



kdecopula: An R Package for the Kernel Estimation of Bivariate Copula Densities

Thomas Nagler

Technische Universität München

Abstract

We describe the R package **kdecopula** (current version 0.9.2), which provides fast implementations of various kernel estimators for the copula density. Due to a variety of available plotting options it is particularly useful for the exploratory analysis of dependence structures. It can be further used for accurate nonparametric estimation of copula densities and resampling. The implementation features spline interpolation of the estimates to allow for fast evaluation of density estimates and integrals thereof. We utilize this for a fast renormalization scheme that ensures that estimates are *bona fide* copula densities and additionally improves the estimators' accuracy. The performance of the methods is illustrated by simulations.

Keywords: dependence, copula, nonparametric, kernel density, exploratory data analysis, R.

1. Introduction

Dependence modeling with copulas has attracted a lot of attention in recent decades. By now, copulas are established tools in many fields of applied statistics, such as finance (Cherubini, Luciano, and Vecchiato 2004), hydrology (Salvadori and De Michele 2007), or machine learning (Elidan 2013).

At the very heart of copula theory is the famous theorem of Sklar (Sklar 1959). It states that any multivariate distribution function can be decomposed into the marginal distributions and a copula, which captures the dependence between variables. Let X and Y be two continuous random variables with joint distribution F and marginal distributions F_X and F_Y . Then, for all (x, y) in the support of the random vector (X, Y) ,

$$F(x, y) = C(F_X(x), F_Y(y)).$$

The copula $C: [0, 1]^2 \rightarrow [0, 1]$ is the bivariate distribution function of the random vector

$(U, V) = (F_X(X), F_Y(Y))$ which has a uniform marginal distribution. If F admits a density, we can also decompose the density f into

$$f(x, y) = c(F_X(x), F_Y(y))f_X(x)f_Y(y), \quad (1)$$

where c , f_X and f_Y are the densities corresponding to C , F_X and F_Y , respectively.

One of the major benefits of copula-based modeling is that inference for marginal distributions can be separated from the modeling of the dependence structure, i.e., the copula. For the estimation of the copula density c , it is most common to take a two-step approach: First, obtain estimates \hat{F}_X, \hat{F}_Y of the marginal distributions. A convenient and flexible way to do this is to use the empirical distribution function as an estimator. Second, define *pseudo-observations* $(\hat{U}, \hat{V}) = (\hat{F}_X(X), \hat{F}_Y(Y))$. The copula density is then estimated as the joint density of (\hat{U}, \hat{V}) .

Often, one assumes a parametric model for the copula density c and estimates its parameters by maximum-likelihood. Although there is a large variety of parametric copula models, they notoriously lack flexibility and bear the risk of misspecification. Nonparametric density estimators remedy these issues. But since copulas live on a bounded support – the unit hypercube – estimators have to be carefully tailored to this problem.

A specific class of nonparametric density estimators are kernel estimators. They are a popular tool for exploratory data analysis and widely used in many disciplines (e.g., Aitken and Lucy 2004; Kie, Matthiopoulos, Fieberg, Powell, Cagnacci, Mitchell, Gaillard, and Moorcroft 2010). The R (R Core Team 2017) package **kdecopula** (Nagler 2018b) implements several bivariate kernel copula density estimators that have been proposed in recent years. In a nutshell, the package provides methods for:

- estimation,
- bandwidth selection,
- simulation,
- visualization.

Package **kdecopula** is available from the Comprehensive R Archive Network (CRAN) at <https://CRAN.R-project.org/package=kdecopula>.

There exist two alternative methods for the kernel estimation of copula densities in R: the function **kcopula** from the **ks** package (Duong 2018) implements an estimator that glues together two independent kernel estimates for the center and boundary of the unit square; the function **npcopula** from the **np** package (Hayfield and Racine 2008) derives the copula density from an estimate of the joint distribution (as proposed by Racine 2015). However, these implementations do not reflect the numerous specialized contributions on the topic. Our package closes this gap by implementing state-of-the-art methods for kernel copula density estimation. The implemented methods are substantially more accurate than existing implementations. Additionally, the package provides a normalization algorithm which ensures that estimates are *bona fide* copula densities and further improves the accuracy.

Apart from kernel estimators, Schellhase (2014) implemented nonparametric copula density estimators based on penalized likelihood estimation in the **pencopula** package (using B -splines

or Bernstein polynomials). The extension package **penDvine** (Schellhase 2015) provides a convenient version with automatic bandwidth selection. A comparison with our implementation will show that these estimators are only competitive when the dependence is weak. The author is not aware of any software implementations for nonparametric copula density estimation outside of R.

In Section 2, we give a review of kernel copula density estimators and point to the relevant literature. Section 3 describes the functionality of the package and gives examples for its use. In Section 4, we give background on the implementation of the estimators using spline interpolation for fast evaluation and renormalization of the estimates. The statistical accuracy of the estimators in this package and other nonparametric copula density estimators (see previous paragraphs) is compared in Section 5. A summary is given in Section 6.

2. Kernel estimators of the copula density: A review

This section will review different approaches to kernel estimation of the copula density. As is common in the literature, we focus on the bivariate case.

Assume we have *iid* observations (U_i, V_i) , $i = 1, \dots, n$, from a bivariate copula C and are interested in the estimation of the corresponding density $c(u, v)$. One could apply the usual kernel density estimator to this problem:

$$\hat{c}_n(u, v) = \frac{1}{n} \sum_{i=1}^n K_{b_n}(u - U_i) K_{b_n}(v - V_i), \quad (u, v) \in [0, 1]^2,$$

where we used the notation $K_b(\cdot) = K(\cdot/b)/b$. The kernel function K is typically assumed to be a symmetric, bounded probability density function on \mathbb{R}^2 and $b_n > 0$ is the smoothing or bandwidth parameter. There is a problem, however. The estimator will put a considerable amount of probability mass outside of the unit square. This implies that \hat{c}_n is not a density function on $[0, 1]^2$, because it does not integrate to one. The estimator will additionally suffer from severe bias at the boundaries (see, e.g., Charpentier, Fermanian, and Scaillet 2006). Three different approaches to tackle this problem have emerged. All three techniques arose initially in the context of univariate kernel density estimation on the unit line. The following sections explain the ideas behind them (in the context of copulas) and give references for more detailed accounts.

2.1. The mirror-reflection method

An intuitive way of adapting \hat{c}_n to make sure that it is a density on $[0, 1]^2$ is the following: gather all probability mass that was put outside of the unit square, and redistribute it back to $[0, 1]^2$. This is the idea behind the *mirror-reflection technique*, which was proposed for copula density estimation by Gijbels and Mielniczuk (1990). As indicated by the name, all data are reflected at the corners and edges of the boundary region. The augmented data set containing all reflections is given by

$$(\tilde{U}_{ik}, \tilde{V}_{ik})_{k=1, \dots, 9} = \left\{ (U_i, V_i), (-U_i, V_i), (U_i, -V_i), (-U_i, -V_i), (U_i, 2 - V_i), \right. \\ \left. (-U_i, 2 - V_i), (2 - U_i, V_i), (2 - U_i, -V_i), (2 - U_i, 2 - V_i) \right\}.$$

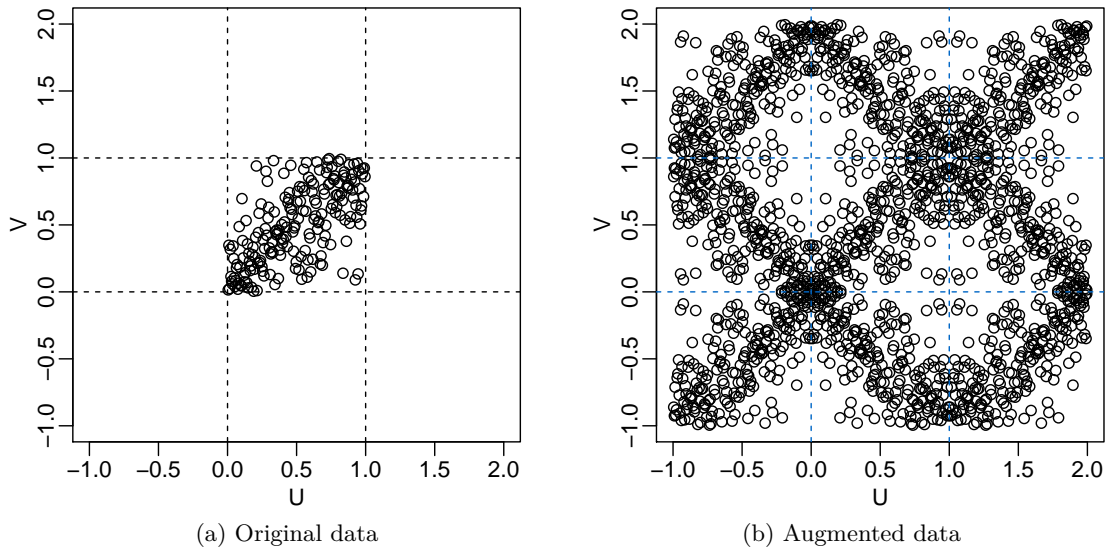


Figure 1: The data augmentation process. The set in (b) is obtained by reflecting all original data points w.r.t. to all corners and edges.

A visualization of the augmented data set is given in Figure 1. The *mirror-reflection estimator* is then defined as the usual kernel density estimator on the augmented data:

$$\hat{c}_n^{(\text{MR})}(u, v) = \frac{1}{n} \sum_{i=1}^n \sum_{k=1}^9 K_{b_n}(u - \tilde{U}_{ik}) K_{b_n}(v - \tilde{V}_{ik}), \quad (u, v) \in [0, 1]^2.$$

By reflecting all data points at the corners and edges also the probability mass outside of the unit square gets reflected back to the interior. As a result, the estimator now integrates to one. A detailed analysis of the asymptotic properties and a method for automatic bandwidth selection are given in Nagler (2014).

2.2. The beta kernel method

A second approach is to use kernels whose support matches the support of the density we want to estimate, and vary the shape of those kernels depending on the point where the density shall be estimated. This is achieved by so-called *boundary kernels*, and *beta kernels* are one instance. An estimator of the copula density based on this idea was proposed by Charpentier *et al.* (2006):

$$c^{(\beta)}(u, v) = \frac{1}{n} \sum_{i=1}^n \beta\left(U_i; \frac{u}{b_n} + 1, \frac{1-u}{b_n} + 1\right) \beta\left(V_i; \frac{v}{b_n} + 1, \frac{1-v}{b_n} + 1\right), \quad (u, v) \in [0, 1]^2,$$

where $\beta(\cdot; p, q)$ is the density of a Beta(p, q)-distributed random variable. We refer to Nagler (2014) for details on asymptotics and bandwidth selection.

2.3. The transformation method

A third approach is inspired by the early work of Devroye and Györfi (1985) and was introduced to kernel copula density estimation by Charpentier *et al.* (2006). The simple idea is

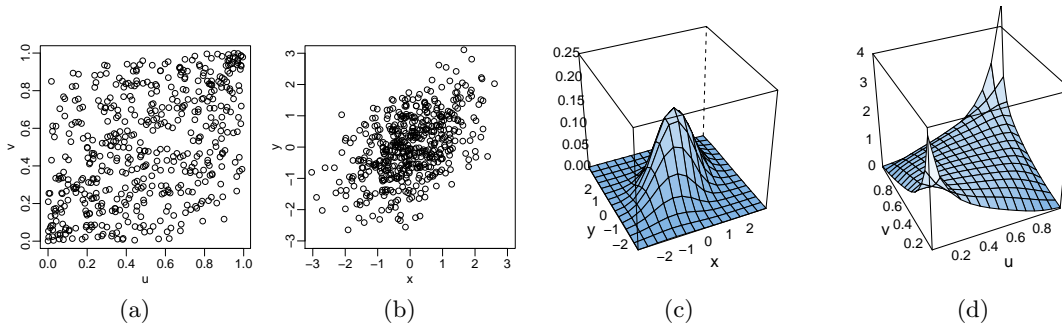


Figure 2: The transformation estimator. (a) copula sample; (b) transformed sample; (c) kernel density estimate for transformed sample; (d) kernel estimate of the copula density.

to transform the data so that it is supported on the full \mathbb{R}^2 (instead of the unit cube). On this transformed domain, standard kernel techniques can be used to estimate the density. An adequate back-transformation then yields an estimate of the copula density. For the transformation, the inverse of the standard normal *cdf* is most commonly used since it is known that kernel estimators tend to do well for Gaussian random variables.

Denote Φ as the standard Gaussian *cdf* and ϕ as its first order derivative. Then $(X_i, Y_i) = (\Phi^{-1}(U_i), \Phi^{-1}(V_i))$ is a random vector with Gaussian margins and copula C . By Sklar's theorem, the corresponding density f can be written as

$$f(x, y) = c(\Phi(x), \Phi(y))\phi(x)\phi(y). \quad (2)$$

This density can be easily estimated by a standard kernel estimator. From such an estimator \hat{f}_n , we can derive an estimator for the copula density c by isolating c in (2):

$$\hat{c}_n^{(T)}(u, v) = \frac{\hat{f}_n(\Phi^{-1}(u), \Phi^{-1}(v))}{\phi(\Phi^{-1}(u))\phi(\Phi^{-1}(v))}, \quad (u, v) \in [0, 1]^2. \quad (3)$$

This procedure is illustrated in Figure 2. The left panel shows the original data from the copula density c ; next to it we see the transformed data after the inverse Gaussian *cdf* has been applied. The third plot shows a kernel estimate of the density of the transformed data; and finally, the fourth plot shows the corresponding kernel estimate of the copula density.

The most natural choice for f_n is the conventional kernel density estimator. More recently, [Geenens, Charpentier, and Paindaveine \(2017\)](#) proposed to use a local likelihood estimator with nearest-neighbor bandwidths instead. Another recent extension was introduced by [Wen and Wu \(2015\)](#) who suggested to taper the back-transformation in the tails by increasing the variance of the Gaussian densities in the denominator of (3). For more details, we refer to the original papers.

3. The package's functionality

In the following, we describe the most important functions provided by the package. All functions either produce or take objects of the S3 class 'kdecopula' for which several methods are available.

3.1. Estimation and bandwidth selection: `kdecop`

At the core of the `kdecopula` package is the function `kdecop`, which estimates the copula density from data. The only mandatory input is an $n \times 2$ matrix of copula data, i.e., data with standard uniform margins. Such data is usually obtained in a first step by applying the empirical marginal *cdfs* to the data. This is equivalent to a rank transformation as shown below. The following lines of code load the package and a data set available in the package. The data is transformed to uniform margins in the third line, and the last line fits the kernel estimator.

```
R> library("kdecopula")
R> data("wdbc", package = "kdecopula")
R> uv <- apply(wdbc[, c(2, 8)], 2, rank) / (nrow(wdbc) + 1)
R> fit <- kdecop(uv)
```

The output of the function `kdecop` is an object of class ‘`kdecopula`’ that contains all information collected during the estimation process. The most relevant information can be summarized as follows.

```
R> summary(fit)
```

```
Kernel copula density estimate (tau = 0.47)
-----
Variables:    mean radius -- mean concavity
Observations: 569
Method:       Transformation local likelihood, log-quadratic
              (nearest-neighbor, 'TLL2nn')
Bandwidth:    alpha = 0.3519647
              B = matrix(c(0.71, 0.7, -0.7, 0.71), 2, 2)
---
logLik: 201.22    AIC: -367.97    cAIC: -366.83    BIC: -293.11
Effective number of parameters: 17.23
```

Summary statistics such as *AIC* or the *effective number of parameters/degrees of freedom* can be accessed via the usual generic functions.

```
R> logLik(fit)
```

```
'log Lik.' 201.2196 (df=17.23373)
```

```
R> AIC(fit)
```

```
[1] -367.9718
```

The function `kdecop` provides all estimation methods mentioned in Section 2. The estimation method can be specified via the `method` argument, e.g., `kdecop(..., method = "MR")`. For each method, we implemented an automatic bandwidth selection procedure. Below we list all implemented methods including a reference to the bandwidth selection procedure used:

- MR** The mirror-reflection estimator of [Gijbels and Mielniczuk \(1990\)](#). Smoothing parameters are selected by minimizing the AMISE using the Frank copula as the reference copula (see [Nagler 2014](#), Section 3.2.4).
- beta** The beta kernel estimator of [Charpentier *et al.* \(2006\)](#). Smoothing parameters are selected by minimizing the AMISE using the Frank copula as the reference copula (see, [Nagler 2014](#), Section 3.3.3).
- T** The transformation estimator of [Charpentier *et al.* \(2006\)](#), but allowing for a bandwidth matrix and not just one parameter. The bandwidth matrix is set by a rule of thumb which is the normal reference rule on the transformed domain (see, [Nagler 2014](#), Section 3.4.4):

$$B_{\text{rot}} = n^{-1/6} \widehat{\Sigma}_{\mathbf{Z}}^{1/2},$$

where $\widehat{\Sigma}_{\mathbf{Z}}$ is the empirical covariance matrix of $\Phi^{-1}(U_i)$ and $\Phi^{-1}(V_i)$, $i = 1, \dots, n$.

- TLL1, TLL2, TLL1nn, TLL2nn** (default): The transformation local likelihood estimator of [Geenens *et al.* \(2017\)](#). TLL1 approximates the log-density linearly; TLL2 by quadratic polynomials. The ***nn** versions use nearest-neighbor bandwidths instead of fixed ones. For fixed-bandwidth versions, the bandwidth matrix is set by the rule of thumb

$$B_{\text{rot}} = 3n^{-1/(4q^*+2)} \widehat{\Sigma}_{\mathbf{Z}}^{1/2}, \quad q^* = 1 + \lfloor q/2 \rfloor,$$

where q is the degree of the polynomial, $\widehat{\Sigma}_{\mathbf{Z}}$ is the empirical covariance matrix of $\Phi^{-1}(U_i)$ and $\Phi^{-1}(V_i)$, $i = 1, \dots, n$. This rule of thumb is similar to the normal reference rule, but ensures that the bandwidth matrix vanishes at the mean-square optimal rate. For nearest-neighbor methods, smoothing parameters are selected based on univariate least-squares cross-validation on the first principal component in the transformed domain (see [Geenens *et al.* 2017](#), Section 4). Local likelihood fitting is implemented via the **locfit** package ([Loader 2013](#)).

- TTCV, TTPI** Tapered transformation estimator of [Wen and Wu \(2015\)](#)¹. Smoothing parameters are selected in the transformed domain by profile cross-validation (TTCV, see [Wen and Wu 2015](#), Section 4.2) or plug-in minimization of the AMISE (TTPI, see [Wen and Wu 2015](#), Section 4.1).

It is possible to specify the bandwidths manually using the **bw** argument of **kdecop**, although we recommend against it. If it is necessary to manually make an estimate more or less smooth, we advise to use the bandwidth multiplier argument **kdecop**(..., **mult** = 1). Values larger than one will make the estimate smoother; values less than one make the estimate less smooth.

3.2. Working with the estimated density: (d/p/r)kdecop

In analogy to the usual (d/p/r) prefixes for distribution families in R, we provide (d/p/r) versions for the **kdecop** family. The functions **dkdecop** and **pkdecop** can be used to evaluate the density and *cdf* of a ‘**kdecopula**’ object, respectively.

¹The implementation of the tapered transformation estimators was kindly provided by Kuangyu Wen.

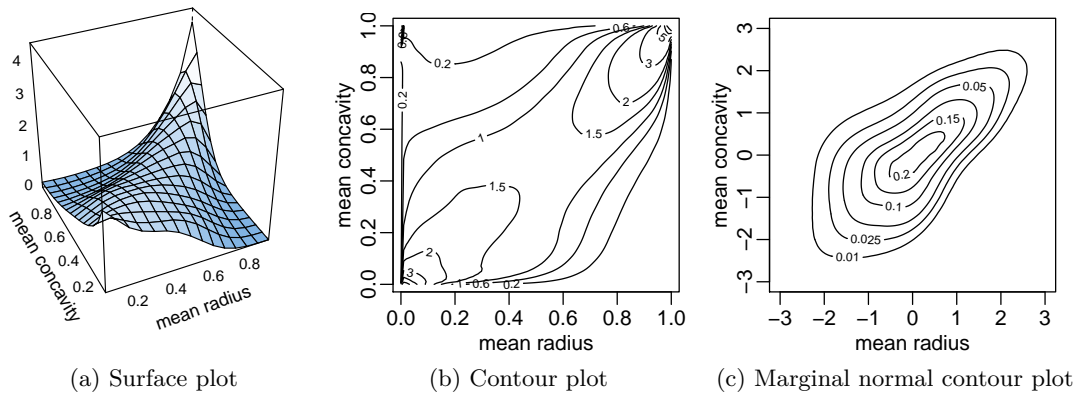


Figure 3: Three ways to visualize a copula density.

```
R> dkdecop(c(0.1, 0.2), fit)
```

```
[1] 1.691764
```

```
R> pkdecop(cbind(c(0.1, 0.9), c(0.1, 0.9)), fit)
```

```
[1] 0.0327257 0.8505370
```

The `rkdecop` function simulates data from the estimated density. This can be done in two ways: (a) using pseudo-random numbers based on `runif`, (b) using quasi-random numbers based on `ghalton` from the `qrng` package (Hofert and Lemieux 2016).

```
R> pseudo <- rkdecop(500, fit)
```

```
R> quasi <- rkdecop(500, fit, quasi = TRUE)
```

3.3. Visualization: The plot and contour generics

For many people, the most interesting feature is probably to make exploratory plots. There are three common ways to visualize a copula density: (a) a surface (or perspective) plot of the copula density, (b) a contour plot of the copula density, (c) a contour plot of the copula density when combined with standard normal margins. The following three lines of code produce the plots shown in Figure 3. Optionally, further arguments can be passed to improve the aesthetics.

```
R> plot(fit)
```

```
R> contour(fit, margins = "unif")
```

```
R> contour(fit)
```

In the author's experience, the most useful plot is (c), the marginal normal contour plot. Copula densities usually explode at some corners of the unit square. As a result, it is not possible to reasonably visualize the raw density (a) on the whole $[0, 1]^2$. The `plot` method for 'kdecopula' objects therefore restricts the displayed area to $[0.01, 0.99]^2$. Obviously, this

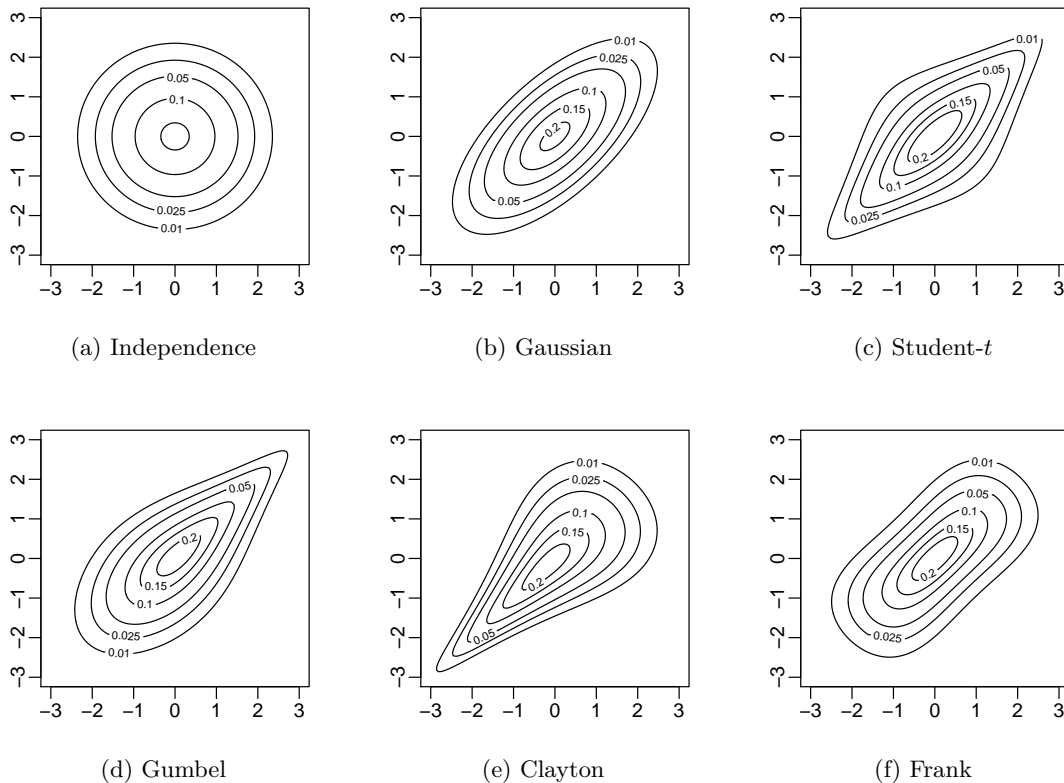


Figure 4: Marginal normal contour plots for several parametric copula families.

hides some information in the tails. This can be problematic because the tails are often of particular interest in copula models. But even on the restricted domain the copula density often attains values larger than 20 when there is strong dependence in the tails. On this scale, the shape of the copula density close to the center of the unit square is difficult to assess. Similarly, the contours of the raw density (b) are inappropriate to reflect the tail behavior because contour lines near to the corners become too close to be visually distinguishable. The marginal normal contour plot overcomes these issues by transforming the margins such that the transformed density is bounded. It has the additional advantage that it allows for an intuitive interpretation that is relative to the Gaussian copula as explained in the following paragraph.

If the true copula is the Independence copula, the contours are perfect circles (see Figure 4a). This is obviously not the case for the estimated density in Figure 3. Figure 4b shows a Gaussian copula with Kendall's τ set to the estimated τ from the data. A Gaussian copula combined with Gaussian margins results in a bivariate Gaussian density and its contours are ellipses. Since most statisticians are familiar with this kind of distribution it seems natural to use this as a benchmark when interpreting the marginal normal contour plot for other copulas. The next plot, Figure 4c, shows the Student- t copula ($df = 3$). Here the contours look like a diamond due to the higher density values in the tails (i.e., the corners of the square). This reflects that – in contrast to the Gaussian copula – the Student- t copula exhibits *tail dependence* (e.g., Joe 2014), a concept that is very important in the modeling of risks. In general, a spiky shape in the corners of the contours is an indication of tail dependence in

the respective corner. This can be observed again in Figures 4d and 4e, where the Gumbel and Clayton copula are shown, respectively. The Gumbel copula is asymmetric and features upper tail dependence only. This is reflected by a spiky shape in the upper right corner and a flatter shape in the lower left corner. For the Clayton copula it is the other way around. Finally the Frank copula has no tail dependence and has lighter tails than the Gaussian, which corresponds to a more flat shape of the contours.

Going back to the estimated density in Figure 3, we see a rather flat shape in the lower left corner and a more spiky shape in the upper right corner. This would indicate that there is no lower, but upper tail dependence. Hence, the Gumbel copula is the most appropriate fit choosing from the parametric families in Figure 4. However, we also observe some asymmetry with respect to the main diagonal. This is not reflected by any of the parametric models under consideration.

3.4. Dependence measures

It is often useful to summarize the dependence in a single number, a *dependence measure*. Many of the popular dependence measures are functionals of the copula. In fact, this property is required by Rényi's axioms, see Schweizer and Wolff (1981). Such measures can be calculated for copula density estimates fitted with `kdecop()`. For example, Kendall's τ can be expressed as

$$\tau = 4 \int_{[0,1]^2} C(u, v) c(u, v) du dv - 1. \quad (4)$$

There are several ways to calculate this measure for a copula density estimate. A straightforward way is to solve the integral numerically with the `cubature` package (Narasimhan and Johnson 2016):

```
R> library("cubature")
R> f <- function(u) pkdecop(u, fit) * dkdecop(u, fit)
R> int <- adaptIntegrate(f, lowerLimit = c(0, 0), upperLimit = c(1, 1))
R> 4 * int$integral - 1
```

```
[1] 0.4666633
```

This method is quite slow because two numerical integrals (one is in `pkdecop`) are nested. A second way to compute τ utilizes a stochastic interpretation of (4). Note that $\tau = 4E[C(U, V)] - 1$, where (U, V) is a random vector with density c . The expectation can then be approximated by Monte Carlo integration:

```
R> uv_qmc <- rkdecop(10^4, fit, quasi = TRUE)
R> 4 * mean(pkdecop(uv_qmc, fit)) - 1
```

```
[1] 0.4685433
```

The third way is to compute the sample version of Kendall's τ on simulated data. This method is asymptotically equivalent to the other methods, but simpler and faster:

```
R> cor(uv_qmc[, 1], uv_qmc[, 2], method = "kendall")
```

```
[1] 0.4718781
```

We see that the three methods lead to different results, but are close. The two Monte Carlo methods have the drawback that their result is random. A large number of samples need to be drawn to get an accuracy of multiple digits, but only the first one or two digits are useful for interpretation of the measure.

Other copula-based measures of dependence can be computed in a similar way. The `dep_measures()` function automates these computations for Kendall's τ , Spearman's ρ , Blomqvist's β , Gini's γ , van der Waerden's coefficient, mutual information, and Linfoot's correlation coefficient. For the definition of these measures, see, e.g., [Nelsen \(2006\)](#), [Genest and Verret \(2005\)](#), and [Joe \(1989\)](#).

```
R> dep_measures(fit)
```

kendall	spearman	blomqvist	gini	vd_waerden	minfo	linfoot
0.4708465	0.6527604	0.4980365	0.5368616	0.6561498	0.3311651	0.6959538

Package **kdecopula** can be used similarly to detect other dependence properties, like positive quadrant dependence (see [Gijbels, Omelka, and Sznajder 2010](#); [Racine 2015](#)).

4. Implementation based on spline interpolation

Typically, the evaluation of a kernel density estimate requires going back to the original data. As a result, the computational effort increases with the sample size. We avoid that issue by evaluating the actual density estimate only once on a fixed number of grid points. For further evaluations we use cubic spline interpolation between the values on this grid. This way, the density can be evaluated efficiently – independently of the sample size. It has the additional advantage that analytical expressions for integrals of the (interpolated) density estimate become available. We make use of that fact to implement a fast renormalization scheme that ensures that the density estimate is close to a *bona fide* copula density.

4.1. Evaluating the estimate by cubic spline interpolation

Recall that the support of a copula density is the unit cube $[0, 1]^2$. Let $m \in \mathbb{N}$ and define a finite set of points $p_j \in [0, 1]$ such that $0 \leq p_1 < \dots < p_m \leq 1$. Then, the set

$$\mathcal{P}_m = \{(u_j, v_k) : (j, k) \in \{1, \dots, m\}^2\} = \{(p_j, p_k) : (j, k) \in \{1, \dots, m\}^2\}$$

defines a symmetric grid on the unit cube. Cubic splines are piecewise cubic polynomials that can be used to approximate or interpolate some function between points on a grid. We will show how cubic spline interpolation can be used to approximate a copula density estimate. We explain in detail how a one-dimensional cubic spline interpolation is constructed when one of the coordinates is fixed. The two-dimensional case is a straightforward extension and only sketched.

The one-dimensional case

Let us first fix v_k and assume that the values of an estimate $\widehat{c}(\cdot, v_k): [0, 1] \rightarrow \mathbb{R}_+$ are available on the grid points u_j , $1 \leq j \leq m$. We want to interpolate the function $\widehat{c}(\cdot, v_k)$ at another point u^* , where $u_j < u^* < u_{j+1}$ for some $j \in \{2, \dots, m-2\}$. We define the interpolated curve segment $\tilde{c}^{j,j+1}(\cdot, v_k): [u_j, u_{j+1}] \rightarrow \mathbb{R}$ as some cubic polynomial

$$\tilde{c}^{j,j+1}(u, v_k) = a_0^{j,j+1} + a_1^{j,j+1}(u - u_j) + a_2^{j,j+1}(u - u_j)^2 + a_3^{j,j+1}(u - u_j)^3.$$

A cubic polynomial defined on a closed interval is fully determined by its function values and first derivatives at the boundary points. Define $\tilde{c}_1^{j,j+1}$ as the partial derivative of $\tilde{c}^{j,j+1}$ w.r.t. its first argument. After some simple algebraic manipulations, we find that the coefficients of a cubic spline approximation can be written as

$$\begin{aligned} a_0^{j,j+1} &= \tilde{c}^{j,j+1}(u_j, v_k), \\ a_1^{j,j+1} &= \tilde{c}_1^{j,j+1}(u_j, v_k), \\ a_2^{j,j+1} &= -3\tilde{c}^{j,j+1}(u_j, v_k) + 3\tilde{c}^{j,j+1}(u_{j+1}, v_k) - 2\tilde{c}_1^{j,j+1}(u_j, v_k) - \tilde{c}_1^{j,j+1}(u_{j+1}, v_k), \\ a_3^{j,j+1} &= 2\tilde{c}_1^{j,j+1}(u_j, v_k) - 2\tilde{c}_1^{j,j+1}(u_{j+1}, v_k) + \tilde{c}_1^{j,j+1}(u_j, v_k) + \tilde{c}_1^{j,j+1}(u_{j+1}, v_k). \end{aligned}$$

Now we replace $\tilde{c}^{j,j+1}(u_j, v_k)$ and $\tilde{c}^{j,j+1}(u_{j+1}, v_k)$ by the known values $\widehat{c}(u_j, v_k)$ and $\widehat{c}(u_{j+1}, v_k)$. Similarly, we want to replace the derivatives $\tilde{c}_1^{j,j+1}(u_j, v_k)$ and $\tilde{c}_1^{j,j+1}(u_{j+1}, v_k)$ by $\widehat{c}_1(u_j, v_k)$ and $\widehat{c}_1(u_{j+1}, v_k)$. These are unknown, but can be approximated by a finite difference scheme. We set

$$\begin{aligned} \tilde{c}_1^{j,j+1}(u_j, v_k) &= \frac{\widehat{c}(u_j, v_k) - \widehat{c}(u_{j-1}, v_k)}{u_j - u_{j-1}} - \frac{\widehat{c}(u_{j+1}, v_k) - \widehat{c}(u_{j-1}, v_k)}{u_{j+1} - u_{j-1}} + \frac{\widehat{c}(u_{j+1}, v_k) - \widehat{c}(u_j, v_k)}{u_{j+1} - u_j}, \\ \tilde{c}_1^{j,j+1}(u_{j+1}, v_k) &= \frac{\widehat{c}(u_{j+1}, v_k) - \widehat{c}(u_j, v_k)}{u_{j+1} - u_j} - \frac{\widehat{c}(u_{j+2}, v_k) - \widehat{c}(u_j, v_k)}{u_{j+2} - u_j} + \frac{\widehat{c}(u_{j+2}, v_k) - \widehat{c}(u_{j+1}, v_k)}{u_{j+2} - u_{j+1}}. \end{aligned}$$

Note that these can only be computed for $2 \leq j \leq m-2$, since four distinct values $\widehat{c}(u_{j+j^*}, v_k)$, $j^* = -1, 0, 1, 2$, show up in the above formulas. For fixed v_k and some $u \in [u_2, u_{k-1}]$, the spline approximation $\tilde{c}(u, v_k)$ of the function $\widehat{c}(u, v_k)$ can then be written as

$$\tilde{c}(u, v_k) = \sum_{j=2}^{m-2} \tilde{c}^{j,j+1}(u, v_k) \mathbb{1}_{[u_j, u_{j+1})}(u).$$

We extend this to allow for the full range ($u \in [0, 1]$) by extrapolating the ‘‘outer’’ two polynomials at the borders, i.e.,

$$\tilde{c}(u, v_k) = \sum_{j=2}^{m-2} \tilde{c}^{j,j+1}(u, v_k) \mathbb{1}_{A_j}(u), \quad \text{where } A_j = \begin{cases} [0, u_3) & \text{for } j = 2, \\ [u_j, u_{j+1}) & \text{for } 3 \leq j \leq k-3, \\ [u_{m-2}, 1] & \text{for } j = m-2. \end{cases}$$

The advantage of cubic spline interpolation is that it is easy to compute. In particular, the computational effort only depends on m , the number of knots. Additionally, the above

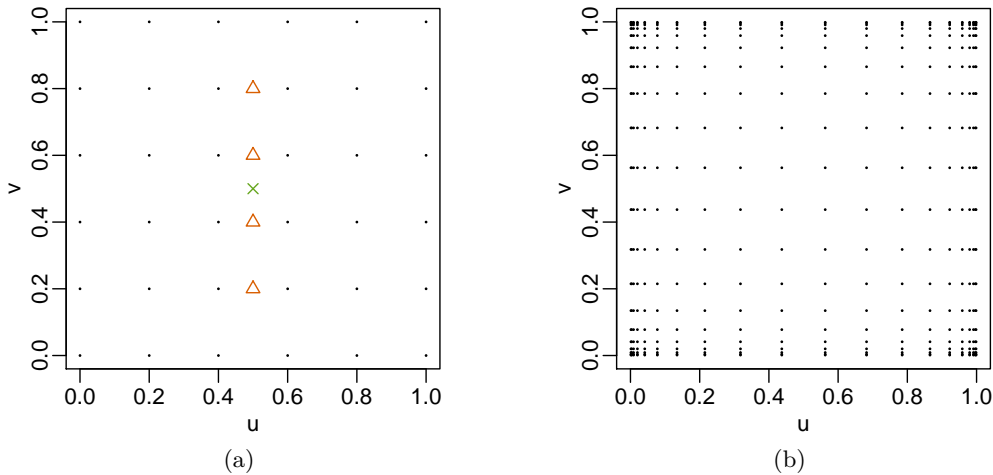


Figure 5: (a) Visualization of a two-dimensional interpolation as the sequence of two one-dimensional interpolations. (b) A grid with 20 knots that is equally spaced after transformation with the inverse Gaussian *cdf*.

approximation allows to write integrals as a sum of quartic polynomials, which can be computed equally fast. This will prove advantageous in Section 4.2, where we use such integrals to renormalize the copula density estimates.

The two-dimensional case

Bivariate functions can be approximated similarly by a sequence of two one-dimensional interpolations. We will illustrate this by a small example and omit any further details (for more, cf., [Habermann and Kindermann 2007](#)). Figure 5a shows the unit cube with grid points $(u_j, v_k) = (j/5, k/5)$, $j, k = 0, \dots, 5$, indicated as dots. Assume we know all function values $\tilde{c}(u_j, v_k)$ on this grid and want to approximate the function at the point $(0.5, 0.5)$ indicated by a cross. We first do four one-dimensional (horizontal) interpolations $\tilde{c}(0.5, v_k)$, $j = 1, 2, 3, 4$ (triangles). Note that all values that are required to calculate the spline coefficients are known. Another one-dimensional (vertical) interpolation based on the four new values $\tilde{c}(0.5, v_k)$ gives us the final interpolated value $\tilde{c}(0.5, 0.5)$.

Choice of grid

In Figure 5a we show a grid that has equal spacings between grid points. This seems natural, but in our context we found it more appropriate to use a grid that is equally spaced after a transformation by the inverse Gaussian *cdf*. Figure 5b depicts such a grid with 20 knots. It is constructed by placing 20 equidistant knots on the line segment $[-3, 3]$ and then applying the Gaussian *cdf* to them. The two-dimensional set-product of these 20 points yields the final two-dimensional grid.

We see that the grid points are more sparse in the center of the unit square and concentrate towards the boundaries and corners. This choice takes into account that for copula densities the areas near the corners are most important. In those areas, copula densities often explode while being rather flat in the center. This allows us to keep the approximation errors in the important areas small. A nice side-effect is that the marginal normal contour plots described

in the last section can be visualized more nicely. The number of knots can be controlled by the `knots` argument of `kdecop` and defaults to 30. A smaller number reduces computation time, but comes at the cost of a larger approximation error.

4.2. Renormalization of the density estimate

We now introduce the idea of iterative renormalization of kernel copula density estimators. Recall that by the definition of a copula, the marginal densities have to be uniform, i.e.,

$$\int_0^1 c(u, s) ds = \int_0^1 c(s, u) ds = 1, \quad \text{for all } u \in [0, 1]. \quad (5)$$

This property is of particular importance, when other functionals of the density are of interest. For example, assume that we integrate the density estimate to obtain an estimate for the corresponding conditional *cdf*, $C(v|u) = \int_0^v c(u, s) ds$. This is a common task in vine copula models which gained a lot of popularity following the seminal paper of Aas, Czado, Frigessi, and Bakken (2009). If the estimated density does not satisfy the uniform margins property, the estimate of the conditional *cdf* may exceed unity, which makes it problematic. The lack of uniform margins was mostly ignored in the literature, although kernel estimates usually do not satisfy the *uniform margins property* (5).

Now let $\hat{c}(u, v)$ be a consistent kernel estimator of $c(u, v)$ for all $(u, v) \in [0, 1]^2$. From Sklar's theorem for density functions (1), we know that dividing a bivariate density by the product of its marginal densities yields a copula density. Hence, a simple way to adjust the estimator is to divide the initial estimate by $\int_0^1 \hat{c}(u, s) ds \int_0^1 \hat{c}(s, v) ds$. The renormalized estimator writes

$$\hat{c}^*(u, v) = \frac{\hat{c}(u, v)}{\int_0^1 \hat{c}(u, s) ds \int_0^1 \hat{c}(s, v) ds} \quad (6)$$

and is asymptotically equivalent to the original one under mild conditions. For sophisticated kernel estimators – such as the beta kernel or local likelihood estimators – the two integrals have to be computed numerically. Conveniently, the spline approximations introduced in the previous section allow for fast computation of the integrals in (5).

The proposed renormalization procedure can be split into two steps.

1. Find a spline approximation of the initial estimate that is defined by its values on a finite grid.
2. Renormalize the approximated density values on this grid by dividing by the (approximated) marginal densities (see Equation 6).

The resulting estimate will typically be closer to a *bona fide* copula density. However, the renormalization is only carried out on a finite number of grid points. Apart from that grid, the renormalized estimate typically still does not satisfy the uniform margins property. But we can simply repeat the two steps above until a satisfactory result is achieved. Our experience suggests that a very small number of iterations is sufficient. The number of iterations can be set by the `renorm.iter` argument of `kdecop` and defaults to three.

The renormalization will turn out to have two benefits: functionals of kernel estimates will show the desired behavior and, additionally, the estimates are more accurate (see Section 5).

For the latter there is an intuitive interpretation. By ensuring that margins are uniform, we incorporate additional information about the true density. This reduces the set of plausible estimates and increases the probability of being “close” to the true one.

5. Comparison of methods

We compare all estimators implemented in this package in a simulation study. Additionally, we include results for other nonparametric copula density estimators available in R.

5.1. Setup

We consider the following estimators:

- MR, beta, T, TLL1, TLL2, TLL1nn, TLL2nn TTPI, TTCV: These estimators are provided by the `kdecop` function presented in this paper (cf., Section 2).
- np: The kernel estimator provided by the `npcopula` function in the `np` package (Hayfield and Racine 2008).
- ks: The kernel estimator provided by the `kcopula.de` function in the `ks` package (Duong 2018).
- bern, bspl: The penalized Bernstein polynomial and B -spline estimators provided by the `paircopula` function in the `penDvine` package (Schellhase 2015).

Default settings are used for all implementations. In particular, smoothing parameters are selected automatically.

As a performance measure, we use the *integrated absolute error (IAE)*

$$\text{IAE}[\hat{c}(u, v)] = \int_{[0,1]} \int_{[0,1]} |\hat{c}(u, v) - c(u, v)| dudv,$$

where we estimate the integrals as the mean over the grid $(j/101, k/101)$, $j, k = 1, \dots, 100$.

We study several target copula densities and two sample sizes ($n = 200, n = 1000$). We use three families from Figure 4 (Independence, Gaussian, Gumbel) and three copulas derived from multi-modal mixtures of bivariate Gaussian distributions (marginal normal contour plots are shown in Figure 6). For the Gaussian and Gumbel copulas we create scenarios with weak and strong dependence (Kendall’s τ of 0.3 and 0.7). For mixture families we use the package `GMCM` (Bilgrau, Eriksen, Rasmussen, Johnsen, Dybkær, and Bøgsted 2016).

5.2. Error analysis

Figure 7 shows the mean of the IAE over 100 replications for various scenarios. The figure is divided into 8 panels, each representing a different target density. Each panel shows results for the nine estimators implemented in package `kdecopula` (left of the dotted line) and four estimators implemented in other packages on CRAN. Each estimator is represented by a bar which is split into two parts. The solid part indicates the mean IAE for $n = 1000$, the transparent part indicates the mean IAE for $n = 200$. As expected, all estimators improve when the sample size increases.

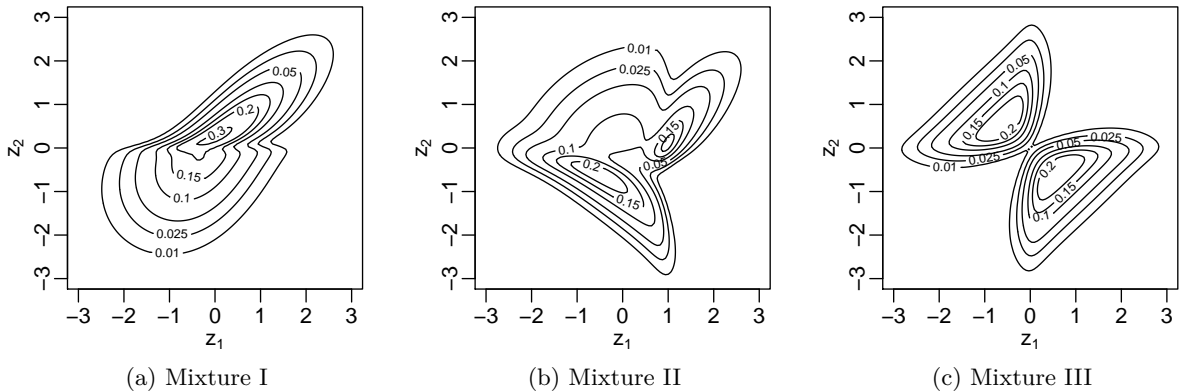


Figure 6: Marginal normal contour plots of copulas derived from mixtures of multivariate Gaussian distributions.

Two scenarios, the Independence and Mixture III copulas, show somewhat different results from the others. We shall discuss them in more detail later. In all other scenarios, TLL2nn is the top performer and TLL2 is a close second. In general, the transformation methods seem to work better than the MR and beta methods. The estimators from the **np** and **ks** packages are often among the least accurate. The methods from the **penDvine** package work rather well when there is weak dependence, but struggle when dependence is strong. The relative performance is largely consistent across the two sample sizes.

The density of the Independence copula is constant. The two estimators from package **penDvine** perform best for this target. The reason is that their basis function formulation can reproduce constants, so they behave similar to a parametric estimator in this case. The Independence copula is an example where the transformation approach is suboptimal. A constant density can be estimated easily because there is no curvature. But after transformation, the target density is a bivariate Gaussian which is more difficult to estimate.

The second scenario standing out is where Mixture III is the target density. The *B*-spline estimator from package **penDvine** is most accurate, followed by TTCV and TTPI. This scenario is the only one where the other transformation methods are consistently outperformed by non-**kdecopula** methods. The issue here is how the bandwidth matrix is selected. Methods T, TLL1, TLL2, TLL1nn, and TLL2nn all use a bandwidth matrix that is proportional to the square root of the empirical covariance. This is often a good choice because it “stretches” the kernels in a way that resembles the shape of the data. For Mixture III, the overall correlation is negative. The correlation in each mixture component, however, is strongly positive. In this case the shape of the kernels does not reflect the local orientation of the data. TTCV, TTPI, and non-**kdecopula** methods do better in this scenario because they do not rely on such a rule of thumb, but use more sophisticated criteria. This comes at a cost in terms of speed, as we shall see in Section 5.4.

5.3. The effect of renormalization

In Section 4.2 we claimed that the renormalization algorithm implemented in this package improves the performance of the estimators. The results presented in Figure 7 are based on default settings, i.e., three iterations of the algorithm. Table 1 show the relative reduction of

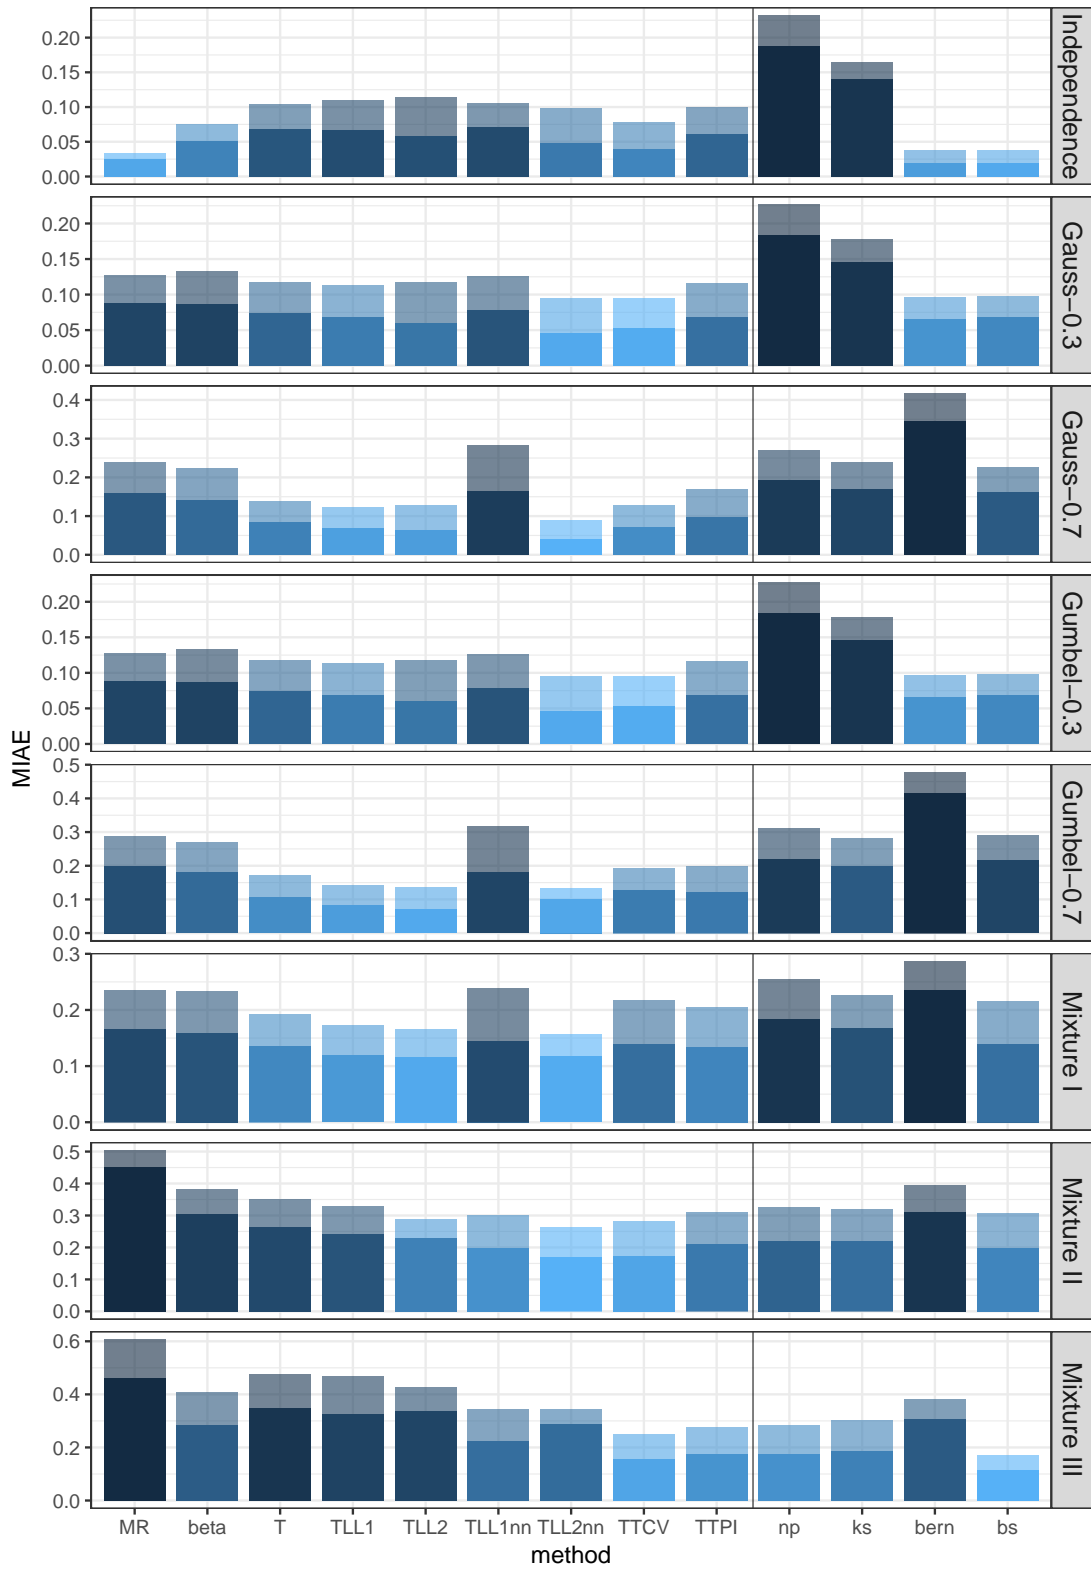


Figure 7: Mean of the integrated absolute error achieved by all methods in different scenarios. Each bar is split in two parts, where the larger corresponds to $n = 200$ and the smaller to $n = 1000$.

MR	beta	T	TLL1	TLL2	TLL1nn	TLL2nn	TTCV	TTPI
13%	15%	25%	39%	21%	28%	29%	24%	22%

Table 1: The effect of renormalization. The numbers indicate by how much the IAE could be improved after three iterations of the renormalization algorithm (rounded to the next integer).

n	MR	beta	T	TLL2	TLL2nn	TTCV	TTPI	np	ks	bern	bs
200	0.11	0.26	0.07	0.31	0.62	3.76	0.35	3.03	4.61	1.88	2.10
1000	0.37	0.99	0.11	0.68	1.30	58.10	13.63	37.84	41.96	3.08	3.45

Table 2: Computation time (in seconds) required for estimating the density and evaluating it on a 100×100 grid.

IAE compared to non-normalized estimators averaged over all scenarios.

We observe that the performance has improved for all estimators. The average gain ranges between 13% and 39%. This contributed significantly to the good performance observed in Figure 7. In fact, the estimator that could be improved the second most is TLL2nn, the top performer in our initial study.

5.4. Computation time

Some design choices in package **kdecopula** prioritize speed over accuracy. One example is the use of spline interpolation, another is the choice of bandwidth selection methods. Table 2 shows the computation time for most methods. Methods TLL1 and TLL1nn are omitted to make the table more compact; they are slightly faster than their “nn” variants.

All **kdecopula** methods except TTCV and TTPI are much faster than estimators from other packages. The difference gets larger for larger sample sizes. Note that TTCV, TTPI, np, and ks are much slower when $n = 1000$. This is caused by them using sophisticated bandwidth selection techniques that have a large complexity with respect to sample size.

6. Summary and extensions

We have described the R package **kdecopula**, which implements several cutting-edge kernel estimation techniques for copula densities. The package allows for automatic selection of the smoothing parameter and resampling. Several plotting options make it particularly useful for the exploratory analysis of copula data. Its abilities have been illustrated by small code examples.

The implementation utilizes spline interpolation for fast evaluation and renormalization of the density estimates. Simulations show that the implementations in this package perform best among available methods for nonparametric copula density estimation. A contributing factor for the good performance is the renormalization of the estimators, which was shown to notably improve the accuracy.

Compared with the functionality provided by the **np** package, package **kdecopula** lacks two features:

1. It does not provide functionality for estimating copula densities when the data contain

discrete variables. A problem is that the copula is not unique when margins are discrete, but there are infinitely many copulas that, when combined with the marginal distributions, lead to the same joint distribution (see, e.g., Genest and Nešlehová 2007). It is unclear which of these copulas shall be estimated and how to justify the choice. Additionally, estimating a copula from discrete data necessarily involves modeling of the marginal distributions, which is deliberately avoided in package **kdecopula**.

2. It does not allow for more than two variables. One major issue is that package **kdecopula** uses interpolation to evaluate and renormalize the estimators. In more than two dimensions the number of grid points explodes rapidly and renders the interpolation approach infeasible.

A **kdecopula**-based solution for both points is the **kdevine** package (Nagler 2017). It implements a kernel estimator of general multivariate densities based on vine copulas (Nagler and Czado 2016), which use marginal densities and bivariate copulas as building blocks. *Continuous convolution* (Nagler 2018a) is used to handle discrete variables, which induces copulas similar to the multilinear copula (see Genest and Nešlehová 2007; Genest, Nešlehová, and Rémillard 2014).

Acknowledgments

This work was partially supported by the German Research Foundation (DFG grant CZ 86/5-1). The author wants to thank two anonymous referees for many remarks and suggestions that considerably improved the article and software.

References

- Aas K, Czado C, Frigessi A, Bakken H (2009). “Pair-Copula Constructions of Multiple Dependence.” *Insurance: Mathematics and Economics*, **44**(2), 182–198. doi:10.1016/j.insmatheco.2007.02.001.
- Aitken CGG, Lucy D (2004). “Evaluation of Trace Evidence in the Form of Multivariate Data.” *Journal of the Royal Statistical Society C*, **53**(1), 109–122. doi:10.1046/j.0035-9254.2003.05271.x.
- Bilgrau AE, Eriksen PS, Rasmussen JG, Johnsen HE, Dybkær K, Bøgsted M (2016). “GMCM: Unsupervised Clustering and Meta-Analysis Using Gaussian Mixture Copula Models.” *Journal of Statistical Software*, **70**(2), 1–23. doi:10.18637/jss.v070.i02.
- Charpentier A, Fermanian JD, Scaillet O (2006). “The Estimation of Copulas: Theory and Practice.” In J Rank (ed.), *Copulas: From Theory to Application in Finance*. Risk Books.
- Cherubini U, Luciano E, Vecchiato W (2004). *Copula Methods in Finance*. Wiley Finance Series. John Wiley & Sons. doi:10.1002/9781118673331.
- Devroye L, Györfi L (1985). *Nonparametric Density Estimation: The L_1 View*. John Wiley & Sons.

- Duong T (2018). *ks: Kernel Smoothing*. R package version 1.11.0, URL <https://CRAN.R-project.org/package=ks>.
- Elidan G (2013). “Copulas in Machine Learning.” In P Jaworski, F Durante, WK Härdle (eds.), *Copulae in Mathematical and Quantitative Finance*, volume 213 of *Lecture Notes in Statistics*, pp. 39–60. Springer-Verlag.
- Geenens G, Charpentier A, Paindaveine D (2017). “Probit Transformation for Nonparametric Kernel Estimation of the Copula Density.” *Bernoulli*, **23**(3), 1848–1873. doi:10.3150/15-bej798.
- Genest C, Nešlehová J (2007). “A Primer on Copulas for Count Data.” *ASTIN Bulletin*, **37**(02), 475–515. doi:10.2143/ast.37.2.2024077.
- Genest C, Nešlehová JG, Rémillard B (2014). “On the Empirical Multilinear Copula Process for Count Data.” *Bernoulli*, **20**(3), 1344–1371. doi:10.3150/13-bej524.
- Genest C, Verret F (2005). “Locally Most Powerful Rank Tests of Independence for Copula Models.” *Journal of Nonparametric Statistics*, **17**(5), 521–539. doi:10.1080/10485250500038926.
- Gijbels I, Mielniczuk J (1990). “Estimating the Density of a Copula Function.” *Communications in Statistics – Theory and Methods*, **19**(2), 445–464. doi:10.1080/03610929008830212.
- Gijbels I, Omelka M, Sznajder D (2010). “Positive Quadrant Dependence Tests for Copulas.” *Canadian Journal of Statistics*, **38**(4), 555–581. doi:10.1002/cjs.10088.
- Habermann C, Kindermann F (2007). “Multidimensional Spline Interpolation: Theory and Applications.” *Computational Economics*, **30**(2), 153–169. doi:10.1007/s10614-007-9092-4.
- Hayfield T, Racine JS (2008). “Nonparametric Econometrics: The **np** Package.” *Journal of Statistical Software*, **27**(5), 1–32. doi:10.18637/jss.v027.i05.
- Hofert M, Lemieux C (2016). *qrng: (Randomized) Quasi-Random Number Generators*. R package version 0.0-3, URL <https://CRAN.R-project.org/package=qrng>.
- Joe H (1989). “Relative Entropy Measures of Multivariate Dependence.” *Journal of the American Statistical Association*, **84**(405), 157–164. doi:10.2307/2289859.
- Joe H (2014). *Dependence Modeling with Copulas*. Chapman & Hall/CRC, Boca Raton.
- Kie JG, Matthiopoulos J, Fieberg J, Powell RA, Cagnacci F, Mitchell MS, Gaillard JM, Moorcroft PR (2010). “The Home-Range Concept: Are Traditional Estimators Still Relevant with Modern Telemetry Technology?” *Philosophical Transactions of the Royal Society B*, **365**(1550), 2221–2231. doi:10.1098/rstb.2010.0093.
- Loader C (2013). *locfit: Local Regression, Likelihood and Density Estimation*. R package version 1.5-9.1, URL <http://CRAN.R-project.org/package=locfit>.
- Nagler T (2014). *Kernel Methods for Vine Copula Estimation*. Master’s thesis, Technische Universität München.

- Nagler T (2017). **kdevine**: *Multivariate Kernel Density Estimation with Vine Copulas*. R package version 0.4.1, URL <https://CRAN.R-project.org/package=kdevine>.
- Nagler T (2018a). “A Generic Approach to Nonparametric Function Estimation with Mixed Data.” *Statistics & Probability Letters*, **137**, 326–330. doi:10.1016/j.spl.2018.02.040.
- Nagler T (2018b). **kdecopula**: *Kernel Smoothing for Bivariate Copula Densities*. R package version 0.9.2, URL <https://CRAN.R-project.org/package=kdecopula>.
- Nagler T, Czado C (2016). “Evading the Curse of Dimensionality in Nonparametric Density Estimation with Simplified Vine Copulas.” *Journal of Multivariate Analysis*, **151**, 69–89. doi:10.1016/j.jmva.2016.07.003.
- Narasimhan B, Johnson SG (2016). **cubature**: *Adaptive Multivariate Integration over Hypercubes*. R package version 1.3-6, URL <https://CRAN.R-project.org/package=cubature>.
- Nelsen RB (2006). *An Introduction to Copulas*. Springer Series in Statistics. Springer-Verlag. doi:10.1007/0-387-28678-0.
- Racine JS (2015). “Mixed Data Kernel Copulas.” *Empirical Economics*, **48**(1), 37–59. doi:10.1007/s00181-015-0913-3.
- R Core Team (2017). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.
- Salvadori G, De Michele C (2007). “On the Use of Copulas in Hydrology: Theory and Practice.” *Journal of Hydrologic Engineering*, **12**(4), 369–380. doi:10.1061/(asce)1084-0699(2007)12:4(369).
- Schellhase C (2014). **penCopula**: *Flexible Copula Density Estimation with Penalized Hierarchical B-Splines*. R package version 0.3.5, URL <https://CRAN.R-project.org/package=penCopula>.
- Schellhase C (2015). **penDvine**: *Flexible Pair-Copula Estimation in D-Vines Using Bivariate Penalized Splines*. R package version 0.2.4, URL <https://CRAN.R-project.org/package=penDvine>.
- Schweizer B, Wolff EF (1981). “On Nonparametric Measures of Dependence for Random Variables.” *The Annals of Statistics*, **9**(4), 879–885. doi:10.1214/aos/1176345528.
- Sklar A (1959). *Fonctions de répartition à n dimensions et leurs marges*. Université Paris 8.
- Wen K, Wu X (2015). “Transformation-Kernel Estimation of the Copula Density.” Working paper, URL <http://agecon2.tamu.edu/people/faculty/wu-ximing/agecon2/public/copula.pdf>.

Affiliation:

Thomas Nagler

Technische Universität München

Zentrum Mathematik

Lehrstuhl für Mathematische Statistik

Boltzmannstraße 3,

85748 Garching, Germany

E-mail: thomas.nagler@tum.de

URL: <https://www.statistics.ma.tum.de/personen/mitarbeiter/thomas-nagler/>