**University of
Zurich**^{UZH}

**Zurich Open Repository and
Archive**

University of Zurich
University Library
Strickhofstrasse 39
CH-8057 Zurich
www.zora.uzh.ch

Year: 2022

Multivariate conditional transformation models

Klein, Nadja ; Hothorn, Torsten ; Barbanti, Luisa ; Kneib, Thomas

Abstract: Regression models describing the joint distribution of multivariate responses conditional on covariate information have become an important aspect of contemporary regression analysis. However, a limitation of such models are the rather simplistic assumptions often made, for example, a constant dependence structure not varying with covariates or the restriction to linear dependence between the responses. We propose a general framework for multivariate conditional transformation models that overcomes these limitations and describes the entire distribution in a tractable and interpretable yet flexible way conditional on nonlinear effects of covariates. The framework can be embedded into likelihood-based inference, including results on asymptotic normality, and allows the dependence structure to vary with covariates. In addition, it scales well-beyond bivariate response situations, which were the main focus of most earlier investigations. We illustrate the benefits in a trivariate analysis of childhood undernutrition and demonstrate empirically that complex truly multivariate data-generating processes can be inferred from observations.

DOI: <https://doi.org/10.1111/sjos.12501>

Posted at the Zurich Open Repository and Archive, University of Zurich

ZORA URL: <https://www.zora.uzh.ch/197553>

Journal Article

Published Version

Originally published at:

Klein, Nadja; Hothorn, Torsten; Barbanti, Luisa; Kneib, Thomas (2022). Multivariate conditional transformation models. *Scandinavian Journal of Statistics*, 49(1):116-142.

DOI: <https://doi.org/10.1111/sjos.12501>

Multivariate Conditional Transformation Models

Nadja Klein^{1*}, Torsten Hothorn², Luisa Barbanti² and Thomas Kneib³

¹Humboldt-Universität zu Berlin, ²Universität Zürich,

³Georg-August-Universität Göttingen

Abstract

Regression models describing the joint distribution of multivariate response variables conditional on covariate information have become an important aspect of contemporary regression analysis. However, a limitation of such models is that they often rely on rather simplistic assumptions, *e.g.* a constant dependency structure that is not allowed to vary with the covariates or the restriction to linear dependence between the responses only. We propose a general framework for multivariate conditional transformation models that overcomes these limitations and describes the entire distribution in a tractable and interpretable yet flexible way conditional on nonlinear effects of covariates. The framework can be embedded into likelihood-based inference, including results on asymptotic normality, and allows the dependence structure to vary with covariates. In addition, the framework scales well beyond bivariate response situations, which were the main focus of most earlier investigations. We illustrate the application of multivariate conditional transformation models in a trivariate analysis of childhood undernutrition and demonstrate empirically that our approach can be beneficial compared to existing benchmarks such that complex truly multivariate data-generating processes can be inferred from observations.

Key words: Constrained optimization; copula; marginal distributions; multivariate regression; most likely transformations; normalizing flows; seemingly unrelated regression.

* Correspondence should be directed to Prof. Dr. Nadja Klein at Humboldt Universität zu Berlin, Unter den Linden 6, 10099 Berlin. Email: nadja.klein@hu-berlin.de.

1 Introduction

In a broad sense, regression models describe the distribution of a response conditional on a set of covariates. Such models are a versatile tool to understand how changes in the covariates propagate to changes in the distribution of the response. Distributional and multivariate regression models have received much interest during the last decade. Rather than focusing on the conditional mean, distributional regression strives to describe relevant features of the complete conditional distribution of a usually univariate response by flexible functions of the covariates. Multivariate regression models employ covariates to express the joint conditional distribution of a multivariate response. Known for half a century, transformation models have recently received renewed interest in statistics as an important technique for distributional regression and, under the term normalizing flows, in machine learning (Papamakarios et al., 2019) for modelling high-dimensional responses in an unconditional way. The core idea of flows or transformation models is to apply a data-driven transformation to the response such that the transformed variable is standard normal or follows some other convenient distribution. In this paper, we propose a framework of multivariate conditional transformation models (MCTMs) that apply this principle to define a novel class of multivariate distributional regression models. We review relevant developments in multivariate distributional regression first before highlighting some special features of the new method.

The most prominent multivariate regression model is seemingly unrelated regression (SUR), which uses a vector of correlated normal error terms to combine several linear model regression specifications with a common correlation structure that does not depend on any of the covariates (Zellner, 1962). By construction, the model is restricted to capture linear dependencies. Lang et al. (2003) extended SUR models by replacing the frequently used linear predictor with a structured additive predictor, while retaining the assumption that the linear correlation structure does not depend on the covariates. Multivariate probit models use a latent SUR model for a multivariate set of latent utilities that, via thresholds, are transformed to the observed binary response vector (Heckman, 1978). The approach of Klein, Kneib, Klasen and Lang (2015) embeds bivariate SUR-type specifications into generalised additive models for location, scale and shape (GAMLSS, Rigby and Stasinopoulos, 2005) by allowing all distribution parameters, including the correlations, to be related to additive predictors.

Beyond SUR-type models, copulas provide a flexible approach to the construction of multivariate distributions and regression models. As a major advantage, the model-building process is conveniently decomposed into the specification of the marginals and the selection of an appropriate copula function that defines the dependence structure (see Joe, 1997; Nelson, 2006, for reviews on copula models and their properties).

There is a rich literature on conditional copula modelling. To name just a few, Veraverbeke et al. (2011) use kernel methods to estimate the copula parameter after having determined the marginal distributions empirically. Assuming that the marginal distributions are known, a copula can be fitted based on a local likelihood using the approach of Acar et al. (2011). Bayesian inference in bivariate conditional copula models with homoscedastic Gaussian marginals has been proposed in Sabeti et al. (2014) and Levi and Craiu (2018). Analogous to GAMLSS, bivariate copula models with parametric marginal distributions, one-parameter copulas, as well as joint semiparametric specifications for the predictors of all parameters of both marginal and copula models have been developed by Marra and Radice (2017) in a penalised likelihood framework and by Klein and Kneib (2016) using a Bayesian approach. Following these lines, Marra and Radice (2020) recently extended the framework to copula link-based survival models, while Sun and Ding (2019) develop a copula-based semiparametric regression method for bivariate data under general interval censoring. Alternatives to simultaneous estimation are two-step procedures that first estimate the marginals and then the copula given the marginals and have been proposed by *e.g.* Vatter and Chavez-Demoulin (2015) and Yee (2015) for parametric marginal distributions and bivariate one-parameter and copulas. However, these approaches are mostly limited to the bivariate case. Vatter and Nagler (2018) recently proposed a sequential method for conditional pair-copula constructions.

Nonparametric attempts to simultaneously study multivariate response variables have been reported in the context of multivariate quantiles. Because no natural ordering exists beyond univariate settings, definitions of multivariate quantiles are challenging and there has been considerable debate regarding their desirable properties (see Serfling, 2002, for an introduction to the different definitions). For example, one group of approaches draws on the concept of data depths (see for example Mosler, 2010), utilising options for multivariate depth functions based on distances such as Mahalanobis and Oja depths, weighted distances or on half-spaces. However, potential quantile crossings need further investigations

to ensure a coherent model for the joint distribution, because single quantiles only relate to local properties of a response to covariates. For more information on depth functions and multivariate quantiles, we refer the reader to Chernozhukov et al. (2017) and Carlier et al. (2016, 2017).

MCTMs constitute a novel and coherent approach to multivariate regression analysis which is in many aspects different to existing approaches in copula or nonparametric regression. In particular, this framework makes six important contributions, none of which are available simultaneously in any existing method to regression for multivariate responses:

- i. MCTMs allow for direct estimation and inference of the entire multivariate conditional cumulative distribution function (CDF) $F_{\mathbf{Y}}(\mathbf{y} \mid \mathbf{x}) = \mathbb{P}(\mathbf{Y} \leq \mathbf{y} \mid \mathbf{x})$ of a J -dimensional response vector $\mathbf{Y}$ given covariate information $\mathbf{x}$ under rather weak assumptions. A key feature of MCTMs is that they extend likelihood-based inference in univariate conditional transformation models (CTMs, Hothorn et al., 2018) to the multivariate situation in a natural way.
- ii. MCTMs can capture nonlinear aspects of covariates on all aspects of the distribution, *e.g.* marginal moments, marginal and joint quantiles, dependence structures etc. As in the case of copulas, a feature of the model specification process is that joint distributions are constructed by their decomposition into marginals and the dependence structure. Most existing approaches assume a constant dependence structure not varying over the covariate space.
- iii. Model estimation can be performed simultaneously for all model components, thus avoiding the need for two-step estimators that are commonly applied in most copula-based approaches.
- iv. Theoretical results on optimality properties, such as consistency and asymptotic normality are available, building on the achievements in univariate CTMs.
- v. Unlike multivariate GAMLSS, MCTMs neither require strong parametric assumptions nor separate the model estimation process into local properties, as in multivariate quantile regression.
- vi. The method scales well to situations beyond the bivariate case $J = 2$ and readily allows for the determination of both the marginal distributions of subsets of the response vector and the conditional distributions of some response elements, given the others. MCTMs are not equivalent to copulas, however, Gaussian copulas (Song, 2000) with

arbitrary marginal distributions are treated as a special case in this paper. Both the marginal distributions and the correlation parameters of the copula can depend on covariates when such a copula model is specified by means of an MCTM.

The paper is structured as follows: Section 2 provides details on the specification of multivariate transformation models for the unconditional case of absolutely continuous responses. Likelihood-based inference and optimality properties are derived in Section 3, along with an illustration on multivariate density estimation with highly non-Gaussian marginal distributions. Section 4 considers how multivariate conditional transformation models may depend on covariates, and the approach is illustrated by a trivariate analysis of childhood undernutrition indicators. Section 5 presents simulation-based empirical evidence on the performance of MCTMs, including examples with up to 10 response dimensions. Finally, Section 6 proposes directions for future research.

2 Multivariate Transformation Models

2.1 Basic Model Setup

First, *unconditional* transformation models are developed for the joint multivariate distribution of a J -dimensional, absolutely continuous random vector $\mathbf{Y} = (Y_1, \dots, Y_J)^\top \in \mathbb{R}^J$ with density $f_{\mathbf{Y}}(\mathbf{y})$ and CDF $F_{\mathbf{Y}}(\mathbf{y}) = \mathbb{P}(\mathbf{Y} \leq \mathbf{y})$. These unconditional models are then extended to the regression case in Section 4.

The key component of multivariate transformation models is an unknown, bijective, strictly monotonically increasing transformation function $h : \mathbb{R}^J \rightarrow \mathbb{R}^J$. This function maps the vector $\mathbf{Y}$, whose distribution is unknown and shall be estimated from data, to a set of J independent and identically distributed, absolutely continuous random variables $Z_j \sim \mathbb{P}_Z, j = 1, \dots, J$ with an a priori defined distribution $\mathbb{P}_Z$, such that

$$h(\mathbf{Y}) = (h_1(\mathbf{Y}), \dots, h_J(\mathbf{Y}))^\top \stackrel{d}{=} (Z_1, \dots, Z_J)^\top = \mathbf{Z} \in \mathbb{R}^J.$$

For an absolutely continuous distribution $\mathbb{P}_Z$ with log-concave density f_Z , it can easily be shown that a unique, monotonically increasing transformation function h exists for arbitrary, absolutely continuous distributions of $\mathbf{Y}$ (Hothorn et al., 2018). Thus, the model class is effectively limited only by the flexibility of the specific choice of h in an actual model. As a default for $\mathbb{P}_Z$, we consider the standard normal distribution, *i.e.* $Z_j \sim \mathcal{N}(0, 1)$ with $\mathbb{P}_Z = \mathcal{N}(0, 1)$ and thus $f_Z = \phi_{0,1}$ and $F_Z = \Phi_{0,1}$. Enforcing independent standard

normality of the transformed response $h(\mathbf{Y})$ is computationally attractive and allows the dependence structure to be described by a Gaussian copula (see Section 2.3). Alternative choices for $\mathbb{P}_Z$ are discussed in Section 2.6.

Under this transformation model, the task of estimating the distribution of $\mathbf{Y}$ simplifies to the task of estimating h . Because h is strictly monotonically increasing in each element, it has a positive definite Jacobian, *i.e.*

$$\left| \frac{\partial h(\mathbf{y})}{\partial \mathbf{y}} \right| > 0. \quad (1)$$

The density of $\mathbf{Y}$ implied by the transformation model is then

$$f_{\mathbf{Y}}(\mathbf{y}) = \left[\prod_{j=1}^J f_Z(h_j(\mathbf{y})) \right] \cdot \left| \frac{\partial h(\mathbf{y})}{\partial \mathbf{y}} \right|.$$

This form of an unconditional multivariate density with $\mathbb{P}_Z = \mathcal{N}(0, 1)$ is called a normalizing flow in machine learning (Papamakarios et al., 2019). However, in this generality, the model is cumbersome in terms of both interpretation and tractability. Thus, in the following, we introduce simplified parameterisations of h that lead to interpretable models.

2.2 Models with Recursive Structure

In a first step, we impose a triangular structure on the transformation function h by assuming

$$h_j(\mathbf{y}) = h_j(y_1, \dots, y_J) = h_j(y_1, \dots, y_j)$$

i.e. the j th component of the transformation function depends only on the first j elements of its argument $\mathbf{y}$. Consequently, this model formulation depends inherently on the ordering of the elements in $\mathbf{Y}$. Because any multivariate distribution can be factored into a sequence of conditional distributions, the triangular structure does not pose a limitation in the representation of J -dimensional continuous distributions, as long as the transformation functions are appropriately chosen. Furthermore, the triangular structure of h considerably simplifies the determinant of the Jacobian (1), which reduces to

$$\left| \frac{\partial h(\mathbf{y})}{\partial \mathbf{y}} \right| = \prod_{j=1}^J \left| \frac{\partial h_j(y_1, \dots, y_j)}{\partial y_j} \right|.$$

In a second step, we assume that the triangularly structured transformation functions are linear combinations of marginal transformation functions $\tilde{h}_j : \mathbb{R} \rightarrow \mathbb{R}$, *i.e.*

$$h_j(y_1, \dots, y_j) = \lambda_{j1} \tilde{h}_1(y_1) + \dots + \lambda_{jj} \tilde{h}_j(y_j)$$

where each $\tilde{h}_j$ increases strictly monotonically and $\lambda_{jj} > 0$ for all $j = 1, \dots, J$ to ensure the bijectivity of h . Because the last coefficient, λ_{jj} , cannot be separated from the marginal transformation function $\tilde{h}_j(y_j)$, we use the restriction $\lambda_{jj} \equiv 1$. Thus, our parameterisation of the transformation function h finally reads

$$h_j(y_1, \dots, y_j) = \lambda_{j1}\tilde{h}_1(y_1) + \dots + \lambda_{j,j-1}\tilde{h}_{j-1}(y_{j-1}) + \tilde{h}_j(y_j). \quad (2)$$

Each of the marginal transformation functions $\tilde{h}_j(y_j)$ includes an intercept, such that no additional intercept term can be inserted in (2). The Jacobian of h now further simplifies to

$$\left| \frac{\partial h(\mathbf{y})}{\partial \mathbf{y}} \right| = \prod_{j=1}^J \frac{\partial \tilde{h}_j(y_j)}{\partial y_j} > 0,$$

and the model-based density function for $\mathbf{Y}$ is therefore

$$f_{\mathbf{Y}}(\mathbf{y}) = \prod_{j=1}^J f_Z \left(\lambda_{j1}\tilde{h}_1(y_1) + \dots + \lambda_{j,j-1}\tilde{h}_{j-1}(y_{j-1}) + \tilde{h}_j(y_j) \right) \frac{\partial \tilde{h}_j(y_j)}{\partial y_j}.$$

Summarising the model specification, our multivariate transformation model is characterised by a set of marginal transformations $\tilde{h}_j(y_j)$, $j = 1, \dots, J$, each applying to only a single component of the vector $\mathbf{Y}$, and by a lower triangular $(J \times J)$ matrix of transformation coefficients

$$\mathbf{\Lambda} = \begin{pmatrix} 1 & & & & 0 \\ \lambda_{21} & 1 & & & \\ \lambda_{31} & \lambda_{32} & 1 & & \\ \vdots & \vdots & & \ddots & \\ \lambda_{J1} & \lambda_{J2} & \dots & \lambda_{J,J-1} & 1 \end{pmatrix}.$$

Under the standard normal reference distribution $\mathbb{P}_Z = \mathcal{N}(0, 1)$, the coefficients in $\mathbf{\Lambda}$ characterise the dependence structure via a Gaussian copula, while the marginal transformation functions $\tilde{h}_j$ allow the generation of arbitrary marginal distributions for the components of $\mathbf{Y}$. Furthermore, the entries of $\mathbf{\Lambda}$ have the interpretation of entries in the inverse precision matrix of the correlation matrix between the marginally transformed components of $\mathbf{Y}$, as we derive in the following.

2.3 Relation to Gaussian Copula Models

The relationship between multivariate transformation models and Gaussian copulas can be made more precise by defining random variables $\tilde{Z}_j = \tilde{h}_j(Y_j)$. Under a standard normal reference distribution $\mathbb{P}_Z = \mathcal{N}(0, 1)$, the vector $\tilde{\mathbf{Z}} = (\tilde{Z}_1, \dots, \tilde{Z}_J)^\top$ follows a zero mean

multivariate normal distribution $\tilde{\mathbf{Z}} \sim N_J(\mathbf{0}_J, \mathbf{\Sigma})$ with covariance matrix $\mathbf{\Sigma} = \mathbf{\Lambda}^{-1} \mathbf{\Lambda}^{-\top}$. As a consequence, the elements of $\tilde{\mathbf{Z}}$ are marginally normally distributed as $\tilde{Z}_j \sim N(0, \sigma_j^2)$, where the variances σ_j^2 can be determined from the diagonal elements of $\mathbf{\Sigma}$.

For the transformation functions $\tilde{h}_j$, the explicit representation

$$\tilde{h}_j(Y_j) = \Phi_{0, \sigma_j^2}^{-1}(F_j(Y_j)) = \tilde{Z}_j \quad (3)$$

is obtained, where $F_j(\cdot)$ is the univariate marginal CDF of Y_j . In summary,

$$\begin{aligned} \mathbb{P}(\mathbf{Y} \leq \mathbf{y}) &= \mathbb{P}(\tilde{\mathbf{Z}} \leq \tilde{\mathbf{z}}) = \Phi_{\mathbf{0}, \mathbf{\Sigma}}(\tilde{\mathbf{z}}) = \Phi_{\mathbf{0}, \mathbf{\Sigma}} \left[\Phi_{0, \sigma_1^2}^{-1} \{F_1(y_1)\}, \dots, \Phi_{0, \sigma_J^2}^{-1} \{F_J(y_J)\} \right] \\ &= \Phi_{\mathbf{0}, \mathbf{\Sigma}} \left(\tilde{h}_1(y_1), \dots, \tilde{h}_J(y_J) \right). \end{aligned}$$

and therefore the CDF of $\mathbf{Y}$ has exactly the same structure as a Gaussian copula, except that our representation relies on a different parameterisation of $\mathbf{\Sigma}$ through $\mathbf{\Sigma} = \mathbf{\Lambda}^{-1} \mathbf{\Lambda}^{-\top}$ rather than a covariance matrix with unit diagonal. This is compensated for by the inclusion of univariate Gaussians with variances different from those acting on the marginals. However, because

$$\tilde{Z}_j / \sigma_j \sim N(0, 1) \text{ and } \Phi_{0, \sigma_j^2}(\tilde{z}_j) = \Phi_{0, 1}(\tilde{z}_j / \sigma_j) \quad (4)$$

unconditional MCTMs with $\mathbb{P}_Z = N(0, 1)$ are equivalent to a Gaussian copula with flexible marginal distributions. This is no longer the case of the conditional case of regression considered in Section 4, where our approach can capture nonlinear aspects of covariates on all aspects of the distribution, *e.g.* marginal moments, marginal and joint quantiles, dependence structures etc.

2.4 Some Properties of the Dependence Structure

We highlight some properties of our MCTM that can be of interest in applied studies.

- The transformed vector $\tilde{\mathbf{Z}}$ is jointly multivariate normally distributed such that pairwise dependencies are restricted to linear dependence through linear correlations. Importantly however, $\mathbf{Y}$ is allowed to have a nonlinear dependence structure due to the inverse marginal transformation functions $Y_j = \tilde{h}_j^{-1}(\tilde{Z}_j)$. We illustrate this feature in Figure 1 which shows a bivariate scatterplot with one marginally normally and one marginally gamma distributed response component for the vector $(Y_1, Y_2)^\top$ (on the left) together with the bivariate normally distributed variables $\tilde{\mathbf{Z}}$ (on the right).
- Assuming $\mathbb{P}_Z = N(0, 1)$ as reference distribution, the entries in $\mathbf{\Lambda}$ determine the conditional independence structure between the transformed responses $\tilde{Z}_j$ and therefore,

Figure 1: Bivariate illustration of the relation between the distributions of $\tilde{\mathbf{Z}}$ (right) and $\mathbf{Y}$ (left). The distribution of $\mathbf{Y}$ is constructed from a Gaussian copula with correlation parameter $\rho = 0.9$, standard normally distributed Y_1 and gamma distributed Y_2 .

implicitly, also the observed responses Y_j as it is for a multivariate Gaussian distribution, see the Appendix Part A.1 for details.

- Rather than looking at linear correlations, common measures of dependence in the context of multivariate modelling are Spearman's rho ρ^S , Kendall's tau τ^K and lower/upper quantile dependence λ^L/λ^U . These can be computed in closed form using the results known for a Gaussian copula, see again the Appendix Part A.1 for formulas. One appealing property of these measures is that they are invariant with respect to monotonic transformations of the marginals and we will use the ρ^S later in our trivariate application on childhood undernutrition.

2.5 Model-Implied Marginal and Conditional Distributions

The relationship of unconditional transformation models and Gaussian copulas can now be employed to facilitate the derivation of model-implied marginal and conditional distributions. The univariate marginal distributions of elements Y_j are given by

$$F_j(y_j) = \mathbb{P}(Y_j \leq y_j) = \Phi_{\mathbf{0}, \Sigma} \left(\infty, \dots, \infty, \tilde{h}_j(y_j), \infty, \dots, \infty \right) = \Phi_{0, \sigma_j^2} \left(\tilde{h}_j(y_j) \right)$$

but more general versions (*i.e.* marginals of a subvector of $\mathbf{Y}$) and conditional distributions are also easily obtained, see the Appendix Part A.2.

Finally, using the marginal CDFs and densities, the marginal quantiles or moments can be derived. The latter can be computed by solving simple univariate numerical integrals, for example:

$$\mathbb{E}(Y_j) = \int F_j^{-1}(\Phi_{0, \sigma_j^2}(\tilde{z})) \phi_{0, \sigma_j^2}(\tilde{z}) d\tilde{z},$$

for the marginal mean.

2.6 Alternative Reference Distributions

Although we have discussed our model specification in the context of a normal reference distribution and a Gaussian copula, these choices can be readily modified. In particular, if a reference distribution $\mathbb{P}_Z \neq N(0, 1)$ is chosen, the transformation function has to be modified to

$$h_j(y_1, \dots, y_j) = \sum_{j=1}^{j-1} \lambda_{jj} \Phi_{0, \sigma_j^2}^{-1} \left[F_Z \left\{ \tilde{h}_j(y_j) \right\} \right] + \Phi_{0, \sigma_j^2}^{-1} \left[F_Z \left\{ \tilde{h}_j(y_j) \right\} \right].$$

We therefore obtain J independent random variables $\Phi_{0, \sigma_j^2}^{-1} \left\{ F_Z \left[\tilde{h}_j(Y_j) \right] \right\} \sim N(0, \sigma_j^2)$. The model then implies marginal distributions

$$F_j(y_j) = \Phi_{\mathbf{0}, \Sigma} \left(\infty, \dots, \infty, \Phi_{0, \sigma_j^2}^{-1} \left\{ F_Z \left[\tilde{h}_j(y_j) \right] \right\}, \infty, \dots, \infty \right) = F_Z \left(\tilde{h}_j(y_j) \right).$$

Attractive alternative choices for the reference distribution are $F_Z^{-1} = \text{logit}$ and $F_Z^{-1} = \text{cloglog}$, because regression coefficients can be interpreted as log-odds ratios and log-hazard ratios (in fact, in the latter case, the marginal model is then a Cox proportional hazards model), respectively.

Extensions beyond the Gaussian copula structure are also conceivable when the linear combination of marginal transformations is replaced by nonlinear specifications. However, those types of models easily lead to identification problems and do not provide direct links to existing parametric copula classes. Accordingly, we leave this topic for future research.

3 Transformation Analysis

This section defines the maximum likelihood estimator and establishes its consistency and asymptotic normality based on suitable parameterisations of the marginal transformation functions $\tilde{h}_j$. It closes with an illustration on bivariate density estimation for highly non-Gaussian data.

3.1 Parameterisation of the Transformation Functions

Following Hothorn et al. (2018), the marginal transformation functions $\tilde{h}_j(y_j)$ are parameterised as linear combinations of the basis-transformed argument y_j , such that $\tilde{h}_j(y_j) = \mathbf{a}_j(y_j)^\top \boldsymbol{\vartheta}_j$ is monotonically increasing. The P_j -dimensional basis functions $\mathbf{a}_j : \mathbb{R} \rightarrow \mathbb{R}^{P_j}$ with basis coefficients $\boldsymbol{\vartheta}_j$ and corresponding derivative $\tilde{h}'_j(y_j) = \mathbf{a}'_j(y_j)^\top \boldsymbol{\vartheta}_j > 0$ are problem-specific, see Hothorn et al. (2018) for suitable choices in different applications. Because

marginal transformation functions $\tilde{h}_j(y_j)$ and therefore also the plug-in estimators $\hat{F}_j$ of the marginal CDF should be smooth with respect to y_j , in principle any polynomial or spline-based basis is a suitable choice for $\mathbf{a}_j$. The empirical results of Sections 4 and 5 rely on Bernstein polynomials of order M ; suitable choices of this parameter are discussed in Section 5.1. The basis functions $\mathbf{a}_j(y_j)$ are then densities of beta distributions, a choice that is computationally appealing because strict monotonicity can be formulated as a set of linear constraints on the components of the parameters $\boldsymbol{\vartheta}_j$, see Curtis and Ghosh (2011); Farouki (2012) for details. Furthermore, Bernstein polynomials of sufficiently large order M can uniformly approximate any function over an interval as a result of the Weierstrass approximation theorem. Hothorn et al. (2018) investigate the choice of M for univariate CTMs.

3.2 Inference

In the following, we denote the set of parameters describing all marginal transformation functions $\tilde{h}_j, j = 1, \dots, J$ as $\boldsymbol{\vartheta} = (\boldsymbol{\vartheta}_1^\top, \dots, \boldsymbol{\vartheta}_J^\top)^\top \in \mathbb{R}^{\sum_{j=1}^J P_j}$, while $\boldsymbol{\lambda}$ contains all unknown elements of $\boldsymbol{\Lambda}$, such that $\boldsymbol{\theta} = (\boldsymbol{\vartheta}^\top, \boldsymbol{\lambda}^\top)^\top$ comprises all unknown model parameters. The parameter space is denoted as $\Theta = \{\boldsymbol{\theta} | h \in \mathcal{H}\}$, where

$$\mathcal{H} = \left\{ h : \mathbb{R}^J \rightarrow \mathbb{R}^J \mid h \text{ as in (2), } h \text{ strictly monotonically increasing} \right\}$$

is the space of all strictly monotonic triangular transformation functions. Consequently, the problem of estimating the unknown transformation function h , and thus the unknown distribution function $F_{\mathbf{Y}}$, reduces to the problem of estimating the parameter vector $\boldsymbol{\theta}$. With the construction of multivariate transformation models, this is conveniently achieved using likelihood-based inference. For $\mathbb{P}_Z = N(0, 1)$, the log-likelihood contribution of a given datum $\mathbf{y}_i = (y_{i1}, \dots, y_{iJ})^\top \in \mathbb{R}^J, i = 1, \dots, n$ is

$$\ell_i(\boldsymbol{\theta}) = -\frac{1}{2} \sum_{j=1}^J \left(\sum_{j=1}^{j-1} \lambda_{jj} \mathbf{a}_j(y_{ij})^\top \boldsymbol{\vartheta}_j + \mathbf{a}_j(y_{ij})^\top \boldsymbol{\vartheta}_j \right)^2 + \log(\mathbf{a}'_j(y_{ij})^\top \boldsymbol{\vartheta}_j)$$

with corresponding score contributions

$$\frac{\partial \ell_i(\boldsymbol{\theta})}{\partial \boldsymbol{\vartheta}_k} = \sum_{j=k}^J - \left(\sum_{j=1}^{j-1} \lambda_{jj} \mathbf{a}_j(y_{ij})^\top \boldsymbol{\vartheta}_j + \mathbf{a}_j(y_{ij})^\top \boldsymbol{\vartheta}_j \right) \lambda_{jk} \mathbf{a}_k(y_{ik}) + \frac{\mathbf{a}'_k(y_{ik})}{\mathbf{a}'_k(y_{ik})^\top \boldsymbol{\vartheta}_k} \quad (5)$$

$$\frac{\partial \ell_i(\boldsymbol{\theta})}{\partial \lambda_{\tilde{k}k}} = - \left(\sum_{j=1}^{\tilde{k}-1} \lambda_{\tilde{k}j} \mathbf{a}_j(y_{ij})^\top \boldsymbol{\vartheta}_j + \mathbf{a}_{\tilde{k}}(y_{i\tilde{k}})^\top \boldsymbol{\vartheta}_{\tilde{k}} \right) \mathbf{a}_k(y_{ik})^\top \boldsymbol{\vartheta}_k \quad (6)$$

for $k = 1, \dots, J$, $1 \leq k < \tilde{k} \leq J$ (and zero otherwise) and with $\lambda_{jj} \equiv 1$. We furthermore define $\mathcal{F}_i(\boldsymbol{\theta}) = -\frac{\partial^2 \ell_i(\boldsymbol{\theta})}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^\top}$ as the i th contribution to the observed Fisher information. Explicit expressions for the entries are given in Appendix Part B. Despite the estimation of a fairly complex multivariate distribution with a Gaussian copula dependence structure and arbitrary marginals, the log-likelihood contributions have a very simple form. In addition, the log-concavity of f_Z ensures the concavity of the log-likelihood and thus the existence and uniqueness of the estimated transformation function $\hat{h}$.

Definition 1. (*Maximum likelihood estimator.*) *The maximum likelihood estimator (MLE) for the parameters of a multivariate transformation model is given by*

$$\hat{\boldsymbol{\theta}}_n = \arg \max_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} \sum_{i=1}^n \ell_i(\boldsymbol{\theta}). \quad (7)$$

Based on the maximum likelihood estimator $\hat{\boldsymbol{\theta}}_n$, maximum likelihood estimators for the marginal and joint CDFs are also obtained, by plugging in $\hat{\boldsymbol{\theta}}_n$. Specifically, the estimated marginal CDFs are given by $\hat{F}_j(y_j) = \Phi_{0, \hat{\sigma}_j^2}(\mathbf{a}_j(y_j)^\top \hat{\boldsymbol{\vartheta}}_j)$, where $\hat{\sigma}_j^2$ is the j th diagonal entry of $\hat{\boldsymbol{\Sigma}}$. The estimated joint CDF reads

$$\hat{F}_{\mathbf{Y}}(\mathbf{y}) = \Phi_{0, \hat{\boldsymbol{\Sigma}}}(\mathbf{a}_1(y_1)^\top \hat{\boldsymbol{\vartheta}}_1, \dots, \mathbf{a}_J(y_J)^\top \hat{\boldsymbol{\vartheta}}_J).$$

3.3 Parametric Inference

In this section, we discuss likelihood-based inference and establish asymptotic results for multivariate transformation models based the theoretical results derived in Hothorn et al. (2018) for univariate conditional transformation models. Assume $\mathbf{Y}_1, \dots, \mathbf{Y}_n \stackrel{\text{i.i.d.}}{\sim} F_{\mathbf{Y}, \boldsymbol{\theta}_0}$ where $\boldsymbol{\theta}_0$ denotes the true parameter vector, then the following assumptions are made:

(A1) The parameter space $\boldsymbol{\Theta}$ is compact.

(A2) $\mathbb{E}_{\boldsymbol{\theta}_0}[\sup_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} \log\{f_{\mathbf{Y}}(\mathbf{Y}|\boldsymbol{\theta})\}] < \infty$, where

$$\sup_{\{\boldsymbol{\theta} \mid |\boldsymbol{\theta} - \boldsymbol{\theta}_0| > \epsilon\}} \mathbb{E}_{\boldsymbol{\theta}_0}[\log\{f_{\mathbf{Y}}(\mathbf{Y}|\boldsymbol{\theta})\}] < \mathbb{E}_{\boldsymbol{\theta}_0}[\log\{f_{\mathbf{Y}}(\mathbf{Y}|\boldsymbol{\theta}_0)\}].$$

(A3)

$$\mathbb{E}_{\boldsymbol{\theta}_0} \left(\sup_{\boldsymbol{\theta}} \left\| \frac{\partial \log(f_{\mathbf{Y}}(\mathbf{Y}|\boldsymbol{\theta}))}{\partial \boldsymbol{\theta}} \right\|^2 \right) < \infty.$$

(A4) $\mathbb{E}_{\boldsymbol{\theta}_0}(\mathcal{F}(\boldsymbol{\theta}))$ is nonsingular.

(A5) $0 < f_Z < \infty$, $\sup |f'_Z| < \infty$, $\sup |f''_Z| < \infty$.

Remark 1. Assumption (A1) is made for convenience, and relaxations of such a condition are given in Vaart (1998). The assumptions in (A2) are rather weak: the first one holds if the functions $\mathbf{a}$ are not arbitrarily ill-posed, and the second one holds if the function $\mathbb{E}_{\theta_0}[\sup_{\theta \in \Theta} [\log\{f_{\mathbf{Y}}(\mathbf{Y}|\theta)\}]]$ is strictly convex in θ (if the assumption would not hold, we would still have convergence to the set $\mathbb{E}_{\theta_0}[\sup_{\theta \in \Theta} [\log\{f_{\mathbf{Y}}(\mathbf{Y}|\theta)\}]]$). Assumptions (A3)–(A5) are needed to derive the asymptotic distribution.

Corollary 1. *Assuming (A1)–(A2), the sequence of estimators $\hat{\theta}_n$ converges in probability $\hat{\theta}_n \xrightarrow{\mathbb{P}} \theta_0$ for $n \rightarrow \infty$.*

The proof of Corollary 1 follows from Theorem 5.8 of Vaart (1998).

Corollary 2. *Assuming (A1)–(A5), the sequence of estimators $\sqrt{n}(\hat{\theta}_n - \theta_0)$ is asymptotically normally distributed with covariance matrix*

$$\left(\mathbb{E}_{\theta_0} \left(-\frac{\partial^2 \log(f_{\mathbf{Y}}(\mathbf{Y}|\theta))}{\partial \theta \partial \theta^\top} \right) \right)^{-1}. \quad (8)$$

Proof. By further assumption, $\sqrt{f_{\mathbf{Y}}}$ is continuously differentiable in θ for all $\mathbf{y}$ and

$$\mathbb{E}_{\theta_0} \left(\left[\frac{\partial \log(f_{\mathbf{Y}}(\mathbf{Y}|\theta))}{\partial \theta} \right] \left[\frac{\partial \log(f_{\mathbf{Y}}(\mathbf{Y}|\theta))}{\partial \theta} \right]^\top \right)$$

is continuous in θ , due to (5),(6). Thus, $F_{\mathbf{Y},\theta_0}$ is differentiable in quadratic mean by Lemma 7.6 of Vaart (1998). Based on Assumptions (A3)–(A5) and Corollary 1, the claim hence follows from Theorem 5.39 of Vaart (1998). $\square$

Remark 2. Similar as in the univariate case (Hothorn et al., 2018), Corollaries 1, 2 also extend to the conditional regression models considered in Section 4.

3.4 Parametric Bootstrap

The asymptotic results allow, at least in principle, the derivation of confidence intervals also for transformed model parameters, by using the Delta-rule. However, many quantities of practical interest, such as the correlation matrix of the implied Gaussian copula or the densities of marginal distributions, are indeed highly nonlinear functions of the parameter vector θ . Accordingly, a parametric bootstrap is a more promising alternative.

Because the proposed multivariate transformation models allow a direct evaluation of estimated joint CDFs, drawing bootstrap samples from the joint distribution is straightforward. For $\mathbb{P}_Z = N(0, 1)$ and with estimated marginal transformation functions $\hat{h}_j(y_j) =$

Algorithm 1: Parametric bootstrap

Given $\hat{\mathbf{h}} = (\hat{h}_1(y_1), \dots, \hat{h}_J(y_J))^\top$ and $\hat{\mathbf{\Lambda}}$:

1: **for** $b = 1, \dots, B$ **do**

2: Generate

$$\mathbf{Z}_{1b}, \dots, \mathbf{Z}_{nb}, \quad b = 1, \dots, B, \quad \mathbf{Z}_{ib} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_J), \quad i = 1, \dots, n.$$

3: Compute

$$\tilde{\mathbf{Z}}_{1b}, \dots, \tilde{\mathbf{Z}}_{nb}, \quad \tilde{\mathbf{Z}}_{ib} = \hat{\mathbf{\Lambda}}^{-1} \mathbf{Z}_{ib}.$$

4: Compute recursively

$$\mathbf{Y}_{1b}, \dots, \mathbf{Y}_{nb}, \quad Y_{ijb} = \hat{h}_j^{-1}(\tilde{Z}_{ijb}), \quad j = 1, \dots, J.$$

5: Re-fit the model to obtain $\hat{h}_{(b)}, \hat{\mathbf{\Sigma}}_{(b)}$.

6: Compute summaries using $\hat{h}_{(1)}, \dots, \hat{h}_{(B)}, \hat{\mathbf{\Sigma}}_{(1)}, \dots, \hat{\mathbf{\Sigma}}_{(B)}$

7: **end for**

$\mathbf{a}_j(y_j)^\top \hat{\boldsymbol{\vartheta}}_j, j = 1, \dots, J$ and estimated covariance matrix $\hat{\mathbf{\Sigma}} = \hat{\mathbf{\Lambda}}^{-1} \hat{\mathbf{\Lambda}}^{-\top}$, the parametric bootstrap can be implemented by Algorithm 1. The inverse $\hat{h}_j^{-1}$ exists because the estimated marginal distribution function is strictly monotonically increasing. For simple basis functions $\mathbf{a}_j$ (for example, linear functions), the inverse can be computed analytically. For more complex basis functions, numerical inversion has to be applied.

3.5 Illustration: Bivariate Density Estimation

Unconditional MCTMs can be employed for multivariate density estimation. For the famous 1920s cars data (Ezekiel, 1930) consisting of speed and distance needed to stop for 50 cars, the bivariate distribution was estimated from an unconditional bivariate transformation model with $\mathbb{P}_Z = \mathcal{N}(0, 1)$, order $M = 6$ of Bernstein polynomials for the two transformation functions, and a constant parameter $\lambda \in \mathbb{R}$. The model is equivalent to a Gaussian copula, however, the marginal distributions are highly non-Gaussian and were estimated by maximum likelihood simultaneously with the correlation parameter λ . The fit of the bivariate density contours and marginal densities is given in Figure 2, which shows that the dependence between speed and distance is clearly nonlinear. We obtained $\hat{\lambda} = -1.633$

For the marginal (with respect to the response y_j) conditional (given covariates $\mathbf{x}$) transformation functions, this leads to a parameterisation

$$\tilde{h}_j(y_j | \mathbf{x}) = \mathbf{c}_j(y_j, \mathbf{x})^\top \boldsymbol{\vartheta}_j$$

where the basis functions $\mathbf{c}_j(y_j, \mathbf{x})$, in general, depend on both element y_j of the response and the covariates $\mathbf{x}$. These can, for example, be constructed as a composition of the basis functions $\mathbf{a}_j(y_j)$ for only y_j from the previous section combined with a basis $\mathbf{b}_j(\mathbf{x})$ depending exclusively on $\mathbf{x}$. Specifically, a purely additive model results from $\mathbf{c}_j = (\mathbf{a}_j^\top, \mathbf{b}_j^\top)^\top$, and a flexible interaction from the tensor product $\mathbf{c}_j = (\mathbf{a}_j^\top \otimes \mathbf{b}_j^\top)^\top$. Response-varying coefficients in distributional regression, or time-varying effects in survival analysis, correspond to a basis $\mathbf{c}_j = (\mathbf{a}_j^\top \otimes (1, \mathbf{x}^\top)^\top)^\top$, see also Section 4.2. A simple linear transformation model for the marginal conditional distribution can be parameterised as

$$\mathbf{c}_j(y_j, \mathbf{x})^\top \boldsymbol{\vartheta}_j = \mathbf{a}_j(y_j)^\top \boldsymbol{\vartheta}_{j,1} - \mathbf{x}^\top \boldsymbol{\beta}_j,$$

with parameters $\boldsymbol{\vartheta}_j = (\boldsymbol{\vartheta}_{j,1}^\top, \boldsymbol{\beta}_j^\top)^\top$. The model restricts the impact of the covariates to a linear shift $\mathbf{x}^\top \boldsymbol{\beta}_j$. For arbitrary choices of $\mathbb{P}_Z$, the marginal distribution, given covariates $\mathbf{X} = \mathbf{x}$, is then a marginal linear transformation model

$$\begin{aligned} \mathbb{P}(Y_j \leq y_j | \mathbf{X} = \mathbf{x}) &= \Phi_{0, \sigma_j^2} \left(\tilde{h}_j(y_j | \mathbf{x}) \right) = \Phi_{0, \sigma_j^2} \left(\Phi_{0, \sigma_j^2}^{-1} (F_Z(\mathbf{a}_j(y_j)^\top \boldsymbol{\vartheta}_{j,1} - \mathbf{x}^\top \boldsymbol{\beta}_j)) \right) \\ &= F_Z(\mathbf{a}_j(y_j)^\top \boldsymbol{\vartheta}_{j,1} - \mathbf{x}^\top \boldsymbol{\beta}_j) \end{aligned}$$

and, consequently, the regression coefficients $\boldsymbol{\beta}_j$ can be directly interpreted as marginal log-odds ratios ($F_Z^{-1} = \text{logit}$) or log-hazard ratios ($F_Z^{-1} = \text{cloglog}$; this is a marginal Cox model). Details of the parameterisations $\mathbf{c}_j(y_j, \mathbf{x})^\top \boldsymbol{\vartheta}_j$ for the marginal transformation functions and a discussion of the practical aspects in different areas of application are provided in Hothorn et al. (2018).

For practical applications, an important and attractive feature of multivariate transformation models for multivariate regression is the possible dependency of $\boldsymbol{\Lambda}$ on covariates $\mathbf{x}$. Thus, the dependence structure of $\mathbf{Y}$ potentially changes as a function of $\mathbf{x}$, if suggested by the data. This feature is implemented by covariate-dependent coefficients of $\boldsymbol{\Lambda}(\mathbf{x})$. A simple linear model of the form

$$\lambda_{jj}(\mathbf{x}) = \alpha_{jj} + \mathbf{x}^\top \boldsymbol{\gamma}_{jj}, \quad 1 \leq j < j \leq J$$

is one option. The case $\boldsymbol{\gamma}_{jj} = \mathbf{0}$ implies that the correlation between Y_j and Y_j does not depend on $\mathbf{x}$. More complex forms of additive models would also be conceivable. Of course,

the number of parameters grows quadratically in J , such that models that are too complex may require additional penalisation terms in the likelihood.

4.2 Application: Trivariate Conditional Transformation Models for Undernutrition in India

To illustrate several practical aspects of the parameterisation and interpretation of MCTMs, we present a trivariate analysis of undernutrition in India in the following. Childhood undernutrition is among the most urgent problems in developing and transition countries. A rich database available from Demographic and Health Surveys (DHS, <https://dhsprogram.com/>) provides nationally representative information about the health and nutritional status of populations in many of those countries. Here we use data from India that were collected in 1998. Overall, the data set comprised 24,316 observations, after pre-processing of the data. For the latter, we use the same steps as in Fahrmeir and Kneib (2011), see the documentation available at <http://www.smoothingbook.org> together with further details on the pre-processing steps. We used three indicators, *stunting*, *wasting* and *underweight*, as the trivariate response vector, where *stunting* refers to stunted growth, measured as an insufficient height of a child with respect to age, while *wasting* and *underweight* refer to insufficient weight for height and insufficient weight for age, respectively. Hence *stunting* is an indicator of chronic undernutrition, *wasting* reflects acute undernutrition and *underweight* reflects both. Our aim was to model the joint distribution of *stunting*, *wasting* and *underweight* conditional upon the age of the child. To the best of our knowledge, there is no implementation available that could estimate the dependence structure and the marginal distributions nonparametrically and conditional on covariates beyond a trivariate normal distribution (which is implemented in the R add-on package (Yee, 2020)).

Model Specification. The focus of our analysis was the variation in the trivariate undernutrition process with respect to the age of the child (in months). To be flexible in the marginal distributions and the dependence structure, we specify response-varying marginal models for $\tilde{h}_j(y_j \mid \text{age})$ of the form

$$\tilde{h}_j(y_j \mid \text{age}) = \mathbf{a}_j(y_j)^\top \boldsymbol{\vartheta}_{j,1} - \mathbf{a}_j(y_j)^\top \boldsymbol{\beta}_j \times \text{age}, \quad j \in \{\text{stunting}, \text{wasting}, \text{underweight}\},$$

while the coefficients of Λ are parameterised through

$$\lambda_{jj}(\text{age}) = \mathbf{b}(\text{age})^\top \boldsymbol{\gamma}_{jj}, \quad j < j \in \{\text{stunting}, \text{wasting}, \text{underweight}\}.$$

Figure 3: Undernutrition. Estimated marginal conditional CDFs $F_j(y_j | \text{age})$, (first row) and marginal densities $f_j(y_j | \text{age})$ (second row), $j \in \{\text{stunting}, \text{wasting}, \text{underweight}\}$ for selected ages in months.

We choose the normal reference distribution $\mathbb{P}_Z = N(0, 1)$ and the basis functions $\mathbf{a}_j$ and $\mathbf{b}(\text{age})$ are Bernstein polynomials of order six (Section 5.1 gives a rational for choosing this default). Furthermore, the parameters $\boldsymbol{\vartheta}_{j,1}$ were estimated under the constraint $\mathbf{D}\boldsymbol{\vartheta}_{j,1} > \mathbf{0}$, where $\mathbf{D}$ is a difference matrix. This leads to monotonically increasing estimated marginal transformation functions $\tilde{h}_j$ (Hothorn et al., 2018). No such shape constraint was applied to functions of age, *i.e.* the parameters $\boldsymbol{\beta}_j$ and $\boldsymbol{\gamma}_{jj}$ were estimated unconstrained.

Results for Marginal Distributions. Figure 3 depicts the estimated marginal conditional CDFs $F_j(y_j | \text{age})$ (first row) and marginal densities $f_j(y_j | \text{age})$ (second row), with the different colours indicating the ages of the children. Clearly, the shapes of the marginals differ for the three indicators, where the differences are mostly restricted to a simple shift effect for stunting, while varying amounts of asymmetry are present for stunting and even more complex changes in the shape of the distribution are identified for underweight.

Results for the Dependence Structure. Figure 4 depicts the conditional rank correlations ρ^S between stunting, wasting and underweight as functions of age along with the point estimates and 95% confidence intervals obtained from $B = 1,000$ parametrically drawn bootstrap samples (see Algorithm 1 in Section 3.4). The rank correlation between stunting

Figure 4: Undernutrition. Spearman’s rho $\rho_{12}^S(\text{age}) = \rho^S(y_{\text{stunting}}, y_{\text{wasting}} | \text{age})$, $\rho_{13}^S(\text{age}) = \rho^S(y_{\text{stunting}}, y_{\text{underweight}} | \text{age})$ and $\rho_{23}^S(\text{age}) = \rho^S(y_{\text{wasting}}, y_{\text{underweight}} | \text{age})$. Shown are the maximum likelihood estimates (solid red line), the bootstrapped mean estimate and the 95% bootstrapped confidence intervals (dashed black lines).

and wasting is initially negative around -0.4 for young children and then approaches zero with increasing age of the child. This finding is in line with the study of Klein, Kneib, Klasen and Lang (2015), who reported the results of a bivariate analysis based on normal and t distributions. In our study, the remaining dependencies were positive and the variation in the rank correlation over age was stronger for the relationship between wasting and underweight compared to stunting and underweight, which varied only between 0.6 and 0.75 whereas the variations of the remaining rank correlations explained by the age of the child were more substantial.

The parametric bootstrap requires re-estimation of the model B times which can be time consuming without parallelisation. A faster alternative in cases where n is large is to draw samples $\hat{\theta}_{(b)}$, $b = 1, \dots, B$ from the asymptotic normal distribution derived in Corollary 2 with mean vector equal to the MLE from (7) and covariance matrix (8). These samples can then be used to compute $\hat{h}_{(1)}, \dots, \hat{h}_{(B)}$, $\hat{\Sigma}_{(1)}, \dots, \hat{\Sigma}_{(B)}$. Because in our example n is rather large we compared both versions of confidence intervals and found them to be very similar (compare Figure A in Appendix Part C).

5 Empirical Evaluation

In this section, we provide empirical evidence on the performance of our MCTMs. In Section 5.1 we demonstrate that the performance of our flexible transformation model was highly competitive relative to two parametric and correctly specified alternatives. Hence, our model is useful not only when parametric assumptions about the marginals are questionable. In Section 5.2, a trivariate example demonstrates that our model is also applicable

to situations beyond the bivariate case, as underpinned by five and ten-dimensional illustrations in Section 5.3. To the best of our knowledge, there are currently no directly competing models that allow for a similar flexibility.

5.1 Bivariate Simulation

Simulation Design. We simulated $R = 100$ data sets of size $n = 1,000$, following a method similar to that used in the parametric bootstrap procedure:

1. Covariate values x were simulated as i.i.d. variables, where $x \sim \text{U}[-0.9, 0.9]$.
2. The latent variables $\tilde{\mathbf{z}}_{ir} \in \mathbb{R}^2$ were generated as

$$\tilde{\mathbf{z}}_{ir} = \mathbf{\Lambda}_i^{-1} \mathbf{z}_{ir}, \quad i = 1, \dots, n; r = 1, \dots, R$$

with

$$\mathbf{z}_{ir} \sim \text{N}(\mathbf{0}, \mathbf{I}_2) \text{ and } \mathbf{\Lambda}_i = \begin{pmatrix} 1 & 0 \\ x_i^2 & 1 \end{pmatrix},$$

such that

$$\text{Cov}(\tilde{z}_{i1}, \tilde{z}_{i2} | x_i) \equiv \mathbf{\Sigma}_i(x_i) = \begin{pmatrix} 1 & -x_i^2 \\ -x_i^2 & 1 + x_i^4 \end{pmatrix}.$$

3. From the latent variables, the observed responses were computed as

$$\mathbf{y}_{ir} = [F_1^{-1}\{\Phi_{0,1}(\tilde{z}_{ir,1})\}, F_2^{-1}\{\Phi_{0,\sigma_{i2}^2}(\tilde{z}_{ir,2})\}]^\top,$$

where $\sigma_{i2}^2 = 1 + x_i^4$ and F_1 and F_2 are the CDFs of two Dagum distributions with parameters $a_1 = \exp(2), b_1 = \exp(1), p_1 = \exp(1.3)$ and $a_2 = \exp(1.8), b_2 = \exp(0), p_2 = \exp(0.9)$, respectively. Note that the CDF of an unconditional Dagum distribution (Kleiber, 1996) reads

$$F(y) = \left(1 + \left(\frac{y}{b}\right)^{-a}\right)^{-p}, \quad \text{for } y > 0, a > 0, b > 0, p > 0.$$

This model specification is equivalent to a Gaussian copula model with Dagum marginals, but by its construction, the first marginal is independent of the covariate x , while the scale parameter b_2 of the second marginal varies as a function of x .

As competitors for MCTMs, we considered Bayesian structured additive distributional regression models (Klein, Kneib, Lang and Sohn, 2015), as implemented in the software package BayesX (Belitz et al., 2015), and vector generalised additive models (VGAM, Yee, 2015), as implemented in the corresponding R add-on package (Yee, 2020). For VGAM

and BayesX, we employed the true specification, *i.e.* a Gaussian copula with correlation parameter $\rho(x_i) = -\lambda(x_i)/\sqrt{1+\lambda(x_i)^2}$ and Dagum marginals, in which the parameter b_2 of the second marginal depends on x but the first marginal as well as the parameters a_2 and p_2 did not. For BayesX, both the predictor for b_2 and the correlation parameter ρ of the Gaussian copula were specified using cubic B-splines with 20 inner knots on an equidistant grid in the range of x with a second-order random walk prior (following suggested default values by Lang and Brezger, 2004); the other parameters of the marginals were estimated as constants.

Because **VGAM** does not allow for simultaneous estimation of the marginals and the dependence structure, we first estimated the Dagum margins with constant parameters a_1, b_1, p_1, a_2, p_2 and covariate-dependent parameters b_2 . The copula predictor was then estimated with plugged-in estimates of the margins, using cubic B-splines according to the **sm.ps** function of the package.

For the multivariate transformation models (denoted as **MCTM-6/6**), we employed Bernstein polynomials of order six (as in Section 4.2) for both the transformation functions ($\tilde{h}_1$ and $\tilde{h}_2$) and the parameter λ . Because of the monotonicity constraints on $\tilde{h}_1$ and $\tilde{h}_2$, the order of the corresponding Bernstein polynomials can be larger without decreasing model performance (Hothorn et al., 2018; Hothorn, 2020). In contrast, too large values of a Bernstein polynomial for λ will result in overly erratic estimates with negative impact on model performance. We demonstrate this effect empirically by two additional MCTMs with order $M = 3$ for λ and with orders $M = 6$ (**MCTM-6/3**), and 12 (**MCTM-12/3**), for the transformation functions $\tilde{h}_1$ and $\tilde{h}_2$.

Measures of Performance. Let $\hat{\lambda}^{(r)}(x)$ be the estimate of the lower triangular element of $\mathbf{\Lambda}^{(r)}$ obtained from data replicate $r = 1, \dots, R$. To evaluate the performance of the three competing methods, we investigated the function estimates $\hat{\lambda}^{(r)}(x)$ relative to $\lambda(x)$ as well as the root mean squared errors $\text{RMSE}(\lambda, \hat{\lambda}^{(r)}) = \sqrt{(\lambda(x_g) - \hat{\lambda}^{(r)}(x_g))^2}$ on a grid of length $G = 100$ within the range of x .

Results. Figure 5 shows the estimates for $\lambda(x) = x^2$ of the 100 simulated data sets for BayesX (first panel), **VGAM** (second panel), and **MCTM** (last three panels). All three models reproduced the general functional form correctly. However, BayesX yielded the most reasonable smoothing properties, while **VGAM** has the wiggliest curves. Although BayesX

Figure 5: Bivariate simulation: Function estimates $\hat{\lambda}(x)$ for the effect $\lambda(x) = x^2$ on the correlation parameter. The black line is the true function and the grey lines are the estimates of the $R = 100$ replicates.

Figure 6: Bivariate simulation: $\text{RMSE}(\lambda(x), \hat{\lambda}(x))$ for BayesX, VGAM and MCTM with Bernstein bases of order 6, 6, and 12 for the marginals and 6, 3, and 3 for $\lambda(x)$ respectively.

and VGAM employed the correct model specification in terms of the parametric distribution assumption for the marginal distributions and the correlation parameter, the performance of MCTM is competitive in terms of the RMSE (Figure 6) without the requirement to either estimate the marginal distributions in a first step and plug the empirical copula data in to obtain the dependence structure (as for VGAM) or to specify predictors for parametric marginal distributions (as for BayesX). Both requirements are restrictive in practice because typically it is impossible to pick the ‘correct’ parametric distribution that exactly matches the marginal distributions of the underlying random variables.

A larger value of $M = 12$ for the transformation functions did not lead to degraded performance, however, a less flexible parameterisation of λ was better able to recover the

quadratic function.

5.2 Trivariate Simulation

Simulation Design. We employed a similar setting as in the previous section, with $R = 100$ data sets of size $n = 1,000$, $x \stackrel{\text{i.i.d.}}{\sim} \text{U}[-0.9, 0.9]$, but Steps 2 and 3 of the simulation design in Section 5.1 were extended to three dimensions. The latent variables $\tilde{\mathbf{z}}_{ir} \in \mathbb{R}^3$ were generated as

$$\tilde{\mathbf{z}}_{ir} = \mathbf{\Lambda}_i^{-1} \mathbf{z}_{ir}, \quad i = 1, \dots, n; r = 1, \dots, R$$

with

$$\mathbf{z}_{ir} \sim \text{N}(\mathbf{0}, \mathbf{I}_3) \text{ and } \mathbf{\Lambda}_i = \begin{pmatrix} 1 & 0 & 0 \\ x_{ir}^2 & 1 & 0 \\ -x_{ir} & x_{ir}^3 - x_{ir} & 1 \end{pmatrix}.$$

Consequently,

$$\text{Cov}(\tilde{z}_{i1}, \tilde{z}_{i2}, \tilde{z}_{i3} | x_i) \equiv \mathbf{\Sigma}_i(x_i) = \begin{pmatrix} 1 & \sigma_{i12} & \sigma_{i13} \\ \sigma_{i12} & \sigma_{i2}^2 & \sigma_{i23} \\ \sigma_{i13} & \sigma_{i23} & \sigma_{i3}^2 \end{pmatrix},$$

with

$$\begin{aligned} \sigma_{i2}^2 &= \lambda_{21}(x_i)^2 + 1 \\ \sigma_{i3}^2 &= (\lambda_{21}(x_i)\lambda_{32}(x_i) - \lambda_{31}(x_i))^2 + \lambda_{32}(x_i)^2 + 1 \\ \sigma_{i12} &= -\lambda_{21}(x_i), \\ \sigma_{i13} &= \lambda_{21}(x_i)\lambda_{32}(x_i) - \lambda_{31}(x_i), \\ \sigma_{i23} &= -\lambda_{21}(x_i)^2\lambda_{32}(x_i) + \lambda_{21}(x_i)\lambda_{31}(x_i) - \lambda_{32}(x_i). \end{aligned}$$

To compute $\mathbf{y}_{ir}$ in Step 3, we additionally chose F_3 to be the CDF of another Dagum distribution with parameters $a_3 = \exp(1.5)$, $b_3 = \exp(-0.9)$ and $p_3 = \exp(1)$, such that

$$\mathbf{y}_{ir} = [F_1^{-1}\{\Phi_{0,1}(\tilde{z}_{ir,1})\}, F_2^{-1}\{\Phi_{0,\sigma_{i2}^2}(\tilde{z}_{ir,2})\}, F_3^{-1}\{\Phi_{0,\sigma_{i3}^2}(\tilde{z}_{ir,3})\}]^\top, \quad i = 1, \dots, n; r = 1, \dots, R.$$

Note that the marginals of y_2 and y_3 (or more precisely their marginal parameters b_2 and b_3) depend on the covariate x .

Results. Figure 7 shows the function estimates for the three parameters $\lambda_{21}(x_i)$, $\lambda_{31}(x_i)$ and $\lambda_{32}(x_i)$. The grey lines indicate overall convincing results for all replicates, without any problematic outliers even though the estimation errors increased with the increasing

complexity of the functional form. We omit the RMSE plot because it yielded qualitatively the same results.

Figure 7: Trivariate simulation. Function estimates of $\lambda_{21}(x)$ (left), $\lambda_{31}(x)$ (middle) and $\lambda_{32}(x)$ (right). The black line is the true function and the grey lines are the estimates of $R = 100$ replicates.

5.3 Higher-Dimensional Responses

To investigate the performance of MCTMs in higher response dimensions, we conducted an experiment with 5- and 10-dimensional responses. Specifically, we employ the settings from Section 5.2 but add 2 and 7 marginally Dagum distributed components assuming independence, *i.e.* $\lambda_{ij}(x) = 0$ for $i > j, i > 3$. The function estimates for the three parameters $\lambda_{21}(x_i)$, $\lambda_{31}(x_i)$ and $\lambda_{32}(x_i)$ is qualitatively similar to the results in Figure 7, see Figure 8, while the true zero components of Λ are identified correctly (not shown in the Figure). Furthermore the scale of the RMSE does not increase but is for all replicates and all $\lambda_{ij}(x) < 0.2$ as before, so that we omit the additional plot. Overall, these results indicate very satisfying results even for high-dimensional response situations.

6 Summary and Discussion

Renewed interest in transformation models (Manuguerra and Heller, 2010; McLain and Ghosh, 2013; Chernozhukov et al., 2013; Hothorn et al., 2014; Liu et al., 2017; Hothorn et al., 2018; Garcia et al., 2019) has been motivated by the combination of model flexibility, parameter interpretability, and broad applicability of this class of regression models. Rather than assuming a specific distribution of the response, transformation models rely on a suitable transformation of the response into an a priori defined reference distribution. The problem of directly estimating a distribution is replaced by the problem of estimating this

Figure 8: 5-dimensional simulation (upper row) and 10-dimensional simulation (lower row). Each column shows function estimates of $\lambda_{21}(x)$ (left), $\lambda_{31}(x)$ (middle) and $\lambda_{32}(x)$ (right). The black line is the true function and the grey lines are the estimates of $R = 100$ replicates.

transformation function. However, in many cases, conceptually and computationally simple solutions to this problem exist (Hothorn, 2020).

The MCTMs introduced herein apply this core principle to multivariate regression. Similar technical approaches have been used in discriminant analysis (Lin and Jeon, 2003, refer to transnormal models), quantile regression (Fan et al., 2016), receiver-operating characteristic curve analysis (Lyu et al., 2019), and are ubiquitous in neural networks as flows, but the generality of multivariate transformation models for regression purposes had yet to be fully developed. MCTMs enjoy the same flexibility, parameter interpretability and broad applicability as their univariate counterparts. The models are highly adaptive and, in our simulation experiments, performed akin to parametric models that exactly matched the data-generating process. While the parametric models are often restricted to bivariate responses, our MCTMs work well far beyond that as illustrated empirically for up to ten dimensions. Appropriate model parameterisations allow both the marginal distributions and the joint distribution to depend on covariates. An important application of this new

model class is the estimation of conditional dependencies, while accounting for covariate effects in the marginal distributions.

Conceptually, our framework carries over to multivariate random vectors that are discrete or censored. In particular, both discrete and censored data can be interpreted as incomplete information $\underline{\mathbf{y}}_i < \mathbf{y}_i \leq \bar{\mathbf{y}}_i$, where, instead of exact observations $\mathbf{y}_i \in \mathbb{R}^J$, only the upper and lower boundaries $\underline{\mathbf{y}}_i$ and $\bar{\mathbf{y}}_i$ respectively, are observed. For discrete data, the underlying rationale would be that discrete realisations are obtained by discretisation from an underlying continuous process. For censored data, the interval boundaries result from the censoring mechanism, and random right censoring, left censoring as well as interval censoring can be handled by appropriate choices of $\underline{\mathbf{y}}_i$ and $\bar{\mathbf{y}}_i$.

For $\mathbb{P}_Z = N(0, 1)$, the log-likelihood contributions are then given by

$$\ell_i(\boldsymbol{\theta}) = \log \left(\int_{\tilde{h}(\underline{\mathbf{y}}_i)}^{\tilde{h}(\bar{\mathbf{y}}_i)} \phi_{\mathbf{0}, \boldsymbol{\Sigma}}(\tilde{\mathbf{z}}) d\tilde{\mathbf{z}} \right),$$

where $\tilde{h}(\mathbf{y}) = (\tilde{h}_1(y_1), \dots, \tilde{h}_J(y_J))$ and $\boldsymbol{\Sigma} = \boldsymbol{\Lambda}^{-1} \boldsymbol{\Lambda}^{-\top}$. Numerical approximations need to be applied in evaluating these likelihood contributions. For $J > 2$, the quasi-Monte-Carlo algorithm by Genz (1992) seems especially appropriate because it relies on the Cholesky factor $\boldsymbol{\Lambda}$ of the precision matrix rather than the covariance matrix $\boldsymbol{\Sigma}$. Nonetheless, the simplicity and explicit structure of the score contributions from the previous sections do not carry over to these more general cases, and the numerical evaluation of the log-likelihood becomes more demanding. We will investigate these challenges in some future work.

More complex models, for example, models featuring additive or spatial effects are conceptually easy to integrate if present in the data because they only require the addition of suitable penalty terms to the log-likelihood. In addition, the analytic expressions for score and Fisher information functions presented herein apply only to a standard normal reference distribution; adaptations to the general case beyond linear dependence structures are still needed.

Computational Details

A reference implementation of conditional and unconditional multivariate transformation models is available in package **tram** (Hothorn et al., 2020). Augmented Lagrangian Minimization implemented in the `auglag()` function of package **alabama** (Varadhan, 2015) was used for optimising the log-likelihood, with starting values obtained from marginal transformation models.

Source code for the reproduction of the empirical results presented in Sections 4 and 5 is distributed as part of this package; the two illustrations can be executed from within R

```
install.packages("tram")
library("tram")
example(mmlt)
demo("undernutrition")
```

Empirical results were obtained using R (version 4.0.2., R Core Team, 2020), a developer version of BayesX (Belitz et al., 2015), **tram** (version 0.5-1, Hothorn et al., 2020), and **VGAM** (version 1.1-3, Yee, 2020). Source code for simulations is available from

```
system.file("simulations", package="tram")
```

Acknowledgements

The authors thank two referees and an associate editor for helpful comments that improved the manuscript. The authors gratefully acknowledge funding through the Emmy Noether grant KL 3037/1-1 (Nadja Klein) from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation), and SNF grant 200021-184603 from the Swiss National Science Foundation (Torsten Hothorn).

Part A Some Properties of the MCTM

Part A.1 Dependence Structure

Let $\mathbf{P} = \mathbf{\Sigma}^{-1} = \mathbf{\Lambda}^\top \mathbf{\Lambda}$ be the precision matrix of the distribution of $\tilde{\mathbf{Z}}$, and $\mathbf{Y}_{-jj}$ and $\tilde{\mathbf{Z}}_{-jj}$ denote the vectors of the observed and transformed responses, excluding elements j and j . Let furthermore $\mathbf{R} = \mathbf{S}\mathbf{\Sigma}\mathbf{S}$, $\mathbf{S} = \text{diag}(\sigma_1^{-1}, \dots, \sigma_J^{-1})$, be the corresponding correlation matrix and C_{Ga}^{jj} the bivariate sub-Gaussian copula of components j and j and correlation matrix entry $R[j, j]$.

Conditional Independence. *The entries in $\mathbf{\Lambda}$ determine the conditional independence structure between the transformed responses $\tilde{Z}_j$ (and therefore, implicitly, also the observed responses Y_j), i.e.*

$$Y_j \perp Y_j \mid \mathbf{Y}_{-jj} \iff \tilde{Z}_j \perp \tilde{Z}_j \mid \tilde{\mathbf{Z}}_{-jj} \iff \mathbf{P}[j, j] = 0$$

Proof. Since $\mathbb{P}(\mathbf{Y} \leq \mathbf{y}) = \mathbb{P}(\tilde{\mathbf{Z}} \leq \tilde{\mathbf{z}})$ holds, the dependence structure of $\mathbf{Y}$ is that of $\tilde{\mathbf{Z}}$. However, $\tilde{\mathbf{Z}} \sim N_J(\mathbf{0}_J, \Sigma)$ with $\Sigma = \Lambda^{-1} \Lambda^{-\top}$. Hence, the result is a direct consequence of Theorem 2.2 of Rue and Held (2005) $\square$

Measures of Dependence.

(i) For $q \in (0, 1)$ and $U_j = F_j(\tilde{h}_j^{-1}(\tilde{Z}_j))$, $U_j = F_j(\tilde{h}_j^{-1}(\tilde{Z}_j))$, the lower and upper quantile dependence are

$$\begin{aligned}\lambda_{jj}^L(q|\Lambda) &\equiv \Pr(U_j < q | U_j < q) = \frac{C_{Ga}^{jj}(q, q)}{q}, \text{ and} \\ \lambda_{jj}^U(q|\Lambda) &\equiv \Pr(U_j > q | U_j > q) = \frac{1 - 2q + C_{Ga}^{jj}(q, q)}{1 - q}.\end{aligned}$$

(ii) The lower and upper extremal tail dependence

$$\lambda_{jj}^L = \lim_{q \downarrow 0} \lambda_{jj}^L(q|\Lambda) = 0, \text{ and } \lambda_{jj}^U = \lim_{q \uparrow 1} \lambda_{jj}^U(q|\Lambda) = 0.$$

(iii) Spearman's rho is

$$\rho_{jj}^S(\Lambda) = \frac{6}{\pi} \arcsin(R[j, j]/2),$$

where $R[j, j]$ is as defined above and is a function of Λ .

(iv) Kendall's tau is

$$\tau_{jj}^K(\Lambda) = \frac{2}{\pi} \arcsin(R[j, j]).$$

Proof. The proofs for (i) and (ii) can be found in Coles et al. (1999), and for (iii) and (iv) in Fang et al. (2002). $\square$

Part A.2 Details on Marginal and Conditional Distributions

Assume that the random vector $\mathbf{Y}$ is partitioned into $\mathbf{Y}_{\mathcal{I}}$ and $\mathbf{Y}_{\mathcal{I}^c}$, where $\mathcal{I} \subset \{1, \dots, J\}$ is a non-empty set of I indices $j_1, \dots, j_I$ and $\mathcal{I}^c = \{1, \dots, J\} \setminus \mathcal{I}$ is its complement, consisting of all indices not contained in $\mathcal{I}$. The vectors $\tilde{\mathbf{Z}}_{\mathcal{I}}$ and $\tilde{\mathbf{Z}}_{\mathcal{I}^c}$ can be similarly defined. The marginal distribution of $\tilde{\mathbf{Z}}_{\mathcal{I}}$ is then given by $N_I(\mathbf{0}_I, \Sigma_{\mathcal{I}})$, where $\Sigma_{\mathcal{I}}$ is the submatrix of Σ containing the elements related to the subset $\mathcal{I}$. We can therefore deduce both the marginal CDF and the density of $\mathbf{Y}_{\mathcal{I}}$ as

$$\mathbb{P}(\mathbf{Y}_{\mathcal{I}} \leq \mathbf{y}_{\mathcal{I}}) = \mathbb{P}(\tilde{\mathbf{Z}}_{\mathcal{I}} \leq \tilde{\mathbf{z}}_{\mathcal{I}}) = \Phi_{\mathbf{0}_I, \Sigma_{\mathcal{I}}}(\tilde{\mathbf{z}}_{\mathcal{I}})$$

where $\tilde{z}_{j_i} = \tilde{h}_{j_i}(y_{j_i})$ and

$$f_{\mathbf{Y}_{\mathcal{I}}}(\mathbf{y}_{\mathcal{I}}) = \phi_{\mathbf{0}_I, \Sigma_{\mathcal{I}}}(\tilde{\mathbf{z}}_{\mathcal{I}}) \prod_{j_i \in \mathcal{I}} \frac{\partial \tilde{h}_{j_i}(y_{j_i})}{\partial y_{j_i}}. \quad (9)$$

We use a similar process for the conditional distribution of $\mathbf{Y}_{\mathcal{I}}$, given $\mathbf{Y}_{\mathcal{I}^c} = \mathbf{y}_{\mathcal{I}^c}$, and note that

$$\mathbb{P}(\mathbf{Y}_{\mathcal{I}} \leq \mathbf{y}_{\mathcal{I}} \mid \mathbf{Y}_{\mathcal{I}^c} = \mathbf{y}_{\mathcal{I}^c}) = \mathbb{P}(\tilde{\mathbf{Z}}_{\mathcal{I}} \leq \tilde{\mathbf{z}}_{\mathcal{I}} \mid \mathbf{Y}_{\mathcal{I}^c} = \mathbf{y}_{\mathcal{I}^c}) = \mathbb{P}(\tilde{\mathbf{Z}}_{\mathcal{I}} \leq \tilde{\mathbf{z}}_{\mathcal{I}} \mid \tilde{\mathbf{Z}}_{\mathcal{I}^c} = \tilde{\mathbf{z}}_{\mathcal{I}^c}).$$

From the rules for multivariate normal distributions we furthermore obtain

$$\tilde{\mathbf{Z}}_{\mathcal{I}} \mid \tilde{\mathbf{Z}}_{\mathcal{I}^c} = \tilde{\mathbf{z}}_{\mathcal{I}^c} \sim \mathcal{N}_I(\boldsymbol{\mu}_{\mathcal{I}|\mathcal{I}^c}, \boldsymbol{\Sigma}_{\mathcal{I}|\mathcal{I}^c})$$

with

$$\boldsymbol{\mu}_{\mathcal{I}|\mathcal{I}^c} = \boldsymbol{\Sigma}_{\mathcal{I},\mathcal{I}^c} \boldsymbol{\Sigma}_{\mathcal{I}^c}^{-1} \tilde{\mathbf{z}}_{\mathcal{I}^c} \text{ and } \boldsymbol{\Sigma}_{\mathcal{I}|\mathcal{I}^c} = \boldsymbol{\Sigma}_{\mathcal{I}} - \boldsymbol{\Sigma}_{\mathcal{I},\mathcal{I}^c} \boldsymbol{\Sigma}_{\mathcal{I}^c}^{-1} \boldsymbol{\Sigma}_{\mathcal{I},\mathcal{I}^c}^{\top}$$

where $\boldsymbol{\Sigma}_{\mathcal{I}^c}$ and $\boldsymbol{\Sigma}_{\mathcal{I},\mathcal{I}^c}$ denote the sub-blocks of $\boldsymbol{\Sigma}$ corresponding to the respective index sets. Therefore, the conditional density of $\mathbf{Y}_{\mathcal{I}}$ given $\mathbf{Y}_{\mathcal{I}^c}$ is

$$f_{\mathbf{Y}_{\mathcal{I}}|\mathbf{Y}_{\mathcal{I}^c}}(\mathbf{y}_{\mathcal{I}}) = \phi_{\boldsymbol{\mu}_{\mathcal{I}|\mathcal{I}^c}, \boldsymbol{\Sigma}_{\mathcal{I}|\mathcal{I}^c}}(\tilde{\mathbf{z}}_{\mathcal{I}}) \prod_{j_i \in \mathcal{I}} \frac{\partial \tilde{h}_{j_i}(y_{j_i})}{\partial y_{j_i}}.$$

Part B Observed Fisher Information

Let $\mathcal{F}_i(\boldsymbol{\theta}) = -\frac{\partial^2 l_i(\boldsymbol{\theta})}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^{\top}} = -\frac{\partial^2 l_i(\boldsymbol{\theta})}{\partial (\boldsymbol{\theta}^{\top}, \boldsymbol{\lambda}^{\top})^{\top} \partial (\boldsymbol{\theta}^{\top}, \boldsymbol{\lambda}^{\top})}$. Then the elements of the observed Fisher information are given by

$$\begin{aligned} \mathcal{F}_i(\lambda_{\tilde{k}k}, \lambda_{\tilde{l}l}) &= -\frac{\partial^2 l_i(\boldsymbol{\theta})}{\partial \lambda_{\tilde{k}k} \partial \lambda_{\tilde{l}l}} \\ &= \begin{cases} \mathbf{a}_k(y_{ik})^{\top} \boldsymbol{\vartheta}_k \mathbf{a}_l(y_{il})^{\top} \boldsymbol{\vartheta}_l & \text{if } \tilde{k} = \tilde{l} \text{ and } l < \tilde{k} \\ 0 & \text{otherwise,} \end{cases} \end{aligned} \quad (10)$$

$$\begin{aligned} \mathcal{F}_i(\lambda_{\tilde{k}k}, \boldsymbol{\vartheta}_l) &= -\frac{\partial^2 l_i(\boldsymbol{\theta})}{\partial \lambda_{\tilde{k}k} \partial \boldsymbol{\vartheta}_l^{\top}} \\ &= \begin{cases} \lambda_{\tilde{k}l} \mathbf{a}_l(y_{il}) \mathbf{a}_k(y_{ik})^{\top} \boldsymbol{\vartheta}_k & \text{if } l < \tilde{k}, l \neq k \\ \mathbf{a}_{\tilde{k}}(y_{i\tilde{k}}) \mathbf{a}_k(y_{ik})^{\top} \boldsymbol{\vartheta}_k & \text{if } \tilde{k} = l \\ \left(\sum_{j=1}^{\tilde{k}-1} \lambda_{\tilde{k}j} \mathbf{a}_j(y_{ij})^{\top} \boldsymbol{\vartheta}_j + \mathbf{a}_{\tilde{k}}(y_{i\tilde{k}})^{\top} \boldsymbol{\vartheta}_{\tilde{k}} \right) \mathbf{a}_k(y_{ik}) & \\ \quad + \lambda_{\tilde{k}k} \mathbf{a}_k(y_{ik}) \mathbf{a}_k(y_{ik})^{\top} \boldsymbol{\vartheta}_k & \text{if } k = l \\ 0 & \text{otherwise,} \end{cases} \end{aligned} \quad (11)$$

Figure A: Undernutrition. Spearman's rho $\rho_{12}^S(\text{age}) = \rho^S(y_{\text{stunting}}, y_{\text{wasting}}|\text{age})$, $\rho_{13}^S(\text{age}) = \rho^S(y_{\text{stunting}}, y_{\text{underweight}}|\text{age})$ and $\rho_{23}^S(\text{age}) = \rho^S(y_{\text{wasting}}, y_{\text{underweight}}|\text{age})$. Shown are the maximum likelihood estimates (solid red line), the bootstrapped mean estimate and the 95% bootstrapped confidence intervals (dashed black lines), as well as mean estimate and 95% confidence intervals (dashed blue lines) based on the asymptotic distribution of $\boldsymbol{\theta}$.

and

$$\begin{aligned} \mathcal{F}_i(\boldsymbol{\vartheta}_{\tilde{k}}, \boldsymbol{\vartheta}_k) &= -\frac{\partial^2 l_i(\boldsymbol{\theta})}{\partial \boldsymbol{\vartheta}_{\tilde{k}} \partial \boldsymbol{\vartheta}_k^\top} \\ &= \begin{cases} \sum_{j=k}^J \lambda_{jk}^2 \mathbf{a}_k(y_{ik}) \mathbf{a}_k(y_{ik})^\top + \frac{\mathbf{a}'_k(y_{ik}) \mathbf{a}'_k(y_{ik})^\top}{(\mathbf{a}'_k(y_{ik})^\top \boldsymbol{\vartheta}_k)^2} & \text{if } k = \tilde{k} \\ \sum_{j=\tilde{k}}^J \lambda_{jk} \mathbf{a}'_k(y_{ik}) \lambda_{j\tilde{k}} \mathbf{a}'_{\tilde{k}}(y_{i\tilde{k}})^\top & \text{if } k \neq \tilde{k} \end{cases} \end{aligned} \quad (12)$$

Part C Fast Alternative to Parametric Bootstrap

As discussed in Section 3.4, a fast alternative to the parametric bootstrap procedure for cases in which n is large is to draw samples $\hat{\boldsymbol{\theta}}_{(b)}$, $b = 1, \dots, B$ from the asymptotic normal distribution of $\boldsymbol{\theta}$ from Corollary 2. We here present the Figure A showing the 95% confidence intervals of the conditional rank correlations ρ^S between stunting, wasting and underweight as functions of age of the parametric bootstrap and the fast alternative.

References

- Acar, E. F., Craiu, R. V. and Yao, F. (2011). Dependence calibration in conditional copulas: A nonparametric approach, *Biometrics* **67**: 445–453.
- Belitz, C., Brezger, A., Klein, N., Kneib, T., Lang, S. and Umlauf, N. (2015). BayesX - Software for Bayesian inference in structured additive regression models. Version 3.0.2. Anonymous, read-only SVN access to the BayesX source code is available via <https://svn.gwdg.de/svn/bayesx/> with user name "anonymous" and empty password. Full details at www.bayesx.org.

- Carlier, G., Chernozhukov, V. and Galichon, A. (2016). Vector quantile regression: An optimal transport approach, *The Annals of Statistics* **44**(3): 1165–1192.
- Carlier, G., Chernozhukov, V. and Galichon, A. (2017). Vector quantile regression beyond the specified case, *Journal of Multivariate Analysis* **161**: 96–102.
- Chernozhukov, V., Fernández-Val, I. and Melly, B. (2013). Inference on counterfactual distributions, *Econometrica* **81**(6): 2205–2268.
- Chernozhukov, V., Galichon, A., Hallin, M. and Henry, M. (2017). Monge-Kantorovich depth, quantiles, ranks and signs, *The Annals of Statistics* **45**(1): 223–256.
- Coles, S., Heffernan, J. and Tawn, J. (1999). Dependence measures for extreme value analyses, *Extremes* **2**: 339–365.
- Curtis, S. M. and Ghosh, S. K. (2011). A variable selection approach to monotonic regression with Bernstein polynomials, *Journal of Applied Statistics* **38**(5): 961–976.
- Ezekiel, M. (1930). *Methods of Correlation Analysis*, Wiley, New York, U.S.A.
- Fahrmeir, L. and Kneib, T. (2011). *Bayesian Smoothing and Regression for Longitudinal, Spatial and Event History Data*, Oxford University Press, New York.
- Fan, J., Xue, L. and Zou, H. (2016). Multitask quantile regression under the transnormal model, *Journal of the American Statistical Association* **111**(516): 1726–1735.
- Fang, H.-B., Fang, K.-T. and Kotz, S. (2002). The meta-elliptical distributions with given marginals, *Journal of Multivariate Analysis* **82**: 1–16.
- Farouki, R. T. (2012). The Bernstein polynomial basis: A centennial retrospective, *Computer Aided Geometric Design* **29**(6): 379–419.
- Garcia, T. P., Marder, K. and Wang, Y. (2019). Time-varying proportional odds model for mega-analysis of clustered event times, *Biostatistics* **20**(1): 129–146.
- Genz, A. (1992). Numerical computation of multivariate normal probabilities, *Journal of Computational and Graphical Statistics* **1**: 141–149.
- Heckman, J. (1978). Dummy endogenous variables in a simultaneous equation system, *Econometrica* **46**: 931–959.

- Hothorn, T. (2020). Most likely transformations: The mlt package, *Journal of Statistical Software* **92**(1): 1–68.
- Hothorn, T., Barbanti, L., Klein, N., Ripley, B., Venables, B. and Bates, D. M. (2020). *tram: Transformation Models*. R package version 0.5-1.
URL: <http://CRAN.R-project.org/package=tram>
- Hothorn, T., Kneib, T. and Bühlmann, P. (2014). Conditional transformation models, *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* **76**: 3–27.
- Hothorn, T., Möst, L. and Bühlmann, P. (2018). Most likely transformations, *Scandinavian Journal of Statistics* **45**(1): 110–134.
- Joe, H. (1997). *Multivariate Models and Dependence Concepts*, Chapman & Hall/CRC, London, UK.
- Kleiber, C. (1996). Dagum vs. Singh-Maddala income distributions, *Economics Letters* **57**: 39–44.
- Klein, N. and Kneib, T. (2016). Simultaneous inference in structured additive conditional copula regression models: A unifying Bayesian approach, *Statistics and Computing* **26**: 841–860.
- Klein, N., Kneib, T., Klasen, S. and Lang, S. (2015). Bayesian structured additive distributional regression for multivariate responses, *Journal of the Royal Statistical Society. Series C (Applied Statistics)* **64**: 569–591.
- Klein, N., Kneib, T., Lang, S. and Sohn, A. (2015). Bayesian structured additive distributional regression with an application to regional income inequality in Germany, *The Annals of Applied Statistics* **9**: 1024–1052.
- Lang, S., Adebayo, S. B., Fahrmeir, L. and Steiner, W. J. (2003). Bayesian geoadditive seemingly unrelated regression, *Computational Statistics* **18**: 263–292.
- Lang, S. and Brezger, A. (2004). Bayesian P-splines, *Journal of Computational and Graphical Statistics* **13**: 183–212.
- Levi, E. and Craiu, R. V. (2018). Bayesian inference for conditional copulas using Gaussian process single index models, *Computational Statistics & Data Analysis* **122**: 115–134.
- Lin, Y. and Jeon, Y. (2003). Discriminant analysis through a semiparametric model, *Biometrika* **90**(2): 379–392.

- Liu, Q., Shepherd, B. E., Li, C. and Harrell, F. E. (2017). Modeling continuous response variables using ordinal regression, *Statistics in Medicine* **36**(27): 4316–4335.
- Lyu, T., Ying, Z. and Zhang, H. (2019). A new semiparametric transformation approach to disease diagnosis with multiple biomarkers, *Statistics in Medicine* **38**(8): 1386–1398.
- Manuguerra, M. and Heller, G. Z. (2010). Ordinal regression models for continuous scales, *International Journal of Biostatistics* **6**(1): 14.
- Marra, G. and Radice, R. (2017). Bivariate copula additive models for location, scale and shape, *Computational Statistics & Data Analysis* **112**: 99–113.
- Marra, G. and Radice, R. (2020). Copula link-based additive models for right-censored event time data, *Journal of the American Statistical Association* **115**(530): 886–895.
- McLain, A. C. and Ghosh, S. K. (2013). Efficient sieve maximum likelihood estimation of time-transformation models, *Journal of Statistical Theory and Practice* **7**(2): 285–303.
- Mosler, K. (2010). Depth statistics, in C. Becker, R. Fried and S. Kuhnt (eds), *Robustness and Complex Data Structures*, Springer-Verlag, Heidelberg, Germany, pp. 17–34.
- Nelson, R. (2006). *An Introduction to Copulas*, 2 edn, Springer-Verlag, New York, U.S.A.
- Papamakarios, G., Nalisnick, E., Rezende, D. J., Mohamed, S. and Lakshminarayanan, B. (2019). Normalizing flows for probabilistic modeling and inference. arXiv:1912.02762.
- R Core Team (2020). *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.
- Rigby, R. A. and Stasinopoulos, D. M. (2005). Generalized additive models for location, scale and shape (with discussion), *Journal of the Royal Statistical Society. Series C (Applied Statistics)* **54**: 507–554.
- Rue, H. and Held, L. (2005). *Gaussian Markov Random Fields*, Chapman & Hall/CRC, New York/Boca Raton.
- Sabeti, A., Wei, M. and Craiu, R. V. (2014). Additive models for conditional copulas, *Stat* **3**(1): 300–312.
- Serfling, R. (2002). Quantile functions for multivariate analysis: approaches and applications, *Statistica Neerlandica* **56**: 214–232.

- Song, P. (2000). Multivariate dispersion models generated from Gaussian copula, *Scandinavian Journal of Statistics* **27**(2): 305–320.
- Sun, T. and Ding, Y. (2019). Copula-based semiparametric regression method for bivariate data under general interval censoring, *Biostatistics* . in press.
- Vaart, A. W. v. d. (1998). *Asymptotic Statistics*, Cambridge University Press, Cambridge, UK.
- Varadhan, R. (2015). *alabama: Constrained Nonlinear Optimization*. R package version 2015.3-1.
URL: <https://CRAN.R-project.org/package=alabama>
- Vatter, T. and Chavez-Demoulin, V. (2015). Generalized additive models for conditional dependence structures, *Journal of Multivariate Analysis* **141**: 147–167.
- Vatter, T. and Nagler, T. (2018). Generalized additive models for pair-copula constructions, *Journal of Computational and Graphical Statistics* **27**(4): 715–727.
- Veraverbeke, N., Omelka, M. and Gijbels, I. (2011). Estimation of a conditional copula and association measures, *Scandinavian Journal of Statistics* **38**(4): 766–780.
- Yee, T. W. (2015). *Vector Generalized Linear and Additive Models: With an Implementation in R*, Springer-Verlag, New York, U.S.A.
- Yee, T. W. (2020). *VGAM: Vector Generalized Linear and Additive Models*. R package version 1.1-3.
URL: <http://CRAN.R-project.org/package=VGAM>
- Zellner, A. (1962). An efficient method of estimating seemingly unrelated regression equations and tests for aggregation bias, *Journal of the American Statistical Association* **57**: 348–368.