal variances.

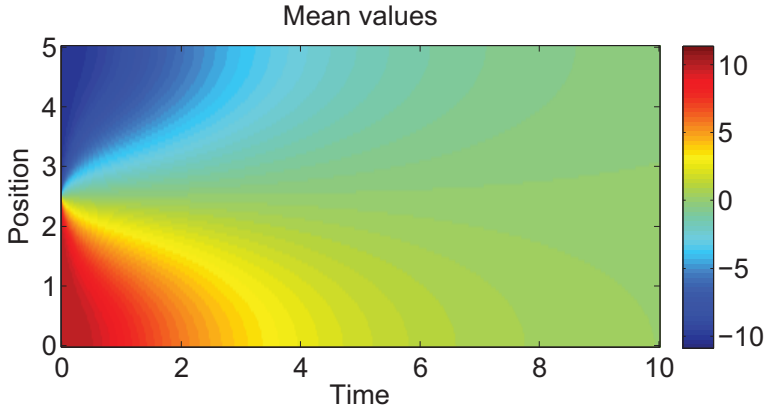

(b) Mean values.

Figure 4.3: Mean values and marginal variances as a function of time and position for the conditional distribution with $\sigma = 1$ in the stochastic heat equation and the initial condition specified in Equation (4.14).

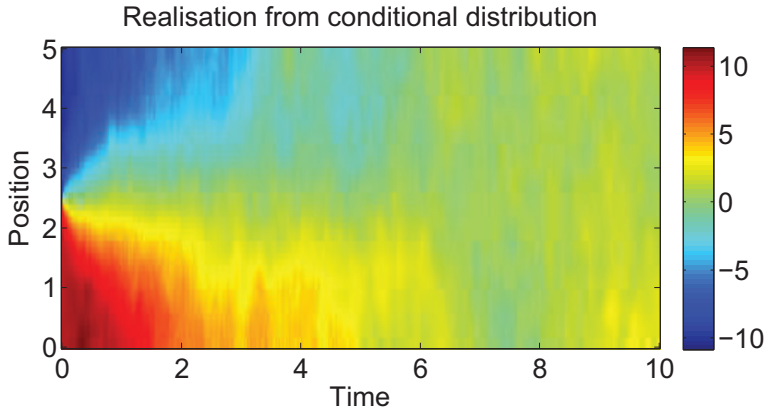

(a) A realization from the conditional distribution.

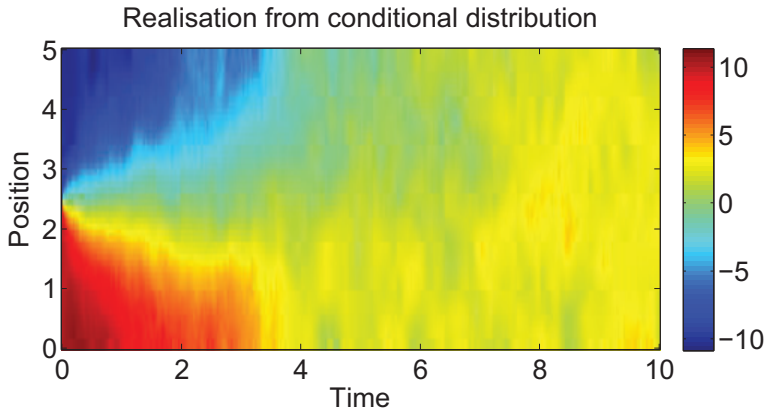

(b) A realization from the conditional distribution.

Figure 4.4: Mean values and marginal variances as a function of time and position for the conditional distribution with $\sigma = 1$ in the stochastic heat equation and the initial condition specified in Equation (4.14).

The time discretization by an implicit Euler scheme appears to give good results for the joint distribution when the SPDE in Equation (4.4) is solved with the Neumann boundary conditions of zero flux at the spatial boundaries for all times.

Chapter 5

Conclusion

The use of Gaussian Markov random field approximations for the solution of stochastic partial differential equation gives joint distributions with sparse precision matrices. These kinds of distributions can be handled well by numerical tools for sparse matrices. This means that such operations as simulations and calculation of marginal variances can be done more efficiently than when using the true precision matrix, or the covariance matrix of the solution which often will be dense.

The construction of the approximation as a Gaussian multivariate distribution means that the solutions can be used easily as a part of a greater model, for example the use of second-order random walk as a prior model when smoothing time-series.

This report shows that this type of construction is also possible for non-stationary stochastic partial differential equations, by the use of a spatial discretization to create a system of stochastic differential equations only in time and solving this system by a time discretization. This type of construction shows good results and merits further study.

Bibliography

[1] Yanqing Chen, Timothy A. Davis, William W. Hager, and Sivasankaran Rajamanickam. Algorithm 887: Cholmod, supernodal sparse cholesky factorization and update/downdate. *ACM Trans. Math. Softw.*, 35(3):1–14, 2008.

[2] R. Eymard, T. Gallouët, and R. Herbin. Finite volume methods. *Handbook of numerical analysis*, 7:713–1018, 2000.

[3] A.E. Gelfand, M. Fuentes, P. Diggle, and P. Guttorp. *Handbook of spatial statistics*. CRC, 2010.

[4] Gaël Guennebaud, Benoît Jacob, et al. Eigen v3. http://eigen.tuxfamily.org, 2010.

[5] F. Lindgren, J. Lindström, and H. Rue. An explicit link between Gaussian fields and Gaussian Markov random fields: The SPDE approach. *Preprint Statistics*, 5, 2010.

[6] F. Lindgren and H. Rue. On the Second-Order Random Walk Model for Irregular Locations. *Scandinavian journal of statistics*, 35(4):691–700, 2008.

[7] H. Rue and L. Held. *Gaussian Markov random fields: theory and applications*. Chapman & Hall, 2005.