




Peer Community Journal

Section: Statistics & Machine Learning

Research article

Published
2026-10-07

Cite as
Torsten Hothorn (2026) *On
Nonparanormal Likelihoods*,
Peer Community Journal,
6: e105.

Correspondence
Torsten.Hothorn@R-project.org

Peer-review
Peer reviewed and
recommended by
PCI Statistics & Machine
Learning,
[https://doi.org/10.24072/pci.
statml.100138](https://doi.org/10.24072/pci.statml.100138)



This article is licensed
under the Creative Commons
Attribution 4.0 License.

On Nonparanormal Likelihoods

Torsten Hothorn ¹

Volume 6 (2026), article e105

<https://doi.org/10.24072/pcjournal.814>

Abstract

Nonparanormal models describe the joint distribution of multivariate responses via latent Gaussian, and thus parametric, copulae while allowing flexible nonparametric marginals. Some aspects of such distributions, for example conditional independence, are formulated parametrically. Other features, such as marginal distributions, can be formulated non- or semiparametrically. Such models are attractive when multivariate normality is questionable but interpretability paramount. Most estimation procedures perform two steps, first estimating the nonparametric part. The copula parameters come second, treating the marginal estimates as known. This is sufficient for some applications. For other applications, e.g. when a semiparametric margin features parameters of interest or when standard errors are important, a simultaneous estimation of all parameters might be more advantageous. We present suitable parameterisations of nonparanormal models, possibly including semiparametric effects, and define four novel nonparanormal log-likelihood functions. In general, the corresponding one-step optimisation problems are shown to be non-convex. In some cases, however, biconvex problems emerge. Several convex approximations are discussed. From a low-level computational point of view, the core contribution is the score function for multivariate normal log-probabilities computed via Genz' procedure. As a demonstration for the versatility of the theoretical and computational framework, we present a series of nonparanormal models for transformation discriminant analysis when some biomarkers are subject to limit-of-detection problems. Possible empirical gains of full maximum likelihood estimation compared to two-step approaches are illustrated in a simulation study targeting semiparametric efficient polychoric correlation analysis where a theoretical benchmark is available.

¹Institut für Epidemiologie, Biostatistik und Prävention Universität Zürich Hirschengraben 84, CH-8001 Zürich, Switzerland

1. Introduction

The multivariate normal distribution comes with a high potential for addiction due to its covariance and precision matrix containing information about marginal and conditional independence, respectively. The fact that many foundations of classical and contemporary multivariate statistics, such as linear or quadratic discriminant analysis, graphical models, or structural equation models, have been defined in terms of this distribution can be explained by these favourable properties. However, normality is the exception rather than the rule in the real world. As an alternative to a full normality detox, statisticians may sacrifice marginal normality while retaining joint normality on some latent scale. This idea has been popularised under different terms, for example as “nonparanormal models” (Liu et al., 2009) or “coordinatewise Gaussianisation” (Mai et al., 2023), but its roots go deeper. For multivariate ordinal variables, Jöreskog, 1994 suggested the estimation of “polychoric correlations” defined by a latent bivariate normal distribution coupled with marginal ordinal probit models. Corresponding concepts in multivariate analysis are discussed in Joe, 1997 and similar principles have been applied in semiparametric copula estimation (Joe, 2005; Klaassen and Wellner, 1997), where marginal parameters are estimated first, followed by a second step of estimating the copula parameters *conditionally* on margins. The rank likelihood (Hermes et al., 2024; Hoff, 2007) does not condition on marginal ranks but treats the marginal distributions as nuisance parameters and focuses on the sole estimation of Gaussian copula parameters.

While such ideas have been very successfully applied for the estimation of dependency structures also in high-dimensional multivariate data, more complex models for both marginal and joint distributions are necessary in many applications. Very much in the spirit of Chen et al., 2006, simultaneous likelihood estimation of and inference for marginal *and* copula parameters in more complex models is our main interest here. The motivation comes from a wide range of applications of the nonparanormal model where the application of a “normalise and forget” scheme is not adequate. For general discrete (Popovic et al., 2018) or mixes of continuous and discrete variables (potentially allowing missing observations in some responses, Christoffersen et al., 2021; Göbner et al., 2024; Pritikin et al., 2018), rank-based approaches are more difficult to justify. Access to the full likelihood covering all model parameters is required in semiparametric discriminant analysis (Mai and Zou, 2015) or for analysing multivariate interval-censored survival data (Ding and Sun, 2022). Most interesting are nonparanormal models where the marginal distributions, for example in multivariate regression models (e.g. in multivariate GLMs or other linear models, Lesaffre and Kaufmann, 1992; Nikoloulopoulos, 2023; Song et al., 2009), or the copula parameters (e.g. in time-varying graphical models, Lu et al., 2018) feature parameters capturing covariate effects. The most striking example necessitating the joint estimation of marginal and copula parameters is a model class for survival analysis under dependent censoring (Deresa and Keilegom, 2023). Here, marginally estimated distributions for time-to-event and time-to-censoring are biased and only the joint model leads to properly identified and estimable parameters. In Section 7, a selection of special applications and models with potential benefits from the contributions to likelihood-based inference made in this work is discussed.

We proceed by suggesting parameterisations of the nonparanormal model for discrete, continuous, and mixed discrete-continuous multivariate responses and derive the several variants of nonparanormal log-likelihoods, including “empirical” or “nonparametric” and “parametric” flavours, and the corresponding score functions. In general, maximum likelihood estimation in this

model class is shown to be non-convex. We discuss convex approximations, which might be useful at least for the computation of starting values. The theory and computational framework presented here allows implementation of a rather general likelihood estimation toolbox for many interesting applications. For some areas in multivariate analysis discussed in Section 7 and beyond, this toolbox helps to setup likelihood and score functions. Coupled with suitable optimisation routines, joint maximum likelihood estimation of all model parameters becomes computationally feasible in a range of relevant problems. As an illustration highlighting the practical potential of this framework, we introduce novel forms of discriminant analyses for evaluating the diagnosis of hepatocellular carcinoma based on only partially observed non-normal biomarker data. Regarding the statistical properties of maximum nonparanormal likelihood estimation, it is demonstrated empirically that copula parameters obtained from optimising the nonparanormal log-likelihood attain the semiparametric efficiency bound derived by Klaassen and Wellner, 1997. A more in-depth analysis of theoretical properties of these estimators is beyond the scope of this paper.

2. The Nonparanormal Model

We jointly observe J response variables $\mathbf{Y} = (Y_1, \dots, Y_J)^\top$ from at least ordered sample spaces $\mathcal{Y}_j \in \mathcal{Y}_j, j = 1, \dots, J$. The nonparanormal (NPN) model $\mathbf{Y} \sim \text{NPN}(\mathbf{h}, \boldsymbol{\Sigma})$ features J monotonically non-decreasing transformation functions $\mathbf{h} = (h_1, \dots, h_J)^\top$, one for each dimension $h_j : \mathcal{Y}_j \rightarrow \mathbb{R}$ and, in addition, a positive semidefinite $J \times J$ covariance matrix $\boldsymbol{\Sigma}$ such that the joint cumulative distribution function can be written in terms of normal probabilities $\mathbb{P}(\mathbf{Y} \leq \mathbf{y}) = \Phi_{\boldsymbol{\Sigma}}(\mathbf{h}(\mathbf{y}))$, where $\Phi_{\boldsymbol{\Sigma}}$ is the joint cumulative distribution function of $\mathbf{N}_J(\mathbf{0}, \boldsymbol{\Sigma})$. In case all elements of $\mathbf{Y}$ are continuous, each h_j is bijective and one typically (Liu et al., 2009) writes $\text{NPN}(\mathbf{h}^{-1}, \boldsymbol{\Sigma})$ for the absolutely continuous distribution of $\mathbf{Y} = \mathbf{h}^{-1}(\mathbf{Z})$ generated by a latent multivariate normal variable $\mathbf{Z} = \mathbf{h}(\mathbf{Y}) \sim \mathbf{N}_J(\mathbf{0}, \boldsymbol{\Sigma})$. The copula parameters $\boldsymbol{\Sigma}$ are identified as a consequence of the uniqueness of the copula in Sklar's theorem. We allow more general sample spaces for binary, ordered, count, or otherwise discrete variables and mixed continuous-discrete variables and thus neither require $\mathcal{Y}_j = \mathbb{R}$ for $j = 1, \dots, J$ nor the existence of $\mathbf{h}^{-1}$. In this discrete situation, identifiability issues arise with two-step copula estimators, however, maximum likelihood estimation has been demonstrated to work in special cases (Section 6.2., Genest and Nešlehová, 2007).

The model is invariant with respect to rescaling, that is

$$\text{NPN}(\mathbf{h}, \boldsymbol{\Sigma}) = \text{NPN}(\text{diag}(\mathbf{d})^{-1}\mathbf{h}, \text{diag}(\mathbf{d})\boldsymbol{\Sigma}\text{diag}(\mathbf{d}))$$

for all $\mathbf{d} = (d_1, \dots, d_J)^\top \in \mathbb{R}^J$ with $d_j > 0$ for all $j = 1, \dots, J$ and $\text{diag}(\mathbf{d})$ the $J \times J$ diagonal matrix. Thus, identifiability constraints on $\boldsymbol{\Sigma}$ are needed. One option is to require $\Sigma_{jj} \equiv 1$ for $j = 1, \dots, J$ leading to the interpretation of h_j as probit-transformed marginal distribution function $\mathbb{P}(Y_j \leq y_j) = \Phi(h_j(y_j))$ for all $y_j \in \mathcal{Y}_j, j = 1, \dots, J$.

Alternatively, we write $\boldsymbol{\Sigma} = \boldsymbol{\Omega}^{-1}\boldsymbol{\Omega}^{-\top}$ in terms of the inverse lower triangular Cholesky factor $\boldsymbol{\Omega}^{-1}$ of the covariance matrix $\boldsymbol{\Sigma}$ and require $\Omega_{jj} \equiv 1$ for $j = 1, \dots, J$. This implies $\Sigma_{11} \equiv 1$ and $\Sigma_{jj} \geq 1$ for $j = 2, \dots, J$ and we define $\Phi_{\boldsymbol{\Omega}} := \Phi_{\boldsymbol{\Omega}^{-1}\boldsymbol{\Omega}^{-\top}}$. In the absolutely continuous case with $\mathbf{Y} \in \mathbb{R}^J$, the model $\mathbf{Y} \sim \text{NPN}(\mathbf{h}, \boldsymbol{\Omega}^{-1}\boldsymbol{\Omega}^{-\top})$ is identical to a structural equation model defined by a series of additive transformation models beginning with the marginal model

$$\mathbb{P}(Y_1 \leq y_1) = \Phi(h_1(y_1))$$

and proceeding with all $J - 1$ induced additive regression models, that is, the conditional models

$$\mathbb{P}(Y_j \leq y_j \mid Y_1 = y_1, \dots, Y_{j-1} = y_{j-1}) = \Phi \left(\sum_{j=1}^j \Omega_{jj} h_j(y_j) \right), \quad j = 2, \dots, J$$

for any $(y_1, \dots, y_J)^\top \in \mathbb{R}^J$. In other words, the rows of Ω contain regression coefficients for the conditional distribution of Y_j given its predecessors and the transformation $h_j(Y_j)$ is *conditionally* standard normal. For exclusively binary outcomes $\mathcal{Y}_j = \{0, 1\}, j = 1, \dots, J$ we have $\mathbb{P}(Y_j = 0 \mid j = 1, \dots, J) = \Phi_{\Sigma}(\theta)$, where $\theta = (\theta_1, \dots, \theta_J)^\top \in \mathbb{R}^J$ with $\theta_j = \sqrt{\Sigma_{jj}} \Phi^{-1}(\mathbb{P}(Y_j = 0))$.

In the presence of covariates $\mathbf{X} = \mathbf{x} \in \mathcal{X}$, one can characterise the model via the conditional joint cumulative distribution function

$$(1) \quad \mathbb{P}(\mathbf{Y} \leq \mathbf{y} \mid \mathbf{X} = \mathbf{x}) = \Phi_{\Sigma(\mathbf{x})}(\mathbf{h}(\mathbf{y} \mid \mathbf{x}))$$

where the covariates impact the transformation functions $\mathbf{h}(\mathbf{y} \mid \mathbf{x})$, the covariance matrix $\Sigma(\mathbf{x})$, or both. For example, simple probit linear transformation models

$$(2) \quad h_j(y_j \mid \mathbf{x}) = h_j(y_j) - \mathbf{x}^\top \boldsymbol{\beta}_j, \quad j \in \{1, \dots, J\}$$

feature linear covariate effects. More elaborate choices include transformation functions $h_j(y_j \mid \mathbf{x})$ of the form $\Phi^{-1}(F_j(h_j(y_j) - \mathbf{x}^\top \boldsymbol{\beta}_j))$, where $F_j: \mathbb{R} \rightarrow [0, 1]$ denotes an absolutely continuous distribution function with log-concave density. For example, a marginal Weibull model can be formulated via the inverse complementary log-log ($F_j = \text{cloglog}^{-1}$) and a log-linear function $h_j(y_j)$ (see Table 1 in Hothorn et al., 2018). Also the joint distribution might change with $\mathbf{x}$, for example via linear models for the off-diagonal elements of the inverse Cholesky factor

$$(3) \quad \Omega_{jj}(\mathbf{x}) = \begin{cases} 1 & 1 \leq j = J \leq J \\ \alpha_{jj} + \mathbf{x}^\top \boldsymbol{\gamma}_{jj} & 1 \leq j < j \leq J. \end{cases}$$

In the context of multivariate transformation models, such a parameterisation has been proposed by Klein et al., 2022. For multivariate normal distributions, the same idea was applied by Barratt and Boyd, 2023. For the sake of notational simplicity, we will consider the unconditional case in Sections 3 to 6 and comment on such conditional extensions in Section 7.

3. Parameterisation

The term “nonparanormal” insinuates a combination of nonparametrically parameterised marginal distributions with a parametric Gaussian copula. As a gold standard, we therefore first derive the “nonparanormal” log-likelihood via nonparametric margins from N independent samples $\mathbf{Y}_1, \dots, \mathbf{Y}_N \sim \text{NPN}(\mathbf{h}, \Sigma)$ with realisations $\mathbf{Y}_i = (Y_{i1}, \dots, Y_{iJ})^\top$ whose elements were drawn from the corresponding sample space $Y_{ij} \in \mathcal{Y}_j, i = 1, \dots, N; j = 1, \dots, J$.

In the model $\text{NPN}(\mathbf{h}, \Sigma)$, we first parameterise the $j = 1, \dots, J$ transformation functions h_j . In each dimension $j = 1, \dots, J$, we consider the “empirical” sample space given by the ordered unique realisations $\mathbf{v}_j = \{v_{j1}, \dots, v_{jK(j)}\} \subseteq \mathcal{Y}_j$ with $v_{j,k-1} < v_{j,k}$ for $k = 2, \dots, K(j)$. For the i th observation in the j th variable, write $r(i, j) \in \{1, \dots, K(j)\}$ such that $Y_{ij} = v_{j,r(i,j)}$. In the absence of ties, $r(i, j)$ is the rank of the i th observation in the sample $Y_{1j}, \dots, Y_{Nj}$. We can now parameterise the transformation function h_j as a step function $h_j(v_{jk}) = \theta_{jk} \in \mathbb{R}$ for $k = 0, \dots, K(j)$ with values $\theta_{j0} \equiv -\infty$ and $\theta_{jK(j)} \equiv \infty$ at the boundaries. The j th marginal parameter vector

$\boldsymbol{\theta}_j = (\theta_{j1}, \dots, \theta_{j,K(j)-1})^\top \in \mathbb{R}^{K(j)-1}$ comes with a monotonicity constraint $\mathbf{D}_j \boldsymbol{\theta}_j \geq \mathbf{0}_{K(j)-2}$ defined by the $(K(j)-2) \times (K(j)-1)$ first order difference matrix $\mathbf{D}_j$. Finally, we collect all marginal parameters in the parameter vector $\boldsymbol{\theta} = (\boldsymbol{\theta}_1^\top, \dots, \boldsymbol{\theta}_J^\top)^\top \in \mathbb{R}^{\sum_{j=1}^J (K(j)-1)}$ fully specifying $\mathbf{h}$. Because the number of parameters $K(j)$ grows with the sample size whenever the sample space is infinite, the situation here is typically referred to as being “nonparametric”, although the vector $\boldsymbol{\theta}$ contains parameters, in fact, potentially many of them. The connection to the “parametric” situation will be made in the next section.

Second, we parameterise the inverse Cholesky factor $\boldsymbol{\Omega}$ of the covariance matrix $\boldsymbol{\Sigma} = \boldsymbol{\Omega}^{-1} \boldsymbol{\Omega}^{-\top}$ by defining a lower triangular unit matrix $\boldsymbol{\Lambda} = \boldsymbol{\Lambda}(\boldsymbol{\lambda})$ in terms of its unconstrained lower triangular elements $\boldsymbol{\lambda} = (\lambda_{21}, \lambda_{31}, \dots, \lambda_{J,J-1})^\top \in \mathbb{R}^{J(J-1)/2}$. The first option to ensure parameter identifiability is to write $\boldsymbol{\Omega} = \boldsymbol{\Omega}^{(1)}(\boldsymbol{\lambda}) = \boldsymbol{\Lambda}$ to obtain $\Omega_{jj} \equiv 1$ for $j = 1, \dots, J$ was already mentioned. As a second option, we can write $\boldsymbol{\Omega} = \boldsymbol{\Omega}^{(2)}(\boldsymbol{\lambda}) = \boldsymbol{\Lambda} \text{diag}(\boldsymbol{\Lambda}^{-1} \boldsymbol{\Lambda}^{-\top})^{1/2}$ such that $\boldsymbol{\Sigma} = \boldsymbol{\Omega}^{-1} \boldsymbol{\Omega}^{-\top}$ is equal to $\text{diag}(\boldsymbol{\Lambda}^{-1} \boldsymbol{\Lambda}^{-\top})^{-1/2} \boldsymbol{\Lambda}^{-1} \boldsymbol{\Lambda}^{-\top} \text{diag}(\boldsymbol{\Lambda}^{-1} \boldsymbol{\Lambda}^{-\top})^{-1/2}$ ensuring the condition $\Sigma_{jj} \equiv 1, j = 1, \dots, J$. Consequently, $h_j(Y_j)$ is now *marginally* standard normal in the absolute continuous case. The transformation to conditional standard normality arises from $\boldsymbol{\Omega}_{jj}^{(2)} h_j(Y_j)$. In the following we refer to these two options as $\boldsymbol{\Omega}^{(s)}$ for $s = 1, 2$. In either case, $\boldsymbol{\Sigma}$ is parameterised in terms of the $J(J-1)/2$ lower triangular parameters $\boldsymbol{\lambda}$ of $\boldsymbol{\Lambda}$ and is, for all values of $\boldsymbol{\lambda}$, symmetric and positive semidefinite. Furthermore, if $\lambda_{jj} = 0$, then Y_j and Y_j are conditionally independent given the remaining variables in $Y_1, \dots, Y_j$ in both parameterisations (the reverse might not be true when some margins are discrete, see Genest and Nešlehová, 2007, Section 4.1.).

4. Nonparanormal Log-likelihoods

In the following, we use the term “likelihood” in its broadest sense, including Fishers original *definition* of it as being proportional to the probability of the observed data and *approximations* thereof in terms of absolutely continuous densities (Lindsey, 1999, Section 5). This point of view allows us to choose between “empirical” or “nonparametric” and “parametric” variants (or combinations thereof) motivated by practical necessity and computational convenience.

Before deriving the joint log-likelihood for all J variables, we consider the likelihood for $\boldsymbol{\theta}_j$, that is, the parameters defining the j th marginal distribution. For the absolutely continuous case recall that the empirical or nonparametric log-likelihood (Owen, 2001) given by

$$(4) \quad \ell_j(\boldsymbol{\theta}_j) = \sum_{i=1}^N \log(\Phi(\theta_{j,r(i,j)}) - \Phi(\theta_{j,r(i,j)-1})) = \sum_{i=1}^N \log \left(\int_{\theta_{j,r(i,j)-1}}^{\theta_{j,r(i,j)}} \phi(z) dz \right)$$

leads to a convex problem whose analytical solution $\Phi(\hat{\theta}_{j,r(i,j)}) = r(i,j)/N$ is identical to the empirical cumulative distribution function evaluated at $\theta_{j,r(i,j)}$. From this perspective, many “nonparametric” estimators can be motivated as maximum *empirical* likelihood estimators and we will follow this path with the nonparanormal log-likelihood to be introduced in Formula (6).

Furthermore, assume we had directly observed the latent multivariate normal variables $\mathbf{Z}_i \sim N_J(\mathbf{0}, \boldsymbol{\Omega}^{-1} \boldsymbol{\Omega}^{-\top})$ with absolute continuous density

$$(5) \quad \phi(\mathbf{z} | \boldsymbol{\Omega}) = \exp \left(-\frac{J}{2} \log(2\pi) + \sum_{j=1}^J \log(\Omega_{jj}) - \frac{1}{2} \|\boldsymbol{\Omega} \mathbf{z}\|_2^2 \right)$$

$\text{diag}(\boldsymbol{\Omega}) > \mathbf{0}, \boldsymbol{\Omega} \in \mathbb{R}^{J \times J}$ lower triangular

for $i = 1, \dots, N$. Then, the parametric negative log-likelihood $-\tilde{\ell}_J^{(0)}(\Omega) = -\sum_{i=1}^N \tilde{\ell}_{J,i}^{(0)}(\Omega)$ with

$$\tilde{\ell}_{J,i}^{(0)}(\Omega) = \log(\phi(\mathbf{Z}_i | \Omega)) \propto -\frac{1}{2} \|\Omega \mathbf{Z}_i\|_2^2 + \sum_{j=1}^J \log(\Omega_{jj})$$

is convex in Ω (Barratt and Boyd, 2023). In this section, we leverage both principles to define a novel log-likelihood for the NPN model.

4.1. The Nonparanormal Log-likelihood

The nonparanormal log-likelihood for all J variables is a direct extension of the bivariate log-likelihood for ordinal data proposed by Jöreskog, 1994. By replacing the univariate standard normal density ϕ in the nonparametric log-likelihood (4) with the J -dimensional density $\phi(\mathbf{z} | \Omega)$ of $N_J(\mathbf{0}, \Omega^{-1} \Omega^{-\top})$ while keeping the integration limits for the j th dimension in a J -dimensional integral, we define the nonparanormal log-likelihood by $\ell_J^{(s)}(\theta, \lambda) = \sum_{i=1}^N \ell_{J,i}^{(s)}(\theta, \lambda)$ with

$$\begin{aligned} (6) \quad \ell_{J,i}^{(s)}(\theta, \lambda) &= \log \left(\mathbb{P}_{\theta, \lambda} \left(\bigcap_{j=1}^J \theta_{j,r(i,j)-1} < h_j(Y_{ij}) \leq \theta_{j,r(i,j)} \right) \right) \\ &= \log \left(\mathbb{P}_{\Omega^{(s)}(\lambda)} (h(\mathbf{Y}_i) \in \mathcal{B}_i(\theta)) \right) \\ &= \log \left(\int_{\mathcal{B}_i(\theta)} \phi(\mathbf{z} | \Omega^{(s)}(\lambda)) d\mathbf{z} \right), \quad s \in \{1, 2\} \end{aligned}$$

where $\mathcal{B}_i(\theta) = \{\mathbf{z} \in \mathbb{R}^J \mid \theta_{j,r(i,j)-1} < z_j \leq \theta_{j,r(i,j)}; j = 1, \dots, J\}$. We refer to this nonparanormal log-likelihood $\ell_J^{(s)}(\theta, \lambda)$ as “NPN log-likelihood”.

4.2. The Smooth Nonparanormal Log-likelihood

The number of parameters $K(j)$ grows with N for infinite sample spaces $\mathcal{Y}_j$ and one might want to reduce the number of parameters in such cases. For $\mathcal{Y}_j \subseteq \mathbb{R}$, define $\theta_{jk} = h_j(v_{jk} | \boldsymbol{\vartheta}_j) = \mathbf{a}_j(v_{jk})^\top \boldsymbol{\vartheta}_j$ in terms of a spline basis $\mathbf{a}_j : \mathcal{Y}_j \rightarrow \mathbb{R}^{P(j)}$ and corresponding coefficients $\boldsymbol{\vartheta}_j \in \mathbb{R}^{P(j)}$, potentially under some constraint $\mathbf{D}_j \boldsymbol{\vartheta}_j \geq \mathbf{0}$. Typically, $P(j) < K(j)$. For finite discrete sample spaces $\mathcal{Y}_j$, we use the same notation with $\vartheta_{jk} = \theta_{jk}$ and $h_j(v_{jk} | \boldsymbol{\vartheta}_j) = \mathbf{e}_{K(j)}(k)^\top \boldsymbol{\vartheta}_j$, where $\mathbf{a}_j(v_{jk}) = \mathbf{e}_{K(j)}(k)$ denotes the unit vector of length $K(j)$ with non-zero element k and $P(j) = K(j)$. For a fixed sample size, the empirical sample space discussed in Section 3 is finite discrete and we can thus establish a one-to-one connection between discrete and nonparametrically parameterised continuous models. Motivations for and examples of such parameterisations can be found in Hothorn et al., 2018. Let $\boldsymbol{\theta}_j = \boldsymbol{\theta}_j(\boldsymbol{\vartheta}_j)$ and $\boldsymbol{\theta}(\boldsymbol{\vartheta}) = (\boldsymbol{\theta}_1(\boldsymbol{\vartheta}_1)^\top, \dots, \boldsymbol{\theta}_J(\boldsymbol{\vartheta}_J)^\top)^\top$ for $\boldsymbol{\vartheta} = (\boldsymbol{\vartheta}_1, \dots, \boldsymbol{\vartheta}_J)^\top$ and define the log-likelihood $\ell_J^{(s)}(\boldsymbol{\vartheta}, \lambda) := \ell_J^{(s)}(\boldsymbol{\theta}(\boldsymbol{\vartheta}), \lambda)$. Because the bases $\mathbf{a}_j$ and thus the transformations $\mathbf{a}_j(y_j)^\top \boldsymbol{\vartheta}_j$ are smooth in y_j , we refer to the log-likelihood $\ell_J^{(s)}(\boldsymbol{\vartheta}, \lambda)$ as “smooth NPN log-likelihood”, allowing us to estimate a fully parametric model featuring a fixed number of parameters by maximising an empirical likelihood.

4.3. The Flow Nonparanormal Log-likelihood

When all response variables are absolutely continuous, we can approximate the smooth NPN log-likelihood defined through log-probabilities by the corresponding multivariate log-densities.

The multivariate density function giving rise to the multivariate distribution function of the non-paranormal model

$$\mathbb{P}(\mathbf{Y} \leq \mathbf{y}) = \Phi_{\Omega}(\mathbf{h}(\mathbf{y})) = \int_{-\infty}^{\mathbf{h}(\mathbf{y})} \phi(\mathbf{z} | \Omega) d\mathbf{z} = \int_{-\infty}^{\mathbf{y}} \phi(\mathbf{h}(\tilde{\mathbf{y}}) | \Omega) \det(\mathbf{h}'(\tilde{\mathbf{y}})) d\tilde{\mathbf{y}}$$

motivates the approximate log-likelihood

$$\begin{aligned} \tilde{\ell}_J^{(s)}(\boldsymbol{\vartheta}, \boldsymbol{\lambda}) &= \sum_{i=1}^N \log \left(\phi \left((h_1(Y_{i1} | \boldsymbol{\vartheta}_1), \dots, h_J(Y_{iJ} | \boldsymbol{\vartheta}_J))^{\top} | \Omega^{(s)}(\boldsymbol{\lambda}) \right) \right) + \\ &\quad \sum_{j=1}^J \log(h'_j(Y_{ij} | \boldsymbol{\vartheta}_j)), \quad s \in \{1, 2\} \end{aligned}$$

where $h_j(Y_{ij} | \boldsymbol{\vartheta}_j) = \mathbf{a}_j(Y_{ij})^{\top} \boldsymbol{\vartheta}_j$ and $h'_j(Y_{ij} | \boldsymbol{\vartheta}_j) = \mathbf{a}'_j(Y_{ij})^{\top} \boldsymbol{\vartheta}_j$ (Hothorn et al., 2018). Because $\Omega \mathbf{h}(\mathbf{Y}) \sim N_J(\mathbf{0}, \mathbf{I})$ is a simple normalising flow (Papamakarios et al., 2021), we use the term “flow NPN log-likelihood” for $\tilde{\ell}_J^{(s)}(\boldsymbol{\vartheta}, \boldsymbol{\lambda})$.

4.4. The Mixed Nonparanormal Log-likelihood

If some response variables are discrete and some absolutely continuous, one can approximate the absolutely continuous parts by the corresponding flow NPN log-likelihood in a mixed continuous-discrete log-likelihood. Without loss of generality, assume that the first $1 \leq j < J$ variables Y_j are absolutely continuous and the remaining $J - j$ variables are discrete. We first partition the inverse Cholesky factor

$$\Omega = \begin{pmatrix} \Omega_A & \mathbf{0} \\ \Omega_B & \Omega_C \end{pmatrix}$$

with the continuous $\Omega_A \in \mathbb{R}^{j \times j}$ and discrete $\Omega_C \in \mathbb{R}^{(J-j) \times (J-j)}$ parts being lower triangular and the full matrix $\Omega_B \in \mathbb{R}^{(J-j) \times j}$ representing the interplay between continuous and discrete variables. We then obtain $(Y_1, \dots, Y_J)^{\top} \sim \text{NPN}((h_1, \dots, h_J)^{\top}, \Omega_A^{-1} \Omega_A^{-\top})$, a NPN model for the continuous part with flow NPN log-likelihood $\tilde{\ell}_j^{(s)}(\boldsymbol{\vartheta}_1, \dots, \boldsymbol{\vartheta}_j, \lambda_{21}, \dots, \lambda_{J,j-1})$, and the conditional distribution of discrete given continuous variables

$$Y_{j+1}, \dots, Y_J | Y_1 = y_1, \dots, Y_j = y_j \sim \text{NPN}((h_{j+1}, \dots, h_J)^{\top} - \boldsymbol{\mu}, \Omega_C^{-1} \Omega_C^{-\top})$$

with $\boldsymbol{\mu} = -\Omega_C^{-1} \Omega_B (h_1(y_1), \dots, h_j(y_j))^{\top} \in \mathbb{R}^{J-j}$, that is, a NPN model for the discrete part given the realisations of the continuous variables. The log-likelihood contribution of all variables is then the sum of $\tilde{\ell}_{j,i}^{(s)}(\boldsymbol{\vartheta}_1, \dots, \boldsymbol{\vartheta}_j, \lambda_{21}, \dots, \lambda_{J,j-1})$ and the term

$$\begin{aligned} \ell_{j,i|J}^{(s)}(\boldsymbol{\vartheta}_1, \dots, \boldsymbol{\vartheta}_j, \boldsymbol{\theta}_{j+1}, \dots, \boldsymbol{\theta}_J, \lambda_{j+1,1}, \dots, \lambda_{J,j-1}) = \\ \log \left(\int_{\mathcal{B}_i(\boldsymbol{\theta}_{j+1}, \dots, \boldsymbol{\theta}_J)} \phi(\mathbf{z} - \boldsymbol{\mu} | \Omega_C^{(s)}) d\mathbf{z} \right), \quad s \in \{1, 2\} \end{aligned}$$

where $\mathcal{B}_i(\boldsymbol{\theta}_{j+1}, \dots, \boldsymbol{\theta}_J) = \{\mathbf{z} \in \mathbb{R}^{J-j} | \theta_{j,r(i,j)-1} < z_j \leq \theta_{j,r(i,j)}; j = j+1, \dots, J\}$. Here, $\boldsymbol{\mu} = \boldsymbol{\mu}(\boldsymbol{\vartheta}_1, \dots, \boldsymbol{\vartheta}_j, \lambda_{j+1,1}, \dots, \lambda_{J,j-1}) = -\Omega_C^{(s)-1} \Omega_B^{(s)} (h_1(y_1 | \boldsymbol{\vartheta}_1), \dots, h_j(y_j | \boldsymbol{\vartheta}_j))^{\top} \in \mathbb{R}^{J-j}$ depends on $\Omega_B^{(s)}$ and $\Omega_C^{(s)}$ which, in turn, depend on $(\lambda_{j+1,1}, \dots, \lambda_{J,j-1})$.

In total, we can write the log-likelihood function for all, discrete and continuous, response variables as

$$\tilde{\ell}_{j|j}^{(s)}(\boldsymbol{\vartheta}_1, \dots, \boldsymbol{\vartheta}_j, \boldsymbol{\theta}_j, \dots, \boldsymbol{\theta}_J, \boldsymbol{\lambda}) = \sum_{i=1}^N \tilde{\ell}_{j,i}^{(s)}(\boldsymbol{\vartheta}_1, \dots, \boldsymbol{\vartheta}_j, \lambda_{21}, \dots, \lambda_{j,j-1}) + \ell_{j,i|j}^{(s)}(\boldsymbol{\vartheta}_1, \dots, \boldsymbol{\vartheta}_j, \boldsymbol{\theta}_{j+1}, \dots, \boldsymbol{\theta}_J, \lambda_{j+1,1}, \dots, \lambda_{j,j-1})$$

and we refer to this form of the log-likelihood as “mixed NPN log-likelihood”.

In summary, we defined the NPN log-likelihood $\ell_j^{(s)}(\boldsymbol{\theta}, \boldsymbol{\lambda})$ and three approximations thereof. The smooth NPN log-likelihood $\ell_j^{(s)}(\boldsymbol{\theta}(\boldsymbol{\vartheta}), \boldsymbol{\lambda})$ computes log-probabilities based on smooth transformations, the flow NPN log-likelihood $\tilde{\ell}_j^{(s)}(\boldsymbol{\vartheta}, \boldsymbol{\lambda})$ for absolute continuous responses approximates log-probabilities by log-densities, and the mixed NPN log-likelihood, given by the term $\tilde{\ell}_{j|j}^{(s)}(\boldsymbol{\vartheta}_1, \dots, \boldsymbol{\vartheta}_j, \boldsymbol{\theta}_j, \dots, \boldsymbol{\theta}_J, \boldsymbol{\lambda})$, mixes the flow NPN log-likelihood defined by log-densities of the marginal distribution for $j = 1, \dots, J$ with log-probabilities for the remaining elements (which again can be in form of a NPN log-likelihood or a smooth NPN log-likelihood). Each of these log-likelihoods can be coupled with either constraint $s = 1$ (unit diagonal in $\boldsymbol{\Omega}$) or $s = 2$ ($\boldsymbol{\Sigma}$ being a correlation matrix). Ways to enhance these log-likelihoods to covariate effects in $\mathbf{h}$ or $\boldsymbol{\Omega}$ are discussed in Section 7.

4.5. Evaluation of Log-likelihood and Score Functions

Computing the flow NPN log-likelihood involves simple matrix multiplications whose gradient with respect to $\boldsymbol{\vartheta}$ is

$$\frac{\partial \tilde{\ell}_{j,i}^{(s)}(\boldsymbol{\vartheta}, \boldsymbol{\lambda})}{\partial \boldsymbol{\vartheta}_j} = -\frac{1}{2} \boldsymbol{\Omega}^{(s)}(\boldsymbol{\lambda})^\top \boldsymbol{\Omega}^{(s)}(\boldsymbol{\lambda}) (\mathbf{a}_1(Y_{i1})^\top \boldsymbol{\vartheta}_1, \dots, \mathbf{a}_J(Y_{iJ})^\top \boldsymbol{\vartheta}_J)^\top \mathbf{a}_j(Y_{ij})^\top + \frac{\mathbf{a}'(Y_{ij})^\top}{\mathbf{a}_j'(Y_{ij})^\top \boldsymbol{\vartheta}_j}$$

both for $s = 1$ and $s = 2$. For $s = 1$, the score function with respect to $\boldsymbol{\lambda}$ is

$$\frac{\partial \tilde{\ell}_{j,i}^{(1)}(\boldsymbol{\vartheta}, \boldsymbol{\lambda})}{\partial \lambda_{jj}} = -\left(\boldsymbol{\Lambda}(\mathbf{a}_1(Y_{i1})^\top \boldsymbol{\vartheta}_1, \dots, \mathbf{a}_J(Y_{iJ})^\top \boldsymbol{\vartheta}_J)^\top (\mathbf{a}_1(Y_{i1})^\top \boldsymbol{\vartheta}_1, \dots, \mathbf{a}_J(Y_{iJ})^\top \boldsymbol{\vartheta}_J)^\top \right)_{jj}$$

for $1 < j < J$. The case of $s = 2$ is more elaborate, the score function is derived in Hothorn, 2026b listed in the supplementary material, see Section 10.

The NPN log-likelihood, smooth NPN log-likelihood, and mixed NPN log-likelihood require the evaluation of multivariate normal probabilities over boxes $\mathcal{B}_i$ and the algorithm by Genz, 1992 has been widely applied to approximate such probabilities by quasi-Monte-Carlo integration. The only attempt to also approximate the score function for these log-probabilities we are aware of was described for the special case of binary outcomes by Christoffersen et al., 2021.

In this simplest form, for $\mathcal{B}_i = \{\mathbf{z} \in \mathbb{R}^J \mid \underline{\mathbf{b}} < \mathbf{z} \leq \bar{\mathbf{b}}\}$, the probability defining the NPN log-likelihood contribution is approximated as

$$\begin{aligned} \exp\left(\ell_{j,i}^{(1)}(\boldsymbol{\theta}, \boldsymbol{\lambda})\right) &\approx \mathbb{E}_{\mathbf{W}} \prod_{j=1}^J (e_j(\mathbf{W}) - d_j(\mathbf{W})) \quad \text{with} \\ d_j(\mathbf{W}) &= \Phi_1\left(\underline{b}_j - \sum_{j=1}^{j-1} \boldsymbol{\Lambda}_{jj}^{-1} \Phi^{-1}(d_j + W_j(e_j(\mathbf{W}) - d_j(\mathbf{W})))\right); \quad d_1(\mathbf{W}) = \Phi(\underline{b}_1) \\ e_j(\mathbf{W}) &= \Phi_1\left(\bar{b}_j - \sum_{j=1}^{j-1} \boldsymbol{\Lambda}_{jj}^{-1} \Phi^{-1}(d_j + W_j(e_j(\mathbf{W}) - d_j(\mathbf{W})))\right); \quad e_1(\mathbf{W}) = \Phi(\bar{b}_1) \end{aligned}$$

and the expectation is over $\mathbf{W} = (W_1, \dots, W_{J-1})^\top \in \mathbb{R}^{J-1}$, $W_j \sim \text{U}(0, 1)$ whose elements are independent. The expectation in turn is approximated by the mean over M independent draws of $\mathbf{W}$. The question arises how many samples M are required when searching for the maximiser of the sample log-likelihood

$$N^{-1} \sum_{i=1}^N \ell_{j,i}^{(1)}(\boldsymbol{\theta}, \boldsymbol{\lambda}).$$

Conceptually, because $\mathbf{Y}_1, \dots, \mathbf{Y}_N$ were assumed to be independent, the variance of this mean of log-probabilities $\ell_{j,i}^{(1)}(\boldsymbol{\theta}, \boldsymbol{\lambda})$ converges to zero as the number of observations N increases, regardless of M . For finite samples, of course, one should assess the impact of M on the maximum likelihood estimators for $\boldsymbol{\theta}$ and $\boldsymbol{\lambda}$ and choose M large enough. Graz et al., 2026 reports extensive simulation results indicating that, depending on context, maximum likelihood estimation may be to stabilise with values as small as $M = 250$ or $M = 500$.

Instead of approximating both the log-likelihood and the corresponding score function, we propose to approximate the log-likelihood by Genz' method in a first step and, in a second step, to derive the exact score function of this approximation rather than an approximate score function of the true log-likelihood. Conditionally on M realisations of $\mathbf{W}$, the score function with respect to $\underline{\mathbf{b}}$ and $\bar{\mathbf{b}}$ and the score function with respect to the lower off-diagonal elements of $\boldsymbol{\Lambda}^{-1}$ can be computed deterministically by the chain-rule. Scores with respect to $\boldsymbol{\Lambda}$ are then given by $-\boldsymbol{\Lambda}^{-\top} \otimes \boldsymbol{\Lambda}^{-1}$. A modular re-implementation of Genz, 1992 algorithm and its score function, also for $s = 2$, is given in Hothorn, 2026b as listed in Section 10.

5. Properties and Convex Approximations

Unfortunately, neither of these nonparanormal log-likelihoods leads to a convex optimisation problem. We study the properties of each of the four log-likelihoods in the following theorems whose proofs are given in Appendix A. These disappointing results, however, lead to some insights allowing to suggest some convex approximations to these problems.

We first consider the flow NPN log-likelihood $\tilde{\ell}_j^{(s)}(\boldsymbol{\vartheta}, \boldsymbol{\lambda})$, which only involves multivariate normal densities, transformation functions, and derivatives thereof.

Theorem 1. Minimising $-\tilde{\ell}_j^{(s)}(\boldsymbol{\vartheta}, \boldsymbol{\lambda})$ subject to $\mathbf{D}_j \boldsymbol{\vartheta}_j \geq \mathbf{0}$ for $j = 1, \dots, J$ is a biconvex problem in $\boldsymbol{\vartheta} \in \mathbb{R}^{\sum_{j=1}^J (P(j)-1)}$ and $\boldsymbol{\lambda} \in \mathbb{R}^{J(J-1)/2}$ for $s = 1, 2$.

The negative NPN log-likelihood, defined by log-probabilities, is not necessarily convex in $\boldsymbol{\lambda}$.

Theorem 2. Minimising $-\ell_j^{(s)}(\boldsymbol{\theta}, \boldsymbol{\lambda})$ subject to $\mathbf{D}_j \boldsymbol{\theta}_j \geq \mathbf{0}$ for $j = 1, \dots, J$ is a convex problem in $\boldsymbol{\theta} \in \mathbb{R}^{\sum_{j=1}^J (K(j)-1)}$ for given $\boldsymbol{\lambda} \in \mathbb{R}^{J(J-1)/2}$ for $s = 1, 2$. It is not necessarily convex in $\boldsymbol{\lambda} \in \mathbb{R}^{J(J-1)/2}$.

Consequently, the negative smooth NPN log-likelihood is also not necessarily convex in $\boldsymbol{\lambda}$. Under independence (that is, for $\boldsymbol{\lambda} = \mathbf{0}$), the NPN log-likelihood $\ell_j^{(s)}(\boldsymbol{\theta}, \mathbf{0}) = \sum_{j=1}^J \ell_j(\boldsymbol{\theta}_j)$ is the sum of J marginal empirical log-likelihoods and therefore $\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_J$ are orthogonal. This property is lost whenever $\boldsymbol{\lambda} \neq \mathbf{0}$ and therefore $\ell_j^{(s)}(\boldsymbol{\theta}, \boldsymbol{\lambda})$ has to be maximised with respect to both $\boldsymbol{\theta}$ and $\boldsymbol{\lambda}$ simultaneously for proper estimation and uncertainty assessment. We conclude with a statement about the mixed NPN log-likelihood.

Corollary 1. Minimising $-\tilde{\ell}_{j|j}^{(s)}(\boldsymbol{\vartheta}_1, \dots, \boldsymbol{\vartheta}_j, \boldsymbol{\theta}_j, \dots, \boldsymbol{\theta}_J, \boldsymbol{\lambda})$ subject to $\mathbf{D}_j \boldsymbol{\vartheta}_j \geq \mathbf{0}$ for $j = 1, \dots, J-1$ and $\mathbf{D}_J \boldsymbol{\theta}_J \geq \mathbf{0}$ is a biconvex problem for $s = 1$ and $j = J-1$, that is in $(\boldsymbol{\vartheta}_1, \dots, \boldsymbol{\vartheta}_{J-1}, \boldsymbol{\theta}_J) \in \mathbb{R}^{K(J)-1 + \sum_{j=1}^{J-1} (P(j)-1)}$ and $\boldsymbol{\lambda} \in \mathbb{R}^{J(J-1)/2}$ for $s = 1, 2$.

Given the malign nature of the optimisation problems involved, we discuss three convex approximations which, at the very least, help to derive good starting values:

- (1) Minimise $-\ell_j(\theta_j)$, or $-\ell_j(\theta_j; \vartheta_j)$, in (4), and obtain the empirical marginal estimate $\hat{\theta}_j$ for all $j = 1, \dots, J$ and get $\hat{\theta} = (\hat{\theta}_1, \dots, \hat{\theta}_J)^\top$. One typically tries to avoid $\theta_{j,N} = \infty$ by changing the estimator to the normal score $r(i, j)(N+1)^{-1} = \Phi(\hat{\theta}_{j,r(i,j)})$ (or a winsorised version, Mai et al., 2023) when estimating θ_j (this problem is not present when a smoothly parameterised model is given by $\hat{\theta}_j$). Define $\hat{\mathbf{Z}}_i = (\hat{\theta}_{1,r(i,1)}, \dots, \hat{\theta}_{J,r(i,J)})^\top$ and minimise $-\tilde{\ell}_j^{(0)}(\Omega^{(2)}(\lambda))$ with respect to λ . This maintains the interpretation of $\Phi(\hat{\theta}_{j,r(i,j)})$ as j th marginal distribution function evaluated at Y_{ij} .
- (2) For the flow NPN log-likelihood, an iterative version with alternating estimation of λ or ϑ , that is, switching between the two target functions $\tilde{\ell}_j^{(s)}(\hat{\vartheta}, \lambda)$ and $\tilde{\ell}_j^{(s)}(\vartheta, \hat{\lambda})$, is known as alternate convex search (ACS) which, under certain conditions, might converge (Gorski et al., 2007).
- (3) (a) Solve the convex problem (in ϑ_1) and minimise $-\tilde{\ell}_1^{(s)}(\vartheta_1)$ subject to $\mathbf{D}_1 \vartheta_1 \geq \mathbf{0}$.
 (b) Solve the convex problem (in ϑ_2 and λ_{21}) and minimise $-\tilde{\ell}_2^{(s)}(\hat{\vartheta}_1, \vartheta_2, \lambda_{21})$ subject to $\mathbf{D}_2 \vartheta_2 \geq \mathbf{0}$.
 (c) Solve the convex problem (in ϑ_3 and λ_3) and minimise $-\tilde{\ell}_3^{(s)}(\hat{\vartheta}_1, \hat{\vartheta}_2, \vartheta_3, \hat{\lambda}_{21}, \lambda_3)$ subject to $\mathbf{D}_3 \vartheta_3 \geq \mathbf{0}$.
 (d) Repeat until $j = J$. Solve the convex problem (in ϑ_J and λ_J) and minimise

$$-\tilde{\ell}_J^{(s)}(\hat{\vartheta}_1, \hat{\vartheta}_2, \dots, \hat{\vartheta}_{J-1}, \vartheta_J, \hat{\lambda}_{21}, \dots, \hat{\lambda}_{(J-1)}, \lambda_J)$$
 subject to $\mathbf{D}_J \vartheta_J \geq \mathbf{0}$.

The approaches in 1. are variants of the two-step maximum pseudo likelihood estimator, where the parameters of the marginal models are fitted first. Treating these parameters as known, the copula parameters are then estimated in a second step. Two-step maximum pseudo likelihood estimators are used as starting values for the simultaneous maximisation of NPN log-likelihood functions with respect to all model parameters in Section 6.

The sequential approximation 3. fits a series of linear transformation models to regressors $\hat{h}_1, \dots, \hat{h}_{j-1}$, where only ϑ_j and the j th row λ_j of Λ are updated. This works for $s = 1, 2$ and also allows penalisation of the Λ parameters for high(er)-dimensional data as suggested for normal models by Khare et al., 2019. Variants 2. and 3. could also be combined with the smooth and mixed NPN log-likelihood.

6. Empirical Comparisons

The theoretical and computational framework presented in Sections 2–5 is too broad to be empirically evaluated in an exhaustive way. We therefore focus on one application highlighting the versatility of the framework and a simple simulation setup to illustrate potential practical merits. A more comprehensive picture of the empirical performance of maximum nonparanormal likelihood estimation emerges from simulation results published elsewhere: Klein et al., 2022 investigate the flow NPN log-likelihood with covariate-dependent Λ , Graz et al., 2026 estimate and evaluate multivariate count models obtained by maximisation of the NPN log-likelihood, and the mixed NPN log-likelihood was evaluated for the estimation of time-dependent receiver operating characteristic curves (Sewak et al., 2026a) and in covariate adjustment procedures (Dandl and Hothorn, 2026).

6.1. Transformation Discriminant Analysis

We discuss a discrimination function for hepatocellular carcinoma (HCC) diagnosis based on four biomarkers (DKK: Dickkopf-1, OPN: osteopontin, PIV: protein induced by vitamin K absence or antagonist-II, and AFP: alpha-fetoprotein). Based on data reported from a retrospective case-control study of 403 subjects by Jang et al., 2016, Sewak et al., 2026b proposed the log-likelihood ratio function of a transformation discriminant analysis (TDA) model as optimal discrimination function. The empirical biomarker distributions presented for HCC cases and non-HCC controls in Figure 1 show that PIV and AFP readings are affected by a limit-of-detection problem. For these subjects, it is only known that PIV (or AFP) is larger than a specific detection limit, in other words, these observations are right-censored.

For the analysis of this dataset, we show how one can put together different models playing with different parameterisations of the marginal and copula parameters and then picking the ‘right’ likelihood based on properties of the observations (the mixed NPN log-likelihood because of right-censoring). The same case is made by the R code snippet presented in Section 9, demonstrating how the theoretical concepts translate into software.

We fit three models to the data. First, a classical linear discriminant analysis (LDA) assuming a linear transformation function (and thus a linear basis function $\mathbf{a}_j$) for each of the four biomarkers and a common covariance, resulting in a joint normal distribution of the biomarker values with class-specific means. Second, we replace the linear transformation functions with potentially non-linear ones (2) featuring a location term differentiating between classes. As a third option, we introduce a scale term such that the marginal variability may differ between classes. In all models, we restrict our attention to a common correlation matrix. With $x = 1$ for HCC case and $x = 0$ for a non-HCC control and $\mathbf{Y} = (Y_{\text{AFP}}, Y_{\text{DKK}}, Y_{\text{OPN}}, Y_{\text{PIV}})^{\top} \in \mathbb{R}^4$ the LDA model is equivalent to $\mathbf{Y} \sim N_J(\Omega^{-1}\boldsymbol{\eta}(x), \Omega^{-1}\Omega^{-\top})$. Because this problem is convex in both $\boldsymbol{\eta}(x)_j = \beta_j x$ and Ω (Section 5.2.1. in Barratt and Boyd, 2023), we use a convex solver as a benchmark for later method comparison.

All NPN models feature variants of the transformation function implementing a location-scale model (Siegfried et al., 2023) with $h_j(y_j | x) = \mathbf{a}_j(y_j)^{\top} \boldsymbol{\vartheta}_j \exp(\xi_j x) - \beta_j x$. The LDA model can be formulated by choosing linear bases $\mathbf{a}_j^{\top} = (1, y_j)$ and location-only part ($\xi_j = 0$ for all $j \in \{\text{DKK}, \text{OPN}, \text{PIV}, \text{AFP}\}$), however, this parameterisation leads to a non-convex optimisation problem when minimising the negative flow NPN log-likelihood $-\tilde{\ell}^{(2)}(\boldsymbol{\vartheta}, \boldsymbol{\lambda})$ simultaneously in all model parameters $\boldsymbol{\vartheta} = (\boldsymbol{\vartheta}_{\text{DKK}}^{\top}, \beta_{\text{DKK}}, \dots, \boldsymbol{\vartheta}_{\text{AFP}}^{\top}, \beta_{\text{AFP}})^{\top}$ and $\boldsymbol{\lambda}$.

The log-likelihoods (normal negative convex and flow NPN log-likelihood) obtained by both optimisation routines are equivalent (-2373.492), this also applies to the log-likelihood ratios. After dividing each column of $\hat{\Omega}$ obtained from the convex solver with the square-root of the corresponding diagonal element, the estimated values of $\boldsymbol{\lambda}$ are also identical, see Table 1. This means that the simultaneous one-step non-convex optimisation of the flow NPN log-likelihood resulted in the same model as the optimisation of the corresponding convex negative normal log-likelihood.

The location-only transformation discriminant analysis model (ITDA) is obtained from more flexible basis functions; we use Bernstein polynomial bases $\mathbf{a}_j^{\top} \in \mathbb{R}^7$ of order 6. The additional 20 parameters introducing non-linear transformations improve the flow NPN log-likelihood to -2106.183 , at the expense of higher computing times (median 2.523 instead of 0.706 seconds).




Figure 1 – Case-control study for Hepatocellular carcinoma (HCC) by Jang et al., 2016: Empirical cumulative distribution functions (CDFs) for log-transformed biomarkers alpha-fetoprotein (AFP), protein induced by vitamin K absence or antagonist-II (PIV), osteopontin (OPN), and Dickkopf-1 (DKK) in HCC cases and non-HCC controls.

Table 1 – HCC discriminant analysis: Linear discriminant analysis (LDA) and transformation discriminant analysis (location-only: ITDA, location-scale: lsTDA) fitted by minimisation of the convex negative log-likelihood of a multivariate normal (for LDA only), flow NPN log-likelihood, and mixed NPN log-likelihood, the latter taking limit-of-detection problems into account. Maximum likelihood estimates for λ parameters and standard errors (via inverse observed Hessians) are given. The in-sample log-likelihoods and median computing times are reported in the bottom rows.

λ	LDA			ITDA		lsTDA			
	convex	flow		flow		flow		mixed	
	$\hat{\lambda}$	$\hat{\lambda}$	SE($\hat{\lambda}$)	$\hat{\lambda}$	SE($\hat{\lambda}$)	$\hat{\lambda}$	SE($\hat{\lambda}$)	$\hat{\lambda}$	SE($\hat{\lambda}$)
OPN,DKK	−0.180	−0.180	0.051	−0.107	0.050	−0.099	0.050	−0.099	0.050
PIV,DKK	−0.330	−0.330	0.053	−0.296	0.053	−0.300	0.053	−0.300	0.053
PIV,OPN	−0.319	−0.319	0.052	−0.344	0.053	−0.311	0.053	−0.311	0.053
AFP,DKK	0.019	0.019	0.053	0.032	0.053	0.041	0.053	0.041	0.053
AFP,OPN	−0.083	−0.083	0.053	−0.171	0.054	−0.186	0.054	−0.186	0.054
AFP,PIV	−1.945	−1.945	0.109	−1.428	0.099	−1.345	0.097	−1.345	0.097
log-Lik	−2373.492	−2373.492		−2106.183		−2069.909		−5268.329	
time (sec)	0.000		0.706		2.523		3.143		24.341

The maximum-likelihood estimates $\hat{\lambda}$ are similar, as are the corresponding standard errors obtained from the inverse Hessian.

A location-scale version of the above model introduces four additional scale parameters ξ_j . Again, an improvement in the flow NPN log-likelihood was observed (−2069.909), the computing time increased only marginally (to median 3.143 seconds).

The four models estimated via minimisation of the convex negative normal or flow NPN log-likelihood ignored the fact that PIV or AFT biomarker values could not be observed for 17

Table 2 – Location-scale transformation discriminant analysis (lsTDA) model for HCC: simultaneous optimisation of marginal and copula parameters using the mixed NPN log-likelihood with three approximations: two-step pseudo, alternating, and sequential (Section 5).

λ	lsTDA			
	simultaneous $\hat{\lambda}$	pseudo (1.) $\hat{\lambda}$	alternating (2.) $\hat{\lambda}$	sequential (3.) $\hat{\lambda}$
OPN,DKK	−0.099	−0.094	−0.093	−0.096
PIV,DKK	−0.300	−0.276	−0.285	−0.276
PIV,OPN	−0.311	−0.301	−0.303	−0.306
AFP,DKK	0.041	0.016	0.028	0.015
AFP,OPN	−0.186	−0.207	−0.182	−0.199
AFP,PIV	−1.345	−1.036	−1.255	−1.134
log-Lik	−5268.329	−5295.388	−5269.009	−5282.360
time (sec)	24.341	0.702	50.392	10.890

subjects because the upper limit-of-detection was reached. Technically, these observations are right-censored, necessitating a correction of the log-likelihood contributions for these observations. We implemented such a correction by a mixed NPN log-likelihood combining the flow NPN log-likelihood for DKK and OPN and a smooth NPN log-likelihood with right-censoring for AFP and PIV, technically speaking, the maximisation of

$$\tilde{\ell}_{\text{PIV,AFP|DKK,OPN}}^{(2)}(\boldsymbol{\vartheta}_{\text{DKK}}, \beta_{\text{DKK}}, \xi_{\text{DKK}}, \dots, \boldsymbol{\vartheta}_{\text{AFP}}(\boldsymbol{\vartheta}_{\text{AFP}}), \beta_{\text{AFP}}, \xi_{\text{AFP}}, \boldsymbol{\lambda}).$$

The in-sample flow NPN log-likelihood and mixed NPN log-likelihood values are not directly comparable, however, neither the estimated $\boldsymbol{\lambda}$ parameters nor the corresponding standard errors are affected by this more elaborate estimation, which took much longer to compute (median 24.341 seconds with Halton sequences of length $M = 500$). This result, obtained from the much more elaborate mixed NPN log-likelihood function, is practically important, as it demonstrates that ignoring limits of detection had negligible influence on the fitted model in this application.

The mixed NPN log-likelihood is not even biconvex and it might be interesting to look at the results obtained by the convex approximations discussed in Section 5. The results in Table 2 suggest that alternating between the estimation of marginal and copula parameters (2.) provides a better approximation to the in-sample log-likelihood obtained by simultaneous optimisation of all model parameters compared to the two-step pseudo (1.) or sequential (3.) approaches. However, the small mixed standard error 0.097 of $\hat{\lambda}_{\text{AFP,PIV}}$ (last row and column of Table 1) suggest that the discrepancies among estimates in rows of Table 2 might be practically relevant. In addition, the alternation took twice as much time as the simultaneous parameter estimation. Therefore, one might call into question the practical merits of convex approximations in this setup.

6.2. Polychoric Correlations

For bivariate Gaussian copulas, the semiparametric efficiency bound is known and the performance of several estimators against this theoretical benchmark is studied in this section. We sample $N \in \{10, 20, 50\}$ observations from $\mathbf{Y} = (Y_1, Y_2) \sim \text{NPN}(\mathbf{h}, \boldsymbol{\Sigma}(\rho))$ with $h_1(y_1) = \Phi^{-1}(\chi^2_2(y_1))$ and $h_2 = h_1$, that is, $Y_j \sim \chi^2_2$ for $j = 1, 2$. The latent correlation between both variables is given by $\boldsymbol{\Sigma}(\rho) = ((1, \rho)^\top \mid (\rho, 1)^\top)$ for $\rho \in \{0, 0.1, 0.2, \dots, 0.9\}$. Klaassen and Wellner, 1997 established the semiparametric efficiency bound $(1 - \rho^2)/\sqrt{N}$ for the correlation, that




Figure 2 – Bivariate polychoric correlations: Distribution of 100 estimators of $\rho = 0.5$ (horizontal line) from N bivariate observations measured at different scales: continuous (Cont.), ordinal (five levels, Ord.), and binary (Bin.). For continuous variables, the two-step maximum pseudo likelihood approach (“ML” and “MPL” provide identical estimators, but different standard errors) is shown as competitor, for categorical variables, the composite maximum likelihood (cML) is presented.

is, the standard error of semiparametric efficient estimators $\hat{\rho}$. It is identical to the standard error of the Pearson correlation coefficient computed for the *transformed* bivariate observations $(\Phi^{-1}(\chi_2^2(Y_{1i})), \Phi^{-1}(\chi_2^2(Y_{2i}))), i = 1, \dots, N$, that is, in the oracle case where the marginal χ_2^2 distribution functions are known.

For both responses being absolutely continuous, we estimate ρ by the two-step maximum pseudo likelihood estimator (which is, according to Klaassen and Wellner, 1997, semiparametric




Figure 3 – Bivariate polychoric correlations: Distribution of 100 standard errors for estimators of $\rho = 0.5$ from N bivariate observations measured at different scales: continuous (Cont.), ordinal (five levels, Ord.), and binary (Bin.). For continuous variables, maximum likelihood (ML) and two-step maximum pseudo likelihood (MPL) approaches are shown as competitors, for categorical variables, the composite maximum likelihood (cML) is presented. The horizontal line indicates the semiparametric efficiency bound and red triangles the standard deviation of $\hat{\rho}_1, \dots, \hat{\rho}_{100}$.

efficient in this simple case), with two different ways to obtain standard errors (labelled “ML” and “MPL”, see Section 9).

We compare the performance of this estimator to maximum likelihood-based estimation of ρ using the NPN log-likelihood, smooth NPN log-likelihood, and flow NPN log-likelihood, that is, by employing the transformation $\hat{\rho} = -\hat{\lambda}_{21}/\sqrt{1+\hat{\lambda}_{21}^2}$. In addition, we also report estimators of the corresponding standard errors of $\hat{\rho}$, obtained via the Δ -method for the procedures described in this paper.

To study the performance for non-continuous data, we transform each variable to binary and ordinal (five categories) measurements using random empirical quantiles between 20% and 80% as cut-offs. As a competitor for binary or ordinal variables, we use the composite likelihood (“cML”, Nikoloulopoulos, 2023, which in the binary case is identical to the NPN log-likelihood). In

absence of a direct competitor for mixed continuous-discrete responses, we only report the results obtained via the mixed NPN log-likelihood (that is, a mix of NPN log-likelihood, smooth NPN log-likelihood, or flow NPN log-likelihood for the continuous variable and NPN log-likelihood for the categorical variable). For each combination of measurement scales, we repeat the simulation 100 times.

For a true $\rho = 0.5$, the distribution of the estimators and their standard errors are given in Figures 2 and 3. For continuous variables, all three flavours of the NPN log-likelihood attain the semiparametric efficiency bound for all sample sizes and the corresponding Δ standard errors are very close to the theoretical value. In comparison for small sample sizes, the classical two-step maximum pseudo likelihood copula estimators are slightly more biased, this also applies to their standard errors. For larger sample sizes, the estimated correlations are identical, unlike the corresponding standard errors.

When at least one variable is categorical, the NPN estimators are still unbiased but the variance of the estimators increases slightly. Especially for small sample sizes, the standard errors by the competing procedures under- or over-estimate the true variation, whereas the standard errors obtained from inverting the Hessian of some NPN log-likelihood reflect the variability of the corresponding estimates closer. For larger sample sizes, these differences become very small. For smaller and larger values of ρ in the data generating process, results are given in Appendix B. Especially for very high correlations and when both variables are binary, the estimation performance as well as the quality of the standard errors degrades.

In the multivariate setup with $J = 3, 6, 9$ ordinal outcomes, we compared the polychoric correlations estimated by maximising the composite likelihood with nonparanormal estimates. The data generating process was the same as in the bivariate case, except that we randomly sampled λ from uncorrelated uniform distributions on $[-1, 1]$ for each of the 100 simulation runs. We report the root mean squared error between the true and estimated $J(J-1)/2$ polychoric correlations as well as computing times in Appendix C. With increasing dimension and increasing sample size, the number of convergence failures increases for cML, whereas optimisation of the NPN log-likelihood was always possible. In case of successful convergence, the root mean squared errors were comparable. Availability of the score function for the NPN log-likelihood resulted in faster computing times compared to optimisation of the composite likelihood relying on numerically approximated scores.

7. Application Domains

The range of applications of the NPN model, especially with covariate-adjustment, is quite broad and we sketch possible parameterisations for some interesting applications in this section. We start with probit models for multivariate binary outcomes (Lesaffre and Kaufmann, 1992), that is, $Y_j \in \{0, 1\}$ for all $j = 1, \dots, J$. The marginal distributions are given by $h_j(y_j | \mathbf{x}) = \theta_{j1} + \mathbf{x}^\top \boldsymbol{\beta}_j$, where θ_{j1} is the intercept term for the j th binary response. This concept was later generalised to “Copula regression” via a multivariate GLM formulation (Masarotto and Varin, 2012; Song et al., 2009). For example, marginal binary logistic models feature $h_j(y_j | \mathbf{x}) = \Phi^{-1}(\text{expit}(\theta_{j1} + \mathbf{x}^\top \boldsymbol{\beta}_j))$ as marginal transformation functions. The notion of “polychoric correlations” (Jöreskog, 1994) was extended to multivariate proportional-odds models (e.g. Hirk et al., 2019). For ordered sample spaces $\mathcal{Y}_j = \{v_{j1} < \dots < v_{jK(j)}\}$, a marginal proportional-odds model corresponds to the transformation $h_j(y_j | \mathbf{x}) = \Phi^{-1}(\text{expit}(\theta_{jk} + \mathbf{x}^\top \boldsymbol{\beta}_j))$. In all these models, the

NPN log-likelihood $\ell(\boldsymbol{\theta}, \boldsymbol{\lambda})$ can be maximised simultaneously in all model parameters. This also allows likelihood inference for contrasts of marginal parameters, for example when the hypothesis $\beta_{j1} = 0$ for all $j = 1, \dots, J$ is of interest.

For counts $Y_j \in \mathbb{N}$, Siegfried and Hothorn, 2020 suggested marginal proportional-odds models $h_j(y_j | \mathbf{x}) = \Phi^{-1}(\text{expit}(\mathbf{a}_j(\lfloor y_j \rfloor)^\top \boldsymbol{\vartheta}_j + \mathbf{x}^\top \boldsymbol{\beta}_j))$ which, for multiple count outcomes, can be estimated jointly by maximising the smooth NPN log-likelihood $\ell(\boldsymbol{\theta}(\boldsymbol{\vartheta}), \boldsymbol{\lambda})$ (Graz et al., 2026). This also generalises the bivariate count models by Niehaus et al., 2024. For continuous outcomes, Mai and Zou, 2015 and Sewak et al., 2026b studied transformation discriminant analysis models, where marginal transformations include shift and possibly scale effects differentiating between two (or more) classes, see also Section 6.1 for a worked example involving both the flow NPN log-likelihood and mixed NPN log-likelihood. The latter likelihood is relevant for the estimation of graphical models (Göbler et al., 2024) or structural equation models (Pritikin et al., 2018) for mixed outcomes, for the generation of synthetic data from such models (e.g. for missing value imputation, Christoffersen et al., 2021; Zhao and Udell, 2020), and for prognostic covariate adjustment under non-collapsibility (Dandl and Hothorn, 2026).

In survival analysis, multivariate survival times can be analysed by NPN models with Cox-type margins. Independent censoring requires the application of the mixed NPN log-likelihood. Such models have been suggested for the case $J = 2$ (Ding and Sun, 2022; Marra and Radice, 2020), NPN models also allow $J > 2$. A very important topic are recently suggested models for dependent censoring. In addition to some time to event of interest $T > 0$, one also observes drop-out times $C > 0$ and times of administrative independent censoring $A > 0$. For each subject, only $\min(T, C, A)$ can be observed. Czado and Van Keilegom, 2023 and Deresa and Keilegom, 2023 proved that the parameters of a suitably defined NPN model are identified even under this partial information. More specifically, with $h_T(t | \mathbf{x}) = \Phi^{-1}(\text{cloglog}^{-1}(\mathbf{a}_T(t)^\top \boldsymbol{\vartheta}_T + \mathbf{x}^\top \boldsymbol{\beta}_T))$ and $h_C(c | \mathbf{x}) = \Phi^{-1}(\text{cloglog}^{-1}(\mathbf{a}_C(c)^\top \boldsymbol{\vartheta}_C + \mathbf{x}^\top \boldsymbol{\beta}_C))$ one can estimate the latent correlation $\rho = -\lambda_{21}/\sqrt{1+\lambda_{21}^2}$ if $\mathbf{a}_C(c)^\top = (1, \log(c))$, that is, when the marginal drop-out time follows a Weibull model. The marginal time to event might even follow a Cox proportional hazards model, for example with marginal log-baseline cumulative hazard function $\mathbf{a}_T(t)^\top \boldsymbol{\vartheta}_T$ parameterised in terms of a polynomial in Bernstein form. Both marginal distributions ensure that covariate effects are interpretable as marginal log-hazard ratios. The log-likelihood for an observed event time (i.e. $T = t, C > t$) is the mixed NPN log-likelihood $\tilde{\ell}_{C|T}(\boldsymbol{\vartheta}_T, \boldsymbol{\theta}_C(\boldsymbol{\vartheta}_C), \lambda_{21})$. For a drop-out (i.e. $T > c, C = c$), the log-likelihood is the mixed NPN log-likelihood $\tilde{\ell}_{T|C}(\boldsymbol{\vartheta}_C, \boldsymbol{\theta}_T(\boldsymbol{\vartheta}_T), \lambda_{21})$. Administratively censored subjects (i.e. $T > a, C > a$) further add the NPN log-likelihood $\tilde{\ell}_2(\boldsymbol{\theta}_T(\boldsymbol{\vartheta}_T), \boldsymbol{\theta}_C(\boldsymbol{\vartheta}_C), \lambda_{21})$, an illustration is given by Siegfried et al., 2026.

Finally, all the models above can be coupled with covariate-dependent copula parameters (3) as explained in Klein et al., 2022 and Barratt and Boyd, 2023, for example when estimating time-varying graphical models (Lu et al., 2018). Unfortunately, and unlike models with constant $\boldsymbol{\lambda}$ parameters, such models are in general not invariant to the order in which responses enter the model.

8. Discussion

Given the plethora of inference procedures for many special cases of the NPN model, one might wonder in which cases optimisation of the NPN log-likelihood, or any of the approximations discussed in this paper, is beneficial. From a methodological point of view, the NPN

log-likelihood provides a benchmark against which other approximations, for example the composite likelihood in multivariate regression models (Nikoloulopoulos, 2023), can be evaluated. We present a simple version of such a benchmark comparison in Section 6, comparing the two-step pseudo and composite maximum likelihood approaches to several flavours of the NPN log-likelihood, both in terms of their estimation accuracy and corresponding variability assessment. The exercise shows that NPN log-likelihood estimators exhibit the variability of semiparametric efficient estimators for at least ordered response variables. For small samples, the estimation accuracy of NPN log-likelihood estimators is better, compared to two-step estimators. Practically even more relevant is the improved accuracy of standard errors obtained from the NPN log-likelihood approaches. On these grounds, the application of standard maximum-likelihood inference procedures to more complex models, for example when one is interested in supplementing bootstrap procedures by simple Wald tests and confidence intervals in dependent censoring models (Deresa and Keilegom, 2023), seems promising.

An important contribution is the ability to estimate models when the response types are mixed, that is, some variables can be considered continuous while others are clearly discrete. The same applies to missing values in some of the response variables. The NPN log-likelihood allows a straightforward handling of observations missing at random. We simply use the datum $(-\infty, \infty)$ when computing the contribution of the j th, missing, covariate to the likelihood. Imprecise measurements can be handled via interval-censoring.

From a more theoretical point of view, the consistency of the two-step pseudo maximum likelihood approach, based on normal or winsorised scores, in combination with the graphical lasso was recently demonstrated in ultra-high dimensions by Mai et al., 2023. So far, such a result is only available when all responses are absolutely continuous and in the absence of any additional parameters in the marginal or joint distributions. The non-convexity of the negative nonparametric log-likelihoods studied here renders them unattractive for penalisation approaches in higher dimensions. However, the contribution might still be useful for the estimation of graphical models for non-normal and potentially discrete responses in high-dimensions. Following Xue and Zou, 2012 or Suggala et al., 2017, bivariate NPN models could be employed to estimate the polychoric correlations ρ_{jj} . The matrix $(\hat{\rho}_{jj})_{1 \leq j < j' \leq J}$ with $\hat{\rho}_{jj} \equiv 1$ can then replace the sample covariance matrix in a graphical lasso, neighbourhood Dantzig selector, or CLIME. Xue and Zou, 2012 demonstrated that their “rank-based” versions are consistent with the same rates of convergence as the original versions based on the sample covariance matrix of normal data.

Efficiency results on such two-step estimators are sparse. Klaassen and Wellner, 1997 demonstrated efficiency of the correlation parameter in a bivariate Gaussian copula, and we utilised this ground truth in the simulation experiments in Section 6. Even in this simple case, the two-step marginal distributions are inefficient. The flow NPN log-likelihood is conceptually very similar to the semiparametric efficient sieve maximum-likelihood estimators studied by Chen et al., 2006. The main difference lies in their choice of a sieve approximation for marginal densities whereas we utilise polynomials in Bernstein form $\mathbf{a}_j(y_j)^\top \boldsymbol{\vartheta}_j$ to approximate marginal transformation functions. If one allows the number of basis functions $K(j)$ in $\mathbf{a}_j$, and therefore the number of coefficients $\boldsymbol{\vartheta}_j$, to depend on the same size N , the sieve space proposed and analysed by McLain and Ghosh, 2013 emerges. Chen et al., 2006 also proved that semiparametric efficiency carries over to models where some of the marginal distributions are fully parametric. These results make the NPN model, its smooth parameterisation and the corresponding flow NPN log-likelihood,

especially when coupled with the ACS optimisation method, promising candidates for future research.

9. Data, Script and Code Availability

All computations were performed using R version 4.6.1 (R Core Team, 2026). All flavours of nonparanormal log-likelihoods were computed using infrastructure in package **mlt** (Hothorn, 2026a) based on algorithms for the evaluation of multivariate normals in **mvtnorm** (Genz et al., 2026). The convex parameterisation of the LDA model was estimated by package **CVXR** (Fu et al., 2026). Maximum pseudo and composite likelihood estimates of polychoric correlations in Section 6.2 were computed using packages **copula** (Hofert et al., 2026), with variant “ML” corresponding to `fitCopula(..., type = "ml")` and variant “MPL” to `fitCopula(..., type = "mpl")`, and **mvord** (Hirk et al., 2026), respectively. A low-level interface to different forms of the nonparanormal log-likelihoods is available from package **mlt** (Hothorn, 2026a) via the `mlt::mmult` function. Package **tram** provides a high-level interface as `tram::Mmult`.

The location-scale transformation discriminant analysis model for HCC diagnosis under limits-of-detection was estimated by the following code

```
> library("tram")
> ### run demo("npr") from tram package for full reproducibility
>
> ### marginal location-scale models
> mDKK <- BoxCox(
+   DKK ~                                     ### probit, h(DKK) via Bernstein
+   x                                         ### location non-HCC / HCC
+   | x,                                     ### scale non-HCC / HCC
+   data = HCC)
> mOPN <- BoxCox(OPN ~ x | x, data = HCC)
> mPIV <- BoxCox(R(
+   Surv(PIV, event = PIV < PIVm),          ### right censoring
+   as.R.interval = TRUE) ~                 ### empirical likelihood
+   x | x,                                   ### location-scale
+   data = HCC)
> mAFP <- BoxCox(R(Surv(AFP, event = AFP < AFPm), as.R.interval = TRUE) ~
+   x | x, data = HCC)
> ### joint estimation of marginal and Gaussian copula parameters, s = 2
> ### location-scale transformation discriminant analysis
> m <- Mmult(mDKK, mOPN, mPIV, mAFP, data = HCC)
> ### marginal parameters
> coef(m, type = "marginal")
> ### copula parameter: Lambda
> coef(m, type = "Lambda")
> ### standard errors for all parameters
> sqrt(diag(vcov(m)))
> ### convex approximations
> ## pseudo
```

```
> mm <- Mmlt(mDKK, mOPN, mPIV, mAFP, data = HCC, domargins = FALSE)
> ## sequential
> ms <- Mmlt(mDKK, mOPN, mPIV, mAFP, data = HCC, sequentialfit = TRUE)
```

The complete R code necessary to reproduce empirical results discussed in Section 6 is provided in

- **mvtnorm**: [10.32614/CRAN.package.mvtnorm](https://cran.r-project.org/web/packages/mvtnorm/index.html)
- **mlt**: [10.32614/CRAN.package.mlt](https://cran.r-project.org/web/packages/mlt/index.html)
- **tram**: [10.32614/CRAN.package.tram](https://cran.r-project.org/web/packages/tram/index.html), including data and code for Section 6.1 in `tram/inst/npn_empeval` and for Section 6.2 in `tram/inst/npn_simulation`.

10. Supplementary Information

- Appendix A, Proofs: Proof of Theorem 1 and Theorem 2.
- Appendix B, Polychoric Correlations: Additional simulation results for the bi- and multi-variate case.
- Supplementary Material, [10.32614/CRAN.package.mvtnorm](https://cran.r-project.org/web/packages/mvtnorm/index.html): **Likelihood and Score Functions**: A modular re-implementation of Genz, 1992 algorithm tailored to the evaluation of the different nonparanormal log-likelihoods. Implementation aspects of mixed continuous and discrete normal log-likelihoods are discussed in Chapter 5. Log-likelihoods for the case $s = 2$ are described in Chapter 6. The chain-rule to derive scores with respect to Λ^{-1} is given in Section 3.2.

Funding

Financial support by Swiss National Science Foundation N° 200021_219384 is acknowledged.

Conflict of Interest Disclosure

The author declares he has no conflict of interest relating to the content of this article.

Acknowledgements

Preprint version 4 of this article has been peer-reviewed and recommended by Peer Community In Statistics and Machine Learning (<https://doi.org/10.24072/pci.statml.100138>; Chiquet, 2026).

Appendix A. Proofs

Theorem 1.

Proof. As a function of $\mathbf{z} \in \mathbb{R}^J$, $\|\Omega^{(s)}(\boldsymbol{\lambda})\mathbf{z}\|_2^2$ is convex for $s = 1, 2$ and each fixed $\boldsymbol{\lambda} \in \mathbb{R}^{J(J-1)/2}$. The argument is completed noting that $\mathbf{z}_j = \mathbf{a}_j(y_j)^\top \boldsymbol{\vartheta}_j$ is linear in $\boldsymbol{\vartheta}_j$.

As a function of $\boldsymbol{\lambda} \in \mathbb{R}^{J(J-1)/2}$, $\|\Omega^{(1)}(\boldsymbol{\lambda})\mathbf{z}\|_2^2 = \|\boldsymbol{\Lambda}\mathbf{z}\|_2^2$ is convex for each fixed $\mathbf{z} \in \mathbb{R}^J$. For $s = 2$, we follow Khare et al., 2019 and write

$$-\tilde{\ell}_{J,i}^{(0)}(\Omega^{(2)}) = - \left(-\frac{1}{2} \|\Omega^{(2)}(\boldsymbol{\lambda})\mathbf{z}\|_2^2 + \sum_{j=1}^J \log(\Omega_{jj}^{(2)}) \right) = \sum_{j=1}^J \frac{1}{2} \|\Omega^{(2)}(\boldsymbol{\lambda})_{j\cdot}\mathbf{z}\|_2^2 - \log(\Omega_{jj}^{(2)})$$

as a sum of J independent terms.

For $j = 1$, $\Omega^{(2)}(\boldsymbol{\lambda})_{11} = 1$. For $j = 2$, we add the constraint $\Omega^{(2)}(\boldsymbol{\lambda})_{22} = \|\lambda_{2\cdot}\|_2 = \sqrt{1 + \lambda_{21}^2}$. This constraint is convex in λ_{21} and can be relaxed to the convex inequality constraint

$$\|\lambda_{2\cdot}\|_2 - \Omega^{(2)}(\boldsymbol{\lambda})_{22} \leq 0.$$

For $j > 3$, we write $\Omega^{(2)}(\boldsymbol{\lambda})_{j,j} = \lambda_{j,j} \Omega^{(2)}(\boldsymbol{\lambda})_{j,j}$ and note that $\Omega^{(2)}(\boldsymbol{\lambda})_{j,j}$ only depends on $\lambda_{11}, \dots, \lambda_{j,j-1}$ for $j = 1, \dots, j-1$. Thus, the constraint $\|\lambda_{j\cdot}\|_2 - \Omega^{(2)}(\boldsymbol{\lambda})_{jj} \leq 0$ is a convex inequality constraint. $\square$

Theorem 2.

Proof. The normal density $\phi(\mathbf{z} | \Omega)$ is log-concave in $\mathbf{z} \in \mathbb{R}^J$ for fixed Ω and probabilities thereof are again log-concave (Prékopa, 1973).

In a counter-example with $J = 2$ and $s = 1$, we have

$$\begin{aligned} \exp(\ell_{J,i}^{(s)}(\boldsymbol{\theta}, \boldsymbol{\lambda})) &= \int_{\theta_{1,r(i,1)-1}}^{\theta_{1,r(i,1)}} \int_{\theta_{2,r(i,2)-1}}^{\theta_{2,r(i,2)}} \underbrace{\exp\left(-\frac{1}{2}z_1^2\right) \exp\left(-\frac{1}{2}(z_2 + \lambda z_1)^2\right)}_{\psi(\lambda)} dz_2 dz_1 \\ \frac{\partial^2 \psi(\lambda)}{\partial^2 \lambda} &= z_1^2 ((z_2 + \lambda z_1)^2 - 1) \exp\left(-\frac{1}{2}z_1^2\right) \exp\left(-\frac{1}{2}(z_2 + \lambda z_1)^2\right) \end{aligned}$$

The integral of the latter expression is larger than zero for example for $\lambda = 0$ and $\theta_{2,r(i,2)-1} > 1$ and thus not concave in λ for all configurations of $\boldsymbol{\theta}$. Similar issues have been noted by Lesaffre and Kaufmann, 1992 in the multivariate probit model. $\square$

Corollary 1 follows from Theorem 1, noting that $\Omega_C = 1$ and

$$\begin{aligned} &\Phi\left(\theta_{J,r(i,J)} + (\lambda_{J,1}, \dots, \lambda_{J,J-1})^\top (h_1(y_1 | \boldsymbol{\vartheta}_1), \dots, h_J(y_{J-1} | \boldsymbol{\vartheta}_{J-1}))\right) - \\ &\Phi\left(\theta_{J,r(i,J)-1} + (\lambda_{J,1}, \dots, \lambda_{J,J-1})^\top (h_1(y_1 | \boldsymbol{\vartheta}_1), \dots, h_J(y_{J-1} | \boldsymbol{\vartheta}_{J-1}))\right) \end{aligned}$$

is log-concave (a probability of a log-concave density and linearity in $\lambda_{J,1}, \dots, \lambda_{J,J-1}$)^T.

Appendix B. Bivariate Polychoric Correlations

Additional simulation results for correlations $\rho \neq 0.5$.




Figure 4 – Bivariate polychoric correlations: Distribution of 100 estimators of $\rho = 0$.




Figure 5 – Bivariate polychoric correlations: Distribution of 100 standard errors for estimators of $\rho = 0$.




Figure 6 – Bivariate polychoric correlations: Distribution of 100 estimators of $\rho = 0.1$.




Figure 7 – Bivariate polychoric correlations: Distribution of 100 standard errors for estimators of $\rho = 0.1$.




Figure 8 – Bivariate polychoric correlations: Distribution of 100 estimators of $\rho = 0.2$.




Figure 9 – Bivariate polychoric correlations: Distribution of 100 standard errors for estimators of $\rho = 0.2$.




Figure 10 – Bivariate polychoric correlations: Distribution of 100 estimators of $\rho = 0.3$.




Figure 11 – Bivariate polychoric correlations: Distribution of 100 standard errors for estimators of $\rho = 0.3$.




Figure 12 – Bivariate polychoric correlations: Distribution of 100 estimators of $\rho = 0.4$.




Figure 13 – Bivariate polychoric correlations: Distribution of 100 standard errors for estimators of $\rho = 0.4$.




Figure 14 – Bivariate polychoric correlations: Distribution of 100 estimators of $\rho = 0.6$.




Figure 15 – Bivariate polychoric correlations: Distribution of 100 standard errors for estimators of $\rho = 0.6$.




Figure 16 – Bivariate polychoric correlations: Distribution of 100 estimators of $\rho = 0.7$.




Figure 17 – Bivariate polychoric correlations: Distribution of 100 standard errors for estimators of $\rho = 0.7$.




Figure 18 – Bivariate polychoric correlations: Distribution of 100 estimators of $\rho = 0.8$.




Figure 19 – Bivariate polychoric correlations: Distribution of 100 standard errors for estimators of $\rho = 0.8$.




Figure 20 – Bivariate polychoric correlations: Distribution of 100 estimators of $\rho = 0.9$.




Figure 21 – Bivariate polychoric correlations: Distribution of 100 standard errors for estimators of $\rho = 0.9$.

Appendix C. Multivariate Polychoric Correlations

Simulation results for multivariate ($J > 2$) polychoric correlations.




Figure 22 – Multivariate polychoric correlations: Distribution of 100 root-mean squared standard errors for estimators of polychoric correlations ρ maximising the composite (cML, implemented in `mvord::mvord`) and nonparanormal (NPN, implemented in `tram::Mmlt`) log-likelihoods. The width of the boxes is proportional to the square root of the number of successful optimisations out of 100 runs.

Table 3 – Multivariate polychoric correlations: Number of successful optimisations in 100 simulation runs optimising composite (cML, implemented in `mvord::mvord`) and nonparanormal (NPN, implemented in `tram::Mmlt`) log-likelihoods.

		N		
Method	J	N=10	N=30	N=50
cML	J=9	67	47	31
	J=6	79	99	99
	J=3	94	99	99
NPN	J=9	100	100	100
	J=6	100	100	100
	J=3	100	100	100




Figure 23 – Multivariate polychoric correlations: Distribution of running 100 times when estimating polychoric correlations maximising the composite (cML, implemented in `mvord::mvord`) and nonparanormal (NPN, implemented in `tram::Mmlt`) log-likelihoods. The width of the boxes is proportional to the square root of the number of successful optimisations out of 100 runs.

References

- Barratt S, Boyd S (2023). Covariance Prediction via Convex Optimization. *Optimization and Engineering* **24.3**, 2045–2078. <https://doi.org/10.1007/s11081-022-09765-w>.
- Chen X, Fan Y, Tsyrennikov V (2006). Efficient Estimation of Semiparametric Multivariate Copula Models. *Journal of the American Statistical Association* **101.475**, 1228–1240. <https://doi.org/10.1198/016214506000000311>.
- Chiquet J (2026). Making the Joint Likelihood of Gaussian Copula Models with Flexible Margins Computable. *Peer Community in Statistics and Machine Learning*, 100138. <https://doi.org/10.24072/pci.statml.100138>.
- Christoffersen B, Clements M, Humphreys K, Kjellström H (2021). Asymptotically Exact and Fast Gaussian Copula Models for Imputation of Mixed Data Types. In: *Proceedings of The 13th Asian Conference on Machine Learning*. Ed. by Vineeth N. Balasubramanian and Ivor Tsang. Vol. 157. Proceedings of Machine Learning Research. PMLR, pp. 870–885. URL: <https://proceedings.mlr.press/v157/christoffersen21a.html>.
- Czado C, Van Keilegom I (2023). Dependent Censoring Based on Parametric Copulas. *Biometrika* **110.3**, 721–738. <https://doi.org/10.1093/biomet/asac067>.
- Dandl S, Hothorn T (2026). Nonparanormal Adjusted Marginal Inference (with Discussion). *Biometrics* **82.2**, uja065. <https://doi.org/10.1093/biometc/ujag065>.
- Deresa NW, Keilegom IV (2023). Copula Based Cox Proportional Hazards Models for Dependent Censoring. *Journal of the American Statistical Association* **119.546**, 1044–1054. <https://doi.org/10.1080/01621459.2022.2161387>.
- Ding Y, Sun T (2022). Copula Models and Diagnostics for Multivariate Interval-Censored Data. In: *Emerging Topics in Modeling Interval-Censored Survival Data*. Ed. by Jianguo Sun and Ding-Geng Chen. Cham: Springer International Publishing, pp. 141–165. https://doi.org/10.1007/978-3-031-12366-5_8.
- Fu A, Narasimhan B, Diamond S, Miller J (2026). CVXR: Disciplined Convex Optimization. R package version 1.9.2. <https://doi.org/10.32614/CRAN.package.CVXR>.
- Genest C, Nešlehová J (2007). A Primer on Copulas for Count Data. *ASTIN Bulletin* **37.2**, 475–515. <https://doi.org/10.2143/AST.37.2.2024077>.
- Genz A (1992). Numerical Computation of Multivariate Normal Probabilities. *Journal of Computational and Graphical Statistics* **1.2**, 141–149. <https://doi.org/10.1080/10618600.1992.10477010>.
- Genz A, Bretz F, Miwa T, Mi X, Hothorn T (2026). mvtnorm: Multivariate Normal and t Distributions. R package version 1.4-2. <https://doi.org/10.32614/CRAN.package.mvtnorm>.
- Göbner K, Drton M, Mukherjee S, Miloschewski A (2024). High-Dimensional Undirected Graphical Models for Arbitrary Mixed Data. *Electronic Journal of Statistics* **18.1**, 2339–2404. <https://doi.org/10.1214/24-EJS2254>.
- Gorski J, Pfeuffer F, Klamroth K (2007). Biconvex Sets and Optimization with Biconvex Functions: A Survey and Extensions. *Mathematical Methods of Operations Research* **66.3**, 373–407. <https://doi.org/10.1007/s00186-007-0161-1>.
- Graz L, Barbanti L, Brandl R, Hothorn T (2026). Joint Count Transformation Models with Covariate-dependent Correlations. *Ecology*. Accepted for publication 2026-07-06. <https://doi.org/10.48550/arXiv.2201.13095>.

- Hermes S, van Heerwaarden J, Behrouzi P (2024). Copula Graphical Models for Heterogeneous Mixed Data. *Journal of Computational and Graphical Statistics* **33.3**, 991–1005. <https://doi.org/10.1080/10618600.2023.2289545>.
- Hirk R, Hornik K, Vana L (2019). Multivariate Ordinal Regression Models: An Analysis of Corporate Credit Ratings. *Statistical Methods & Applications* **28.3**, 507–539. <https://doi.org/10.1007/s10260-018-00437-7>.
- Hirk R, Hornik K, Vana L (2026). *mvord: Multivariate Ordinal Regression Models*. R package version 1.2.7. <https://doi.org/10.32614/CRAN.package.mvord>.
- Hofert M, Kojadinovic I, Maechler M, Yan J (2026). *copula: Multivariate Dependence with Copulas*. R package version 1.1-7. <https://doi.org/10.32614/CRAN.package.copula>.
- Hoff PD (2007). Extending the Rank Likelihood for Semiparametric Copula Estimation. *The Annals of Applied Statistics* **1.1**, 265–283. <https://doi.org/10.1214/07-AOAS107>.
- Hothorn T (2026a). *mlt: Most Likely Transformations*. R package version 1.8-2. <https://doi.org/10.32614/CRAN.package.mlt>.
- Hothorn T (2026b). *Multivariate Normal Log-likelihoods in the mvtnorm Package*. R package vignette version 1.4-2. <https://doi.org/10.32614/CRAN.package.mvtnorm>.
- Hothorn T, Möst L, Bühlmann P (2018). Most Likely Transformations. *Scandinavian Journal of Statistics* **45.1**, 110–134. <https://doi.org/10.1111/sjos.12291>.
- Jang ES, Jeong SH, Kim JW, Choi YS, Leissner P, Brechot C (2016). Diagnostic Performance of Alpha-Fetoprotein, Protein Induced by Vitamin K Absence, Osteopontin, Dickkopf-1 and Its Combinations for Hepatocellular Carcinoma. *PLOS One* **11.3**, e0151069. <https://doi.org/10.1371/journal.pone.0151069>.
- Joe H (1997). *Multivariate Models and Multivariate Dependence Concepts*. New York, U.S.A.: Chapman and Hall/CRC. <https://doi.org/10.1201/9780367803896>.
- Joe H (2005). Asymptotic Efficiency of the Two-stage Estimation Method for Copula-based Models. *Journal of Multivariate Analysis* **94.2**, 401–419. <https://doi.org/10.1016/j.jmva.2004.06.003>.
- Jöreskog KG (1994). On the Estimation of Polychoric Correlations and Their Asymptotic Covariance Matrix. *Psychometrika* **59.3**, 381–389. <https://doi.org/10.1007/BF02296131>.
- Khare K, Oh SY, Rahman S, Rajaratnam B (2019). A Scalable Sparse Cholesky Based Approach for Learning High-dimensional Covariance Matrices in Ordered Data. *Machine Learning* **108**, 2061–2086. <https://doi.org/10.1007/s10994-019-05810-5>.
- Klaassen CA, Wellner JA (1997). Efficient Estimation in the Bivariate Normal Copula Model: Normal Margins are Least Favourable. *Bernoulli* **3.1**, 55–77. <https://doi.org/10.2307/3318652>.
- Klein N, Hothorn T, Barbanti L, Kneib T (2022). Multivariate Conditional Transformation Models. *Scandinavian Journal of Statistics* **49**, 116–142. <https://doi.org/10.1111/sjos.12501>.
- Lesaffre E, Kaufmann H (1992). Existence and Uniqueness of the Maximum Likelihood Estimator for a Multivariate Probit Model. *Journal of the American Statistical Association* **87.419**, 805–811. <https://doi.org/10.2307/2290218>.
- Lindsey JK (1999). Some Statistical Heresies. *Journal of the Royal Statistical Society: Series D (The Statistician)* **48.1**, 1–40.

- Liu H, Lafferty J, Wasserman L (2009). The Nonparanormal: Semiparametric Estimation of High Dimensional Undirected Graphs. *Journal of Machine Learning Research* **10**.80, 2295–2328. URL: <http://jmlr.org/papers/v10/liu09a.html>.
- Lu J, Kolar M, Liu H (2018). Post-Regularization Inference for Time-Varying Nonparanormal Graphical Models. *Journal of Machine Learning Research* **18**.203, 1–78. URL: <http://jmlr.org/papers/v18/17-145.html>.
- Mai Q, He D, Zou H (2023). Coordinatewise Gaussianization: Theories and Applications. *Journal of the American Statistical Association* **118**.544, 2329–2343. <https://doi.org/10.1080/01621459.2022.2044825>.
- Mai Q, Zou H (2015). Sparse Semiparametric Discriminant Analysis. *Journal of Multivariate Analysis* **135**, 175–188. <https://doi.org/10.1016/j.jmva.2014.12.009>.
- Marra G, Radice R (2020). Copula Link-Based Additive Models for Right-Censored Event Time Data. *Journal of the American Statistical Association* **115**.530, 886–895. <https://doi.org/10.1080/01621459.2019.1593178>.
- Masarotto G, Varin C (2012). Gaussian Copula Marginal Regression. *Electronic Journal of Statistics* **6**, 1517–1549. <https://doi.org/10.1214/12-EJS721>.
- McLain AC, Ghosh SK (2013). Efficient Sieve Maximum Likelihood Estimation of Time-Transformation Models. *Journal of Statistical Theory and Practice* **7**.2, 285–303. <https://doi.org/10.1080/15598608.2013.772835>.
- Niehaus JM, Zhu L, Cook SJ, Jun M (2024). bizicount: Bivariate Zero-Inflated Count Copula Regression Using R. *Journal of Statistical Software* **109**.1, 1–42. <https://doi.org/10.18637/jss.v109.i01>.
- Nikoloulopoulos AK (2023). Efficient and Feasible Inference for High-dimensional Normal Copula Regression Models. *Computational Statistics & Data Analysis* **179**, 107654. <https://doi.org/10.1016/j.csda.2022.107654>.
- Owen AB (2001). *Empirical Likelihood*. New York, U.S.A.: Chapman and Hall/CRC. <https://doi.org/10.1201/9781420036152>.
- Papamakarios G, Nalisnick E, Rezende DJ, Mohamed S, Lakshminarayanan B (2021). Normalizing Flows for Probabilistic Modeling and Inference. *Journal of Machine Learning Research* **22**.57, 1–64. URL: <http://jmlr.org/papers/v22/19-1028.html>.
- Popovic GC, Hui FK, Warton DI (2018). A General Algorithm for Covariance Modeling of Discrete Data. *Journal of Multivariate Analysis* **165**, 86–100. <https://doi.org/10.1016/j.jmva.2017.12.002>.
- Prékopa A (1973). On Logarithmic Concave Measures and Functions. *Acta Scientiarum Mathematicarum* **34**, 335–343.
- Pritikin JN, Brick TR, Neale MC (2018). Multivariate Normal Maximum Likelihood with Both Ordinal and Continuous Variables, and Data Missing at Random. *Behavior Research Methods* **50**.2, 490–500. <https://doi.org/10.3758/s13428-017-1011-6>.
- R Core Team (2026). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. Vienna, Austria. <https://doi.org/10.32614/R.manuals>.
- Sewak A, Inácio V, Wu J, Benatar M, Hothorn T (2026a). Likelihood-based Modeling of Covariate-Specific Time-Dependent ROC Curves. *Statistical Methods in Medical Research* **35**.6, 1325–1346. <https://doi.org/10.1177/09622802261445416>.

- Sewak A, Siegfried S, Hothorn T (2026b). Transformation Discriminant Analysis for Constructing Optimal Biomarker Combinations. *Statistics in Medicine* **45**.15–17, e70665. <https://doi.org/10.1002/sim.70665>.
- Siegfried S, Hothorn T (2020). Count Transformation Models. *Methods in Ecology and Evolution* **11**.7, 818–827. <https://doi.org/10.1111/2041-210X.13383>.
- Siegfried S, Kook L, Hothorn T (2023). Distribution-Free Location-Scale Regression. *The American Statistician* **77**.4, 345–356. <https://doi.org/10.1080/00031305.2023.2203177>.
- Siegfried S, Tamási B, Hothorn T (2026). Smooth Transformation Models for Survival Analysis: A Tutorial Using R. *Statistical Methods in Medical Research* **35**.5, 980–997. <https://doi.org/10.1177/09622802251414595>.
- Song P XK, Li M, Yuan Y (2009). Joint Regression Analysis of Correlated Data Using Gaussian Copulas. *Biometrics* **65**.1, 60–68. <https://doi.org/10.1111/j.1541-0420.2008.01058.x>.
- Suggala AS, Yang E, Ravikumar P (2017). Ordinal Graphical Models: A Tale of Two Approaches. In: *Proceedings of the 34th International Conference on Machine Learning*. Ed. by Doina Precup and Yee Whye Teh. Vol. 70. Proceedings of Machine Learning Research. PMLR, pp. 3260–3269. URL: <https://proceedings.mlr.press/v70/suggala17a.html>.
- Xue L, Zou H (2012). Regularized Rank-based Estimation of High-dimensional Nonparanormal Graphical Models. *The Annals of Statistics* **40**.5, 2541–2571. <https://doi.org/10.1214/12-AOS1041>.
- Zhao Y, Udell M (2020). Missing Value Imputation for Mixed Data via Gaussian Copula. In: *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. New York, NY, USA: Association for Computing Machinery, pp. 636–646. <https://doi.org/10.1145/3394486.3403106>.