

SELECTING BIVARIATE COPULA MODELS USING IMAGE RECOGNITION*

BY

ANDREAS TSANAKAS AND RUI ZHU

ABSTRACT

The choice of a copula model from limited data is a hard but important task. Motivated by the visual patterns that different copula models produce in smoothed density heatmaps, we consider copula model selection as an image recognition problem. We extract image features from heatmaps using the pre-trained AlexNet and present workflows for model selection that combine image features with statistical information. We employ dimension reduction via Principal Component and Linear Discriminant Analyses and use a Support Vector Machine classifier. Simulation studies show that the use of image data improves the accuracy of the copula model selection task, particularly in scenarios where sample sizes and correlations are low. This finding indicates that transfer learning can support statistical procedures of model selection. We demonstrate application of the proposed approach to the joint modelling of weekly returns of the MSCI and RISX indices.

KEYWORDS

Copula, dependence modelling, image recognition, model selection, classification, transfer learning.

1 INTRODUCTION

Copulas are dependence modelling tools of fundamental importance in actuarial and financial risk management (Frees and Valdez 1998; Denuit *et al.* 2006; McNeil *et al.* 2015) and fields beyond (e.g. Genest and Favre 2007). An extensive literature has emerged on specifically actuarial applications of copula

*The authors are grateful to Markus Gesmann, Marius Hofert, George Tzougas, Tim Verdonck, Jing-Hao Xue, and two anonymous reviewers, for constructive feedback.

Astin Bulletin 52(3), 707–734. doi:10.1017/asb.2022.12 © The Author(s), 2022. Published by Cambridge University Press on behalf of The International Actuarial Association. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution, and reproduction in any medium, provided the original work is properly cited.

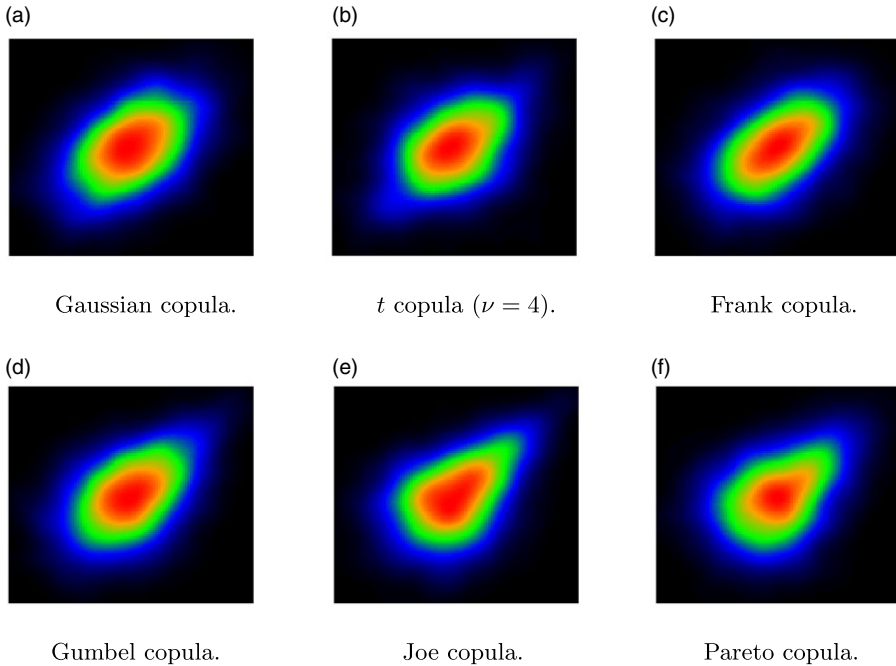


FIGURE 1. Examples of heatmap images of smoothed bivariate densities for different copula models ($n = 2000$, $\tau = 0.3$).

modelling, in credibility (Frees and Wang 2005), stochastic reserving (Shi and Frees 2011; Shi 2014; Abdallah *et al.* 2015), and claims modelling (Czado *et al.* 2012; Shi and Valdez 2014; Hu *et al.* 2021; Tzougas and Pignatelli di Cerchiara 2021). At the same time, the problem of choosing a copula model based on datasets that are of limited size is a non-trivial task, as evidenced in various strands of the literature, indicatively including: seminal work on copula goodness-of-fit (Genest *et al.* 2009); the study of practical problems arising in insurance risk management (Shaw *et al.* 2010); the impact of copula choice on portfolio risk (McNeil *et al.* 2015, Section 11.1.5); and the consideration of dependence uncertainty in a regulatory framework (Embrechts *et al.* 2014). For a historical perspective on the application of copulas to insurance risk modelling, see the interview of Edward (Jed) Frees, by Genest and Scherer (2020).

The different properties of alternative copula models are often visualised by joint density contour plots or heatmaps. For example, in Figure 1, we show heatmap images of smoothed bivariate densities for six well-known copula models, with standard Normal margins. It is obvious that different copula models have heatmaps with different patterns, reflecting for example different degrees of skewness. This observation motivates our research question: *Do images of smoothed joint densities convey useful information that can improve the accuracy of copula model selection procedures?*

Our paper seeks to address this question, in the context of small data sizes and bivariate copula models. Small data sizes, for example less than 250, make the copula model selection problem hard, hence an improvement offered by image data would be welcome. (Conversely, given the asymptotic consistency of statistical procedures, there is less scope for improvement when data sizes are large.) Furthermore, for small data sizes, it is natural to focus on the simplest models, given the likely lack of statistical power to detect more complex model features. In the case where datasets are richer, the problem the modeller faces is not so much one of selecting between different models, but more one of designing a model that is flexible enough to reflect idiosyncratic features of the data, see for example Hofert *et al.* (2021).

Here, we treat bivariate copula density heatmaps as RGB images and exploit the spatial patterns present in the images to aid bivariate copula selection. The copula selection task is treated as an image recognition or classification task: we classify a given copula sample to an element of a model set, based on its density heatmap image.

One vital challenge in image recognition is to obtain good representations of the images that can summarize well their distinct spatial patterns and thus make the recognition task easier; this is known as representation learning. Deep neural models have been demonstrated to be effective for representation learning, especially in the machine vision community (Bengio *et al.* 2013). We utilize a deep convolutional neural network, the AlexNet pre-trained by the ImageNet dataset (Krizhevsky *et al.* 2012), to extract image features with strong representation abilities. This is an example of *transfer learning*, that is, the use of knowledge from addressing a particular problem, to a new task (Pan and Yang 2009; Zhuang *et al.* 2020).

Instead of using the extracted image features to train a classifier directly, we propose three additional amendments on them. First, the AlexNet image features are high dimensional. To avoid potential problems induced by high dimensionality, principal component analysis (PCA) (Wold *et al.* 1987) is applied to reduce the dimensions of the extracted image features. Second, to further enhance the representation ability of the features, summary statistics are concatenated with image features to provide a more complete description of copula samples. Lastly, we aim to make these features more discriminative via linear discriminant analysis (LDA), which projects the concatenated features to a low-dimensional subspace where the observations from the same classes are grouped close together, while those from different classes are pushed apart (Yang and Jin 2006). Hence, the recognition task becomes easier on this discriminative subspace. The features extracted by LDA are used as the final representations of the copula samples to train the classifier. Support vector machine (SVM) is chosen as the classifier for the image recognition task, because it is proven to be effective on various real-world applications (Tzotsos and Argialas 2008; Islam *et al.* 2017; Sheykhmousa *et al.* 2020).

We test the performance of the proposed image recognition approach to copula selection via simulation studies. We consider the six copula models of

Figure 1 and evaluate the classification accuracy of our approach in different scenarios, comparing to the statistical benchmark given by AIC. First, we consider model selection when all training and testing instances arise for copula samples with the same sample size and underlying rank correlation. While this is not a realistic setting, it allows us to explore the performance of image recognition for different problem parameters. We observe that image recognition consistently outperforms AIC, except when the underlying rank correlation is very high. The biggest improvement occurs in those scenarios of low sample size and correlation, where the copula selection problem is the hardest.

Subsequently, we consider a more realistic scenario, where a copula model needs to be selected for data with differing sample sizes and correlations, generated under any rotation of the six bivariate baseline copula models. In this scenario, the heatmap images from the same copula model can present very different patterns, leading to substantial within-class variations. For that reason, we propose to add a first data rotation step, based on sample statistics, with the aim of converting the data to be positively correlated and skewed. Then, in a second step, we apply the image recognition approach to images generated from the rotated data.

We find that this two-step image recognition approach dominates AIC for the copula model selection task, again except in the situation of high correlations. This motivates our final proposal to combine the two-step image recognition approach with AIC. In this combined approach, AIC values (calculated on the rotated data) are concatenated with the image features and statistical features before applying LDA. Experiments show that the combined approach improves on both AIC and image recognition-based model selection.

Finally, we apply the copula selection framework developed in the paper to modelling joint weekly losses of the MSCI World Index and the ICMR (Re)Insurance Specialty Index. The dependence between those two indices is important as it reflects inter-industry diversification and thus (re)insurers' capital availability. We find that both the strength of correlation and the copula family selected vary over time, with the t copula being the most frequent choice.

The paper is organised as follows. In Section 2, we give preliminaries on copula modelling. Section 3 introduces the image recognition approach, with fixed correlation and sample sizes. In Section 4, we discuss the two-step approach to copula model selection for more general datasets. Experimental results are summarised within each section. The real-data application is presented in Section 5. Section 6 presents our concluding remarks.

2 COPULAS

2.1 Copula models and their properties

Consider continuous random variables X, Y with marginal distributions F, G and joint distribution H , on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The *copula* of (X, Y)

is a distribution on $[0, 1]^2$ with uniform marginals, such that

$$H(x, y) = C(F(x), G(y)). \quad (2.1)$$

Denote $U = F(X)$, $V = G(Y)$ and also $\bar{U} = 1 - U$, $\bar{V} = 1 - V$. Then it follows from (2.1) that C is the joint distribution of (U, V) , that is,

$$C(u, v) = \mathbb{P}(U \leq u, V \leq v), \quad (u, v) \in [0, 1]^2. \quad (2.2)$$

Analogously, the joint distribution of (\bar{U}, \bar{V}) is called the *survival copula* of (X, Y) and denoted by \bar{C} .

Definition (2.1) implies a separation of a random vector's marginal behaviour from its dependence structure, which has enabled copulas to be widely employed as multivariate modelling tools. For detailed treatments of copulas, including applications in insurance and financial risk management, see Nelsen (2007), Denuit *et al.* (2006), McNeil *et al.* (2015). We note that the copulas of discontinuous random vectors are not uniquely defined – in such a case the variables U, V as constructed above are not uniform. However, we can always uniquely determine a copula for X, Y via (2.2), with uniform variables U, V constructed via the generalised distributional transform of Rüschendorf and de Valk (1993).¹

In risk management, the specific properties of different copula families are important. Assume that X and Y represent losses, such that high (joint) outcomes are associated with adverse events. Then, beyond considering (rank) correlation measures, it is important to model the extent to which X and Y will jointly achieve high values. A typical way in which the literature considers the propensity of joint extremes is via the coefficients of upper and lower tail dependence (e.g. McNeil *et al.* 2015, Section 7.2.4):

$$\lambda_U(U, V) := \lim_{p \rightarrow 1} \mathbb{P}\left(Y > G^{-1}(p) \mid X > F^{-1}(p)\right) = \lim_{p \rightarrow 1} \frac{\bar{C}(1-p, 1-p)}{1-p},$$

$$\lambda_L(U, V) := \lim_{p \rightarrow 0} \mathbb{P}\left(Y \leq G^{-1}(p) \mid X \leq F^{-1}(p)\right) = \lim_{p \rightarrow 0} \frac{C(p, p)}{p}.$$

Models for which λ_U or λ_L are non-zero are, respectively, called upper or lower tail dependent.

While tail dependence is an asymptotic property, a distinct issue is the skewness or asymmetry of a copula. A copula is radially symmetric if $C(1-u, 1-v) = C(u, v)$. Various measures of bivariate skewness have been proposed by Rosco and Joe (2013). Here we focus on the moment-based measure ζ , defined for $k \in (1, \infty)$ as:

$$\zeta(U, V; k) = \mathbb{E}\left[|U + V - 1|^k \text{sign}(U + V - 1)\right].$$

Implications of the choice of the parameter k are briefly explored in Rosco and Joe (2013).

A further property of bivariate copulas relates to the extent that observations are concentrated in the four corners of the unit box, even when

correlation is low, leading to spider-like pattern. This property, which distinguishes, for example, a t from a Gaussian copula model, is termed *arachnitude* in Shaw *et al.* (2010), see also Androschuck *et al.* (2017), Genest *et al.* (2019). We measure arachnitude in the way proposed by Shaw *et al.* (2010), that is, as

$$\xi(U, V) := \rho \left((U - 0.5)^2, (V - 0.5)^2 \right),$$

where ρ is the Pearson (product-moment) correlation.

In this paper, we consider six bivariate copula models:

1. The *Gaussian copula* is probably the most widely used copula model. It is radially symmetric and tail independent.
2. The *t copula* is radially symmetric but both upper- and lower-tail dependent. It admits a degree of freedom parameter ν ; λ_L, λ_U decrease in ν , while for $\nu \rightarrow \infty$, the *t copula* reduces to a Gaussian.
3. The *Frank copula* is radially symmetric and tail independent.
4. The *Gumbel copula*, commonly used in risk management, is positively skewed and upper tail dependent.
5. The *Joe copula* is positively skewed and upper tail dependent.
6. The *Pareto copula* (or Clayton survival copula), is positively skewed and upper tail dependent.

We do not provide technical detail on these models, as they are all well known and exhaustively discussed in the literature (Denuit *et al.* 2006; Nelsen 2007; McNeil *et al.* 2015). The Gaussian and *t* models are, respectively, the copulas of bivariate Normal and *t* distributions; the remaining four models belong to the family of Archimedean copulas. All models, except the *t copula*, have a single parameter, which can be calibrated to (e.g. Kendall's) rank correlation or estimated by MLE. The *t* model has the degrees of freedom ν as an additional parameter.

The properties of different copula families are illustrated in Figure 1, which shows heatmaps of smoothed bivariate densities, each derived from samples of size $n = 2000$, with underlying Kendall rank correlation $\tau = 0.3$. While copulas are defined as distributions with uniform marginals, here and in the sequel, before producing such plots, we transform the marginals to standard normals, as we find that this transformation enables a better visual inspection of the dependence pattern. The distinct patterns of the different models are visible. At the same time, there is substantial similarity between some of the resulting heatmaps (e.g. Gaussian and *t*; Joe and Pareto), which indicates that selecting the correct model from data is not a trivial task. This is of course even more challenging for smaller datasets. We show heatmaps from the Gaussian and *t* families in Figure 2, comparing plots generated from bivariate samples of sizes $n = 100$ and 2000. It is apparent that with the smaller sample size the observed patterns become substantially noisier.

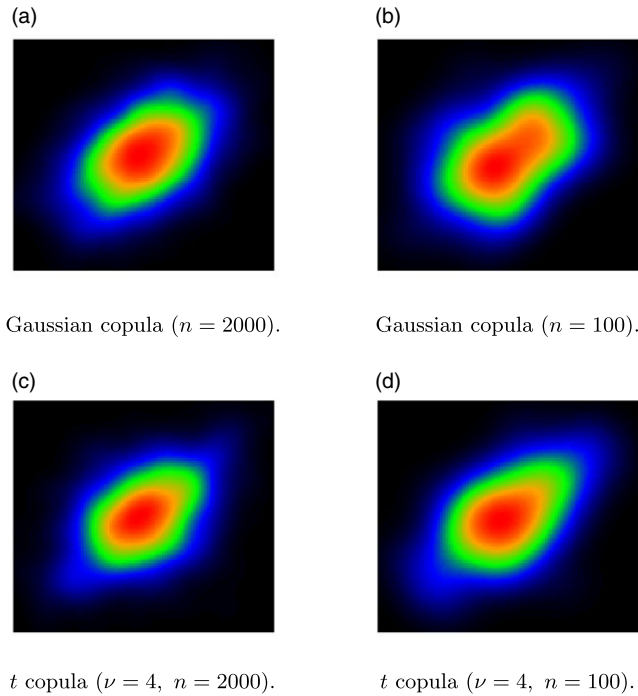


FIGURE 2. Comparison of heatmap images of smoothed bivariate densities for sample sizes $n = 100, 2000$ ($\tau = 0.3$).

2.2 Estimation and model selection

Consider a sample from $(X, Y), (x_1, y_1), \dots, (x_n, y_n)$. Realisations of the random vector (U, V) are not directly observable.² As a result, it is common to construct copula *pseudo-observations* (e.g. Genest *et al.* 2009). Let $r_i(\mathbf{z})$ be the rank of observation z_i in a univariate sample $\mathbf{z} = (z_1, \dots, z_n)$. Then, the pseudo-observations are given by

$$u_i = \frac{r_i(\mathbf{x})}{n+1}, \quad v_i = \frac{r_i(\mathbf{y})}{n+1}.$$

From the pseudo-observations $(u_i, v_i), i = 1, \dots, n$, we can readily estimate skewness and arachnitude, denoting the corresponding estimates by $\hat{\zeta}, \hat{\xi}$. Furthermore, we denote the sample version of Kendall's rank correlation as $\hat{\tau}$.

From the observations $(u_i, v_i), i = 1, \dots, n$, we can readily estimate skewness and arachnitude, denoting the corresponding estimates by $\hat{\zeta}, \hat{\xi}$. Furthermore, we denote the sample version of Kendall's rank correlation as $\hat{\tau}$.

For a parametric family of copulas $C^{(m)}(\cdot; \theta), \theta \in \Theta_m$, the parameters θ can be estimated by maximum likelihood estimation, treating the pseudo-observations as if they are a random sample from $C^{(m)}(\cdot; \theta)$. If we are

considering a family of copula models $\{C^{(m)}, m \in \mathcal{M}\}$, likelihood methods also offer model selection criteria. Let $c^{(m)}$ be the bivariate density corresponding to copula $C^{(m)}$ and $\hat{\theta}^{(m)}$ the (k_m -dimensional) estimate of the corresponding model parameter. The Akaike and Bayes Information criteria are given by

$$\text{AIC}^{(m)} = 2k_m - 2 \sum_{i=1}^n \log c^{(m)}(u_i, v_i; \hat{\theta}^{(m)}),$$

$$\text{BIC}^{(m)} = k_m \log n - 2 \sum_{i=1}^n \log c^{(m)}(u_i, v_i; \hat{\theta}^{(m)}).$$

The selected model is then the one with the lowest $\text{AIC}^{(m)}$ or $\text{BIC}^{(m)}$. A cross-validated log-likelihood criterion is formulated by (Grønneberg and Hjort 2014, Equation (42)); see Jordanger and Tjøstheim (2014) for a simulation study. Other model selection criteria can be constructed using goodness-of-fit statistics; for example Kularatne *et al.* (2021) employ the Cramer-von Mises statistic, the use of which (including variations) in copula goodness-of-fit testing has been thoroughly explored by Genest *et al.* (2009). Bayesian copula selection is discussed in Huard *et al.* (2006).

In this paper, we use as a statistical benchmark for model selection the AIC.³

3 AN IMAGE RECOGNITION-BASED APPROACH TO COPULA MODEL SELECTION

In this section, we introduce a new methodology, which uses image recognition to select a suitable bivariate copula from a sample, by classifying its density heatmap image to one of the six models we consider. The recognition process is designed to incorporate rich information that can well represent the samples and is discriminative to make the classification task easier. The generation of the heatmap images is introduced first, and then the image recognition approach is discussed in detail. Subsequently, experimental results are shown to demonstrate the effectiveness of this approach for copula model selection.

We note that in the present section we apply the classification/copula model selection framework to simulated data with very benign features, with all images in any given dataset derived from samples with the same sample size and underlying correlation. This is of course an unrealistic testing environment, with classes that are more homogeneous than in any practical application. Nonetheless, the setting of this section allows us to evaluate whether image recognition can be effective as a copula selection tool (and under which conditions). The restrictive assumptions of this section are relaxed in the two-step approach of Section 4.

3.1 Generating heatmap images of smoothed bivariate densities

Here we outline how the image datasets are generated, on which classifiers are trained to perform the copula selection task. Each image is a smoothed bivariate density heatmap, generated from a simulated pseudo-sample from (U, V) , drawn from a given copula specification $C^{(m)}(\cdot; \theta)$, $\theta \in \Theta^{(m)}$, $m \in \mathcal{M}$.

For each image dataset that we generate the following hold:

- (a) The dataset contains $R = 20,000$ images.
- (b) Each image is derived from a bivariate sample of size n , drawn from one of the 6 copula families we consider in this paper, $\mathcal{M} = \mathcal{M}_s \cup \mathcal{M}_a$, where $\mathcal{M}_s = \{\text{Gaussian}, t, \text{Frank}\}$ and $\mathcal{M}_a = \{\text{Gumbel}, \text{Joe}, \text{Pareto}\}$ contain radially symmetric and asymmetric (positively skewed) models respectively. Each dataset contains approximately the same number of images from each copula family.
- (c) In each dataset, all images are generated from simulated samples with a fixed sample size $n \in \{100, 150, 200, 250\}$ and (population) Kendall $\tau \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$. Hence we have $4 \times 5 = 20$ datasets, corresponding to different (n, τ) combinations, each containing R images.
- (d) All images are generated from bivariate samples that have positive empirical rank correlation and positive empirical skewness. This means that samples that display negative rank correlation are rejected and not used to produce images in the dataset. Furthermore, samples that display negative sample skewness, when the underlying copula model has positive skewness, are also rejected.⁴
- (e) Images are generated based on pseudo-observations from the copula samples. Furthermore, to generate the heatmap images, we transform pseudo-observations to have a standard normal marginal distribution. This transformation is employed only for image generating purposes and reflects no assumption of normality for the marginal distribution of the underlying data.

The precise process by which images are generated is given in the Supplementary Material (Appendix C, Algorithm 1). All calculations are carried out in **R**. For random number generation we use package *copula* (Yan 2007; Hofert *et al.* 2020). For AIC calculations, we use the package *VineCopula* (Schepsmeier *et al.* 2021).

Joint densities are estimated on a 100×100 grid on $[-3, 3]^2$ using an axis-aligned bivariate normal kernel, by the function `kde2d` of the package *MASS*, with bandwidths set to 1.3 times the values given by the heuristic in (Venables and Ripley 2002, Equation (5.5)). Specifically, for a bivariate sample (x_j, y_j) , $j = 1, \dots, n$, the density estimate at an arbitrary point (x, y) is

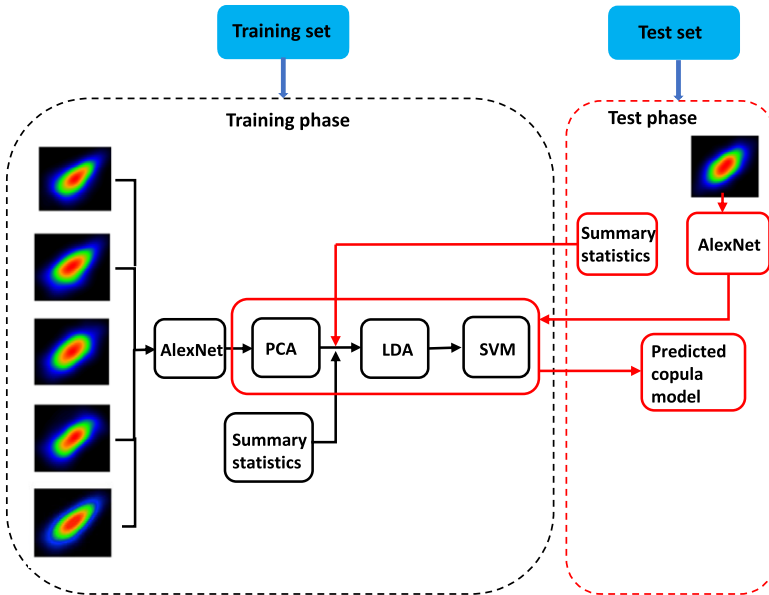


FIGURE 3. The workflow of the image recognition approach for copula model selection.

given by

$$f(x, y) = \sum_{j=1}^n \frac{\varphi((x - x_j)/h_x)\varphi((y - y_j)/h_y)}{h_x h_y},$$

where φ is the standard normal density and h_x, h_y are the bandwidths used. Subsequently, the values of the bivariate density estimate f on the grid are used to generate a heatmap, via the `image` function. The heatmap consists of a 100×100 grid of coloured rectangles with colours corresponding to the values of the density estimate. The heatmap is then saved as a png image file.

3.2 The image recognition approach

Now we present the image recognition approach with the complete workflow shown in Figure 3. Similarly to all classification tasks, the image recognition approach consists of a training phase to extract features and train a classifier, and a test phase to predict a copula model for a test sample. In Figure 3, the training phase is presented by the black flow while the test phase is presented by the red flow.




FIGURE 4. The architecture of the AlexNet.

3.2.1 The training phase

The training phase contains two parts: (1) the representation learning part to extract features with strong representation and discrimination abilities and (2) the classification part to train an effective classifier.

Representation learning. In the representation learning part, two types of features are extracted from the training copula samples: pure image features and statistical features. The pure image features are extracted from the training heatmap images by the pretrained AlexNet that has been trained on the ImageNet dataset with 1000 classes of high-resolution images (Deng *et al.* 2009). Given the complex nature of the images in the ImageNet dataset and the competitive classification performance of AlexNet, we believe that this pretrained network can provide good representations of our heatmap images of relatively simple patterns.

AlexNet is a deep convolutional neural network, consisting of layers with three-dimensional volumes of neurons. The architecture of AlexNet is depicted in Figure 4, with an input layer of RGB images, five convolutional layers and three fully connected layers. The input layer is of dimension $227 \times 227 \times 3$: 227 and 227 are the width and height of the layer, which represent the width and height of the image, while 3 is the depth of the layer, representing the red, green and blue channels of the RGB image. The convolutional layer contains neurons that are only connected to a small local region in the previous layer, each computing a dot product between the filters and the region they are connected to, that is the convolution between the filters and input regions. Let us take the first convolutional layer as an example. The filters of this layer are of dimension $11 \times 11 \times 3$ and each of them captures different features from the input image, for example edge or colour. By sliding each filter across the width and height of the input layer and applying the convolution operation, we can obtain a two-dimensional activation map that contains the responses of the filter at every spatial position. Specifically, we reshape the $11 \times 11 \times 3$ volume of the input and filter to two $11 \times 11 \times 3 = 363$ dimensional column vectors. The dot product of the two vectors is fed to the RELU activation function, $\max\{0, x\}$, to capture the nonlinearity of features. The output from the activation function is called activation and fills one entry in the activation map. This process leads to a volume of $55 \times 55 \times 96$ neurons in the first convolutional




FIGURE 5. Left: The heatmap image of a Gaussian copula sample. Right: The activations of the first convolutional layer of AlexNet for the sample.

layer, with 96 filters. The max pooling layers between the convolutional layers can effectively reduce the number of parameters and control overfitting.

The outputs of the convolutional layers are then fed to fully connected layers to provide the predicted distribution of the class labels, that is the predicted probabilities for each of the 1000 classes considered. Note that we do not use the classification output of AlexNet, as the 1000 classes used are not relevant to our copula selection task; however, the representation ability of the network allows us to use it for feature extraction from our heatmap images. To demonstrate the good representation ability of the pretrained AlexNet, we show the activation maps of a Gaussian copula sample for all 96 filters of the first convolutional layer in Figure 5, with each square presenting the activation map of one filter. The bright pixels reflect high activations, which means that they make substantial contributions to the extracted features. It is obvious that the contour shapes of the Gaussian sample can be well captured by the first convolutional layer.

We input the training heatmap images to the pretrained AlexNet, and extract 4096 features from the second fully connected layer, ‘fc2’ in Figure 4, which is the last feature extraction layer before working out predicted probabilities. This layer provides high-level abstract features that can well represent the images. To be precise, each training heatmap image is represented by a 4096-dimensional vector $\mathbf{x}_i^M \in \mathbb{R}^{4096 \times 1}$ ($i = 1, 2, \dots, N$), where N is the number of training copula samples, and the image features of the whole training set is denoted as $\mathbf{X}^M = [\mathbf{x}_1^M, \mathbf{x}_2^M, \dots, \mathbf{x}_N^M]^T \in \mathbb{R}^{N \times 4096}$. (Note that N here is different to the size of the whole data set R , as the generated images for each dataset will be subsequently split into training and test samples.)

The high number of image features can lead to potential problems in classification, for example the curse of dimensionality and high computational cost. Thus, we reduce the number of image features before training a classifier. This dimension reduction step is achieved by principal component analysis (PCA), which is a widely adopted unsupervised dimension reduction method that can provide low-dimensional yet effective representations of the original high-dimensional data (Wold *et al.* 1987). PCA projects data from the original

feature space to a low-dimensional subspace, spanned by the first few principal components (PCs) that can explain most of the variation in data. This can be achieved by applying the reduced singular value decomposition (SVD) on the column-centred \mathbf{X}^M . Let \mathbf{V} denote the matrix whose columns are the PCs sorted by the singular values in a descending order. In this paper, we set $q = 150$ to explain 99.9% of the variation in \mathbf{X}^M . After PCA, the image features of the training set become

$$\mathbf{X}^{MP} = (\mathbf{X}^M)^c \mathbf{V}_{150} \in \mathbb{R}^{N \times 150}, \quad (3.1)$$

where $(\mathbf{X}^M)^c$ is derived by subtracting the column means of \mathbf{X}^M , $\mathbf{V}_{150} \in \mathbb{R}^{4096 \times 150}$ is \mathbf{V} with the first 150 columns. Thus, the image features now lie in a 150-dimensional PC subspace.

Besides the pure image features extracted from AlexNet, we propose to utilise some additional statistical features to enrich the description of the training copula samples. Three summary statistics, Kendall's rank correlation, skewness and arachnitude are chosen as the statistical features. The statistical features of each sample are denoted as $\mathbf{x}_i^S = (\hat{\tau}_i, \hat{\zeta}_i, \hat{\xi}_i) \in \mathbb{R}^{3 \times 1}$, $i = 1, 2, \dots, N$.

The low-dimensional image features and three statistical features are then concatenated to provide a representation for the i th copula sample, $\mathbf{x}_i = [(\mathbf{x}_i^{MP})^T, (\mathbf{x}_i^S)^T]^T \in \mathbb{R}^{153 \times 1}$, where $\mathbf{x}_i^{MP} \in \mathbb{R}^{150 \times 1}$ is the i th observation in \mathbf{X}^{MP} . The feature matrix to represent all training copula samples is denoted as $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N]^T \in \mathbb{R}^{N \times 153}$. Before simply feeding this feature matrix to a classifier, we extract more compact and discriminative information to reflect the differences between classes better and make the classification process easier. For that purpose, we apply LDA on \mathbf{X} . LDA is a well known supervised dimension reduction method that projects data to a subspace such that between-class variation is maximised while within-class variation is minimised (Yang and Jin 2006). The classification task is easier on this subspace because the instances from the same class are pulled close together while those from different classes are pushed further away. By projecting \mathbf{X} on the linear discriminant subspace, we have

$$\mathbf{X}^P = \mathbf{X}\mathbf{W} \in \mathbb{R}^{N \times (K-1)}, \quad (3.2)$$

where $\mathbf{W} \in \mathbb{R}^{153 \times (K-1)}$ contains the bases of the linear discriminant subspace and K is the number of classes. LDA can provide at most a $(K-1)$ -dimensional subspace. In this paper, six copula models are considered, thus the LDA subspace is at most five-dimensional. Here, we take all five discriminative dimensions provided by LDA.

Classification. Support vector machine (SVM) is chosen as the classifier for its efficiency in many real-world applications (Tzotsos and Argialas 2008; Islam *et al.* 2017; Sheykhmousa *et al.* 2020). The training set to train SVM contains N pairs of observations $\{\mathbf{x}_i^P, m_i\}_{i=1}^N$, where $\mathbf{x}_i^P \in \mathbb{R}^{5 \times 1}$ is the feature vector of the i th observation obtained in the representation learning part and $m_i \in \mathcal{M}$ is the corresponding copula model.

SVM aims to find a separating hyperplane $f(\mathbf{x}) = \phi(\mathbf{x}_i^P)^T \mathbf{w} + b$ for classification by maximising the margin between two classes. The sign of the estimated value of $f(x)$ is used to distinguish the positive and negative classes. Thus, the standard SVM can only be used for binary classification. To apply it in our case with six classes, we adopt the one-versus-one strategy (Hastie *et al.* 2009). We apply SVM to pairs of classes, which means that $\binom{K}{2}$ SVM classifiers are trained. For the test observation \mathbf{x} , we obtain $\binom{K}{2}$ classification results. A majority vote is then applied to these classification results to determine the class of \mathbf{x} , that is the class with the highest vote is selected.

More detail on the calculations relating to PCA, LDA, and SVM is given in the Supplementary Material (Appendix A).

3.2.2 The test phase

In the test phase, given one observed copula sample, we first extract its image features $\mathbf{x}_i^M \in \mathbb{R}^{4096 \times 1}$ from the pretrained AlexNet and project them to the PC subspace constructed in the training phase:

$$\mathbf{x}_i^{MP} = \mathbf{V}_{150}^T (\mathbf{x}_i^M)^c \in \mathbb{R}^{150 \times 1}, \quad (3.3)$$

where $(\mathbf{x}_i^M)^c$ is the centred \mathbf{x}_i^M by the column means of \mathbf{X}^M . The image features are then concatenated with the summary statistics $\mathbf{x}_i^S \in \mathbb{R}^{3 \times 1}$ to form the feature vector of the test copula sample, $\mathbf{x}_i = [(\mathbf{x}_i^M)^T, (\mathbf{x}_i^S)^T]^T \in \mathbb{R}^{153 \times 1}$. We then project this vector to the linear discriminant subspace to obtain the discriminative features for classification:

$$\mathbf{x}_i^P = \mathbf{W}^T \mathbf{x}_i \in \mathbb{R}^{5 \times 1}. \quad (3.4)$$

The copula model is then selected by applying the trained SVM on \mathbf{x}_i^P .

3.3 Classification results on copula samples with fixed n and τ

3.3.1 Experimental settings

For each dataset with fixed n and τ , we randomly select 70% of the $R = 20,000$ images to form the training set, with the remaining images used as a test set; hence the training sample size is $N = 0.7 \times R = 14,000$. For SVM, the radial basis function (RBF) kernel is chosen as the kernel function. The hyperparameters associated with the SVM classifier are tuned by 10-fold cross-validation on the training set. The classification accuracies of the image recognition approach are recorded. Furthermore, we record the accuracy by which the copula model is selected when using AIC as a criterion. To make the results more reliable, the training/test random split process is repeated 20 times. Thus for each combination of n and τ , we record 20 classification accuracies for each copula model selection method.

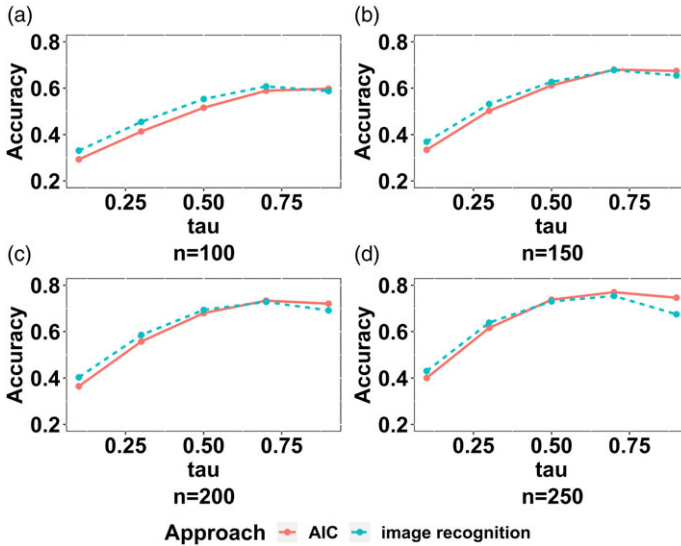


FIGURE 6. Classification accuracies for fixed n and τ . The blue dashed curves represent the accuracies of the image recognition approach while the red solid curves represent those of AIC.

3.3.2 Classification results

The classification results are shown in Figure 6, with each plot presenting the accuracies of the two approaches with a given value of n and all values of τ . For each plot, the horizontal axis represents the values of τ , while the vertical axis represents the classification accuracy. The mean classification accuracies of the image recognition approach are shown by blue dashed curves while those of AIC are shown by red solid curves.

It is clear that, generally, the classification accuracy increases as τ increases for each value of n , and also increases as n increases for each value of τ . This pattern makes sense because the underlying model can be better described by samples with larger n 's. The images generated with larger τ 's are generally easier to classify, as the characteristics of the models are presented clearer when τ is larger. The image recognition approach can beat AIC when τ is not high, that is less than 0.7, and this is more obvious when n is small. This is encouraging, as the improvement provided by the image recognition approach to select copula models occurs for samples that are difficult to classify, that is those with small n and τ . However, when τ is large, the image recognition approach performs apparently worse than AIC; moreover, beyond $\tau = 0.7$, the classification accuracies of the image recognition approach start decreasing. A plausible reason for this is that, when τ is very large, the heatmap images of different copula models exhibit very similar visual patterns and image recognition cannot easily distinguish between them.

To sum up, the preliminary experimental results on copula samples with fixed n and τ demonstrate the potential effectiveness of image recognition to select copula models, especially when n and τ are relatively small.

Finally, we note that the design of the copula selection process proposed in this section entails some important decisions: the marginal distributions used for plotting heatmaps and the dimensions of the PC and LD subspaces. In the Supplementary Material (Appendix B.1), we provide robustness checks to evaluate the extent to which classification performance is impacted by such decisions.

4 SELECTING ROTATED COPULA MODELS

Encouraged by the results of Section 3, we here address the more realistic scenario where n and τ are not fixed. Furthermore, we allow τ to be negative. In particular, for asymmetric copula models, we consider all rotations, such that each model in \mathcal{M}_a has now four distinct versions. We believe that this addresses a realistic modelling scenario, as the classification approach does not assume anything a priori about the sign of either the correlation or the skewness of the underlying model.

In this section, we consider three distinct approaches:

1. **Copula selection with AIC.** This approach requires us to consider an enlarged set of candidate models, which also includes the rotations of the positively correlated and skewed models in \mathcal{M} .
2. **Image recognition with a two-step approach.** In the first step, sample statistics are used to assess the sign of correlation and skewness – essentially trying to detect if elements of \mathcal{M} have been rotated. Then, the pseudo-samples are transformed to have positive correlation and skewness. In the second step, heatmap images are generated from the transformed samples and classified to models in \mathcal{M} .
3. **Combining image recognition with AIC.** The same approach as in 2. is followed, with the AIC of the models in \mathcal{M} added as a feature in the second step of the process. The motivation for this is to not miss out on any information encoded in likelihood-based statistical criteria, which is not present in image features.

The performance of the three copula model selection approaches is assessed on a test set. In contrast to Section 3, for the two-step approaches of this section we cannot validate the result of the classification algorithm on subsets of the training set. In the training set, all copula samples are positively correlated and skewed, to make the within-class variation smaller. However, this is not a realistic testing scenario. Thus, in the test set, we consider copula samples that may have negative or positive correlation and skewness. Before we discuss each

of the three approaches in more depth, we describe the construction of this test set.

4.1 Test set

We produce two test sets of $S = 10,000$ bivariate copula pseudo-samples: one with $n \in [25, 100]$ and the other with $n \in [100, 250]$, to test the classification performance of the copula selection methods on different ranges of sample sizes. These sample size ranges were chosen based on practical considerations. For sample sizes that are large, criteria such as AIC will lead to optimal model choices, which means that there is less scope for improvement from adding image features, as already glanced from Figure 6. If on the other hand sample sizes are very small, the model selection errors will be inescapably high, whichever method is used, and in such a situation a professional would typically prefer to use expert judgement. This leaves us with a ‘small-to-modest’ sample size range, for which improvements to standard model selection methods can be most meaningfully proposed. We consider that sample sizes from 25 to 250 give a suitable such range.

Each pseudo-sample is generated from a copula in \mathcal{M} , with randomly (uniformly) chosen from the corresponding range of n and $\tau \in [0.1, 0.9]$. Before simulating each sample, the underlying copula model is rotated by 0, 90, 180, or 270 degrees counter-clockwise. Hence, we deal with an enlarged model set, denoted by \mathcal{M}' . Let m_r represent a model m in \mathcal{M} , rotated by r degrees, such that $m_0 = m$. For radially symmetric models $m \in \mathcal{M}_s$, we also have $m_{180} = m_0$, $m_{270} = m_{90}$. Then, we let $\mathcal{M}' = \mathcal{M}'_s \cup \mathcal{M}'_a$, where $\mathcal{M}'_s = \{m_r : m \in \mathcal{M}_s, r \in \{0, 90\}\}$ and $\mathcal{M}'_a = \{m_r : m \in \mathcal{M}_a, r \in \{0, 90, 180, 270\}\}$. The data and copula model specification are saved, as well as the rank correlation of the rotated model τ_r .

The process of generating the test set is given in detail in the Supplementary Material (Appendix C, Algorithm 2).

4.2 Copula selection approaches

4.2.1 Copula selection with AIC

For each instance i in the test set, the model in \mathcal{M}' with the smallest AIC is chosen. The classification is successful if the chosen model is identical to the underlying model $m_r \in \mathcal{M}'$.

4.2.2 Image recognition with a two-step approach

When we allow for models with negative correlation and skewness, the classification task becomes harder. One can either assign a different class for each element in the enlarged model space \mathcal{M}' – thus moving from 6 to 18 classes – or within each of the 6 classes in \mathcal{M} accommodate model rotations – resulting in non-homogeneous classes. We address this challenge pragmatically, by

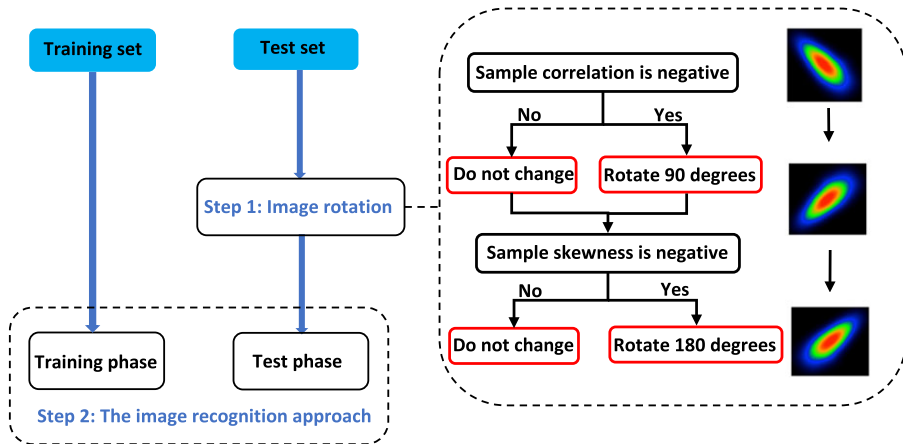


FIGURE 7. The workflow of the two-step approach for copula model selection.

using the sample statistics $\hat{\tau}$, $\hat{\zeta}$ to infer the rotation of the underlying model. Thus, as a first step pseudo-observations are transformed to have positive correlation and skewness. Subsequently, heatmap images are created from the transformed samples. As a second step, these heatmaps are classified, to select one of the 6 copula models in $m \in \mathcal{M}$. Thus in the image recognition stage we avoid the need to consider too many or very heterogeneous classes.

Figure 7 depicts the workflow of the two-step image recognition approach. The two steps are discussed in more detail below.

Step 1 In order to arrive at homogeneous image classes, corresponding to the first step discussed above, samples are pre-processed as follows (the detailed algorithm is given in the Supplementary Material (Appendix C, Algorithm 3)). First we check whether the rank correlation is negative – if so the data are rotated by 90 degrees counter-clockwise. Then skewness is checked – if negative, the data are rotated by another 180 degrees. Let the quantity s represent the degree of data rotation to achieve a positive correlation and skewness; hence $\hat{r} = 360 - s$ is an estimate of the rotation r under which the pseudo-observations (u_j, v_j) were simulated. Sample statistics (including AIC values only used in the combined approach of Section 4.2.3) and images generated from the transformed sample and subsequently exported. The process is illustrated on the right of Figure 7, which shows an example of the heatmap image of a 90-degree rotated Pareto copula.

Finally, we assess whether the first step has been performed successfully, in that the rotation applied to the data before producing the heatmap is consistent with the rotation of the underlying copula model. Radially symmetric and asymmetric models are treated differently, to reflect that for symmetric models $m_0 = m_{180}$, $m_{90} = m_{270}$, which in practice requires us only to test whether the sign of the rank correlation τ_r of the copula model m_r matches that of the pseudo-observations, $\hat{\tau}$.

Step 2 In the second step, images generated by Step 1 are classified to models in \mathcal{M} . The classification approach proceeds analogously with what was discussed in Section 3.2. The only difference is that, when generating a training sample of $R = 20,000$ for each test set rather than using fixed n and τ as in Section 3.1, for each i these are now randomly chosen in the corresponding range of $n \in [25, 100]$ or $n \in [100, 250]$ and $\tau \in [0.1, 0.9]$. That is, predictions on each test set are performed based on the model trained on its corresponding training sample.

Once again we randomly split the $R = 20,000$ to a training set containing with 70% of the samples and a validation set with 30% of the samples. A SVM with RBF kernel is trained based on training sets of size $0.7R = 14,000$. The training/validation split process is repeated 20 times.

Finally, the 2-step approach is applied to classify the samples in the test set. For each test sample, we obtain 20 classification results based on the 20 classifiers trained in the training phase. A majority vote is applied to these results to get the final decision. In other words, the copula model with the highest vote by the 20 results is selected for the test copula sample.

We count a test sample as correctly classified only if both steps in the classification process are successful. In other words, for a sample to be classified correctly we need both to be true:

1. In the first step, the data were rotated in a way consistent with the underlying copula model.⁵
2. In the second step, the classifier identifies the correct copula model out of the six models in \mathcal{M} . In other words, if the model underlying a test set was $m_r \in \mathcal{M}'$, the prediction of the classifier is $m \in \mathcal{M}$.

Given the concurrent use of image-based and statistical features, it is of interest to consider which of those features drive the predictions of each model. A related sensitivity analysis, following the ideas of Pesenti *et al.* (2019), is documented in the Supplementary Material (Appendix B.3).

4.2.3 Combining image recognition with AIC

The analysis of Section 3 has shown that image recognition and statistical approaches may be complementary, each being dominant in different ranges of n and τ . For that reason we propose to combine the two approaches. We adapt the approach of Section 4.2.2, to integrate AIC information for different models in \mathcal{M} in the second step of the process. Specifically:

1. We follow the Step 1 of Section 4.2.2 in exactly the same way.
2. In Step 2, we follow again the approach of 4.2.2, but now adding the AIC values for $l \in \mathcal{M}$ in both the training and testing phases. Specifically, for a training instance i denote the AIC values of

TABLE 1.
TEST CLASSIFICATION ACCURACIES OF AIC, TWO-STEP IMAGE RECOGNITION
APPROACH, AND COMBINING IMAGE RECOGNITION WITH AIC.

	AIC	Image recognition	Combined
$n \in [25, 100]$	0.3557	0.4007	0.4121
$n \in [100, 250]$	0.5688	0.5822	0.6061

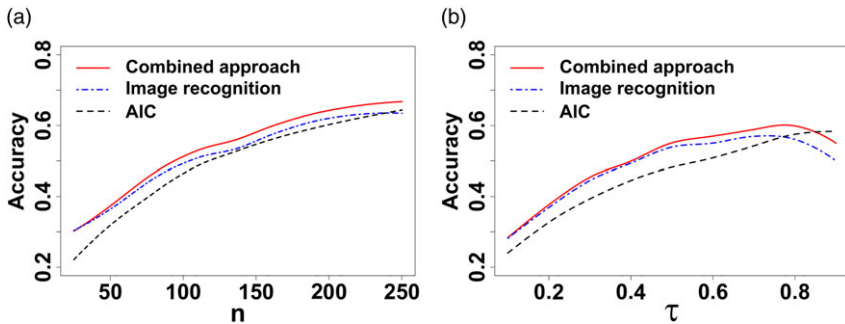


FIGURE 8. Fitted curves of the aggregated classification accuracies of the two test sets for the three approaches against (a) n and (b) τ .

different models by $\mathbf{x}_i^A = \{\text{AIC}_i^{(l)}\}_{l \in \mathcal{M}}$, $i = 1, 2, \dots, N$. The concatenated feature vector for each copula sample then becomes $\mathbf{x}_i = [(\mathbf{x}_i^{MP})^T, (\mathbf{x}_i^S)^T, (\mathbf{x}_i^A)^T]^T \in \mathbb{R}^{159 \times 1}$. The rest of the training phase follows exactly as described in the second step of Section 4.2.2. Similarly, a test sample is represented by $\mathbf{x}_t = [(\mathbf{x}_t^M)^T, (\mathbf{x}_t^S)^T, (\mathbf{x}_t^A)^T]^T \in \mathbb{R}^{159 \times 1}$, where \mathbf{x}_t^A are the AIC values calculated for elements of \mathcal{M} , from the (rotated) data of the test set. \mathbf{x}_t is then projected to the linear discriminant subspace constructed in the training phase and classified by SVM.

4.2.4 Classification results

First we compare the performance of the different copula model selection methods presented in Section 4.2 by calculating their classification accuracies on the two test sets. The classification accuracies of the three approaches show similar patterns for both ranges of n . It is seen in Table 1 that the image recognition approach of Section 4.2.2 outperforms AIC. Furthermore, combining image recognition with AIC information, as in Section 4.2.3, leads to a better accuracy than either of those two approaches.

To gain more insight into the test results, we plot smoothed curves of the average test classification accuracies of the three approaches against n and τ in Figure 8(a) and (b), respectively. Note that here we aggregate the results of both test sets to produce the two plots. The curves are generated by the

TABLE 2.

CONFUSION MATRIX OF THE APPROACH COMBINING IMAGE RECOGNITION AND AIC ON THE TWO TEST SETS.

Predicted \ True	Gaussian	t	Frank	Gumbel	Joe	Pareto
Gaussian	2023	840	453	476	110	163
t	449	1604	200	464	120	96
Frank	503	322	2461	221	122	170
Gumbel	271	344	115	1250	286	314
Joe	58	96	54	276	1785	1405
Pareto	88	60	61	293	857	1059
Accuracy	0.5964	0.4911	0.7359	0.4195	0.5442	0.3302

TABLE 3.

CONFUSION MATRIX OF THE TWO-STEP IMAGE RECOGNITION APPROACH ON THE TWO TEST SETS.

Predicted \ True	Gaussian	t	Frank	Gumbel	Joe	Pareto
Gaussian	1905	928	474	513	117	185
t	517	1498	186	441	118	87
Frank	551	338	2440	230	133	161
Gumbel	273	355	121	1215	275	300
Joe	63	93	59	296	1815	1518
Pareto	83	54	64	285	822	956
Accuracy	0.5616	0.4587	0.7297	0.4077	0.5534	0.2981

function `locfit` of the package `locfit` with the smoothing parameter of 0.5. It is clear that the image recognition approach outperforms AIC, except for high correlations. Furthermore, the combined approach, including the information from both images and AIC, provides better results than simply using the image information in the two-step approach for all values of n and τ . Figure 8(a) shows that the combined approach is the best over all values of n while AIC is the worst, which is consistent with our conclusion in Table 1. However, Figure 8(b) presents a slightly different pattern for τ : AIC performs better than both the combined and two-step approaches when the value of τ is larger than 0.9. This indicates that the heatmap image information does not help copula selection when τ is very high.

In Tables 2 and 3, respectively, we present the confusion matrices of the image recognition and combined approaches by aggregating the results of the two test sets. The classification accuracies for each class are summarised in the bottoms of each table. For simplicity of exposition, we calculate the confusion matrices only for those testing instances where we had $FirstStep = TRUE$ (99.02% of instances). It is apparent that the Frank, Gaussian and Joe models are best predicted. Predictions for underlying Gumbel and t models are less accurate, while for Pareto models the predictions are the worst. Pareto models are confused with Joe by both approaches, due to the similarity between their

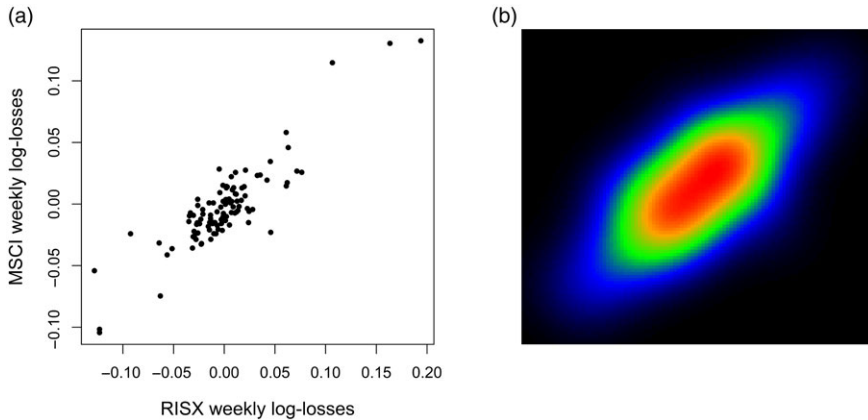


FIGURE 9. (a) Scatter plot of weekly log-losses for the MSCI (ticker: 'NDDUWI') and RISX (ticker: 'RISXNTR') indices (10/01/2020 – 31/12/2021); (b) smoothed bivariate density heatmap (normal marginals).

heatmap images. Comparing Tables 2 and 3, we can see that the combined approach can provide better predictions across all models (with only a small drop in accuracy for Joe), which demonstrates the effectiveness of involving the information provided by AIC.

Based on this analysis, we find that image recognition offers improvements in the accuracy of bivariate copula model selection, when compared to using the AIC benchmark. Furthermore, combining those two is best and we consider this combined approach our preferred model. Further supporting evidence is given in the Supplementary Material (Appendix B.2), where the impact of using purely statistical or image features is explored.

5 REAL DATA APPLICATION

Here, we apply the copula selection framework developed in the paper to modelling (unconditional) joint weekly net total returns – that is, including dividends reinvested net of withholding taxes – of the MSCI World Index (ticker: 'NDDUWI') and the ICMR (Re)Insurance Specialty Index (ticker: 'RISXNTR') (both in USD). The MSCI World Index is a market capitalisation weighted index representing equity performance across developed markets (MSCI 2022). RISX is an equity-based benchmark for the global specialty (re)insurance sector, based on publicly listed companies that underwrite risks in the Lloyd's of London market (Insurance Capital Markets Research 2022). Modelling the dependence between those two indices reflects diversification considerations that are important for investment decisions and for insurance operations. In particular, a lower level of correlation between global equity markets and the (re)insurance industry makes the latter attractive for investors,

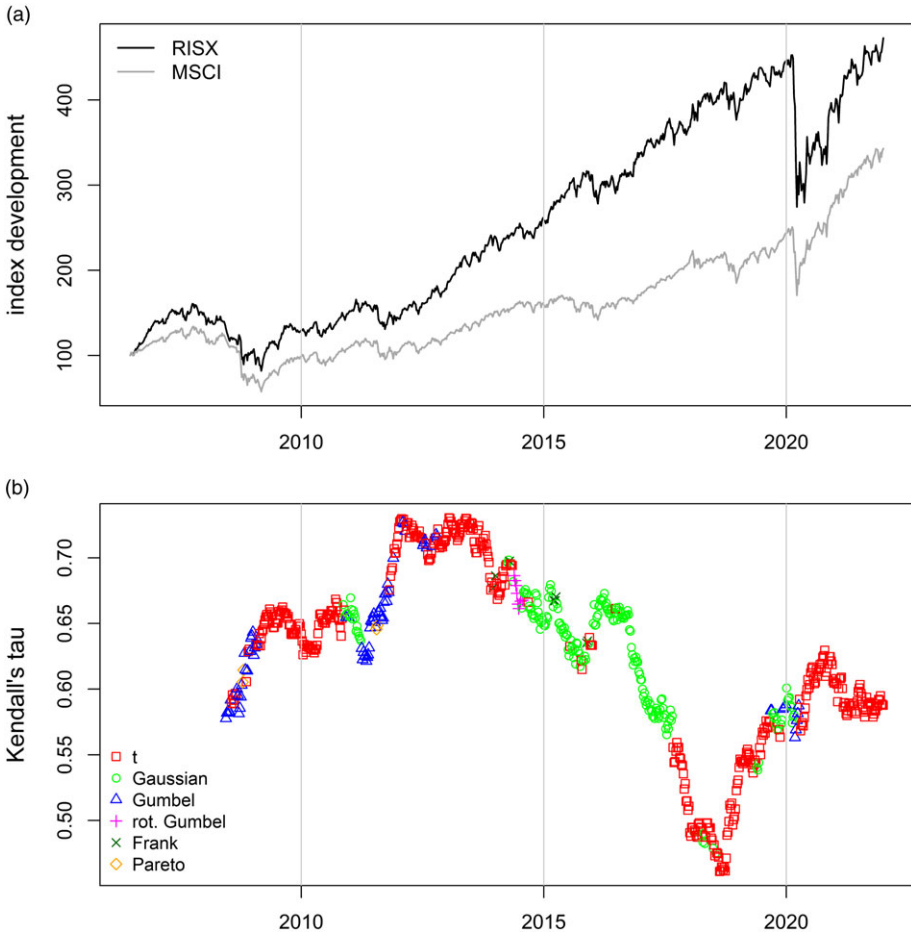


FIGURE 10. (a) Development of the MSCI (ticker: 'NDDUWI') and RISX (ticker: 'RISXNTR') indices (06/2006 – 12/2021); (b) rank correlation between weekly log-losses, with a 2-year rolling window. The colour indicates the copula model selected.

which has a direct impact on insurers' capital availability and, by implication, competition and commercial premium rates.

We consider the dependence structure of the weekly log-losses (negative log-returns) of the two indices. First, we select a copula model for the 104 joint observations from 10/01/2020 – 31/12/2021, hence covering an approximate 2-year period. In Figure 9(a), we show the scatter plot of weekly log-losses from the two indices; in Figure 9(b), as in previous sections, we represent the dependence structure by a smoothed density heatmap, after transforming the marginals to normal. The substantial correlation between the returns ($\hat{\tau} = 0.5878$) is apparent. Furthermore, we have sample skewness and arachnitude of $\hat{\zeta} = 0.0007$ (indicating near symmetry) and $\hat{\xi} = 0.6271$, respectively. Using




FIGURE 11. Predicted probabilities for different types of copula model; darker colours indicate models with more (right) tail risk.

those statistics and Figure 9(b), we select a copula model using the combined approach of Section 4.2.3. A t model is chosen, with fitted degree-of-freedom parameter $\hat{\nu} = 5.064$.

We now examine changes in the dependence structure of weekly log-losses over a longer time period. We use data from 16/06/2006 – 31/12/2021 (811 observations), hence reflecting the impacts of both the 2007–2008 global financial crisis and the 2020(-present) COVID-19 pandemic. First, in Figure 10(a), we plot the development of the two indices over time, normalised to have the same starting point – the joint drop in early 2020 and subsequent recovery are clearly visible. In Figure 10(b), we show the path of Kendall's τ for the weekly log-losses, calculated with a 104-week (2-year) rolling window. It is seen that correlation has changed over time, with a more or less continuous drop during 2011–2019, and a subsequent increase (one cannot attribute this increase to COVID-19, as it had started already before 2020). On the same plot, we also indicate the chosen copula family using the combined approach of Section 4.2.3. It is seen that there are some changes in the selected model over time. The most frequently selected model is the t copula (red squares). Early in the series, the Gumbel copula (blue triangles) is also commonly selected. Furthermore, the Gaussian copula (green circles) seems to often be chosen during periods when the correlation value is on a decreasing trajectory. Other models appear less frequently; it is notable that occasionally the (180-degree-)rotated Gumbel copula (magenta crosses) is chosen, reflecting a benign scenario where the dependence of log-losses is negatively skewed. Overall, this analysis provides

strong evidence for non-stationary of the dependence structure (and not just the correlation) of the two indices.

The SVM-based classifiers allow us to state not just the chosen model, but also predicted probabilities for each model considered. We visualise these predicted probabilities in the stacked area plot of Figure 11. To aid interpretability, we now group the selected copula models in four categories, from the most adverse (darkest) to the least adverse (lightest): positively skewed (Gumbel, Pareto, Joe); symmetric and tail dependent (t); symmetric and tail independent (Gaussian, Frank); and negatively skewed (rotated Gumbel, Pareto, Joe). Hence, during time periods when the dark greys dominate, a more adverse dependence structure is present. From that perspective, it appears that we have adverse dependence regimes during 2008–2014 and from 2019 to the present. The extent to which these trends can be associated to phenomena such as insurance cycles remains a topic for further work.

6 CONCLUSIONS

In this paper, we proposed approaches for selecting copula models by utilising image recognition of the density heatmap images obtained from copula samples. PCA-reduced AlexNet image features and summary statistics were utilised to represent each copula sample and a low-dimensional discriminative projection by LDA used to train an SVM for classification. Experimental results showed that the proposed image recognition approach can provide improved classification performance on copula samples with relatively low sample size and correlation, compared to AIC. When combining image recognition with AIC, performance improves further.

Hence, we can answer our research question in the affirmative: indeed, heatmap images do contain relevant information for copula model selection that can help improve on standard procedures, and we propose workflows to harness this information for model selection. Furthermore, we demonstrate the potential value of transfer learning in statistical applications. In this paper, the knowledge obtained from the domain of natural images is transferred to the different domain of copula heatmap images, to assist in the new classification task of copula selection. This shows the potential of utilising machine learning algorithms that have been trained on a large amount of high-quality data to provide additional useful information for statistical applications, or to help in situations where collecting enough data is not an easy task (Pan and Yang 2009).

The workflows we propose are more complex than evaluating likelihood-based criteria, at least where efficient implementation of the latter is available. We do not argue for replacing well-established criteria such as AIC, but provide a ‘proof of concept’ that image recognition can supplement standard statistical approaches. With this in mind, future work can consider designing an expert system with a broader scope, for example including a wider

range of models and sample sizes, as well as handling multivariate dependence structures. For example, one may consider (partially) representing a copula sample with multivariate dependence structure, via two-dimensional heatmap images generated from pairwise bivariate dependence structures. In this way, each sample is represented by a set of images, which leads to the image set classification task in computer vision (Fukui and Maki 2015).

SUPPLEMENTARY MATERIAL

To view supplementary material for this article, please visit <http://doi.org/10.1017/asb.2022.12>.

NOTES

1 Note though that in applications such as personal lines insurance, where discrete variables have many zeros or variables are categorical, such a transformation will introduce unwelcome variability. Alternatively, one may consider the jitter approach of Denuit and Lambert (2005), see for example Shi and Valdez (2014).

2 In fact, when (continuous) marginal distributions are estimated with reasonable confidence, observations $(u_1, v_1), \dots, (u_n, v_n)$ from (U, V) can be obtained by the probability integral transform. In this paper, we do not make such an assumption.

3 In numerical experiments, we found that AIC consistently outperformed BIC, the Cramer-von Mises statistic and the cross-validated log-likelihood criterion with 10-fold cross-validation.

4 This process takes away from realism, by removing some noise from the generated datasets. Nonetheless, we pursue this strategy for image generation in order to maintain consistency with the two-step approaches of Section 4.

5 Specifically, the variable *FirstStep* in Algorithm 3 (Supplementary Material, Appendix C) has the value *TRUE*.

REFERENCES

- ABDALLAH, A., BOUCHER, J.-P. and COSSETTE, H. (2015) Modeling dependence between loss triangles with hierarchical archimedean copulas. *ASTIN Bulletin*, **45**(3), 577–599.
- ANDROSCHUCK, T., GIBBS, S., KATRAKIS, N., LAU, J., ORAM, S., RADDALL, P., SEMCHYSHYN, L., STEVENSON, D. and WATERS, J. (2017) Simulation-based capital models: Testing, justifying and communicating choices. A report from the life aggregation and simulation techniques working party. *British Actuarial Journal*, **22**(2), 257–335.
- BENGIO, Y., COURVILLE, A. and VINCENT, P. (2013). Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **35**(8), 1798–1828.
- CZADO, C., KASTENMEIER, R., BRECHMANN, E.C. and MIN, A. (2012) A mixed copula model for insurance claims and claim sizes. *Scandinavian Actuarial Journal*, **2012**(4), 278–305.
- DENG, J., DONG, W., SOCHER, R., LI, L.-J., LI, K. and FEI-FEI, L. (2009) ImageNet: A large-scale hierarchical image database. *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255. IEEE.
- DENUIT, M., DHAENE, J., GOOVAERTS, M. and KAAS, R. (2006) *Actuarial Theory for Dependent Risks: Measures, Orders and Models*. New Jersey, US: John Wiley & Sons.
- DENUIT, M. and LAMBERT, P. (2005) Constraints on concordance measures in bivariate discrete data. *Journal of Multivariate Analysis*, **93**(1), 40–57.

- EMBRECHTS, P., PUC CETTI, G., RÜSCHENDORF, L., WANG, R. and BELERAJ, A. (2014) An academic response to Basel 3.5. *Risks*, **2**(1), 25–48.
- FREES, E.W. and VALDEZ, E.A. (1998) Understanding relationships using copulas. *North American Actuarial Journal*, **2**(1), 1–25.
- FREES, E.W. and WANG, P. (2005) Credibility using copulas. *North American Actuarial Journal*, **9**(2), 31–48.
- FUKUI, K. and MAKI, A. (2015) Difference subspace and its generalization for subspace-based methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **37**(11), 2164–2177.
- GENEST, C. and FAVRE, A.-C. (2007) Everything you always wanted to know about copula modeling but were afraid to ask. *Journal of Hydrologic Engineering*, **12**(4), 347–368.
- GENEST, C., NEŠLEHOVÁ, J.G., RÉMILLARD, B. and MURPHY, O.A. (2019) Testing for independence in arbitrary distributions. *Biometrika*, **106**(1), 47–68.
- GENEST, C., RÉMILLARD, B. and BEAUDOIN, D. (2009) Goodness-of-fit tests for copulas: A review and a power study. *Insurance: Mathematics and Economics*, **44**(2), 199–213.
- GENEST, C. and SCHERER, M. (2020) Insurance applications of dependence modeling. *Dependence Modeling*, **8**(1), 93–106.
- GRØNNEBERG, S. and HJORT, N.L. (2014) The copula information criteria. *Scandinavian Journal of Statistics*, **41**(2), 436–459.
- HASTIE, T., TIBSHIRANI, R. and FRIEDMAN, J. (2009) *The Elements of Statistical Learning: Data Mining, Inference and Prediction*. New York: Springer Series in Statistics.
- HOFERT, M., KOJADINOVIC, I., MAECHLER, M. and YAN, J. (2020) *copula: Multivariate Dependence with Copulas*. R package version 1.0-1.
- HOFERT, M., PRASAD, A. and ZHU, M. (2021) Quasi-random sampling for multivariate distributions via generative neural networks. *Journal of Computational and Graphical Statistics*, **30**(3), 647–670.
- HU, S., MURPHY, T.B. and O'HAGAN, A. (2021) mvClaim: An R package for multivariate general insurance claims severity modelling. *Annals of Actuarial Science*, **15**(2), 441–457.
- HUARD, D., EVIN, G. and FAVRE, A.-C. (2006) Bayesian copula selection. *Computational Statistics & Data Analysis*, **51**(2), 809–822.
- Insurance Capital Markets Research (2022) RISX: ICMR (Re)Insurance Specialty Index. <https://risxindex.com/>. Accessed: 2022-01-19.
- ISLAM, M., DINH, A., WAHID, K. and BHOWMIK, P. (2017) Detection of potato diseases using image segmentation and multiclass support vector machine. *2017 IEEE 30th Canadian Conference on Electrical and Computer Engineering (CCECE)*, pp. 1–4. IEEE.
- JORDANGER, L.A. and TJØSTHEIM, D. (2014) Model selection of copulas: AIC versus a cross validation copula information criterion. *Statistics & Probability Letters*, **92**, 249–255.
- KRIZHEVSKY, A., SUTSKEVER, I. and HINTON, G.E. (2012) ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, pp. 1097–1105.
- KULARATNE, T.D., LI, J. and PITT, D. (2021) On the use of Archimedean copulas for insurance modelling. *Annals of Actuarial Science*, **15**(1), 57–81.
- MCNEIL, A.J., FREY, R. and EMBRECHTS, P. (2015) *Quantitative Risk Management: Concepts, Techniques and Tools – revised edition*. New Jersey, US: Princeton University Press.
- MSCI (2022). MSCI developed markets indexes. <https://www.msci.com/our-solutions/indexes/developed-markets>. Accessed: 2022-01-19.
- NELSEN, R.B. (2007) *An Introduction to Copulas*. New York: Springer.
- PAN, S.J. and YANG, Q. (2009) A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, **22**(10), 1345–1359.
- PESENTI, S.M., MILLOSovich, P. and TSANAKAS, A. (2019) Reverse sensitivity testing: What does it take to break the model? *European Journal of Operational Research*, **274**(2), 654–670.
- ROSCO, J. and JOE, H. (2013) Measures of tail asymmetry for bivariate copulas. *Statistical Papers*, **54**(3), 709–726.
- RÜSCHENDORF, L. and DE VALK, V. (1993) On regression representations of stochastic processes. *Stochastic Processes and their Applications*, **46**(2), 183–198.

- SCHEPSMEIER, U., STOEBER, J., BRECHMANN, E. C., GRAELER, B., NAGLER, T., ERHARDT, T., ALMEIDA, C., MIN, A., CZADO, C., HOFMANN, M., KILLICHES, M., JOE, H. and VATTER, T. (2021) Package ‘vinecopula’. *R package version*, 2.4.2.
- SHAW, R., SMITH, A.D. and SPIVAK, G. (2010). Calibration and communication of dependencies with a case study based on market returns. <https://www.actuaries.org.uk/system/files/documents/pdf/b3.pdf>. Presented to the Institute and Faculty of Actuaries, November 2010.
- SHEYKHOUSA, M., MAHDIANPARI, M., GHANBARI, H., MOHAMMADIMANESH, F., GHAMISI, P. and HOMAYOUNI, S. (2020) Support vector machine vs. random forest for remote sensing image classification: A meta-analysis and systematic review. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, **13**, 6308–6325.
- SHI, P. (2014) A copula regression for modeling multivariate loss triangles and quantifying reserving variability. *ASTIN Bulletin*, **44**(1), 85–102.
- SHI, P. and FREES, E.W. (2011) Dependent loss reserving using copulas. *ASTIN Bulletin*, **41**(2), 449–486.
- SHI, P. and VALDEZ, E.A. (2014) Multivariate negative binomial models for insurance claim counts. *Insurance: Mathematics and Economics*, **55**, 18–29.
- TZOTSOS, A. and ARGIALAS, D. (2008) Support vector machine classification for object-based image analysis. In *Object-Based Image Analysis*, pp. 663–677. Springer.
- TZOUGAS, G. and PIGNATELLI DI CERCHIARA, A. (2021) The multivariate mixed negative binomial regression model with an application to insurance a posteriori ratemaking. *Insurance: Mathematics and Economics*, **101**, 602–625.
- VENABLES, W.N. and RIPLEY, B.D. (2002) *Modern Applied Statistics with S*, 4th ed. New York: Springer. ISBN 0-387-95457-0.
- WOLD, S., ESSENSEN, K. and GELADI, P. (1987) Principal component analysis. *Chemometrics and Intelligent Laboratory Systems*, **2**(1–3), 37–52.
- YAN, J. (2007). Enjoy the joy of copulas: With a package copula. *Journal of Statistical Software*, **21**(4), 1–21.
- YANG, L. and JIN, R. (2006) Distance metric learning: A comprehensive survey. *Michigan State University*, **2**(2), 4.
- ZHUANG, F., QI, Z., DUAN, K., XI, D., ZHU, Y., ZHU, H., XIONG, H. and HE, Q. (2020) A comprehensive survey on transfer learning. *Proceedings of the IEEE*, **109**(1), 43–76.

ANDREAS TSANAKAS

*Bayes Business School (formerly Cass),
City, University of London,
London EC1V 0HB, UK
E-Mail: a.tsanakas.1@city.ac.uk*

RUI ZHU (CORRESPONDING AUTHOR)

*Bayes Business School (formerly Cass),
City, University of London,
London EC1V 0HB, UK
E-Mail: rui.zhu@city.ac.uk*