
Copula Processes

Andrew Gordon Wilson*
Department of Engineering
University of Cambridge
agw38@cam.ac.uk

Zoubin Ghahramani†
Department of Engineering
University of Cambridge
zoubin@eng.cam.ac.uk

Abstract

We define a copula process which describes the dependencies between arbitrarily many random variables independently of their marginal distributions. As an example, we develop a stochastic volatility model, Gaussian Copula Process Volatility (GCPV), to predict the latent standard deviations of a sequence of random variables. To make predictions we use Bayesian inference, with the Laplace approximation, and with Markov chain Monte Carlo as an alternative. We find our model can outperform GARCH on simulated and financial data. And unlike GARCH, GCPV can easily handle missing data, incorporate covariates other than time, and model a rich class of covariance structures.

Imagine measuring the distance of a rocket as it leaves Earth, and wanting to know how these measurements correlate with one another. How much does the value of the measurement at fifteen minutes depend on the measurement at five minutes? Once we've learned this correlation structure, suppose we want to compare it to the dependence between measurements of the rocket's velocity. To do this, it is convenient to separate dependence from the marginal distributions of our measurements. At any given time, a rocket's distance from Earth could have a Gamma distribution, while its velocity could have a Gaussian distribution. And separating dependence from marginal distributions is precisely what a copula function does.

While copulas have recently become popular, especially in financial applications [1, 2], as Nelsen [3] writes, "the study of copulas and the role they play in probability, statistics, and stochastic processes is a subject still in its infancy. There are many open problems. . ." Typically only bivariate (and recently trivariate) copulas are being used and studied. In our introductory example, we are interested in learning the correlations in different stochastic processes, and comparing them. It would therefore be useful to have a *copula process*, which can describe the dependencies between arbitrarily many random variables independently of their marginal distributions. We define such a process. And as an example, we develop a stochastic volatility model, *Gaussian Copula Process Volatility* (GCPV). In doing so, we provide a Bayesian framework for the learning the marginal distributions and dependency structure of what we call a *Gaussian copula process*.

The *volatility* of a random variable is its standard deviation. Stochastic volatility models are used to predict the volatilities of a *heteroscedastic* sequence – a sequence of random variables with different variances, like distance measurements of a rocket as it leaves the Earth. As the rocket gets further away, the variance on the measurements increases. Heteroscedasticity is especially important in econometrics; the returns on equity indices, like the S&P 500, or on currency exchanges, are heteroscedastic. Indeed, in 2003, Robert Engle won the Nobel Prize in economics "for methods of analyzing economic time series with time-varying volatility". GARCH [4], a generalized version of Engle's ARCH, is arguably unsurpassed for predicting the volatility of returns on equity indices and currency exchanges [5, 6, 7]. GCPV can outperform GARCH, and is competitive on financial data that especially suits GARCH [8, 9, 10]. Before discussing GCPV, we first introduce copulas and the copula process. For a review of Gaussian processes, see Rasmussen and Williams [11].

*<http://mlg.eng.cam.ac.uk/andrew>

†Also at the machine learning department at Carnegie Mellon University.

1 Copulas

Copulas are important because they separate the dependency structure between random variables from their marginal distributions. Intuitively, we can describe the dependency structure of any multivariate joint distribution $H(x_1, \dots, x_n) = P(X_1 \leq x_1, \dots, X_n \leq x_n)$ through a two step process. First we take each univariate random variable X_i and transform it through its cumulative distribution function (cdf) F_i to get $U_i = F_i(X_i)$, a uniform random variable. We then express the dependencies between these transformed variables through the n -copula $C(u_1, \dots, u_n)$. Formally, an n -copula $C : [0, 1]^n \rightarrow [0, 1]$ is a multivariate cdf with uniform univariate marginals: $C(u_1, u_2, \dots, u_n) = P(U_1 \leq u_1, U_2 \leq u_2, \dots, U_n \leq u_n)$, where U_1, U_2, \dots, U_n are standard uniform random variables. Sklar [12] precisely expressed our intuition in the theorem below.

Theorem 1.1. *Sklar's theorem*

Let H be an n -dimensional distribution function with marginal distribution functions F_1, F_2, \dots, F_n . Then there exists an n -copula C such that for all $(x_1, x_2, \dots, x_n) \in [-\infty, \infty]^n$,

$$H(x_1, x_2, \dots, x_n) = C(F_1(x_1), F_2(x_2), \dots, F_n(x_n)) = C(u_1, u_2, \dots, u_n). \quad (1)$$

If F_1, F_2, \dots, F_n are all continuous then C is unique; otherwise C is uniquely determined on $\text{Range } F_1 \times \text{Range } F_2 \times \dots \times \text{Range } F_n$. Conversely, if C is an n -copula and F_1, F_2, \dots, F_n are distribution functions, then the function H is an n -dimensional distribution function with marginal distribution functions F_1, F_2, \dots, F_n .

As a corollary, if $F_i^{(-1)}(u) = \inf\{x : F_i(x) \geq u\}$, the quasi-inverse of F_i , then for all $u_1, u_2, \dots, u_n \in [0, 1]^n$,

$$C(u_1, u_2, \dots, u_n) = H(F_1^{(-1)}(u_1), F_2^{(-1)}(u_2), \dots, F_n^{(-1)}(u_n)). \quad (2)$$

In other words, (2) can be used to construct a copula. For example, the bivariate Gaussian copula is defined as

$$C(u, v) = \Phi_\rho(\Phi^{-1}(u), \Phi^{-1}(v)), \quad (3)$$

where Φ_ρ is a bivariate Gaussian cdf with correlation coefficient ρ , and Φ is the standard univariate Gaussian cdf. Li [2] popularised the bivariate Gaussian copula, by showing how it could be used to study financial risk and default correlation, using credit derivatives as an example.

By substituting $F(x)$ for u and $G(y)$ for v in equation (3), we have a bivariate distribution $H(x, y)$, with a Gaussian dependency structure, and marginals F and G . Regardless of F and G , the resulting $H(x, y)$ can still be uniquely expressed as a Gaussian copula, so long as F and G are continuous. It is then a copula itself that captures the underlying dependencies between random variables, regardless of their marginal distributions. For this reason, copulas have been called *dependence functions* [13, 14]. Nelsen [3] contains an extensive discussion of copulas.

2 Copula Processes

Imagine choosing a covariance function, and then drawing a sample function at some finite number of points from a Gaussian process. The result is a sample from a collection of Gaussian random variables, with a dependency structure encoded by the specified covariance function. Now, suppose we transform each of these values through a univariate Gaussian cdf, such that we have a sample from a collection of uniform random variables. These uniform random variables also have this underlying Gaussian process dependency structure. One might call the resulting values a draw from a *Gaussian-Uniform Process*. We could subsequently put these values through an inverse beta cdf, to obtain a draw from what could be called a *Gaussian-Beta Process*: the values would be a sample from beta random variables, again with an underlying Gaussian process dependency structure. We could also transform the uniform values with different inverse cdfs, which would give a sample from different random variables, with dependencies encoded by the Gaussian process.

The above procedure is a means to generate samples from arbitrarily many random variables, with arbitrary marginal distributions, and desired dependencies. It is an example of how to use what we call a *copula process* – in this case, a *Gaussian copula process*, since a Gaussian copula describes the dependency structure of a finite number of samples. We now formally define a copula process.

Definition 2.1. *Copula Process*

Let $\{W_t\}$ be a collection of random variables indexed by $t \in \mathcal{T}$, with marginal distribution functions F_t , and let $Q_t = F_t(W_t)$. Further, let μ be a stochastic process measure with marginal distribution functions G_t , and joint distribution function H . Then W_t is copula process distributed with base measure μ , or $W_t \sim \mathcal{CP}(\mu)$, if and only if for all $n \in \mathbb{N}$, $a_i \in \mathbb{R}$,

$$P\left(\bigcap_{i=1}^n \{G_{t_i}^{(-1)}(Q_{t_i}) \leq a_i\}\right) = H_{t_1, t_2, \dots, t_n}(a_1, a_2, \dots, a_n). \quad (4)$$

Each $Q_{t_i} \sim \text{Uniform}(0, 1)$, and $G_{t_i}^{(-1)}$ is the quasi-inverse of G_{t_i} , as previously defined.

Definition 2.2. *Gaussian Copula Process*

W_t is Gaussian copula process distributed if it is copula process distributed and the base measure μ is a Gaussian process. If there is a mapping Ψ such that $\Psi(W_t) \sim \mathcal{GP}(m(t), k(t, t'))$, then we write $W_t \sim \mathcal{GCP}(\Psi, m(t), k(t, t'))$.

For example, if we have $W_t \sim \mathcal{GCP}$ with $m(t) = 0$ and $k(t, t) = 1$, then in the definition of a copula process, $G_t = \Phi$, the standard univariate Gaussian cdf, and H is the usual GP joint distribution function. Supposing this GCP is a *Gaussian-Beta process*, then $\Psi = \Phi^{-1} \circ F_B$, where F_B is a univariate Beta cdf. One could similarly define other copula processes.

We described generally how a copula process can be used to generate samples of arbitrarily many random variables with desired marginals and dependencies. We now develop a specific and practical application of this framework. We introduce a stochastic volatility model, *Gaussian Copula Process Volatility* (GCPV), as an example of how to learn the joint distribution of arbitrarily many random variables, the marginals of these random variables, and to make predictions. To do this, we fit a Gaussian copula process by using a type of Warped Gaussian Process [15]. However, our methodology varies substantially from Snelson et al. [15], since we are doing inference on latent variables as opposed to observations, which is a much greater undertaking that involves approximations, and we are doing so in a different context.

3 Gaussian Copula Process Volatility

Assume we have a sequence of observations $\mathbf{y} = (y_1, \dots, y_n)^\top$ at times $\mathbf{t} = (t_1, \dots, t_n)^\top$. The observations are random variables with different latent standard deviations. We therefore have n unobserved standard deviations, $\sigma_1, \dots, \sigma_n$, and want to learn the correlation structure between these standard deviations, and also to predict the distribution of σ_* at some unrealised time t_* .

We model the standard deviation function as a Gaussian copula process:

$$\sigma_t \sim \mathcal{GCP}(g^{-1}, 0, k(t, t')). \quad (5)$$

Specifically,

$$f(t) \sim \mathcal{GP}(m(t) = 0, k(t, t')) \quad (6)$$

$$\sigma(t) = g(f(t), \boldsymbol{\omega}) \quad (7)$$

$$y(t) \sim \mathcal{N}(0, \sigma^2(t)), \quad (8)$$

where g is a monotonic warping function, parametrized by $\boldsymbol{\omega}$. For each of the observations $\mathbf{y} = (y_1, \dots, y_n)^\top$ we have corresponding GP latent function values $\mathbf{f} = (f_1, \dots, f_n)^\top$, where $\sigma(t_i) = g(f_i, \boldsymbol{\omega})$, using the shorthand f_i to mean $f(t_i)$.

$\sigma_t \sim \mathcal{GCP}$, because any finite sequence $(\sigma_1, \dots, \sigma_p)$ is distributed as a Gaussian copula:

$$\begin{aligned} P(\sigma_1 \leq a_1, \dots, \sigma_p \leq a_p) &= P(g^{-1}(\sigma_1) \leq g^{-1}(a_1), \dots, g^{-1}(\sigma_p) \leq g^{-1}(a_p)) \quad (9) \\ &= \Phi_\Gamma(g^{-1}(a_1), \dots, g^{-1}(a_p)) = \Phi_\Gamma(\Phi^{-1}(F(a_1)), \dots, \Phi^{-1}(F(a_p))) \\ &= \Phi_\Gamma(\Phi^{-1}(u_1), \dots, \Phi^{-1}(u_p)) = C(u_1, \dots, u_p), \end{aligned}$$

where Φ is the standard univariate Gaussian cdf (supposing $k(t, t) = 1$), Φ_Γ is a multivariate Gaussian cdf with covariance matrix $\Gamma_{ij} = \text{cov}(g^{-1}(\sigma_i), g^{-1}(\sigma_j))$, and F is the marginal distribution of

each σ_i . In (5), we have $\Psi = g^{-1}$, because it is g^{-1} which maps σ_i to a GP. The specification in (5) is equivalently expressed by (6) and (7). With GCPV, the form of g is learned so that $g^{-1}(\sigma_i)$ is best modelled by a GP. By learning g , we learn the marginal of each σ : $F(a) = \Phi(g^{-1}(a))$ for $a \in \mathbb{R}$. Recently, a different sort of ‘kernel copula process’ has been used, where the marginals of the variables being modelled are not learned [16].¹ Further, we also consider a more subtle and flexible form of our model, where the function g itself is indexed by time: $g = g_t(f(t), \omega)$. We only assume that the marginal distributions of σ_t are stationary over ‘small’ time periods, and for each of these time periods (5)-(7) hold true. We return to this in the final discussion section.

Here we have assumed that each observation, conditioned on knowing its variance, is normally distributed with zero mean. This is a common assumption in heteroscedastic models. The zero mean and normality assumptions can be relaxed and are not central to this paper.

4 Predictions with GCPV

Ultimately, we wish to infer $p(\sigma(t_*)|\mathbf{y}, \mathbf{z})$, where $\mathbf{z} = \{\boldsymbol{\theta}, \boldsymbol{\omega}\}$, and $\boldsymbol{\theta}$ are the hyperparameters of the GP covariance function. To do this, we sample from

$$p(f_*|\mathbf{y}, \mathbf{z}) = \int p(f_*|\mathbf{f}, \boldsymbol{\theta})p(\mathbf{f}|\mathbf{y}, \mathbf{z})d\mathbf{f} \quad (10)$$

and then transform these samples by g . Letting $(C_f)_{ij} = \delta_{ij}g(f_i, \boldsymbol{\omega})^2$, where δ_{ij} is the Kronecker delta, $K_{ij} = k(t_i, t_j)$, $(\mathbf{k}_*)_i = k(t_*, t_i)$, we have

$$p(\mathbf{f}|\mathbf{y}, \mathbf{z}) = \mathcal{N}(\mathbf{f}; 0, K)\mathcal{N}(\mathbf{y}; 0, C_f)/p(\mathbf{y}|\mathbf{z}), \quad (11)$$

$$p(f_*|\mathbf{f}, \boldsymbol{\theta}) = \mathcal{N}(\mathbf{k}_*^\top K^{-1} \mathbf{f}, k(t_*, t_*) - \mathbf{k}_*^\top K^{-1} \mathbf{k}_*). \quad (12)$$

We also wish to learn \mathbf{z} , which we can do by finding the $\hat{\mathbf{z}}$ that maximizes the marginal likelihood,

$$p(\mathbf{y}|\mathbf{z}) = \int p(\mathbf{y}|\mathbf{f}, \boldsymbol{\omega})p(\mathbf{f}|\boldsymbol{\theta})d\mathbf{f}. \quad (13)$$

Unfortunately, for many functions g , (10) and (13) are intractable. Our methods of dealing with this can be used in very general circumstances, where one has a Gaussian process prior, but an (optionally parametrized) non-Gaussian likelihood. We use the Laplace approximation to estimate $p(\mathbf{f}|\mathbf{y}, \mathbf{z})$ as a Gaussian. Then we can integrate (10) for a Gaussian approximation to $p(f_*|\mathbf{y}, \mathbf{z})$, which we sample from to make predictions of σ_* . Using Laplace, we can also find an expression for an approximate marginal likelihood, which we maximize to determine \mathbf{z} . Once we have found \mathbf{z} with Laplace, we use Markov chain Monte Carlo to sample from $p(f_*|\mathbf{y}, \mathbf{z})$, and compare that to using Laplace to sample from $p(f_*|\mathbf{y}, \mathbf{z})$. In the supplement we relate this discussion to (9).

4.1 Laplace Approximation

The goal is to approximate (11) with a Gaussian, so that we can evaluate (10) and (13) and make predictions. In doing so, we follow Rasmussen and Williams [11] in their treatment of Gaussian process classification, except we use a parametrized likelihood, and modify Newton’s method.

First, consider as an objective function the logarithm of an unnormalized (11):

$$s(\mathbf{f}|\mathbf{y}, \mathbf{z}) = \log p(\mathbf{y}|\mathbf{f}, \boldsymbol{\omega}) + \log p(\mathbf{f}|\boldsymbol{\theta}). \quad (14)$$

The Laplace approximation uses a second order Taylor expansion about the $\hat{\mathbf{f}}$ which maximizes (14), to find an approximate objective $\tilde{s}(\mathbf{f}|\mathbf{y}, \mathbf{z})$. So the first step is to find $\hat{\mathbf{f}}$, for which we use Newton’s method. The Newton update is $\mathbf{f}^{\text{new}} = \mathbf{f} - (\nabla \nabla s(\mathbf{f}))^{-1} \nabla s(\mathbf{f})$. Differentiating (14),

$$\nabla s(\mathbf{f}|\mathbf{y}, \mathbf{z}) = \nabla \log p(\mathbf{y}|\mathbf{f}, \boldsymbol{\omega}) - K^{-1} \mathbf{f} \quad (15)$$

$$\nabla \nabla s(\mathbf{f}|\mathbf{y}, \mathbf{z}) = \nabla \nabla \log p(\mathbf{y}|\mathbf{f}, \boldsymbol{\omega}) - K^{-1} = -W - K^{-1}, \quad (16)$$

where W is the diagonal matrix $-\nabla \nabla \log p(\mathbf{y}|\mathbf{f}, \boldsymbol{\omega})$.

¹Note added in proof: Also, for a very recent related model, see Rodríguez et al. [17].

If the likelihood function $p(\mathbf{y}|\mathbf{f}, \boldsymbol{\omega})$ is not log concave, then W may have negative entries. Vanhatalo et al. [18] found this to be problematic when doing Gaussian process regression with a Student-t likelihood. They instead use an expectation-maximization (EM) algorithm for finding $\hat{\mathbf{f}}$, and iterate ordered rank one Cholesky updates to evaluate the Laplace approximate marginal likelihood. But EM can converge slowly, especially near a local optimum, and each of the rank one updates is vulnerable to numerical instability. With a small modification of Newton’s method, we often get close to quadratic convergence for finding $\hat{\mathbf{f}}$, and can evaluate the Laplace approximate marginal likelihood in a numerically stable fashion, with no approximate Cholesky factors, and optimal computational requirements. Some comments are in the supplementary material but, in short, we use an approximate negative Hessian, $-\nabla\nabla s \approx M + K^{-1}$, which is guaranteed to be positive definite, since M is formed on each iteration by zeroing the negative entries of W . For stability, we reformulate our optimization in terms of $B = I + M^{\frac{1}{2}}KM^{\frac{1}{2}}$, and let $Q = M^{\frac{1}{2}}B^{-1}M^{\frac{1}{2}}$, $\mathbf{b} = M\mathbf{f} + \nabla \log p(\mathbf{y}|\mathbf{f})$, $\mathbf{a} = \mathbf{b} - QK\mathbf{b}$. Since $(K^{-1} + M)^{-1} = K - KQK$, the Newton update becomes $\mathbf{f}^{new} = K\mathbf{a}$. With these updates we find $\hat{\mathbf{f}}$ and get an expression for \tilde{s} which we use to approximate (13) and (11). The approximate marginal likelihood $q(\mathbf{y}|\mathbf{z})$ is given by $\int \exp(\tilde{s})d\mathbf{f}$. Taking its logarithm,

$$\log q(\mathbf{y}|\mathbf{z}) = -\frac{1}{2}\hat{\mathbf{f}}^\top \mathbf{a}_{\hat{\mathbf{f}}} + \log p(\mathbf{y}|\hat{\mathbf{f}}) - \frac{1}{2} \log |B_{\hat{\mathbf{f}}}|, \quad (17)$$

where $B_{\hat{\mathbf{f}}}$ is B evaluated at $\hat{\mathbf{f}}$, and $\mathbf{a}_{\hat{\mathbf{f}}}$ is a numerically stable evaluation of $K^{-1}\hat{\mathbf{f}}$.

To learn the parameters \mathbf{z} , we use conjugate gradient descent to maximize (17) with respect to \mathbf{z} . Since $\hat{\mathbf{f}}$ is a function of \mathbf{z} , we initialize \mathbf{z} , and update $\hat{\mathbf{f}}$ every time we vary \mathbf{z} . Once we have found an optimum $\hat{\mathbf{z}}$, we can make predictions. By exponentiating \tilde{s} , we find a Gaussian approximation to the posterior (11), $q(\mathbf{f}|\mathbf{y}, \mathbf{z}) = \mathcal{N}(\hat{\mathbf{f}}, K - KQK)$. The product of this approximate posterior with $p(f_*|\mathbf{f})$ is Gaussian. Integrating this product, we approximate $p(f_*|\mathbf{y}, \mathbf{z})$ as

$$q(f_*|\mathbf{y}, \mathbf{z}) = \mathcal{N}(\mathbf{k}_*^\top \nabla \log p(\mathbf{y}|\hat{\mathbf{f}}), k(t_*, t_*) - \mathbf{k}_*^\top Q\mathbf{k}_*). \quad (18)$$

Given n training observations, the cost of each Newton iteration is dominated by computing the cholesky decomposition of B , which takes $\mathcal{O}(n^3)$ operations. The objective function typically changes by less than 10^{-6} after 3 iterations. Once Newton’s method has converged, it takes only $\mathcal{O}(1)$ operations to draw from $q(f_*|\mathbf{y}, \mathbf{z})$ and make predictions.

4.2 Markov chain Monte Carlo

We use Markov chain Monte Carlo (MCMC) to sample from (11), so that we can later sample from $p(\sigma_*|\mathbf{y}, \mathbf{z})$ to make predictions. Sampling from (11) is difficult, because the variables \mathbf{f} are strongly coupled by a Gaussian process prior. We use a new technique, Elliptical Slice Sampling [19], and find it extremely effective for this purpose. It was specifically designed to sample from posteriors with correlated Gaussian priors. It has no free parameters, and jointly updates every element of \mathbf{f} . For our setting, it is over 100 times as fast as axis aligned slice sampling with univariate updates.

To make predictions, we take J samples of $p(\mathbf{f}|\mathbf{y}, \mathbf{z})$, $\{\mathbf{f}^1, \dots, \mathbf{f}^J\}$, and then approximate (10) as a mixture of J Gaussians:

$$p(f_*|\mathbf{y}, \mathbf{z}) \approx \frac{1}{J} \sum_{i=1}^J p(f_*|\mathbf{f}^i, \boldsymbol{\theta}). \quad (19)$$

Each of the Gaussians in this mixture have equal weight. So for each sample of $f_*|\mathbf{y}$, we uniformly choose a random $p(f_*|\mathbf{f}^i, \boldsymbol{\theta})$ and draw a sample. In the limit $J \rightarrow \infty$, we are sampling from the exact $p(f_*|\mathbf{y}, \mathbf{z})$. Mapping these samples through g gives samples from $p(\sigma_*|\mathbf{y}, \mathbf{z})$. After one $\mathcal{O}(n^3)$ and one $\mathcal{O}(J)$ operation, a draw from (19) takes $\mathcal{O}(1)$ operations.

4.3 Warping Function

The warping function, g , maps f_i , a GP function value, to σ_i , a standard deviation. Since f_i can take any value in \mathbb{R} , and σ_i can take any non-negative real value, $g: \mathbb{R} \rightarrow \mathbb{R}^+$. For each f_i to correspond to a unique deviation, g must also be one-to-one. We use

$$g(x, \boldsymbol{\omega}) = \sum_{j=1}^K a_j \log[\exp[b_j(x + c_j)] + 1], \quad a_j, b_j > 0. \quad (20)$$

This is monotonic, positive, infinitely differentiable, asymptotic towards zero as $x \rightarrow -\infty$, and tends to $(\sum_{j=1}^K a_j b_j)x$ as $x \rightarrow \infty$. In practice, it is useful to add a small constant to (20), to avoid rare situations where the parameters ω are trained to make g extremely small for certain inputs, at the expense of a good overall fit; this can happen when the parameters ω are learned by optimizing a likelihood. A suitable constant could be one tenth the absolute value of the smallest nonzero observation.

By inferring the parameters of the warping function, or distributions of these parameters, we are learning a transformation which will best model σ_t with a Gaussian process. The more flexible the warping function, the more potential there is to improve the GCPV fit – in other words, the better we can estimate the ‘perfect’ transformation. To test the importance of this flexibility, we also try a simple unparametrized warping function, $g(x) = e^x$. In related work, Goldberg et al. [20] place a GP prior on the log noise level in a standard GP regression model on observations, except for inference they use Gibbs sampling, and a high level of ‘jitter’ for conditioning.

Once g is trained, we can infer the marginal distribution of each σ : $F(a) = \Phi(g^{-1}(a))$, for $a \in \mathbb{R}$. This suggests an alternate way to initialize g : we can initialize F as a mixture of Gaussians, and then map through Φ^{-1} to find g^{-1} . Since mixtures of Gaussians are dense in the set of probability distributions, we could in principle find the ‘perfect’ g using an infinite mixture of Gaussians [21].

5 Experiments

In our experiments, we predict the latent standard deviations σ of observations \mathbf{y} at times t , and also σ_* at unobserved times t_* . To do this, we use two versions of GCPV. The first variant, which we simply refer to as GCPV, uses the warping function (20) with $K = 1$, and squared exponential covariance function, $k(t, t') = A \exp(-(t-t')^2/l^2)$, with $A = 1$. The second variant, which we call GP-EXP, uses the unparametrized warping function e^x , and the same covariance function, except the amplitude A is a trained hyperparameter. The other hyperparameter l is called the *lengthscale* of the covariance function. The greater l , the greater the covariance between σ_t and σ_{t+a} for $a \in \mathbb{R}$. We train hyperparameters by maximizing the Laplace approximate log marginal likelihood (17).

We then sample from $p(f_*|\mathbf{y})$ using the Laplace approximation (18). We also do this using MCMC (19) with $J = 10000$, after discarding a previous 10000 samples of $p(\mathbf{f}|\mathbf{y})$ as burn-in. We pass these samples of $f_*|\mathbf{y}$ through g and g^2 to draw from $p(\sigma_*|\mathbf{y})$ and $p(\sigma_*^2|\mathbf{y})$, and compute the sample mean and variance of $\sigma_*|\mathbf{y}$. We use the sample mean as a point predictor, and the sample variance for error bounds on these predictions, and we use 10000 samples to compute these quantities. For GCPV we use Laplace and MCMC for inference, but for GP-EXP we only use Laplace. We compare predictions to GARCH(1,1), which has been shown in extensive and recent reviews to be competitive with other GARCH variants, and more sophisticated models [5, 6, 7]. GARCH(p, q) specifies $y(t) \sim \mathcal{N}(0, \sigma^2(t))$, and lets the variance be a deterministic function of the past: $\sigma_t^2 = a_0 + \sum_{i=1}^p a_i y_{t-i}^2 + \sum_{j=1}^q b_j \sigma_{t-j}^2$. We use the Matlab Econometrics Toolbox implementation of GARCH, where the parameters a_0, a_i and b_j are estimated using a constrained maximum likelihood.

We make forecasts of volatility, and we predict historical volatility. By ‘historical volatility’ we mean the volatility at observed time points, or between these points. Uncovering historical volatility is important. It could, for instance, be used to study what causes fluctuations in the stock market, or to understand physical systems.

To evaluate our model, we use the Mean Squared Error (MSE) between the true variance, or proxy for the truth, and the predicted variance. Although likelihood has advantages, we are limited in space, and we wish to harmonize with the econometrics literature, and other assessments of volatility models, where MSE is the standard. In a similar assessment of volatility models, Brownlees et al. [7] found that MSE and quasi-likelihood rankings were comparable.

When the true variance is unknown we follow Brownlees et al. [7] and use squared observations as a proxy for the truth, to compare our model to GARCH.² The more observations, the more reliable these performance estimates will be. However, not many observations (e.g. 100) are needed for a stable ranking of competing models; in Brownlees et al. [7], the rankings derived from high frequency squared observations are similar to those derived using daily squared observations.

²Since each observation y is assumed to have zero mean and variance σ^2 , $\mathbb{E}[y^2] = \sigma^2$.

5.1 Simulations

We simulate observations from $\mathcal{N}(0, \sigma^2(t))$, using $\sigma(t) = \sin(t) \cos(t^2) + 1$, at $\mathbf{t} = (0, 0.02, 0.04, \dots, 4)^\top$. We call this data set TRIG. We also simulate using a standard deviation that jumps from 0.1 to 7 and back, at times $\mathbf{t} = (0, 0.1, 0.2, \dots, 6)^\top$. We call this data set JUMP. To forecast, we use all observations up until the current time point, and make 1, 7, and 30 step ahead predictions. So, for example, in TRIG we start by observing $t = 0$, and make forecasts at $t = 0.02, 0.14, 0.60$. Then we observe $t = 0, 0.02$ and make forecasts at $t = 0.04, 0.16, 0.62$, and so on, until all data points have been observed. For historical volatility, we predict the latent σ_t at the observation times, which is safe since we are comparing to the true volatility, which is not used in training; the results are similar if we interpolate. Figure 1 panels a) and b) show the true volatility for TRIG and JUMP respectively, alongside GCPV Laplace, GCPV MCMC, GP-EXP Laplace, and GARCH(1,1) predictions of historical volatility. Table 1 shows the results for forecasting and historical volatility.

In panel a) we see that GCPV more accurately captures the dependencies between σ at different times points than GARCH: if we manually decrease the lengthscale in the GCPV covariance function, we can replicate the erratic GARCH behaviour, which inaccurately suggests that the covariance between σ_t and σ_{t+a} decreases quickly with increases in a . We also see that GCPV with an unparametrized exponential warping function tends to overestimate peaks and underestimate troughs. In panel b), the volatility is extremely difficult to reconstruct or forecast – with no warning it will immediately and dramatically increase or decrease. This behaviour is not suited to a smooth squared exponential covariance function. Nevertheless, GCPV outperforms GARCH, especially in regions of low volatility. We also see this in panel a) for $t \in (1.5, 2)$. GARCH is known to respond slowly to large returns, and to overpredict volatility [22]. In JUMP, the greater the peaks, and the smaller the troughs, the more GARCH suffers, while GCPV is mostly robust to these changes.

5.2 Financial Data

The returns on the daily exchange rate between the Deutschmark (DM) and the Great Britain Pound (GBP) from 1984 to 1992 have become a benchmark for assessing the performance of GARCH models [8, 9, 10]. This exchange data, which we refer to as DMGBP, can be obtained from www.datastream.com, and the returns are calculated as $r_t = \log(P_{t+1}/P_t)$, where P_t is the number of DM to GBP on day t . The returns are assumed to have a zero mean function.

We use a rolling window of the previous 120 days of returns to make 1, 7, and 30 day ahead volatility forecasts, starting at the beginning of January 1988, and ending at the beginning of January 1992 (659 trading days). Every 7 days, we retrain the parameters of GCPV and GARCH. Every time we retrain parameters, we predict historical volatility over the past 120 days. The average MSE for these historical predictions is given in Table 1, although they should be observed with caution; unlike with the simulations, the DMGBP historical predictions are trained using the same data they are assessed on. In Figure 1c), we see that the GARCH one day ahead forecasts are lifted above the GCPV forecasts, but unlike in the simulations, they are now operating on a similar lengthscale. This suggests that GARCH could still be overpredicting volatility, but that GCPV has adapted its estimation of how σ_t and σ_{t+a} correlate with one another. Since GARCH is suited to this financial data set, it is reassuring that GCPV predictions have a similar time varying structure. Overall, GCPV and GARCH are competitive with one another for forecasting currency exchange returns, as seen in Table 1. Moreover, a learned warping function g outperforms an unparametrized one, and a full Laplace solution is comparable to using MCMC for inference, in accuracy and speed. This is also true for the simulations. Therefore we recommend whichever is more convenient to implement.

6 Discussion

We defined a copula process, and as an example, developed a stochastic volatility model, GCPV, which can outperform GARCH. With GCPV, the volatility σ_t is distributed as a Gaussian Copula Process, which separates the modelling of the dependencies between volatilities at different times from their marginal distributions – arguably the most useful property of a copula. Further, GCPV fits the marginals in the Gaussian copula process by learning a warping function. If we had simply chosen an unparametrized exponential warping function, we would incorrectly be assuming that the log

Table 1: MSE for predicting volatility.

Data set	Model	Historical	1 step	7 step	30 step	
TRIG	GCPV (LA)	0.0953	0.588	0.951	1.71	
	GCPV (MCMC)	0.0760	0.622	0.979	1.76	
	GP-EXP	0.193	0.646	1.36	1.15	
	GARCH	0.938	1.04	1.79	5.12	
$\times 10^3$	JUMP	GCPV (LA)	0.588	0.891	1.38	1.35
	GCPV (MCMC)	1.21	0.951	1.37	1.35	
	GP-EXP	1.43	1.76	6.95	14.7	
	GARCH	1.88	1.58	3.43	5.65	
$\times 10^{-9}$	DMGBP	GCPV (LA)	2.43	3.00	3.08	3.17
	GCPV (MCMC)	2.39	3.00	3.08	3.17	
	GP-EXP	2.52	3.20	3.46	5.14	
	GARCH	2.83	3.03	3.12	3.32	




Figure 1: Predicting volatility and learning its marginal pdf. For a) and b), the true volatility, and GCPV (MCMC), GCPV (LA), GP-EXP, and GARCH predictions, are shown respectively by a thick green line, a dashed thick blue line, a dashed black line, a cyan line, and a red line. a) shows predictions of historical volatility for TRIG, where the shade is a 95% confidence interval about GCPV (MCMC) predictions. b) shows predictions of historical volatility for JUMP. In c), a black line and a dashed red line respectively show GCPV (LA) and GARCH one day ahead volatility forecasts for DMGBP. In d), a black line and a dashed blue line respectively show the GCPV learned marginal pdf of σ_t in DMGBP and a Gamma(4.15, 0.00045) pdf.

volatilities are marginally Gaussian distributed. Indeed, for the DMGBP data, we trained the warping function g over a 120 day period, and mapped its inverse through the univariate standard Gaussian cdf Φ , and differenced, to estimate the marginal probability density function (pdf) of σ_t over this period. The learned marginal pdf, shown in Figure 1d), is similar to a Gamma(4.15, 0.00045) distribution. However, in using a rolling window to retrain the parameters of g , we do not assume that the marginals of σ_t are stationary; we have a time changing warping function.

While GARCH is successful, and its simplicity is attractive, our model is also simple and has a number of advantages. We can effortlessly handle missing data, we can easily incorporate covariates other than time (like interest rates) in our covariance function, and we can choose from a rich class of covariance functions – squared exponential, Brownian motion, Matérn, periodic, etc. In fact, the volatility of high frequency intradaily returns on equity indices and currency exchanges is cyclical [23], and GCPV with a periodic covariance function is uniquely well suited to this data. And the parameters of GCPV, like the covariance function lengthscale, or the learned warping function, provide insight into the underlying source of volatility, unlike the parameters of GARCH.

Finally, copulas are rapidly becoming popular in applications, but often only bivariate copulas are being used. With our copula process one can learn the dependencies between arbitrarily many random variables independently of their marginal distributions. We hope the Gaussian Copula Process Volatility model will encourage other applications of copula processes. More generally, we hope our work will help bring together the machine learning and econometrics communities.

Acknowledgments: Thanks to Carl Edward Rasmussen and Ferenc Huszár for helpful conversations. AGW is supported by an NSERC grant.

References

- [1] Paul Embrechts, Alexander McNeil, and Daniel Straumann. Correlation and dependence in risk management: Properties and pitfalls. In *Risk Management: Value at risk and beyond*, pages 176–223. Cambridge University Press, 1999.
- [2] David X. Li. On default correlation: A copula function approach. *Journal of Fixed Income*, 9(4):43–54, 2000.
- [3] Roger B. Nelsen. *An Introduction to Copulas*. Springer Series in Statistics, second edition, 2006.
- [4] Tim Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327, 1986.
- [5] Ser-Huang Poon and Clive W.J. Granger. Practical issues in forecasting volatility. *Financial Analysts Journal*, 61(1):45–56, 2005.
- [6] Peter Reinhard Hansen and Asger Lunde. A forecast comparison of volatility models: Does anything beat a GARCH(1,1). *Journal of Applied Econometrics*, 20(7):873–889, 2005.
- [7] Christian T. Brownlees, Robert F. Engle, and Bryan T. Kelly. A practical guide to volatility forecasting through calm and storm, 2009. Available at SSRN: <http://ssrn.com/abstract=1502915>.
- [8] T. Bollerslev and E. Ghysels. Periodic autoregressive conditional heteroscedasticity. *Journal of Business and Economic Statistics*, 14:139–151, 1996.
- [9] B.D. McCullough and C.G. Renfro. Benchmarks and software standards: A case study of GARCH procedures. *Journal of Economic and Social Measurement*, 25:59–71, 1998.
- [10] C. Brooks, S.P. Burke, and G. Persaud. Benchmarks and the accuracy of GARCH model estimation. *International Journal of Forecasting*, 17:45–56, 2001.
- [11] Carl Edward Rasmussen and Christopher K.I. Williams. *Gaussian processes for Machine Learning*. The MIT Press, 2006.
- [12] Abe Sklar. Fonctions de répartition à n dimensions et leurs marges. *Publ. Inst. Statist. Univ. Paris*, 8: 229–231, 1959.
- [13] P Deheuvels. Caractérisation complète des lois extrêmes multivariés et de la convergence des types extrêmes. *Publications de l’Institut de Statistique de l’Université de Paris*, 23:1–36, 1978.
- [14] G Kimeldorf and A Sampson. Uniform representations of bivariate distributions. *Communications in Statistics*, 4:617–627, 1982.
- [15] Edward Snelson, Carl Edward Rasmussen, and Zoubin Ghahramani. Warped Gaussian Processes. In *NIPS*, 2003.
- [16] Sebastian Jaimungal and Eddie K.H. Ng. Kernel-based Copula processes. In *ECML PKDD*, 2009.
- [17] A. Rodríguez, D.B. Dunson, and A.E. Gelfand. Latent stick-breaking processes. *Journal of the American Statistical Association*, 105(490):647–659, 2010.
- [18] Jarno Vanhatalo, Pasi Jylanki, and Aki Vehtari. Gaussian process regression with Student-t likelihood. In *NIPS*, 2009.
- [19] Iain Murray, Ryan Prescott Adams, and David J.C. MacKay. Elliptical Slice Sampling. In *AISTATS*, 2010.
- [20] Paul W. Goldberg, Christopher K.I. Williams, and Christopher M. Bishop. Regression with input-dependent noise: A Gaussian process treatment. In *NIPS*, 1998.
- [21] Carl Edward Rasmussen. The Infinite Gaussian Mixture Model. In *NIPS*, 2000.
- [22] Ruey S. Tsay. *Analysis of Financial Time Series*. John Wiley & Sons, 2002.
- [23] Torben G. Andersen and Tim Bollerslev. Intraday periodicity and volatility persistence in financial markets. *Journal of Empirical Finance*, 4(2-3):115–158, 1997.