Gaussian Process Conditional Copulas with Applications to Financial Time Series

José Miguel Hernández–Lobato[1],

October 28, 2013

joint work with
James Robert Lloyd[1] and Daniel Hernández-Lobato[2]



[1]Cambridge University.
[2]Universidad Autónoma de Madrid.

Copulas and Multivariate Probabilistic Models

Copulas link univariate marginal distributions into joint multivariate probabilistic models.

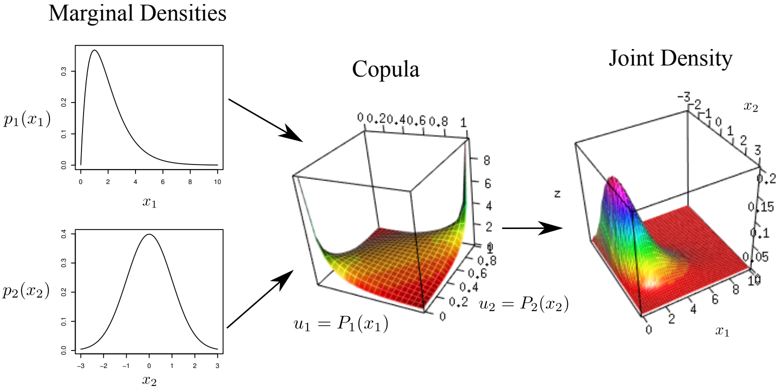

Copulas specify dependencies between the random variables.

Separating Marginals and Dependencies Using Copulas

Sklar's Theorem: any joint distribution can be written in terms of its copula and its univariate marginals [Sklar, 1959].

$$p(x_1, x_2) = \underbrace{c(u_1, u_2)}_{\text{copula}} \underbrace{p_1(x_1)\, p_2(x_2)}_{\text{marginals}},$$

where $u_i := P_i(x_i)$ and P_i denotes the marginal cdf for x_i.

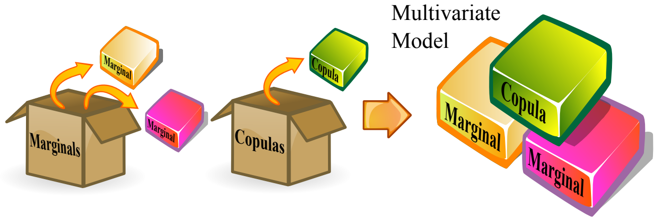

We can use different copulas and different univariate marginals as building blocks for new multivariate models.

Additional Advantages of Copulas

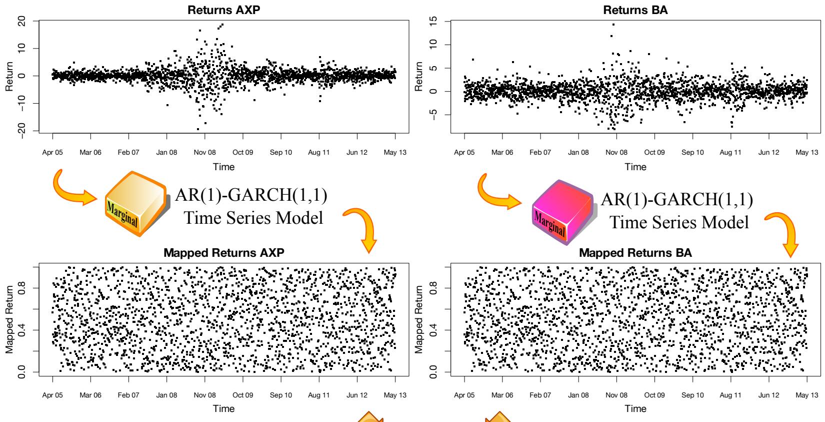

Copulas easily extend univariate models to the multivariate regime.

Copulas simplify the estimation of multivariate models [Joe, 2005].

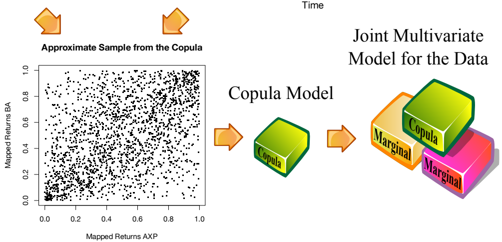

Parametric Copula Models

Any parametric density $p(x_1, x_2)$ has corresponding copula

$$c(u_1, u_2) = \frac{p(x_1, x_2)}{p_1(x_1)p_2(x_2)} \, ,$$

where $x_i = P_i^{-1}(u_i)$. Other parametric copulas are not directly derived from a parametric density, e.g., Archimedean copulas.

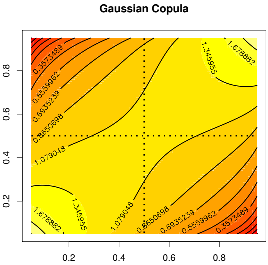

Parameters: ρ

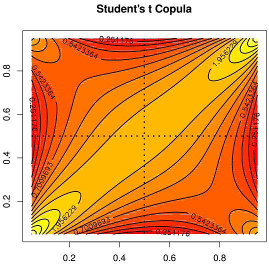

Parameters: ρ, ν

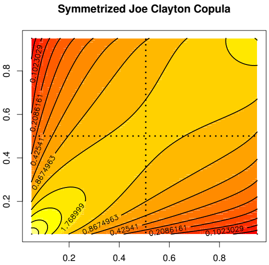

Parameters: τ_U, τ_L

Conditional Copulas

The distribution of x_1 and x_2 may depend on a covariate vector $\mathbf{z}$.

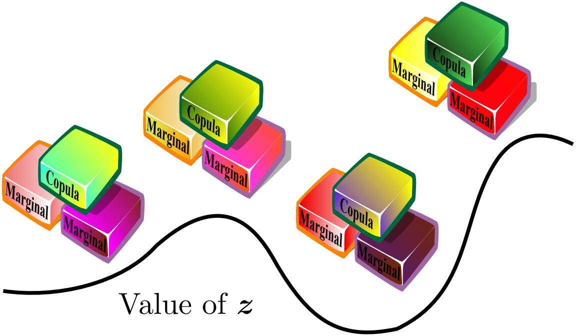

The copula framework can be extended to conditional distributions [Patton, 2006]. In this case,

$$p(x_1, x_2|\mathbf{z}) = \underbrace{c(u_1, u_2|\mathbf{z})}_{\substack{\text{conditional}\\ \text{copula}}} \; \underbrace{p_1(x_1|\mathbf{z})\, p_2(x_2|\mathbf{z})}_{\substack{\text{conditional}\\ \text{marginals}}},$$

where $u_i := P_i(x_i|\mathbf{z})$ and P_i denotes the marginal cdf for x_i given $\mathbf{z}$.

Semiparametric Conditional Copulas

Main Assumption: $c(u_1, u_2|\mathbf{z})$ is given by a parametric copula model $c_{\text{par}}[u_1, u_2|\theta_1(\mathbf{z}), \ldots, \theta_k(\mathbf{z})]$ specified by k scalar parameters $\theta_1, \ldots, \theta_k$ that are functions of $\mathbf{z}$. We select $\theta_i(\mathbf{z}) = \sigma_i[f_i(\mathbf{z})]$, where f_i is a real function and σ_i is a link function that maps $\mathbb{R}$ to a set Θ_i of valid configurations for θ_i.

Example: conditional Student's t copula. $c_{\text{student}}[u_1, u_2|\rho(\mathbf{z}), \nu(\mathbf{z})]$, $\rho(\mathbf{z}) = \sigma_\rho[f_\rho(\mathbf{z})]$, $\nu(\mathbf{z}) = \sigma_\nu[f_\nu(\mathbf{z})]$, $\sigma_\rho(x) = 2\Phi(x) - 1$, $\sigma_\nu(x) = \exp(x)$.

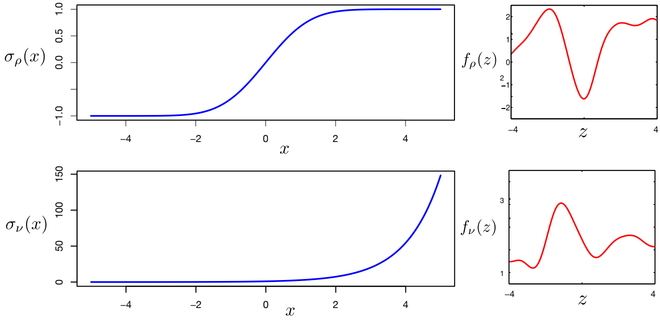

Our objective is to estimate $f_1, \ldots, f_k$ from data.

A Bayesian Approach for the Estimation of $f_1, \ldots, f_k$

Given a dataset formed by $\mathcal{D}_{\mathbf{z}} = \{\mathbf{z}_i\}_{i=1}^n$, $\mathcal{D}_{u_1,u_2} = \{u_{1i}, u_{2i}\}_{i=1}^n$, where $(u_{1i}, u_{2i}) \sim c_{\text{par}}[u_1, u_2 | \theta_1(\mathbf{z}_i), \ldots, \theta_k(\mathbf{z}_i)]$, we place prior distributions on $f_1, \ldots, f_k$ and compute a posterior distribution.

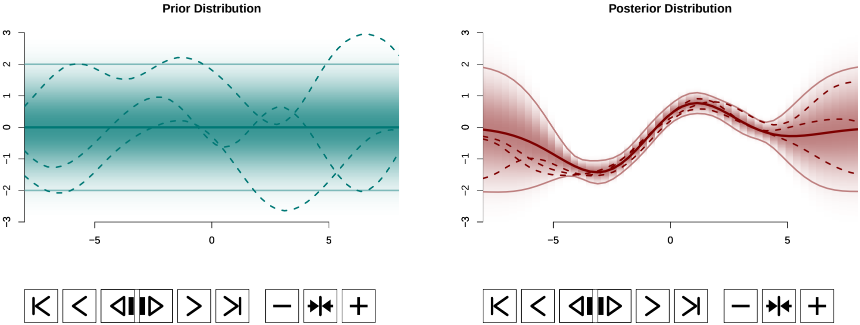

The distribution over functions is going to be a Gaussian Process.

Animations generated using the method desribed in [Henning, 2013]

From Multivariate Gaussians to Gaussian Processes

Let us consider a 20 dimensional vector $\mathbf{f} = (f_1, \ldots, f_{20})^{\mathrm{T}}$. We plot the value of each component f_i against its index i where $\mathbf{f}$ is random with distribution $\mathbf{f} \sim \mathcal{N}(\mathbf{m}, \mathbf{V})$.

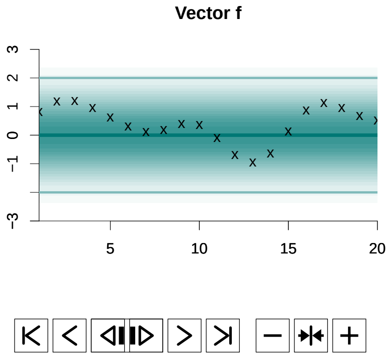

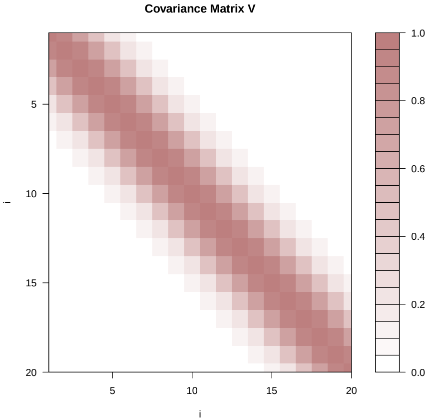

Intuitively, we can understand a GP as a Gaussian distribution on a vector $\mathbf{f}$ whose dimension is infinite and is indexed by numbers in $\mathbb{R}$.

Definition and Basics of Gaussian Processes

Definition:

A Gaussian process $p(f) = \mathcal{P}(f|m,k)$ is a distribution over functions $f : \mathbb{R}^d \to \mathbb{R}$ such that every finite dimensional vector of function values $\mathbf{f} = [f(\mathbf{z}_1), \ldots, f(\mathbf{z}_n)]^\mathrm{T}$ at locations $\mathbf{Z} = [\mathbf{z}_1, \ldots, \mathbf{z}_n]^\mathrm{T}$ follows a multivariate Gaussian distribution $p(\mathbf{f}) = \mathcal{N}(\mathbf{f}|\mathbf{m}_\mathbf{Z}, \mathbf{V}_{\mathbf{Z},\mathbf{Z}})$.

The mean vector and covariance matrix $\mathbf{m}_\mathbf{Z}$ and $\mathbf{V}_{\mathbf{Z},\mathbf{Z}}$ are the result of evaluating the mean function $m : \mathbb{R}^d \to \mathbb{R}$ and the covariance function $k : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$ at the locations $\mathbf{Z}$.

Choices for μ and k:

$m(\mathbf{z}) = \mu$,
$k(\mathbf{z},\mathbf{z}') = \beta \exp\{-(\mathbf{z}-\mathbf{z}')^\mathrm{T} \mathrm{diag}(\boldsymbol{\lambda})(\mathbf{z}-\mathbf{z}')\} + \gamma\delta(\mathbf{z}-\mathbf{z}')$.

Prediction:

Given $\mathbf{f}$, the conditional for the value of f at at new locations $\mathbf{Z}^\star = [\mathbf{z}^\star_{n+1}, \ldots, \mathbf{z}^\star_{n+h}]^\mathrm{T}$ is $p(\mathbf{f}^\star|\mathbf{f}) = \mathcal{N}(\mathbf{f}^\star|\mathbf{m}^\mathrm{pred}, \mathbf{V}^\mathrm{pred})$ with

$$\mathbf{m}^\mathrm{pred} = \mathbf{m}_{\mathbf{Z}^\star} + \mathbf{V}_{\mathbf{Z}^\star,\mathbf{Z}}\mathbf{V}^{-1}_{\mathbf{Z},\mathbf{Z}}(\mathbf{f}-\mathbf{m}_\mathbf{Z})\,,$$
$$\mathbf{V}^\mathrm{pred} = \mathbf{V}_{\mathbf{Z}^\star,\mathbf{Z}^\star} - \mathbf{V}_{\mathbf{Z}^\star,\mathbf{Z}}\mathbf{V}^{-1}_{\mathbf{Z},\mathbf{Z}}\mathbf{V}_{\mathbf{Z},\mathbf{Z}^\star}\,.$$

Sparse Gaussian Processes

Working with GPs has cost $\mathcal{O}(n^3)$.

A low rank covariance approximation reduces the cost to $\mathcal{O}(m^2 n)$. We use the FITC approximation [Snelson, 2006]. The covariance matrix $\mathbf{V}_{\mathbf{Z},\mathbf{Z}}$ is approximated using

$$\mathbf{V}_{\mathbf{Z},\mathbf{Z}} \approx \mathbf{V}_{\mathbf{Z},\bar{\mathbf{Z}}} \mathbf{V}_{\bar{\mathbf{Z}},\bar{\mathbf{Z}}}^{-1} \mathbf{V}_{\bar{\mathbf{Z}},\mathbf{Z}} + \boldsymbol{\Lambda}\,,$$

where $\boldsymbol{\Lambda}$ is diagonal, $\boldsymbol{\Lambda} = \operatorname{diag}(\mathbf{V}_{\mathbf{Z},\mathbf{Z}}) - \operatorname{diag}(\mathbf{V}_{\mathbf{Z},\bar{\mathbf{Z}}} \mathbf{V}_{\bar{\mathbf{Z}},\bar{\mathbf{Z}}}^{-1} \mathbf{V}_{\bar{\mathbf{Z}},\mathbf{Z}})$ and $\bar{\mathbf{Z}}$ is a set of m additional locations or pseudo-inputs.

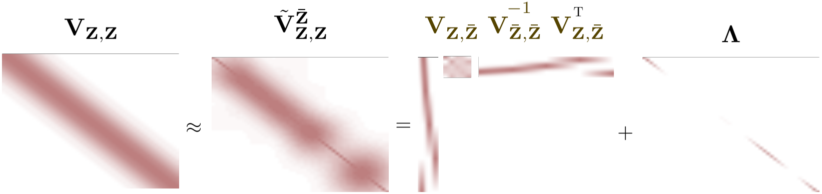

Posterior Distribution and Predictive Distribution

We are given $\mathcal{D}_{\mathbf{z}} = \{\mathbf{z}_i\}_{i=1}^n$ and $\mathcal{D}_{u_1,u_2} = \{u_{1i}, u_{2i}\}_{i=1}^n$. We denote by $\mathbf{f}_j = [f_j(\mathbf{z}_1), \ldots, f_j(\mathbf{z}_n)]^{\mathrm{T}}$ the vector with the evaluation of f_j at $\mathcal{D}_{\mathbf{z}}$. Then,

$$p(\mathbf{f}_1, \ldots, \mathbf{f}_k | \mathcal{D}_{u_1,u_2}, \mathcal{D}_{\mathbf{z}}) = \left[\prod_{i=1}^n c_{\mathrm{par}}\left[u_{1i}, u_{2i} | \sigma_1[f_1(\mathbf{z}_i)], \ldots, \sigma_k[f_k(\mathbf{z}_i)]\right]\right] \left[\prod_{j=1}^k p(\mathbf{f}_j)\right] \left[p(\mathcal{D}_{u_1,u_2}, \mathcal{D}_{\mathbf{z}})\right]^{-1} .$$

We make predictions at $\mathbf{z}^\star$ using $\mathbf{f}^\star = [f_1(\mathbf{z}^\star), \ldots, f_k(\mathbf{z}^\star)]$ and

$$p(u_1^\star, u_2^\star | \mathbf{z}^\star, \mathcal{D}_{\mathbf{z}}, \mathcal{D}_{u_1,u_2}) = \int c_{\mathrm{par}}\left[u_1^\star, u_2^\star | \sigma_1[f_1(\mathbf{z}^\star)], \ldots, \sigma_k[f_k(\mathbf{z}^\star)]\right] p(\mathbf{f}^\star | \mathbf{f}_1, \ldots, \mathbf{f}_k, \mathbf{z}^\star, \mathcal{D}_{\mathbf{z}}) p(\mathbf{f}_1, \ldots, \mathbf{f}_k | \mathcal{D}_{u_1,u_2}, \mathcal{D}_{\mathbf{z}})\, d\mathbf{f}_1 \cdots d\mathbf{f}_k\, d\mathbf{f}^\star .$$

We approximate these distributions using Expectation Propagation [Minka, 2001].

Expectation Propagation

EP approximates the unnormalized joint $p(\mathbf{f}_1, \ldots, \mathbf{f}_k, \mathcal{D}_{u_1,u_2} | \mathcal{D}_{\mathbf{z}})$ with an unnormalized Gaussian $\mathcal{Q}(\mathbf{f}_1, \ldots, \mathbf{f}_k) = k \prod_{j=1}^{k} \mathcal{N}(\mathbf{f}_j | \mathbf{m}_j, \mathbf{V}_j)$.

Copula Likelihood (Exact Factor)

Gaussian Prior

$$p(\mathbf{f}_1, \ldots, \mathbf{f}_k, \mathcal{D}_{U,V} | \mathcal{D}_{\mathbf{Z}}) = \left[\prod_{i=1}^{n} \overbrace{c_{\text{par}}\left[u_{1i}, u_{2i} | \sigma_1(f_{1i}), \ldots, \sigma_k(f_{ki})\right]}^{g_i(f_{1i}, \ldots, f_{ki})}\right] \left[\prod_{j=1}^{k} \mathcal{N}(\mathbf{f}_j | \mathbf{m}_{\mathbf{Z}}^{(j)}, \mathbf{V}_{\mathbf{Z},\mathbf{Z}}^{(j)})\right]$$

EP Projection

Exact

$$\mathcal{Q}(\mathbf{f}_1, \ldots, \mathbf{f}_k) = \left[\prod_{i=1}^{n} \underbrace{\tilde{k}_i \left[\prod_{j=1}^{k} \mathcal{N}(f_{ji}|, \tilde{m}_{ji}, \tilde{v}_{ji})\right]}_{\tilde{g}_i(f_{1i}, \ldots, f_{ki})}\right] \left[\prod_{j=1}^{k} \mathcal{N}(\mathbf{f}_j | \mathbf{m}_{\mathbf{Z}}^{(j)}, \mathbf{V}_{\mathbf{Z},\mathbf{Z}}^{(j)})\right]$$

Approximate Gaussian Factor

EP adjusts $\tilde{g}_i$ by minimizing $\mathrm{KL}[g_i \mathcal{Q}^{\setminus i} || \tilde{g}_i \mathcal{Q}^{\setminus i}]$, where $\mathcal{Q}^{\setminus i}$ is the ratio of $\mathcal{Q}$ and $\tilde{g}_i$. EP with sparse GPs [Naish-Guzman et al. 2008].

An Additional Approximation

Minimizing $\mathrm{KL}[g_i \mathcal{Q}^{\setminus i} || \tilde{g}_i \mathcal{Q}^{\setminus i}]$ requires to compute the first and second moments of $g_i(f_{1i}, \ldots, f_{ki})\mathcal{Q}^{\setminus i}(f_{1i}, \ldots, f_{ki})$ with respect to each f_{ji}.

No analytical solution!
k-dimensional numerical integration infeasible!

Additional approximation:

$$\int f_{1i} g_i(f_{1i}, \ldots, f_{ki}) \mathcal{Q}^{\setminus i}(f_{1i}, \ldots, f_{ki})\, df_{1i} \cdots df_{ki} \approx$$
$$\approx C \times \int f_{1i} g_i(f_{1i}, \bar{f}_{2i}, \ldots, \bar{f}_{ki}) \mathcal{Q}^{\setminus i}(f_{1i}, \bar{f}_{2i}, \ldots, \bar{f}_{ki})\, df_{1i}$$

where $\bar{f}_{1i}, \ldots, \bar{f}_{ki}$ are the means of $f_{1i}, \ldots, f_{ki}$ with respect to $\mathcal{Q}^{\setminus i}$, and C is a constant that approximates the width of the integrand around its maximum in all dimensions except f_{1i}.

Now we only need a 1-dimensional numerical integration!

An Illustrative Example

The approximation is very accurate since g_i is locally very flat.
To illustrate this, we plot for a conditional Student's t copula the normalized versions of

$$q_1^\star(f_{1i}) = \int g_i(f_{1i}, f_{2i}) \mathcal{Q}^{\setminus i}(f_{1i}, f_{2i})\, df_{2i}\,, \quad q_1(f_{1i}) = g_i(f_{1i}, \bar{f}_{2i}) \mathcal{Q}^{\setminus i}(f_{1i}, \bar{f}_{2i})\,,$$

and similarly for $q_2^\star(f_{2i})$ and $q_2(f_{2i})$.

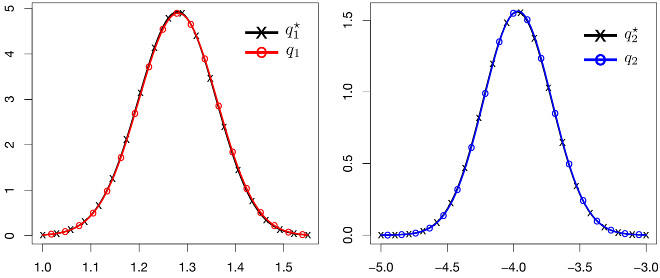

$q_1^\star(f_{1i})$ and $q_1(f_{2i})$ are almost identical. Finding the moments of $q_1^\star(f_{1i})$ requires 2-dimensional numerical integration. Finding the moments of $q_1(f_{2i})$ requires only a 1-dimensional integral.

Experiments

We evaluate the proposed Gaussian Process Conditional Copulas (GPCC) in experiments with synthetic and financial time series.

One-step-ahead rolling-window prediction task where we assume the copula function changes as a function of time.

We consider three parametric copula families for our GPCC:

- Gaussian: GPCC-G.
- Student's t: GPCC-T.
- Symmetrized Joe Clayton: GPCC-SJC.

Copula	Parameters	Transformation	Synthetic parameter function
Gaussian	correlation, τ	$0.99(2\Phi[f(t)]-1)$	$\tau(t) = 0.3 + 0.2\cos(t\pi/125)$
Student's t	correlation, τ	$0.99(2\Phi[f(t)]-1)$	$\tau(t) = 0.3 + 0.2\cos(t\pi/125)$
	degrees of freedom, ν	$1 + 10^6\Phi[g(t)]$	$\nu(t) = 1 + 2(1 + \cos(t\pi/250))$
SJC	upper dependence, τ^U	$0.01 + 0.98\Phi[g(t)]$	$\tau^U(t) = 0.1 + 0.3(1 + \cos(t\pi/125))$
	lower dependence, τ^L	$0.01 + 0.98\Phi[g(t)]$	$\tau^L(t) = 0.1 + 0.3(1 + \cos(t\pi/125 + \pi/2))$

Table: Copula parameters, modeling formulae and parameter functions used to generate synthetic data. Φ is the standard Gaussian cumulative density function f and g are GPs.

Benchmark Methods

We compare with Gaussian, Student's t and symmetrized Joe Clayton copulas with static parameters: CONST-G, CONST-T, CONST-SJC.

Time-varying Correlation Model (TVC). Student's t copula with ν constant and $\rho_t = (1 - \alpha - \beta)\rho + \alpha\epsilon_{t-1} + \beta\rho_{t-1}$, where ϵ_{t-1} is the empirical correlation of the last 10 data points [Jondeau et al. 2006].

Dynamic Symmetrized Joe Clayton Copula (DSJCC). The parameters τ^U and τ^L satisfy

$$\tau_t^U = 0.01 + 0.98\Lambda\left[\omega_U + \alpha_U\varepsilon_{t-1} + \beta_U\tau_{t-1}^U\right],$$
$$\tau_t^L = 0.01 + 0.98\Lambda\left[\omega_L + \alpha_L\varepsilon_{t-1} + \beta_L\tau_{t-1}^L\right],$$

where $\Lambda[\cdot]$ is the logistic function, $\varepsilon_{t-1} = \frac{1}{10}\sum_{j=1}^{10} |u_{t-j} - v_{t-j}|$, (u_t, v_t) is a copula sample at time t [Patton 2006].

Two-state Hidden Markov Model (HMM). At any time the copula is Student's t with different parameters for each state. Two regimes of low/high dependence [Jondeau et al. 2006].

Visualization of Synthetic Time Series

Each time series has length 5001. Rolling-window of size 1000.

Training data sampled from a dynamic Student's t copula.

Red: Evolution of τ. Blue: Evolution of ν.

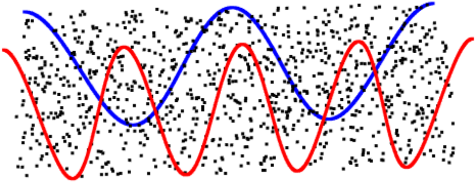

Results on Synthetic Time Series

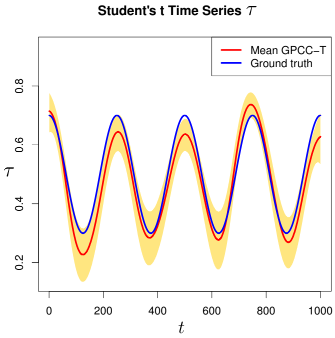

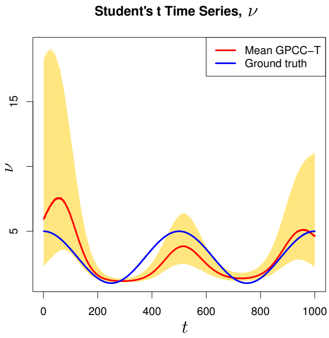

Average Test Log-Likelihood

Method	Gaussian	Student	SJC
GPCC-G	0.3347	0.3879	0.2513
GPCC-T	0.3397	**0.4656**	0.2610
GPCC-SJC	0.3355	0.4132	**0.2771**
HMM	**0.3555**	0.4422	0.2547
TVC	0.3277	0.4273	0.2534
DSJCC	0.3329	0.4096	0.2612
CONST-G	0.3129	0.3201	0.2339
CONST-T	0.3178	0.4218	0.2499
CONST-SJC	0.3002	0.3812	0.2502

Experiments with Foreign Exchange Time Series

Daily returns of nine currency pairs from 02-01-1990 to 15-01-2013 (6011 observations). Each pair includes the Swiss franc (CHF) and another currency.

CHF is a safe haven currency. We expect correlations between CHF and other currencies to have large variability across time.

Data preprocessed using an asymmetric AR(1)-GARCH(1,1) process with non-parametric innovations [Hernández-Lobat et al. 2008].

Code	**Currency Name**
CHF	Swiss Franc
AUD	Australian Dollar
CAD	Canadian Dollar
JPY	Japanese Yen
NOK	Norwegian Krone
SEK	Swedish Krona
EUR	Euro
NZD	New Zeland Dollar
GBP	British Pound

Results on Foreign Exchange Time Series

Method	AUD	CAD	JPY	NOK	SEK	EUR	GBP	NZD
GPCC-G	0.1260	0.0562	**0.1221**	0.4106	0.4132	0.8842	0.2487	0.1045
GPCC-T	**0.1319**	**0.0589**	0.1201	**0.4161**	**0.4192**	**0.8995**	**0.2514**	**0.1079**
GPCC-SJC	0.1168	0.0469	0.1064	0.3941	0.3905	0.8287	0.2404	0.0921
HMM	0.1164	0.0478	0.1009	0.4069	0.3955	0.8700	0.2374	0.0926
TVC	0.1181	0.0524	0.1038	0.3930	0.3878	0.7855	0.2301	0.0974
DSJCC	0.0798	0.0259	0.0891	0.3994	0.3937	0.8335	0.2320	0.0560
CONST-G	0.0925	0.0398	0.0771	0.3413	0.3426	0.6803	0.2085	0.0745
CONST-T	0.1078	0.0463	0.0898	0.3765	0.3760	0.7732	0.2231	0.0875
CONST-SJC	0.1000	0.0425	0.0852	0.3536	0.3544	0.7113	0.2165	0.0796

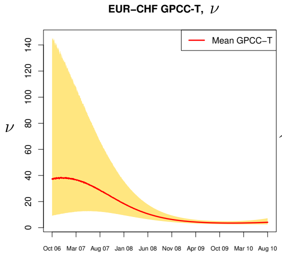

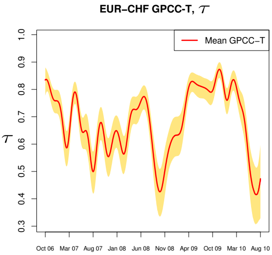

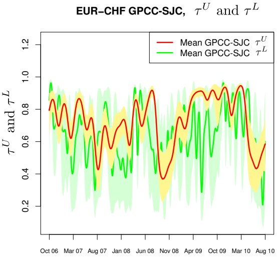

Visualization of Foreign Exchange Time Series

Predictions of GPCC-T on the training data for CHF-EUR.

Red: Evolution of τ. Blue: Evolution of ν.

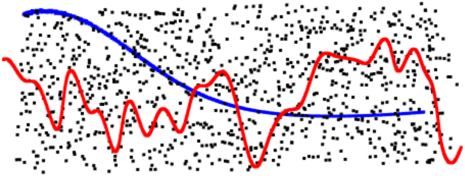

Experiments on Equity Exchange Time Series

Results on 8 pairs of financial time series with equity returns:

Method	HD HON	AXP BA	CNW CSX	ED EIX	HPQ IBM	BARC HSBC	RBS BARC	RBS HSBC
GPCC-G	0.1247	0.1133	0.1450	**0.2072**	0.1536	0.2424	0.3401	0.1860
GPCC-T	**0.1289**	**0.1187**	**0.1499**	0.2059	**0.1591**	0.2486	**0.3501**	**0.1882**
GPCC-SJC	0.1210	0.1095	0.1399	0.1935	0.1462	0.2342	0.3234	0.1753
HMM	0.1260	0.1119	0.1458	0.2040	0.1511	**0.2486**	0.3414	0.1818
TVC	0.1251	0.1119	0.1459	0.2011	0.1511	0.2449	0.3336	0.1823
DSJCC	0.0935	0.0750	0.1196	0.1721	0.1163	0.2188	0.3051	0.1582
CONST-G	0.1162	0.1027	0.1288	0.1962	0.1325	0.2307	0.2979	0.1663
CONST-T	0.1239	0.1091	0.1408	0.2007	0.1481	0.2426	0.3301	0.1775
CONST-SJC	0.1175	0.1046	0.1307	0.1891	0.1373	0.2268	0.2992	0.1639

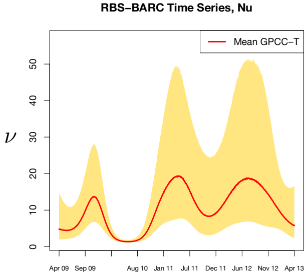

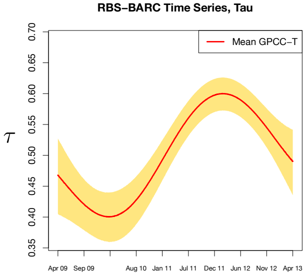

Summary and Conclusions

- We have proposed a framework for conditional dependencies based on parametric copula models with multiple parameters.
- Our approach is based on sparse GPs and expectation propagation for fast approximate inference.
- We have evaluated this framework on synthetic and financial time-series.
- We outperform static copula models and other dynamic copula models.
- Our approach can be easily extended to higher dimensions using vine factorizations [López-Paz et al. 2013].

Vine Factorizations

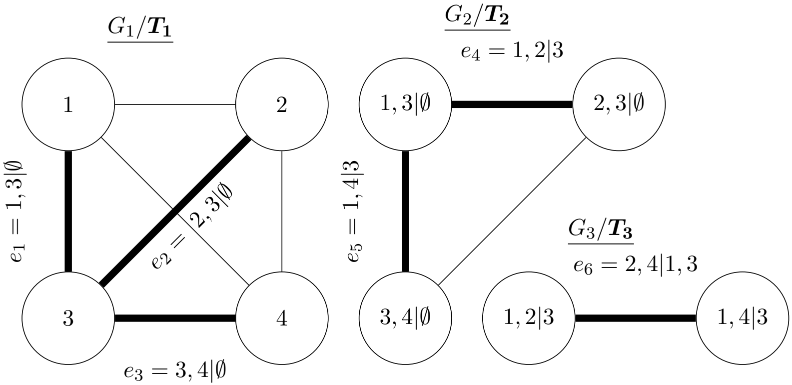

$$c_{1234} = \underbrace{\underbrace{c_{13|\emptyset}}_{e_1}\underbrace{c_{23|\emptyset}}_{e_2}\underbrace{c_{34|\emptyset}}_{e_3}}_{T_1}\underbrace{\underbrace{c_{12|3}}_{e_4}\underbrace{c_{14|3}}_{e_5}}_{T_2}\underbrace{\underbrace{c_{24|13}}_{e_6}}_{T_3}$$

Useful References I

- J. M. Hernndez-Lobato, J. R. Lloyd and D. Hernndez-Lobato. Gaussian Process Conditional Copulas with Applications to Financial Time Series. In Advances in Neural Information Processing Systems 26 (NIPS), 2013.
- A. Sklar. Fonctions de répartition à n dimensions et leurs marges. Publ. Inst. Statis. Univ. Paris, 8(1):229–231, 1959.
- H. Joe. Asymptotic efciency of the two-stage estimation method for copula-based models. J. Multivar. Anal., 94(2):401–419, 2005.
- A. J. Patton. Modelling asymmetric exchange rate dependence. International economic review, 47(2):527–556, 2006.
- P. Hennig. Animating Samples from Gaussian Distributions. Max Planck Institute for Intelligent Systems, Technical Report 8, 2013.
- E. Snelson and Z. Ghahramani. Sparse Gaussian processes using pseudo-inputs. In NIPS, 1257–1264, 2006.

Useful References II

- A. Naish-Guzman and S. B. Holden. The generalized FITC approximation. In NIPS, 1057–1064, 2008.
- E. Jondeau and M. Rockinger. The copula-garch model of conditional dependencies: An international stock market application. Journal of international money and nance, 25(5):827–853, 2006.
- J. M. Hernandez-Lobato, D. Hernández-Lobato, and A. Suárez. GARCH processes with non-parametric innovations for market risk estimation. In ICANN, 718–727, 2007.
- D. López-Paz, J. M. Hernández-Lobato and Z. Ghahramani. Gaussian Process Vine Copulas for Multivariate Dependence, In ICML, 2013.

Thank you for your attention!