



Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Sturma, N;Drton, M;Leung, D

Title:

Testing many constraints in possibly irregular models using incomplete U-statistics

Date:

2024-09-01

Citation:

Sturma, N., Drton, M. & Leung, D. (2024). Testing many constraints in possibly irregular models using incomplete U-statistics. *Journal of the Royal Statistical Society Series B Statistical Methodology*, 86 (4), pp.987-1012. <https://doi.org/10.1093/jrsssb/qkae022>.

Persistent Link:

<https://hdl.handle.net/11343/351738>

License:

[cc-by-nc](#)

Testing many constraints in possibly irregular models using incomplete U -statistics

Nils Sturma¹ , Mathias Drton¹  and Dennis Leung²

¹Munich Center for Machine Learning and Department of Mathematics, School of Computation, Information and Technology, Technical University of Munich, Munich, Germany

²School of Mathematics and Statistics, University of Melbourne, Melbourne, Australia

Address for correspondence: Nils Sturma, Department of Mathematics, School of Computation, Information and Technology, Technical University of Munich, Boltzmannstr. 3, 85748 Garching, Germany. Email: nils.sturma@tum.de

Abstract

We consider the problem of testing a null hypothesis defined by equality and inequality constraints on a statistical parameter. Testing such hypotheses can be challenging because the number of relevant constraints may be on the same order or even larger than the number of observed samples. Moreover, standard distributional approximations may be invalid due to irregularities in the null hypothesis. We propose a general testing methodology that aims to circumvent these difficulties. The constraints are estimated by incomplete U -statistics, and we derive critical values by Gaussian multiplier bootstrap. We show that the bootstrap approximation of incomplete U -statistics is valid for kernels that we call mixed degenerate when the number of combinations used to compute the incomplete U -statistic is of the same order as the sample size. It follows that our test controls type I error even in irregular settings. Furthermore, the bootstrap approximation covers high-dimensional settings making our testing strategy applicable for problems with many constraints. The methodology is applicable, in particular, when the constraints to be tested are polynomials in U -estimable parameters. As an application, we consider goodness-of-fit tests of latent-tree models for multivariate data.

Keywords: Gaussian approximation, high dimensions, incomplete U -statistics, latent-tree model, multiplier bootstrap, non-asymptotic bound

JEL codes: 62F03, 62R01, 62E17

1 Introduction

Let $\{P_\theta : \theta \in \Theta\}$ be a statistical model with parameter space $\Theta \subseteq \mathbb{R}^d$. Given i.i.d. samples X_1, \dots, X_n from an unknown distribution P_θ , we are interested in testing

$$H_0 : \theta \in \Theta_0 \quad \text{vs.} \quad H_1 : \theta \in \Theta \setminus \Theta_0 \quad (1.1)$$

for a subset $\Theta_0 \subseteq \Theta$. In this paper, we consider the situation where the null hypothesis Θ_0 is defined by constraints, that is,

$$\Theta_0 = \{\theta \in \Theta : f_j(\theta) \leq 0 \text{ for all } j = 1, \dots, p\}, \quad (1.2)$$

where each constraint f_j is a function $f_j : \mathbb{R}^d \rightarrow \mathbb{R}$. The description of the null hypothesis also allows for equality constraints since $f_j(\theta) = 0$ can be equivalently described by $f_j(\theta) \leq 0$ and $-f_j(\theta) \leq 0$.

Received: July 18, 2022. Revised: January 17, 2024. Accepted: February 9, 2024

© The Royal Statistical Society 2024.

This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial License (<https://creativecommons.org/licenses/by-nc/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited. For commercial re-use, please contact journals.permissions@oup.com

These types of hypotheses are an important and general class and appear in a variety of statistical problems. Our work is in particular motivated by polynomial hypotheses where each f_j belongs to the ring $\mathbb{R}[\theta_1, \dots, \theta_d]$ of polynomials in the indeterminates $\theta_1, \dots, \theta_d$ with real coefficients. Examples of polynomial hypotheses feature in graphical modelling (Chen et al., 2017, 2014; Shiers et al., 2016; Sullivant et al., 2010), testing causal effects (Pearl, 2009; Spirtes et al., 2000; Steyer, 2005; Strieder et al., 2021), testing sub-determinants, tetrads, pentads and more in factor analysis models (Bollen & Ting, 2000; Drton et al., 2008, 2007; Drton & Xiao, 2016; Dufour et al., 2013; Gaffke et al., 2002; Leung & Drton, 2018; Silva et al., 2006), and in constraint-based causal discovery algorithms (Claassen & Heskes, 2012; Pearl & Verma, 1995; Spirtes et al., 2000).

The standard method for dealing with testing problems like equation (1.1) is the likelihood ratio (LR) test. However, a likelihood function may be multi-modal and difficult to maximize. If the LR test is not suitable, then one might instead make use of the implicit characterization of Θ_0 given by equation (1.2). In Wald-type tests, for example, the strategy is to form estimates of the involved functions f_1, \dots, f_p and aggregate them in a test statistic. Standard Wald tests require the number of restrictions p to be smaller than or equal to the dimension d . However, in many of the above examples, this might not be the case.

Another challenge for the classical LR and Wald test is that the null hypothesis Θ_0 may contain irregular points. For example, when the hypothesis is polynomial, it may contain singularities; a rigorous definition of singularities can be found in Drton (2009, Section 4.1) or Cox et al. (2015, Section 9). At singularities the rank of the Jacobian of the constraints being tested drops and the asymptotic behaviour of the LR test and the Wald-type test can be different than at regular points, resulting in an (also asymptotically) invalid test (Drton, 2009; Drton & Xiao, 2016; Dufour et al., 2013; Gaffke et al., 2002, 1999). In practice, it is unknown whether the true parameter θ is an irregular point, and it is therefore desirable to construct a test statistic for which one can give asymptotic approximations that accommodate and remain valid in irregular settings.

In this work, we propose a testing strategy that aims to cover set-ups where the number of restrictions p can be much larger than the sample size n and where the true parameter may be an irregular point. A precise definition of regular and irregular points is given later. Our method incorporates estimating the constraints f_1, \dots, f_p by an *incomplete* U -statistic. By first considering the commonly used *complete* U -statistic, we now give intuition for how this allows for high dimensionality and irregular points. Complete U -statistics provide an efficient method for unbiased estimation of $f := (f_1, \dots, f_p)$. We assume that $f(\theta)$ is (U -)estimable, i.e. for some integer m there exists a \mathbb{R}^p -valued measurable symmetric function $b(x_1, \dots, x_m)$ such that

$$\mathbb{E}_\theta[b(X_1, \dots, X_m)] = f(\theta) \quad \text{for all } \theta \in \Theta,$$

when X_1, \dots, X_m are i.i.d. with distribution P_θ . The U -statistic with kernel b is the average of $b(X_{i_1}, \dots, X_{i_m})$ over all distinct m -tuples (i_1, \dots, i_m) from $\{1, \dots, n\}$, in formulas

$$U_n = \frac{1}{|I_{n,m}|} \sum_{(i_1, \dots, i_m) \in I_{n,m}} b(X_{i_1}, \dots, X_{i_m}), \quad (1.3)$$

where $I_{n,m} = \{(i_1, \dots, i_m) : 1 \leq i_1 < \dots < i_m \leq n\}$. Due to the form of the null hypothesis Θ_0 , it is natural to define the test statistic as the maximum of the studentized U -statistic, i.e.

$$\max_{1 \leq j \leq p} \sqrt{n} U_{n,j} / \hat{\sigma}_j \quad (1.4)$$

and reject H_0 for ‘large’ values of it. Here, $U_{n,j}$ refers to the j th coordinate of U_n for all $j = 1, \dots, p$ and $\hat{\sigma}_j^2$ is a ‘good’ estimator of the asymptotic variance of $U_{n,j}$.

Example 1.1 As a leading example we consider testing of so-called ‘tetrad constraints’, a problem of particular relevance in factor analysis (Bollen & Ting, 2000; Drton & Xiao, 2016; Hipp & Bollen, 2003; Leung et al., 2016;

Spirtes et al., 2000) that can be traced back to Spearman (1904) and Wishart (1928). For a given symmetric matrix $\Sigma = (\sigma_{uv})$, a tetrad is an off-diagonal 2×2 sub-determinant. An example is $f_j(\Sigma) = \sigma_{uv}\sigma_{wz} - \sigma_{uz}\sigma_{vw}$ with four different indices u, v, w, z . If an l -dimensional normal distribution $N_l(0, \Sigma)$ follows a one-factor analysis model, where it is assumed that all variables are independent conditioned on one hidden factor, then all tetrads that can be formed from the covariance matrix Σ vanish. Thus, for given i.i.d. samples $X_1, \dots, X_n \sim N_l(0, \Sigma)$, one might be interested in testing whether all tetrads vanish simultaneously. It is easy to see that a tetrad $f_j(\Sigma) = \sigma_{uv}\sigma_{wz} - \sigma_{uz}\sigma_{vw}$ is estimable by the kernel

$$h_j(X_1, X_2) = \frac{1}{2} \{ (X_{1u}X_{1v}X_{2w}X_{2z} - X_{1u}X_{1z}X_{2v}X_{2w}) + (X_{2u}X_{2v}X_{1w}X_{1z} - X_{2u}X_{2z}X_{1v}X_{1w}) \}.$$

Interestingly, the resulting U -statistic corresponds to the ‘plug-in’ estimate, i.e. $U_{n,j} = \frac{n}{n-1} f_j(S)$, where $S = \frac{1}{n} \sum_{i=1}^n X_i X_i^T$ is the sample covariance. The plug-in estimate is considered in previous work on testing tetrads; see for example Shiers et al. (2016).

Critical values for the test statistic (1.4) can be derived by bootstrap methods that approximate the sampling distribution of U_n . Crucial to the validity of the bootstrap is the approximation by a Gaussian distribution. Recent progress in high-dimensional central limit theory yields valid Gaussian approximation of U -statistics in settings where $p \gg n$ is allowed. In particular, Chen (2018) and Chen and Kato (2020) derive finite-sample Berry–Esseen-type bounds on the Gaussian approximation in a two-step procedure: In the first step, the centred U -statistic is approximated by the linear component in the Hoeffding decomposition (a.k.a. the Hájek projection)

$$\frac{m}{n} \sum_{i=1}^n (g(X_i) - f(\theta)), \tag{1.5}$$

where $g(x) = \mathbb{E}_\theta[h(x, X_2, \dots, X_m)]$ and the expectation is taken with respect to the unknown distribution P_θ . In the second step, the linear term (1.5) is further approximated by a Gaussian random vector using the ‘classical’ central limit theorem for high-dimensional independent sums (Chernozhukov et al., 2013, 2017a). Crucially, this procedure assumes that the U -statistic is *non-degenerate*, that is, the individual variances $\sigma_{g,\theta,j}^2 = \text{Var}_\theta[g_j(X_1)]$ of the linear component do not vanish. Hence, the Berry–Esseen-type bound on the approximation relies on the standard assumption that the minimum $\underline{\sigma}_{g,\theta} = \min_{1 \leq j \leq p} \sigma_{g,\theta,j}^2$ is bounded away from zero. However, the minimum $\underline{\sigma}_{g,\theta}$ depends on the unknown parameter θ , and there may be certain points in Θ_0 where we have $\underline{\sigma}_{g,\theta} = 0$.

Definition 1.2 We say that a point $\theta \in \Theta$ is *regular* with respect to the kernel h if $\underline{\sigma}_{g,\theta} = \min_{1 \leq j \leq p} \sigma_{g,\theta,j}^2 > 0$. Otherwise, we say that θ is an *irregular* point.

If θ is an irregular point, then the Gaussian approximation of the U -statistic is not valid any more. This is illustrated in Example 1.3. Even if the parameter θ is only ‘close’ to an irregular point, the minimum $\underline{\sigma}_{g,\theta}$ can be very small. In this case, a very large sample size may be required for the Gaussian limiting distribution to provide a good approximation of the U -statistic since convergence is not uniform and the rate depends on the minimum $\underline{\sigma}_{g,\theta}$.

Example 1.3 Recall the kernel h_j of the tetrad $f_j(\Sigma) = \sigma_{uv}\sigma_{wz} - \sigma_{uz}\sigma_{vw}$ from Example 1.1. The corresponding random variable $g_j(X_1)$ in the Hájek projection (1.5) is given by

$$g_j(X_1) = \frac{1}{2} \{ (X_{1u}X_{1v}\sigma_{wz} - X_{1u}X_{1z}\sigma_{vw}) + (\sigma_{uv}X_{1w}X_{1z} - \sigma_{uz}X_{1v}X_{1w}) \}. \tag{1.6}$$

By inspecting equation (1.6), we see that $g_j(X_1)$ is degenerate if $\sigma_{wz} = \sigma_{vw} = \sigma_{uw} = \sigma_{uz} = 0$ and, in general, non-degenerate if at least one of the covariances is non-zero. Thus, the covariance matrices $\Sigma = (\sigma_{uv})$ in the one-factor analysis model that have $\sigma_{wz} = \sigma_{vw} = \sigma_{uw} = \sigma_{uz} = 0$ correspond to irregular points. Moreover, when the covariance matrix is only close to irregular, then the variance $\sigma_{g,\theta,j}^2$ might already be very small.

To accommodate irregularity, we propose to use *randomized incomplete U-statistics* instead of the usual, complete U -statistic from equation (1.3). That is, for a *computational budget parameter* $N \leq |I_{n,m}|$, we randomly choose on average N indices from $I_{n,m}$. Then the incomplete U -statistic is defined as the sample average of $h(X_{i_1}, \dots, X_{i_m})$ taken only over the subset of chosen indices. The test statistic is then formed as in equation (1.4) by replacing U_n with the incomplete counterpart.

The main theoretical contribution of this work is to show that the high-dimensional Gaussian approximation of incomplete U -statistics remains valid under *mixed degeneracy*, that is, for each index j , the variance of the Hájek projection $\sigma_{g,\theta,j}^2$ is allowed to take more or less arbitrary values, including zero. Therefore, our result covers testing hypotheses as in equation (1.2) when the underlying true parameter may be irregular or close to irregular. The setting of mixed degeneracy constitutes a further development of prior results in the literature which prove validity of the Gaussian approximation of incomplete U -statistics in either the fully non-degenerate case or the fully degenerate case (Chen & Kato, 2019; Song et al., 2019). That is, either *all* variances $\sigma_{g,\theta,j}^2$ are bounded away from zero or *all* of them are equal to zero. Our result is intermediate since we allow a different status of degeneracy for each index j . The crucial fact we exploit is that the asymptotic variance in the Gaussian approximation of the incomplete U -statistic is a weighted sum of the variance of the Hájek projection and the variance of the kernel itself. Hence, vanishing of the variance of the Hájek projection need not cause degeneracy of the asymptotic distribution. Indeed, the approximation is valid when we choose the computational budget parameter appropriately, typically of the same order as the sample size, i.e. $N = \mathcal{O}(n)$.

To derive critical values of our test statistic, we further approximate the limiting Gaussian distribution via a data-dependent Gaussian multiplier bootstrap as proposed in Chen and Kato (2019). We show that when choosing N appropriately, the bootstrap approximation remains trustworthy under mixed degeneracy and yields asymptotically valid critical values. The bootstrap is computationally feasible even for a large number of constraints p , as incomplete U -statistics also offer computational advantages over complete U -statistics. The computation of the complete U -statistic (1.3) requires $\mathcal{O}(n^m p)$ operations, which can be challenging for $m \geq 3$, while the incomplete U -statistic only requires $\mathcal{O}(Np)$ operations. We would like to highlight that these computational advantages were the main motivation for consideration of incomplete U -statistics in prior literature; in contrast, our work raises statistical advantages.

Example 1.4 Figure 1 shows histograms of simulated p -values for testing a large number of tetrad constraints in the one-factor analysis model when the true parameter matrix is close to an irregular point. In addition to our proposed strategy of using incomplete U -statistics, we include two other strategies: one that is based on complete U -statistics and another based on the LR test. The method using incomplete U -statistics yields p -values that are only slightly conservative (nearing a uniform distribution), while the other two methods fail drastically.

Our strategy can be applied for general hypotheses (1.2) as long as the functions $f_j(\theta)$ are estimable using a kernel function h . This is an advancement of Leung and Drton (2018), where only tetrads are considered. Leung and Drton (2018) also used a kernel function to estimate the tetrads, but the test statistic is an m -dependent average over the kernels instead of an incomplete U -statistic. The randomized incomplete U -statistic proposed in this work is much more flexible, supports broad application, and yields better results in simulations. Moreover, we show useful theoretical guarantees of incomplete U -statistics that yield an asymptotically valid test even at irregularities.

Remark 1.5 (Non-parametric set-ups). The setting we consider is formulated as pertaining to a parametric model with d -dimensional parameter space. However,

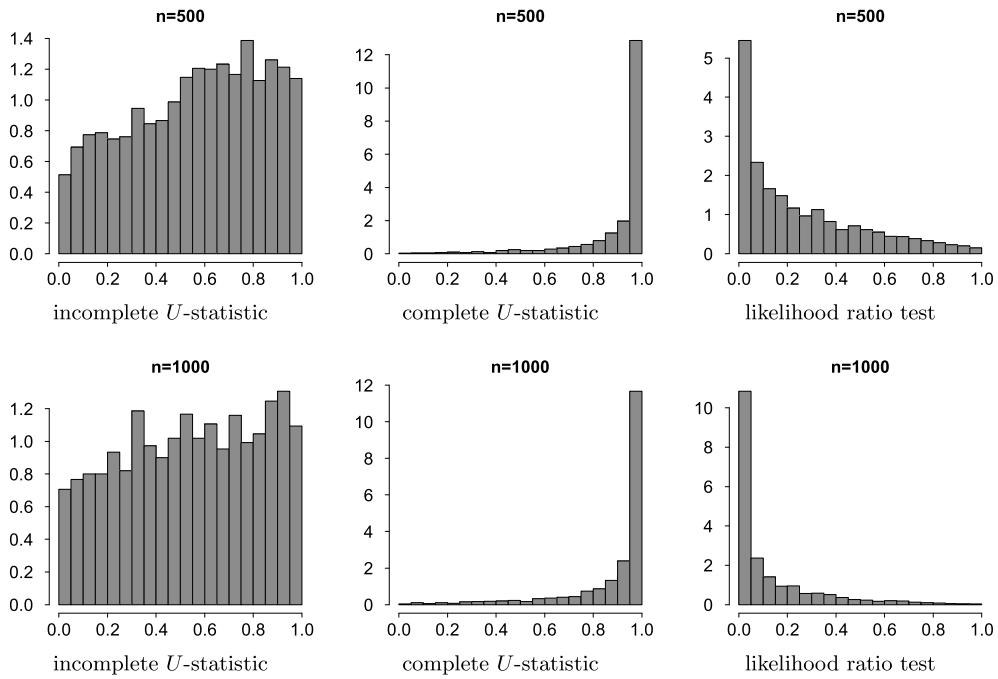


Figure 1. Histograms of 5,000 simulated p -values for simultaneously testing 2,730 tetrad constraints implied by the one-factor model with $l = 15$ observed variables. The computational budget parameter for the incomplete U -statistic is $N = 2n$ and the true covariance matrix is close to an irregular point; for exact parameter values, see Section 5, set-up (b).

our testing methodology applies without change to settings where we test constraints on a d -dimensional parameter $\theta(P)$ of the distributions in a non-parametric model. For example, one could consider testing tetrad constraints in covariance matrices of non-Gaussian distributions.

Remark 1.6 (Comparison to literature on shape restrictions). It is natural to compare our work to recent progress in the literature on shape restrictions that also considers testing equality and inequality constraints; see [Chetverikov et al. \(2018\)](#) for a review. In this line of work, the restricted parameter space $\Theta_0 \subseteq \Theta$ is usually considered to be infinite dimensional, so that more general parameters such as entire function classes are covered. On the other hand, to the best of our knowledge, the methods do not explicitly focus on set-ups with a large amount of constraints and possibly irregular points. For example, in [Chernozhukov et al. \(2023b\)](#), the authors consider a test statistic that minimizes a generalized method of moments objective function over the restricted and the whole parameter space and compares the difference. It is the main goal to study the behaviour of the test statistic in regions near the boundary of the restricted parameter space Θ_0 . The different focus is reflected in the conditions the authors of [Chernozhukov et al. \(2023b\)](#) assume for their theoretical analysis. In particular, they assume that the Jacobian matrix of the equality constraints has full row rank in a neighbourhood of the true parameter. This implies that the maximal number of equality constraints is smaller than the dimension of the parameter space and that algebraic singularities are excluded since the Jacobian is not allowed to drop rank. Thus, the conditions do not allow for many equality constraints and irregular points. Moreover, the test statistic requires an optimization over the restricted and the whole parameter space which requires extra assumptions such as

convexity and compactness and can be difficult to implement in practice. In contrast, our method is optimization free and does not require any assumptions on the parameter space.

1.1 Organization of the paper

In Section 2, we give non-asymptotic Berry–Esseen-type bounds for high-dimensional Gaussian and bootstrap approximation of incomplete U -statistics. Importantly, the incomplete U -statistic is assumed to be of mixed degeneracy. In Section 3, we propose our testing methodology by formally defining the test statistic and showing how to derive critical values. Our results on incomplete U -statistics yield that the test is asymptotically valid and consistent even in irregular settings. In Section 4, we show that our test is applicable to polynomial hypotheses by explaining a general method for constructing a kernel b . In Section 5, we then apply our strategy for testing the goodness-of-fit of latent tree models of which the one-factor analysis model is a special case. In numerical experiments we compare our strategy with the LR test. The [online supplementary material](#) contains additional material such as all technical proofs ([Appendix A, online supplementary material](#)), additional lemmas ([Appendix B, online supplementary material](#)), properties of sub-Weibull random variables ([Appendix C, online supplementary material](#)), and additional simulation results for Gaussian latent-tree models ([Appendix D, online supplementary material](#)). Moreover, we provide a second application of our methodology in [Appendix E, online supplementary material](#), where we test minors in two-factor analysis models.

1.2 Notation

Given $\beta \in (0, \infty)$, we define the function $\psi_\beta(x) = \exp(x^\beta) - 1$ for $x > 0$. A real random variable Y is said to be *sub-Weibull of order β* if $\|Y\|_{\psi_\beta} := \inf\{t > 0 : \mathbb{E}[\psi_\beta(|Y|/t)] \leq 1\}$ is finite. We use the usual convention that $\inf\{\emptyset\} = \infty$. For $\beta \geq 1$ we have that $\|Y\|_{\psi_\beta}$ is a norm, while for $\beta \in (0, 1)$ it is only a quasi-norm. If $\|Y\|_{\psi_1} < \infty$, then Y is called a *sub-Exponential* random variable and if $\|Y\|_{\psi_2} < \infty$, then Y is called *sub-Gaussian*. For a random element Y , let $P|_Y(\cdot)$ and $\mathbb{E}|_Y[\cdot]$ denote the conditional probability and expectation given Y . We denote a sequence of random variables $Y_i, \dots, Y_{i'}$ by $Y_i^{i'}$ for $i \leq i'$ and for a tuple of indices $\iota = (i_1, \dots, i_m)$ we write $Y_\iota = (Y_{i_1}, \dots, Y_{i_m})$.

For $a, b \in \mathbb{R}$, define $a \vee b = \max\{a, b\}$ and $a \wedge b = \min\{a, b\}$. For $a, b \in \mathbb{R}^p$, we write $a \leq b$ if $a_j \leq b_j$ for all $j = 1, \dots, p$, and we write $[a, b]$ for the hyper-rectangle $\prod_{j=1}^p [a_j, b_j]$. If $a \leq b$, then the hyper-rectangle $[a, b]$ is non-empty but if there is at least one index $j \in \{1, \dots, p\}$ such that $a_j > b_j$, then the hyper-rectangle is equal to the empty set. The class of hyper-rectangles in \mathbb{R}^p is denoted by $\mathbb{R}_{\text{re}}^p = \{\prod_{j=1}^p [a_j, b_j] : a_j, b_j \in \mathbb{R} \cup \{-\infty, \infty\}\}$. For a vector $a \in \mathbb{R}^p$ and $r, t \in \mathbb{R}$, we write $ra + t$ for the vector in \mathbb{R}^p with j th component $ra_j + t$. Finally, for a vector $a \in \mathbb{R}^p$ and two integers p_1, p_2 such that $p_1 + p_2 = p$ we write $a = (a^{(1)}, a^{(2)})$, where $a_j^{(1)} = a_j$ for all $j = 1, \dots, p_1$ and $a_j^{(2)} = a_{p_1+j}$ for all $j = 1, \dots, p_2$. Let $\|A\|_\infty = \max_{i,j} |a_{ij}|$ be the element-wise maximum norm of a matrix $A = (a_{ij})$.

2 Incomplete U -statistics under mixed degeneracy

Suppose we are given i.i.d. samples X_1, \dots, X_n from an unknown distribution P in a statistical model. For some integer m , let $b(x_1, \dots, x_m)$ be a fixed \mathbb{R}^p -valued measurable function that is symmetric in its arguments. In this section, we consider the general case of inference on the mean vector $\mathbb{E}[b(X_1, \dots, X_m)] = (\mu_1, \dots, \mu_p)^\top = \mu$, where μ is an arbitrary estimable parameter of the underlying distribution P . In our set-up, where we want to test hypotheses characterized by constraints, the distribution P depends on θ and μ is given by the constraints $f(\theta) = (f_1(\theta), \dots, f_p(\theta))$.

We begin with the formal definition of randomized incomplete U -statistics using similar notation as in [Chen and Kato \(2019\)](#) and [Song et al. \(2019\)](#). Let $N \leq \binom{n}{m}$ be a computational budget parameter and generate i.i.d. Bernoulli random variables $\{Z_\iota : \iota \in I_{n,m}\}$ with success probability $\rho_n = N / \binom{n}{m}$. Then the incomplete U -statistics based on Bernoulli sampling is defined by

$$U'_{n,N} = \frac{1}{N} \sum_{\iota \in I_{n,m}} Z_\iota b(X_\iota), \quad (2.1)$$

where $\hat{N} = \sum_{i \in I_{n,m}} Z_i$ is the number of successes. The variable \hat{N} follows a Binomial distribution with parameters $(|I_{n,m}|, \rho_n)$. Therefore, $\mathbb{E}(\hat{N}) = |I_{n,m}| \rho_n = N$, and the incomplete U -statistic is on average a sum over N objects. Thus, we may view the computational budget N as a sparsity parameter for the incomplete U -statistic. We denote by $\sigma_{b,j}^2 = \mathbb{E}[(h_j(X_1^m) - \mu_j)^2]$ the variance of the j th coordinate of the kernel and by $\sigma_{g,j}^2 = \mathbb{E}[(g_j(X_1) - \mu_j)^2]$ the variance of $g_j(X_1)$; recall that the Hájek projection is given by $g(x) = \mathbb{E}[h(x, X_2, \dots, X_m)]$. Note that $\sigma_{b,j}^2$ and $\sigma_{g,j}^2$ depend on the underlying distribution P as emphasized in the introduction; however, in the rest of the paper, we omit the explicit dependence to simplify notation.

2.1 Gaussian approximation

We will derive non-asymptotic Gaussian approximation error bounds for the incomplete U -statistic $U'_{n,N}$ that allow for mixed degenerate kernels h when choosing the computational budget parameter N appropriately. To state the formal approximation results, we assume $2 \leq m \leq \sqrt{n}$, $n \geq 4$, $p \geq 3$ and $\rho_n = N/|I_{n,m}| < 1/2$. We start by making assumptions on the moment structure of the kernel h before formally introducing mixed degeneracy. Let $\beta \in (0, 1]$ and suppose there exists a constant $D_n \geq 1$ such that:

- (C1) $\mathbb{E}[|h_j(X_1^m) - \mu_j|^{2+l}] \leq \sigma_{b,j}^2 D_n^l$ for all $j = 1, \dots, p$ and $l = 1, 2$.
- (C2) $\|h_j(X_1^m) - \mu_j\|_{\psi_\beta} \leq D_n$ for all $j = 1, \dots, p$.
- (C3) There exists $\underline{\sigma}_n^2 > 0$ such that $\underline{\sigma}_n^2 \leq \min_{1 \leq j \leq p} \sigma_{h,j}^2$.

Assuming conditions similar to (C1)–(C3) is standard in high-dimensional Gaussian approximation theory. Condition (C2) assumes that the kernel h is sub-Weibull. In prior work on Gaussian approximation of high-dimensional U -statistics (Chen, 2018; Chen & Kato, 2019), the authors usually consider sub-Exponential kernels with $\beta = 1$. However, the kernel we propose in Section 4 for testing polynomial hypotheses will typically be sub-Weibull and not sub-Exponential, also see Example 4.2. We discuss important properties of sub-Weibull random variables in Appendix C, online supplementary material. Note that we allow the bound D_n to depend on the sample size n since, in the high-dimensional setting, the distribution P may depend on n . Condition (C1) is of a more technical nature and serves for clear presentation. In principle, it would be possible to omit this assumption, but the resulting error bound would be more complicated. Finally, Condition (C3) requires that the minimal variance of the individual kernels h_j is bounded away from zero, even for large p . Put differently, we only consider kernels h_j such that $h_j(X_1^m)$ is not almost surely constant. It remains to make assumptions with respect to the degeneracy of the Hájek projection, i.e. we formally define mixed degeneracy.

Definition 2.1 Let $p_1, p_2 \in \mathbb{N}_{\geq 0}$ such that $p_1 + p_2 = p$. We say that the kernel h , or also simply the incomplete U -statistic, is *mixed degenerate* for distribution P if the following two conditions are satisfied:

- (C4) There exists $\underline{\sigma}_{g^{(1)}}^2 > 0$ such that $\underline{\sigma}_{g^{(1)}}^2 \leq \min_{1 \leq j \leq p_1} \sigma_{g,j}^2$.
- (C5) There exists $k > 0$ such that $\|g_j(X_1) - \mu_j\|_{\psi_\beta} \leq n^{-k} D_n$ for all $j = p_1 + 1, \dots, p$.

In other words, mixed degeneracy of a kernel requires that each index $j = 1, \dots, p$ either satisfies condition (C4) or (C5). By rearranging the indices, we then find that the first p_1 indices satisfy Condition (C4) and the remaining indices satisfy Condition (C5). Note that whether a kernel is mixed degenerate or not depends on the underlying distribution P from which the samples X_1, \dots, X_n are drawn.

Assuming mixed degenerate kernels is the main difference in comparison to the existing literature on Gaussian approximation of high-dimensional incomplete U -statistics (Chen & Kato, 2019; Song et al., 2019). Usually, either the non-degenerate case where $p_2 = 0$ or the fully degenerate case where $p_1 = 0$ and $\sigma_{g,j}^2 = 0$ for all $j = 1, \dots, p$ are treated. In contrast, we allow for a different status of degeneracy of the Hájek projection for each index j . In particular, Condition (C5) covers degenerate cases where $\sigma_{g,j}^2 = 0$ since zero variance implies that $\|g_j(X_1) - \mu_j\|_{\psi_\beta} = 0$ almost

surely. But our assumptions even allow for more flexibility. For example, it may be the case that the variance $\sigma_{g,j}^2$ decreases with the sample size n . In this case, Condition (C5) is an assumption on the rate of convergence to degeneracy, i.e. the rate is polynomial in n .

Remark 2.2 (Discussion on mixed degeneracy). The notion of mixed degeneracy is particularly interesting for kernels h , where we do not know whether the individual components h_j are degenerate or non-degenerate, for example when the underlying distribution is unknown. If the number of constraints p does *not* grow with the sample size, then mixed degeneracy holds for any distribution P . Indeed, by letting $\underline{\sigma}_{g^{(1)}}^2$ be the minimum of the non-zero variances $\sigma_{g,j}^2$, we see that Condition (C4) is satisfied. All other indices j have $\sigma_{g,j}^2 = 0$ which implies that Condition (C5) is also satisfied. However, we emphasize that mixed degeneracy is a more subtle condition if the number of constraints p is growing.

Remark 2.3 (Parametric families and irregular points). In our testing problem (1.1), we are considering a parametric family of distributions $\{P_\theta, \theta \in \Theta\}$ but the true parameter θ is unknown. In this case, the variances $\sigma_{g,j}^2$ of the individual Hájek projections depend on θ as outlined in the introduction. If the number of constraints p does *not* grow with the sample size, then mixed degeneracy holds uniformly over the whole parameter space Θ , as we have seen in Remark 2.2. However, the rate of convergence of the incomplete U -statistic depends on the minimum $\underline{\sigma}_{g^{(1)}}^2$. If $\underline{\sigma}_{g^{(1)}}^2$ is really small, the index corresponding to the minimum may already satisfy (C5) for large sample sizes n , so $\underline{\sigma}_{g^{(1)}}^2$ can in fact be chosen larger. Therefore, mixed degeneracy is also suitable for points $\theta \in \Theta$ that are close to irregular, see Corollary 2.7 for a precise statement.

In principle, it would be possible to extend our results on more general sequences γ_n converging to zero instead of n^{-k} in Condition (C5), but this would result in more involved error bounds. For simplicity, we also assume $p_1, p_2 \geq 3$, even though one could specify the bounds for arbitrary p_1 and p_2 . Our last technical assumption is similar to Condition (C1) and is also necessary for the sake of clear presentation:

$$(C6) \mathbb{E}[|g_j(X_1) - \mu_j|^{2+l}] \leq \sigma_{g,j}^2 D_n^l \text{ for all } j = 1, \dots, p \text{ and } l = 1, 2.$$

Now, we state our main result that specifies a non-asymptotic error bound on the Gaussian approximation of incomplete U -statistics. We define $a_n = n/N$, $\Gamma_b = \text{Cov}[h(X_1^m)]$ and $\Gamma_g = \text{Cov}[g(X_1)]$. For notational convenience, we further define the quantities

$$\omega_{n,1} = \left(\frac{m^{2/\beta} D_n^2 \log(pn)^{1+6/\beta}}{(\underline{\sigma}_{g^{(1)}}^2 \wedge \underline{\sigma}_h^2 \wedge 1) (n \wedge N)} \right)^{1/6}, \quad \omega_{n,2} = \frac{N^{1/2} m^2 D_n \log(pn)^{1/2+2/\beta}}{\underline{\sigma}_h n^{\min\{1/2+k, 5/6\}}},$$

$$\omega_{n,3} = \left(\frac{Nm^2 D_n^2 \log(p)^2}{(\underline{\sigma}_h^2 \wedge 1) n^{\min\{1+k, m\}}} \right)^{1/3}.$$

Theorem 2.4 Assume Conditions (C1)–(C6) hold. Then there is a constant $C_\beta > 0$ only depending on β such that

$$\sup_{R \in \mathbb{R}^p_{fc}} |P(\sqrt{n}(U'_{n,N} - \mu) \in R) - P(Y \in R)| \leq C_\beta \{\omega_{n,1} + \omega_{n,2} + \omega_{n,3}\},$$

where $Y \sim N_p(0, m^2 \Gamma_g + a_n \Gamma_b)$.

Theorem 2.4 shows that the distribution of $\sqrt{n}(U'_{n,N} - \mu)$ can be approximated by the Gaussian distribution $Y \sim N_p(0, m^2\Gamma_g + \alpha_n\Gamma_b)$ under mixed degeneracy. Since the computational budget parameter N occurs in the numerator of the bound, one has to choose it in proportion to the sample size such that the bound vanishes. In particular, we note that under the regime $N = \mathcal{O}(n)$ the bound vanishes when treating other quantities as constants.

Example 2.5 Let N be of the same order as the sample size. Then α_n can be viewed as a constant and each coordinate of Y is asymptotically non-degenerate. In this case, when we assume $k \geq 1/3$ and treat $m, \underline{\sigma}_{g(1)}, \underline{\sigma}_b$ and D_n as fixed constants, the bound $C_\beta\{\omega_{n,1} + \omega_{n,2} + \omega_{n,3}\}$ vanishes asymptotically under Conditions (C1)–(C6) if the dimension p satisfies $\log(pn)^{3/2+6/\beta} = \mathcal{O}(n)$. On the other hand, one can also choose $N = n^{1+\varepsilon}$ for small enough $\varepsilon > 0$. Then, $\alpha_n \rightarrow 0$ as $n \rightarrow \infty$ and the distribution of Y may become asymptotically degenerate. Nevertheless, the bound still vanishes if we fix all other quantities as before and if the dimension p satisfies $\log(pn)^{\frac{3/2+6/\beta}{1/3-\varepsilon/2}} = \mathcal{O}(n)$.

Remark 2.6 (Order of the kernel). The Berry–Esseen-type bound in Theorem 2.4 depends explicitly on the order of the kernel b . In particular, the result also allows for kernels of diverging order, i.e. increasing m . However, larger m imply worse performance of the Gaussian approximation in terms of the required sample size. Also, if m is increasing with n one has to be careful in choosing the computational budget N ; it has to be chosen smaller to achieve convergence.

The proof of Theorem 2.4 relies on the seminal papers of Chernozhukov et al. (2013) and Chernozhukov et al. (2017a) on Gaussian approximation of high-dimensional independent sums, and it extends the results of Chen and Kato (2019) and Song et al. (2019) on incomplete U -statistics to the mixed degenerate case. Obtaining sharper bounds for the high-dimensional approximation of independent sums than in the original papers is an ongoing area of research, see for example Fang and Koike (2021), Chernozhukov et al. (2022), Lopes (2022) and Chernozhukov et al. (2023a).

The bound in Theorem 2.4 is stated as general as possible. Importantly, it entails that the incomplete U -statistic can be approximated by $N_p(0, m^2\Gamma_g + \alpha_n\Gamma_b)$ even in irregular set-ups of our testing problem (1.1) as long as $\underline{\sigma}_{g(1)}^2 > 0$. However, it might be difficult to read off precise rates for the speed of convergence in the *close to irregular* scenarios. For large n , it is possible that there are close to irregular points such that some indices $j \in \{1, \dots, p\}$ only satisfy mixed degeneracy if one chooses $\underline{\sigma}_{g(1)}^2$ relatively small or k small. On the other hand, under further distributional assumptions, we manage to improve the bound in Theorem 2.4 so that the speed of convergence is completely independent to the irregularity status of the hypothesis.

Corollary 2.7 Assume that X_1, \dots, X_n are i.i.d. samples of a Gaussian distribution and assume that each individual kernel b_1, \dots, b_p is a non-constant polynomial of degree at most $2s$. Suppose that $\mu = \mathbb{E}[b(X_1^m)] = 0$ and that there exists $\underline{\sigma}_b^2 > 0$ such that $\underline{\sigma}_b^2 \leq \min_{1 \leq j \leq p} \sigma_{b,j}^2$. Then there exist $\beta \in (0, 1]$ and $D_n \geq 1$ such that (C1), (C2) and (C6) are satisfied and the kernel b is mixed degenerate. If, additionally, $n \leq N \leq Cn$ for some constant $C > 0$, then it holds that

$$\sup_{R \in \mathbb{R}_c^p} |P(\sqrt{n}U'_{n,N} \in R) - P(Y \in R)| \leq C_{s,m} \frac{(\bar{\sigma}_b^2 \vee 1) \log(pn)^{1/2+2s}}{(\underline{\sigma}_b^2 \wedge 1)n^{1/9}},$$

where $Y \sim N_p(0, m^2\Gamma_g + \alpha_n\Gamma_b)$, $\bar{\sigma}_b^2 = \max_{1 \leq j \leq p} \sigma_{b,j}^2$ and $C_{s,m} > 0$ is a constant only depending on s and m .

In the bound in Corollary 2.7, there is no $\sigma_{g(i)}^2$ and no k showing up anymore, so the convergent speed shows no dependence on how close to irregular the points in the null hypothesis really are. Even though the bound might not be optimal, it completely mathematizes our intuition why incomplete U-statistics are a good choice for a test statistic to guard against irregular points.

Example 2.8 Each coordinate h_j of the proposed kernel in Example 1.1 is a polynomial of degree 4 in Gaussian variables, which implies that Corollary 2.7 is applicable with $s = 2$. See Section 4 for a general method to construct polynomial kernels.

2.2 Bootstrap approximation

Since the covariance matrix $m^2\Gamma_g + \alpha_n\Gamma_b$ of the approximating Gaussian distribution in Theorem 2.4 is typically unknown in statistical applications, we apply a Gaussian multiplier bootstrap. The procedure is exactly the same as in Chen and Kato (2019) and Song et al. (2019) but their error bounds on the approximation require non-degeneracy, that is, there is a constant $c > 0$ such that $\min_{1 \leq j \leq p} \sigma_{g,j}^2 \geq c$. Under mixed degeneracy, this is no longer the case. However, when choosing the computational budget appropriately, we prove that the bootstrap approximation still holds.

The bootstrap is based on the fact that the random vector $Y \sim N_p(0, m^2\Gamma_g + \alpha_n\Gamma_b)$ is a weighted sum of the two independent random vectors $Y_g \sim N_p(0, \Gamma_g)$ and $Y_b \sim N_p(0, \Gamma_b)$, i.e. $Y = mY_g + \sqrt{\alpha_n}Y_b$. Let $\mathcal{D}_n = \{X_1, \dots, X_n\} \cup \{Z_i : i \in I_{n,m}\}$ be the data involved in the definition of the incomplete U-statistic $U'_{n,N}$. We will construct data-dependent random vectors $U^{\#}_{n_1,g}$ and $U^{\#}_{n_1,b}$ such that, given the data \mathcal{D}_n , both vectors are independent and approximate Y_g and Y_b .

To approximate the distribution of Y_b take a collection $\{\zeta'_i : i \in I_{n,m}\}$ of independent $N(0, 1)$ random variables that are also independent of \mathcal{D}_n . Define the multiplier bootstrap

$$U^{\#}_{n,b} = \frac{1}{\sqrt{N}} \sum_{i \in I_{n,m}} \zeta'_i \sqrt{Z_i} (h(X_i) - U'_{n,N}) \tag{2.2}$$

and observe that, conditioned on the data \mathcal{D}_n , the distribution of $U^{\#}_{n,b}$ is Gaussian with mean zero and covariance matrix $\widehat{N}^{-1} \sum_{i \in I_{n,m}} Z_i (h(X_i) - U'_{n,N})(h(X_i) - U'_{n,N})^T$. Intuitively, this covariance matrix should be a good estimator of the true covariance matrix Γ_b and therefore the distribution of $U^{\#}_{n,b}$ should be ‘close’ to the distribution of Y_b .

Approximating $Y_g \sim N_p(0, \Gamma_g)$ is more involved since the Hájek projection (1.5) is in general unknown. Thus, we first construct estimates G_{i_1} of $g(X_{i_1})$ for each i_1 in a chosen subset $S_1 \subseteq \{1, \dots, n\}$ with cardinality $n_1 = |S_1|$. Then, we consider the multiplier bootstrap distribution

$$U^{\#}_{n_1,g} = \frac{1}{\sqrt{n_1}} \sum_{i_1 \in S_1} \zeta_{i_1} (G_{i_1} - \overline{G}), \tag{2.3}$$

where $\{\zeta_{i_1} : i_1 \in S_1\}$ is a collection of independent $N(0, 1)$ random variables that is independent of \mathcal{D}_n and $\{\zeta'_i : i \in I_{n,m}\}$. Here, $\overline{G} = n_1^{-1} \sum_{i_1 \in S_1} G_{i_1}$ denotes the average of the constructed estimates. The exact form of the estimates G_{i_1} is specified later. Similar as above, the distribution of $U^{\#}_{n_1,g}$ given the data \mathcal{D}_n should be ‘close’ to the distribution of Y_g . Combining $U^{\#}_{n,b}$ and $U^{\#}_{n_1,g}$, we obtain the multiplier bootstrap $U^{\#}_{n,n_1} = mU^{\#}_{n_1,g} + \sqrt{\alpha_n}U^{\#}_{n,b}$. It approximates the distribution of Y even under mixed degeneracy, as we verify in our next Lemma. The approximation of $U^{\#}_{n,n_1}$ will depend on the quality of the estimator G_{i_1} that we measure by the quantity

$$\widehat{\Delta}_{g,1} = \max_{1 \leq j \leq p} \frac{1}{n_1} \sum_{i_1 \in S_1} (G_{i_1,j} - g_j(X_{i_1}))^2. \tag{2.4}$$

Moreover, the approximation depends on conditions involving the quantity

$$A_n := A_n(\underline{\sigma}_b, m, \beta, N) = \frac{m^{\max\{2/\beta, 4\}}}{\underline{\sigma}_b^4 \wedge 1} \max\{(N/n)^2, 1\},$$

which is required to be small. In particular, if m does not grow with n and the computational budget N is chosen appropriately, then Assumption (C3) ensures that A_n does not inflate.

Lemma 2.9 Assume the Conditions (C1)–(C3) hold. If

$$\frac{A_n D_n^4 \log(pn)^{3+4/\beta}}{n_1 \wedge N} \leq C_1 n^{-\zeta_1} \tag{2.5}$$

and

$$P\left(A_n D_n^2 \widehat{\Delta}_{g,1} \log(p)^4 > C_1 n^{-\zeta_2}\right) \leq \frac{C_1}{n} \tag{2.6}$$

for some constants $C_1 > 0$ and $\zeta_1, \zeta_2 \in (0, 1)$, then there exists a constant $C > 0$ depending only on β, ζ_1 and C_1 such that with probability at least $1 - C/n$,

$$\sup_{R \in \mathbb{R}_{rc}^p} \left| P|_{\mathcal{D}_n}(U_{n,n_1}^\# \in R) - P(Y \in R) \right| \leq C n^{-(\zeta_1 \wedge \zeta_2)/6}.$$

Note that Lemma 2.9 formally does not require mixed degeneracy, that is, it is completely independent of the status of degeneracy of the kernel. Now, we specify G_{i_1} to be a special case of the divide-and-conquer estimator \hat{g}_{i_1} from Chen and Kato (2019) and Song et al. (2019) that is defined as follows. For each index $i_1 \in S_1$, we partition the remaining indices, $\{1, \dots, n\} \setminus \{i_1\}$, into disjoint subsets $\{S_{2,k}^{(i_1)}, k = 1, \dots, K\}$, each of size $m - 1$, where $K = \lfloor (n - 1)/(m - 1) \rfloor$. Then, we define for each $i_1 \in S_1$ the estimator

$$\hat{g}_{i_1} = \frac{1}{K} \sum_{k=1}^K b(X_{i_1}, X_{S_{2,k}^{(i_1)}}). \tag{2.7}$$

Thus, from now on, we again refer by $U_{n_1,g}^\#$ to the statistics defined in equation (2.3) but with the specialized estimator $G_{i_1} = \hat{g}_{i_1}$ as defined in equation (2.7). The next theorem builds on Lemma 2.9 and verifies that the bootstrap approximation is valid for this specialized estimator.

Theorem 2.10 Assume the Conditions (C1)–(C3) hold and

$$\frac{A_n D_n^4 \log(pn)^{3+4/\beta}}{n_1 \wedge N} \leq C_1 n^{-\zeta} \tag{2.8}$$

for some constants $C_1 > 0, \zeta \in (0, 1)$. Then, for any $v \in (\max\{7/6, 1/\zeta\}, \infty)$, there exists a constant $C > 0$ depending only on β, v, ζ and C_1 such that with probability at least $1 - C/n$,

$$\sup_{R \in \mathbb{R}_{rc}^p} \left| P|_{\mathcal{D}_n}(U_{n,n_1}^\# \in R) - P(Y \in R) \right| \leq C n^{-(\zeta-1/v)/6}.$$

Theorem 2.10 says that we can approximate the asymptotic Gaussian distribution on the hyper-rectangles via the multiplier bootstrap $U_{n,n_1}^\# = m U_{n_1,g}^\# + \sqrt{a_n} U_{n,b}^\#$. As long as the computational budget N is chosen appropriately, for example $N = \mathcal{O}(n)$, this holds independently of the status of degeneracy since we do not require Conditions (C4)–(C6). Crucially, we are able to simulate the distribution of $U_{n,n_1}^\#$ given the data by generating independent $N(0, 1)$ random variables. The complexity of the bootstrap procedure, and therefore the complexity of our testing

methodology, is discussed in Remark 3.2. In a corollary, we combine the Gaussian approximation with the bootstrap approximation.

Corollary 2.11 Assume the Conditions (C1)–(C6) hold. Further assume that for some constants $C_1 > 0$, $\zeta \in (0, 1)$, Condition (2.8) holds and $\omega_{n,1} + \omega_{n,2} + \omega_{n,3} \leq C_1 n^{-\zeta/7}$. Then there exists a constant $C > 0$ depending only on β, ζ and C_1 such that with probability at least $1 - C/n$,

$$\sup_{R \in \mathbb{R}_{re}^p} \left| P(\sqrt{n}(U'_{n,N} - \mu) \in R) - P|_{\mathcal{D}_n}(U^{\#}_{n,n_1} \in R) \right| \leq Cn^{-\zeta/7}.$$

Proof. This follows from Theorems 2.4 and 2.10 with $v = 7/\zeta$. □

2.3 Studentization

Often, the approximate variances of the coordinates of $U'_{n,N}$ are heterogeneous, and it is therefore desirable to studentize the incomplete U -statistic. For $j = 1, \dots, p$, we denote by $\sigma_j^2 = m^2 \sigma_{g,j}^2 + \alpha_n \sigma_{b,j}^2$ the variance of the j th coordinate of the approximating Gaussian random vector Y , that is, σ_j^2 is equal to the diagonal element $m^2 \Gamma_{g,jj} + \alpha_n \Gamma_{b,jj}$ of the approximating covariance matrix. In line with this, we define the empirical variances $\hat{\sigma}_j^2 = m^2 \hat{\sigma}_{g,j}^2 + \alpha_n \hat{\sigma}_{b,j}^2$ to be the diagonal elements of the conditional covariance matrix of the bootstrap distribution $U^{\#}_{n,n_1}$ given the data. Therefore, $\hat{\sigma}_{g,j}^2$ and $\hat{\sigma}_{b,j}^2$ are given by

$$\hat{\sigma}_{g,j}^2 = \frac{1}{n_1} \sum_{i_1 \in \mathcal{S}_1} (\hat{g}_{i_1,j} - \bar{g}_j)^2 \quad \text{and} \quad \hat{\sigma}_{b,j}^2 = \frac{1}{\bar{N}} \sum_{i \in I_{n,m}} Z_i(b_j(X_i) - U'_{n,N,j})^2.$$

Moreover, we define a $p \times p$ diagonal matrix $\hat{\Lambda}$ with diagonal elements $\hat{\Lambda}_{jj} = m^2 \hat{\sigma}_{g,j}^2 + \alpha_n \hat{\sigma}_{b,j}^2$ for all $j = 1, \dots, p$.

Corollary 2.12 Assume the conditions in Corollary 2.11. Then there exists a constant $C > 0$ depending only on β, ζ and C_1 such that with probability at least $1 - C/n$,

$$\sup_{R \in \mathbb{R}_{re}^p} \left| P(\sqrt{n} \hat{\Lambda}^{-1/2}(U'_{n,N} - \mu) \in R) - P|_{\mathcal{D}_n}(\hat{\Lambda}^{-1/2}U^{\#}_{n,n_1} \in R) \right| \leq Cn^{-\zeta/7}.$$

Corollary 2.12 allows us to construct a test for hypotheses of the form (1.2) that asymptotically controls type I error and has power against alternatives outside a small neighbourhood of the null hypothesis.

3 Testing methodology

In this section, we propose our test based on incomplete U -statistics. Recall that X_1, \dots, X_n are i.i.d. samples from a distribution P_θ with parameter $\theta \in \Theta \subseteq \mathbb{R}^d$.

3.1 Test statistic

We want to test null hypotheses $\Theta_0 \subseteq \Theta$ defined by constraints $f(\theta) = (f_1(\theta), \dots, f_p(\theta))$ as specified in equation (1.2). For now, we assume that an \mathbb{R}^p -valued, measurable and symmetric function $b(x_1, \dots, x_m)$ exists such that $\mathbb{E}[b(X_1^m)] = f(\theta)$. For the case of polynomial hypotheses we show a general construction of kernels in Section 4. We define the test statistic \mathcal{T} to be the maximum

of a studentized incomplete U -statistic, that is,

$$\mathcal{T} = \max_{1 \leq j \leq p} \sqrt{n} U'_{n,N,j} / \hat{\sigma}_j,$$

where $\hat{\sigma}_j^2$ is the empirical approximate variance of the j th coordinate of the incomplete U -statistic; recall the definition in Section 2.3. Large values of \mathcal{T} indicate that the null hypothesis is violated and thus it is natural to reject H_0 when \mathcal{T} exceeds a certain critical value.

3.2 Critical value

The idea for construction of a critical value relies on the observation

$$\mathcal{T} \leq \max_{1 \leq j \leq p} \sqrt{n} (U'_{n,N,j} - f_j(\theta)) / \hat{\sigma}_j, \tag{3.1}$$

whenever θ is a point in the null hypothesis Θ_0 . Hence, to make the test size less or equal than α , it is enough to choose the critical value as the $(1 - \alpha)$ -quantile of the distribution of $\max_{1 \leq j \leq p} \sqrt{n} (U'_{n,N,j} - f_j(\theta)) / \hat{\sigma}_j$, which is a centred version of our test statistic \mathcal{T} . Since the distribution of $\max_{1 \leq j \leq p} \sqrt{n} (U'_{n,N,j} - f_j(\theta)) / \hat{\sigma}_j$ is unknown, we approximate it using the Gaussian multiplier bootstrap introduced in Section 2. The distribution function $P(\max_{1 \leq j \leq p} \sqrt{n} (U'_{n,N,j} - f_j(\theta)) / \hat{\sigma}_j \leq t)$ for $t \in \mathbb{R}$ corresponds to the probabilities $P(\sqrt{n} \hat{\Lambda}^{-1/2} (U'_{n,N} - f(\theta)) \in R_t)$, where $R_t = \{x \in \mathbb{R}^p : -\infty \leq x_j \leq t \text{ for all } j = 1, \dots, p\}$ are t -dependent hyper-rectangles. Hence, by Corollary 2.12, the maximum $\max_{1 \leq j \leq p} \sqrt{n} (U'_{n,N,j} - f_j(\theta)) / \hat{\sigma}_j$ can be approximated by the maximum of the studentized Gaussian multiplier statistic $W := \max_{1 \leq j \leq p} U'_{n,m,j} / \hat{\sigma}_j$. In particular, the quantiles of W approximate the quantiles of $\max_{1 \leq j \leq p} \sqrt{n} (U'_{n,N,j} - f_j(\theta)) / \hat{\sigma}_j$ as we verify in the next corollary. For $\alpha \in (0, 1)$ we denote by $c_W(1 - \alpha)$ the conditional $(1 - \alpha)$ -quantile of W given the data \mathcal{D}_n .

Corollary 3.1 Assume the conditions in Corollary 2.11. Then there exists a constant $C > 0$ depending only on β, ζ and C_1 such that

$$\sup_{\alpha \in (0,1)} \left| P\left(\max_{1 \leq j \leq p} \sqrt{n} (U'_{n,N,j} - f_j(\theta)) / \hat{\sigma}_j > c_W(1 - \alpha) \right) - \alpha \right| \leq C n^{-\zeta/7}.$$

Corollary 3.1 says that, under mild regularity conditions involving mixed degeneracy, using $c_W(1 - \alpha)$ as a critical value gives an asymptotically valid test, that is, it asymptotically controls type I error for the significance level α . When the null hypothesis is only defined by equality constraints $f_j(\theta) = 0$, then we have equality in equation (3.1) and Corollary 3.1 implies that our test has asymptotic type I error exactly equal to the chosen level α . When we have $f_j(\theta) < 0$ for certain indices $j = 1, \dots, p$, then the type I error is smaller or equal to the chosen level. As discussed in Remark 2.3, mixed degeneracy also accommodates irregular settings. The computational budget parameter has to be chosen appropriately such that the conditions in Corollary 2.11 are satisfied, typically $N = \mathcal{O}(n)$ is a good choice as we elaborate in Remark 3.3. In practice, we use the empirical version of $c_W(1 - \alpha)$ as a critical value. It is obtained in the following procedure:

- (i) Generate many, say $A = 1,000$, sets of standard normal random variables $\{\xi'_i : i \in I_{n,m}\} \cup \{\xi_i : i_1 \in S_1\}$.
- (ii) Evaluate W for each of these A sets.
- (iii) Take $\hat{c}_W(1 - \alpha)$ to be the $(1 - \alpha)$ -quantile of the resulting A numbers.
- (iv) Reject H_0 if $\mathcal{T} > \hat{c}_W(1 - \alpha)$.

Remark 3.2 (Practical considerations and complexity). There is no need to generate $|I_{n,m}| \approx n^m$ Bernoulli random variables to construct the incomplete U -statistic $U'_{n,N}$, nor to generate the same amount of standard normal random variables ξ'_i . As explained in Chen and Kato (2019, Section 2.1), one

can equivalently generate $\hat{N} \sim \text{Bin}(|I_{n,m}|, \rho_n)$ once and then choose indices $\iota_1, \dots, \iota_{\hat{N}}$ without replacement from $I_{n,m}$. Then we can compute the incomplete U -statistic as

$$U'_{n,N} = \frac{1}{\hat{N}} \sum_{i=1}^{\hat{N}} b(X_{\iota_i}).$$

Similarly, one only needs to generate \hat{N} versions of ζ'_i to construct $U^{\#}_{n,b}$. As another practical guidance, we suggest to compute $n_1 = |S_1| = n$ divide-and-conquer estimators. Choosing n_1 smaller only decreases the accuracy of the bootstrap, which is only useful if the computational cost is otherwise too high. With $n_1 = n$, the overall computational cost of the proposed procedure is $\mathcal{O}(n^2p + A(N+n)p)$ as the estimation step for g can be done outside the bootstrap.

Remark 3.3 (Computational budget parameter N). If one is certain that the underlying parameter θ is regular and *all* random variables $g_j(X_1)$, $j = 1, \dots, p$ are asymptotically non-degenerate, then it may be appealing to choose the computational budget $0 < N < \binom{n}{m}$ arbitrarily since the bootstrap approximation still holds as shown in [Chen and Kato \(2019, Theorem 3.1\)](#) and further refined in [Song et al. \(2019, Theorem 2.4\)](#). However, the rate of convergence depends on $\underline{\sigma}_{g,\theta}^2 = \min_{1 \leq j \leq p} \sigma_{g,\theta,j}^2$ as outlined in the introduction. Thus, if θ is close to an irregular point, where $\underline{\sigma}_{g,\theta}^2$ is small, the rate of convergence can be very slow. In contrast, a lower computational budget parameter N may imply convergence with a reasonable rate, even when the parameter θ is close to irregular. If mixed degeneracy is satisfied, the rate of convergence only depends on $\underline{\sigma}_{g^{(1)}}^2$, which is a fixed constant large enough to achieve fast convergence; recall [Remark 2.3](#) and [Corollary 2.7](#).

On the other hand, consider an underlying parameter θ such that *all* variances $\sigma_{g,\theta,j}^2$ of the Hájek projection are zero, which in particular means that θ is an irregular point. Then, the incomplete U -statistic is fully degenerate and a bootstrap approximation also holds true for larger N ([Chen & Kato, 2019, Theorem 3.3](#)). However, the limiting Gaussian distribution becomes $N_p(0, \Gamma_b)$ under a suitable scaling of the incomplete U -statistic. Choosing the computational budget N lower allows for computing the bootstrap based on the Gaussian approximation $N_p(0, m^2\Gamma_g + \alpha_n\Gamma_b)$, regardless of any potential degeneracy. In particular, we can compute valid critical values without knowing a priori whether or not a point is irregular. Moreover, the degenerate approximation to $N_p(0, \Gamma_b)$ only holds if *all* variances $\sigma_{g,\theta,j}^2$ are zero and not for more general irregular points where only some of them are zero.

To summarize, our approximation results of [Section 2](#) are very useful if it is unknown whether or not the true point is irregular or close to irregular. Mixed degeneracy allows for such set-ups when choosing the computational budget N appropriately. In particular, we recommend to choose the computational budget parameter N of the same order as the sample size n to guard against irregularities. However, this still allows some flexibility in practice. In [Section 5](#), we compare test size and power for different choices of N in numerical experiments.

Remark 3.4 (High dimensionality). We emphasize that bootstrap approximation holds in settings where the number of polynomials p may be much larger than the sample size n , i.e. p may be as large as $\exp(n^c)$ for a constant $c > 0$.

Thus, we can test a very large number of polynomial restrictions simultaneously. Moreover, we do not require any restriction on the correlation structure among the polynomials.

3.3 Power

The following result pertains to the power of the proposed test.

Proposition 3.5 Assume the conditions in Corollary 2.11, and let $\alpha \in (0, 1/2)$. Then there exists a constant $C > 0$ depending only on β, ζ and C_1 such that for every $\epsilon > 0$, whenever

$$\max_{1 \leq j \leq p} (f_j(\theta) / \sigma_j) \geq (1 + \epsilon) \left(1 + Cn^{-3\zeta/7} \log(p)^{-2} \right) \times \frac{\sqrt{2 \log(p)} + \sqrt{2 \log(1/\alpha)}}{\sqrt{n}}, \tag{3.2}$$

we have

$$P(\mathcal{T} > c_W(1 - \alpha)) \geq 1 - \exp\left(-\frac{1}{C} \epsilon^2 \log(p/\alpha)\right) - Cn^{-\zeta/7}.$$

Proposition 3.5 shows that our test is consistent against all alternatives where at least one constraint, normalized by its approximate variance, is violated by a small amount. In other words, the test is consistent against all alternatives excluding the ones in a small neighbourhood of the null set Θ_0 . The size of the neighbourhood shrinks with rate $\sqrt{\log(p)/n}$ as long as $p \rightarrow \infty$ as $n \rightarrow \infty$. A similar result on power is derived in Chernozhukov et al. (2019) for the special case on independent sums, and we refer to their discussion on the rate the neighbourhood converges to zero.

4 Polynomial hypotheses

In this section, we assume that the constraints defining the null hypothesis Θ_0 in equation (1.2) are polynomial. That is, each constraint $f_j \in \mathbb{R}[\theta_1, \dots, \theta_d]$ is a polynomial in the indeterminates $\theta_1, \dots, \theta_d$. We propose a general procedure to find a kernel h such that $h_j(X_1^m)$ is an unbiased estimator of $f_j(\theta)$. For now, let $f \in \mathbb{R}[\theta_1, \dots, \theta_d]$ be a *single* polynomial of total degree s and write

$$f(\theta) = a_0 + \sum_{r=1}^s \sum_{\substack{(i_1, \dots, i_r) \\ i_j \in \{1, \dots, d\}}} a_{(i_1, \dots, i_r)} \theta_{i_1} \cdots \theta_{i_r}$$

with $a_{(i_1, \dots, i_r)} \in \mathbb{R}$ for all multi-indices (i_1, \dots, i_r) . Note that this notation of multivariate polynomials is somewhat inefficient in comparison to the usual multi-index notation since there may appear various indices repeatedly in (i_1, \dots, i_r) . However, the representation is useful to define an estimator of the polynomial. We construct $h(X_1^m)$ by the following three steps:

1. For a fixed integer $\eta \geq 1$, find \mathbb{R} -valued functions $\widehat{\theta}_i$ such that $\widehat{\theta}_i(X_1^\eta)$ is an unbiased estimator of θ_i for all $i = 1, \dots, d$.
2. Let $m = \eta s$ be the order of the kernel, and define the \mathbb{R} -valued function \check{h} via

$$x_1^m \mapsto a_0 + \sum_{r=1}^s \sum_{\substack{(i_1, \dots, i_r) \\ i_j \in \{1, \dots, d\}}} a_{(i_1, \dots, i_r)} \widehat{\theta}_{i_1}(x_1^\eta) \widehat{\theta}_{i_2}(x_{\eta+1}^{2\eta}) \cdots \widehat{\theta}_{i_r}(x_{(r-1)\eta+1}^m),$$

where $x_k^l = (x_k, \dots, x_l)$ for $k < l$. Note that $\check{h}(X_1^m)$ is an unbiased estimator of $f(\theta)$ by the linearity of the expectation and the independence of the samples X_1, \dots, X_m .

3. To get a symmetric kernel h , average over all permutations $\pi \in S_m$ of the set $\{1, \dots, m\}$, that is,

$$h(x_1, \dots, x_m) = \frac{1}{m!} \sum_{\pi \in S_m} \check{h}(x_{\pi(1)}, \dots, x_{\pi(m)}).$$

The construction works for any polynomial as long as we can construct an unbiased estimator of the parameter θ that only involves a small number η of samples. In this case polynomial hypotheses are estimable and our proposed testing methodology is applicable.

Example 4.1 The kernel for estimating the tetrad $f(\Sigma) = \sigma_{uw}\sigma_{wz} - \sigma_{uz}\sigma_{vw}$ in Example 1.1 is also constructed by steps 1) to 3). The total degree of f is $s = 2$ and to estimate an entry σ_{uw} in the covariance matrix Σ we only need $\eta = 1$ sample, i.e. an unbiased estimator of σ_{uw} is given by $\widehat{\sigma}_{uw}(X_1) = X_{1u}X_{1v}$. Thus, the order of the kernel is $m = 2$ and we obtain

$$\check{h}(X_1, X_2) = X_{1u}X_{1v}X_{2w}X_{2z} - X_{1u}X_{1z}X_{2v}X_{2w}.$$

The symmetric kernel is given by

$$h(X_1, X_2) = \frac{1}{2} \{ (X_{1u}X_{1v}X_{2w}X_{2z} - X_{1u}X_{1z}X_{2v}X_{2w}) + (X_{2u}X_{2v}X_{1w}X_{1z} - X_{2u}X_{2z}X_{1v}X_{1w}) \}.$$

The bootstrap approximation established in Section 2 requires that the individual estimators $b_j(X_1^m)$ are sub-Weibull of order $\beta > 0$ for all $j = 1, \dots, p$; recall Condition (C2). If one is able to check that all estimators $\widehat{\theta}_i(X_1^\eta)$, $i = 1, \dots, d$ are sub-Weibull of order γ , then we obtain by Lemma C.3, online supplementary material that the estimator $b_j(X_1, \dots, X_m)$ is sub-Weibull of order $\beta = \gamma/s$, where s is the total degree of the polynomial $f_j(\theta)$. For estimating tetrads, this is illustrated in the next example.

Example 4.2 For a Gaussian random vector $X_i \sim N_l(0, \Sigma)$, it is easy to check that each component X_{iu} , $u = 1, \dots, l$ is sub-Gaussian with $\|X_{iu}\|_{\psi_2} \leq \sqrt{8\sigma_{iu}}/3$. Therefore, $\widehat{\sigma}_{uw}(X_1) = X_{1u}X_{1v}$ is sub-exponential by Lemma C.2, online supplementary material and hence the kernel $h(X_1, X_2)$ for estimating a tetrad in Example 4.1 is sub-Weibull of order $\beta = 1/2$.

Recall our testing problem (1.2) where the null hypothesis $\Theta_0 \subseteq \Theta$ is defined by polynomial constraints. Suppose we have a kernel h such that every point $\theta \in \Theta$ is regular with respect to h and $\sigma_{g,\theta} = \min_{1 \leq j \leq p} \sigma_{g,\theta,j}^2$ is not too small, that is, every point is far away from irregular. Then it is favourable to consider critical values based on the Gaussian approximation for non-degenerate incomplete U -statistics as discussed in Remark 3.3. On the other hand, if all points are irregular with every individual Hájek projection being degenerate, then one should consider critical values based on the limiting distribution of degenerate incomplete U -statistics. For the proposed kernel, we will show in our next result, that parameter spaces typically contain a measure zero subset of irregular points.

Proposition 4.3 Let $f \in \mathbb{R}[\theta_1, \dots, \theta_d]$ be a polynomial of total degree $s \geq 2$ and define the set of indices $D(f) = \{j \in [d] : \theta_j \text{ appears in } f\}$. Further, define $\widehat{\theta}_{i,l}(x) = \mathbb{E}[\widehat{\theta}_i(X_1, \dots, X_{l-1}, x, X_{l+1}, \dots, X_\eta)]$ for all $i = 1, \dots, d$ and $l = 1, \dots, \eta$ and denote the random vector

$$\widehat{\theta}_f(X_1) = \left(\sum_{l=1}^{\eta} \widehat{\theta}_{i,l}(X_1) \right)_{i \in D(f)}.$$

If the covariance matrix $\text{Cov}[\hat{\theta}_f(X_1)]$ is positive definite, then the irregular points $\theta \in \Theta \subseteq \mathbb{R}^d$ with respect to the kernel b form a measure zero set with respect to the Lebesgue measure on \mathbb{R}^d .

Proposition 4.3 is a sufficient condition for checking that almost all points in the parameter space are regular. By inspecting the proof, we see that the set of irregular points is the zero set of a certain non-zero polynomial and thus has Lebesgue measure zero. However, it will typically be non-empty. As we emphasized in the introduction, the Gaussian approximation of complete U -statistics, or incomplete U -statistics with high computational budget, may already be very slow if the true parameter is close to an irregular point. Since the null hypothesis Θ_0 can be defined by a very large amount of constraints p , the sets of irregular points is the union of irregular points given by p sets with Lebesgue measure zero. In other words, there might be many hyper-surfaces with irregular points, such that it is likely that the true parameter is close to an irregular point. Thus, assuming mixed degeneracy is crucial for considering kernels as constructed above.

Example 4.4 Consider again a tetrad $f(\Sigma) = \sigma_{uv}\sigma_{wz} - \sigma_{uz}\sigma_{vw}$ and the corresponding kernel b given in Example 4.2. Since $\hat{\sigma}_{uv}(X_1) = X_{1u}X_{1v}$, we have that $\hat{\theta}_f(X_1) = (X_{1u}X_{1v}, X_{1w}X_{1z}, X_{1u}X_{1z}, X_{1v}X_{1w})^T$. For Gaussian $X_1 \sim N_l(0, \Sigma)$, it is easy to see that $\text{Cov}[\hat{\theta}_f(X_1)]$ is positive definite since no coordinate of $\hat{\theta}_f(X_1)$ is a linear function of the remaining coordinates. It follows by Proposition 4.3 that almost all covariance matrices Σ are regular points in the parameter space given by the cone of positive definite matrices. Covariance matrices with some entries equal to zero correspond to irregular points as we have seen in Example 1.3. Thus, all covariance matrices with those entries having small absolute values are close to irregular points.

5 Testing Gaussian latent-tree models

In this section, we apply our test for assessing the goodness-of-fit of Gaussian latent-tree models. These models are of particular relevance in phylogenetics (Semple & Steel, 2003; Zwiernik, 2016) and the problem of model selection is, for example, considered in Shiers et al. (2016) and Leung and Drton (2018); for a survey see Sung (2009) and Junker and Schreiber (2011). Our methodology is applicable since a full semi-algebraic description of the model is known, that is, all polynomial equalities and inequalities that fully describe the distributions corresponding to a given tree are known. The constraints include tetrads as well as higher order polynomial constraints. However, as we also point out in the introduction, it is challenging to test the large number of constraints simultaneously. For example, in Shiers et al. (2016) only small trees and a subset of constraints is tested. Typical approaches such as the Wald test can only handle $p \leq n$ constraints and the maximum likelihood function is difficult to optimize. For the latter, we implemented an expectation-maximization algorithm, but it is unclear whether it obtains the global maximum; see Section 5.2. Moreover, Drton (2009), Dufour et al. (2013), and Drton and Xiao (2016) show that irregular points lead to different limiting distributions in Wald and LR tests when testing tetrads. Hence, our testing strategy is an approach that targets both challenges in testing Gaussian latent-tree models: a large number of constraints and irregular points in the null hypothesis.

5.1 Model and constraints

We begin by briefly introducing Gaussian latent-tree models. Let $T = (V, E)$ be an undirected tree where V is the set of nodes and E is the set of edges. For $L \subseteq V$ denoting the set of leaves, we assume that every inner node in $V \setminus L$ has minimal degree 3. By Zwiernik (2016, Chapter 8) we have the following parametric representation of Gaussian latent-tree models. Let $PD(l)$ be the cone of symmetric positive definite $l \times l$ matrices.

Proposition 5.1 The Gaussian latent-tree model of a tree $T = (V, E)$ with leaves $L = \{1, \dots, l\} \subseteq V$ is the family of Gaussian distributions $N_l(0, \Sigma)$ such

that Σ is in the set

$$\mathcal{M}(T) = \left\{ \Sigma = (\sigma_{uv}) \in PD(l) : \sigma_{uv} = \sqrt{\omega_u} \sqrt{\omega_v} \prod_{e \in ph_T(u,v)} \rho_e \text{ for } v \neq u, \sigma_{vv} = \omega_v \right\},$$

where $ph_T(u, v)$ denotes the set of *edges* on the unique path from u to v , the vector $(\omega_v)_{v \in L} \in \mathbb{R}_{>0}^l$ contains the variances and $(\rho_e)_{e \in E}$ is an $|E|$ -dimensional vector with $|\rho_e| \in (0, 1)$.

For details on how this parametrization arises from the paradigm of graphical modelling, see [Zwiernik \(2016\)](#) or [Drton et al. \(2017\)](#). We may identify a Gaussian latent tree model with its set of covariance matrices $\mathcal{M}(T)$ and we will simply refer to $\mathcal{M}(T)$ as the model.

From a geometric point of view, the parameter space, i.e. the set of covariance matrices $\mathcal{M}(T)$, is fully understood. It is a semi-algebraic set described by polynomial equalities and inequalities in the entries of the covariance matrix. To state the polynomial constraints, we need the following notation borrowed from [Leung and Drton \(2018\)](#). Recall that in a tree any two nodes are connected by precisely one path. For four distinct leaves $\{u, v, w, z\} \subseteq L$, there are three possible partitions into two subsets of equal size, namely $\{u, v\} \setminus \{w, z\}$, $\{u, w\} \setminus \{v, z\}$, and $\{u, z\} \setminus \{v, w\}$. These three partitions correspond to the three intersection of path pairs

$$ph_T(u, v) \cap ph_T(w, z), \quad ph_T(u, w) \cap ph_T(v, z) \quad \text{and} \quad ph_T(u, z) \cap ph_T(v, w). \quad (5.1)$$

By the structure of a tree, either all of the intersections give an empty set or exactly one is empty while the others are not; cf. [Leung and Drton \(2018\)](#). We let $\mathcal{Q} \subseteq \{\{u, v, w, z\} \subseteq L\}$ be the collections of subsets of size four such that $\{u, v, w, z\} \in \mathcal{Q}$ if exactly one of the intersections is empty. Given $\{u, v, w, z\} \in \mathcal{Q}$, we write $\{u, v\} \setminus \{w, z\} \in \mathcal{Q}$ to indicate that $\{u, v, w, z\}$ belongs to \mathcal{Q} and the paths $ph_T(u, v)$ and $ph_T(w, z)$ have empty intersection.

Proposition 5.2 ([Leung & Drton, 2018](#)). Let $\Sigma = (\sigma_{uv}) \in PD(l)$ be a covariance matrix with no zero entries. Then $\Sigma \in \mathcal{M}(T)$ if and only if Σ satisfies the following constraints:

1. *Inequality constraints:*

- (a) For any $\{u, v, w\} \subseteq L$, $-\sigma_{uv}\sigma_{uw}\sigma_{vw} \leq 0$.
- (b) For any $\{u, v, w\} \subseteq L$, $\sigma_{uv}^2\sigma_{vw}^2 - \sigma_{vw}^2\sigma_{uw}^2 \leq 0$, $\sigma_{uw}^2\sigma_{vw}^2 - \sigma_{uv}^2\sigma_{vw}^2 \leq 0$ and $\sigma_{uv}^2\sigma_{uw}^2 - \sigma_{uv}^2\sigma_{vw}^2 \leq 0$.
- (c) For any $\{u, v\} \setminus \{w, z\} \in \mathcal{Q}$, $\sigma_{uw}^2\sigma_{vz}^2 - \sigma_{uv}^2\sigma_{wz}^2 \leq 0$.

2. *Equality constraints (tetrads):*

- (a) For any $\{u, v\} \setminus \{w, z\} \in \mathcal{Q}$, $\sigma_{uw}\sigma_{vz} - \sigma_{uz}\sigma_{vw} = 0$.
- (b) For any $\{u, v, w, z\} \notin \mathcal{Q}$, $\sigma_{uz}\sigma_{vw} - \sigma_{uw}\sigma_{vz} = \sigma_{uv}\sigma_{wz} - \sigma_{uv}\sigma_{vz} = 0$.

We want to test a Gaussian latent-tree model against the saturated alternative, that is,

$$H_0 : \Sigma \in \mathcal{M}(T) \quad \text{vs.} \quad H_1 : \Sigma \in PD(l) \setminus \mathcal{M}(T)$$

based on i.i.d. samples X_1, \dots, X_n taken from a Gaussian distribution $N_l(0, \Sigma)$. With our methodology from Sections 3 and 4, we can simultaneously test all constraints defining $\mathcal{M}(T)$ given in Proposition 5.2. The number of constraints is $p = 2 \binom{l}{4} + 4 \binom{l}{3}$ which is of order $\mathcal{O}(l^4)$. Even for moderate values of l , say $l = 15$, this is a large number of constraints since $p = 4,550$.

5.2 Simulations

We have published an implementation of our tests for Gaussian latent-tree models in the R package `TestGGM`, available at <https://github.com/NilsSturma/TestGGM>. In the implementation we use $A = 1,000$ sets of Gaussian multipliers to compute the critical value $\hat{c}_W(1 - \alpha)$. The kernel for




Figure 2. Graphical representation of (a) the star tree and (b) the binary caterpillar tree. Solid black dots correspond to leaves (observed variables).

estimating the polynomials is constructed as demonstrated in Example 4.1, where we treated the special case of tetrads. The package `TestGGM` also provides a routine for determining the specific constraints for a given latent-tree input, i.e. for determining the set \mathcal{Q} . For finding a path $ph_T(u, v)$ we use the function `shortest_paths` in the `igraph` package (Csardi & Nepusz, 2006).

We compare our test with the LR test. By Shiers et al. (2016) the dimension of a Gaussian latent-tree model is $|E| + |L|$. Thus, under regularity conditions, the LR test statistic approximates a chi-square distribution with $\binom{l+1}{2} - (|E| + |L|)$ degrees of freedom for $l \geq 4$ (Van der Vaart, 1998).

For the one-factor model, the LR test is implemented by the function `factanal` in the `base` library of R (R Core Team, 2020). Note that it employs a Bartlett correction for better asymptotic approximation. For general trees the implementation shows some of the challenges of the likelihood ratio test. Optimizing the likelihood function on the subspace $\mathcal{M}(T)$ is difficult using the common non-convex methods due to the large number of constraints defining the subspace and the required positive definiteness of the covariance matrices. Therefore we implement an expectation maximization algorithm which was first proposed by Dempster et al. (1977) and is described in Friedman et al. (2002) and Mourad et al. (2013) for the special case of latent-tree models. A drawback of this method is that global optimization is not guaranteed.

In our numerical experiments we consider two different tree structures. One is the star tree, where the model is equal to the one-factor analysis model. Second, we consider the binary caterpillar tree which is extreme in a sense that it has the longest paths possible between leaves. In Figure 2, both structures are illustrated. We fix dimension $l = 15$, sample size $n = 500$ and generate data from both tree structures in three experimental set-ups:

- (a) Star tree: All edge parameters ρ_e are equal to $\sqrt{0.5}$, while we take all variances $\omega_v = 2$.
- (b) Star tree: Let h be the unique, unobserved inner node. The parameters for the edges $\{h, 1\}$ and $\{h, 2\}$ are taken to be 0.998, while all other edge parameters ρ_e are independently generated from $N(0, 0.1)$. The marginal variances are taken to be $\omega_v = 100$ for $v = 1, 2$ and $\omega_v = 1$ for all other leaves $v \in L$.
- (c) Caterpillar tree: The edge parameters ρ_e are taken to be 0.998 except for all edges incident to four selected inner nodes in $V \setminus L$, where the edge parameters are independently generated from $N(0, 0.1)$. All variances ω_v are taken to be 2.

Set-up (a) is a regular problem where the LR test is the gold standard. In contrast, set-ups (b) and (c) are designed to be irregular problems with small covariances; compare with Example 1.3 in the introduction. In set-up (b), parameters are chosen such that the covariance matrix has exactly one off-diagonal entry which is far away from zero while all the remaining off-diagonal entries are close to zero. In this case, the parameters are also close to an algebraic singularity of the parameter space of the star tree (Drton et al., 2007, Proposition 32). The LR test should fail since the LR test statistic does not follow a chi-square distribution at singularities (Drton, 2009). Similarly, parameters in set-up (c) are close to an algebraic singularity since some of the edge parameters are almost zero (Drton et al., 2017) and thus we also expect the LR test to fail. Moreover, set-up (c) is considered to emphasize that the proposed methodology allows for testing arbitrary Gaussian latent tree models.

In Figures 3 and 4, we compare empirical test sizes of our test with the LR test for different, fixed significance levels $\alpha \in (0, 1)$ that we call the ‘nominal level’. In addition, we consider different




Figure 3. Empirical sizes vs. nominal levels for testing tetrads based on 500 experiments. The computational budget parameter N is varied as indicated and empirical sizes of the LR test are also shown. Data are generated from set-up (a) with $(l, n) = (15, 500)$.

computational budget parameters N to see its influence on the behaviour of our test. To better compare test sizes we only consider testing the equality constraints (tetrads) in Proposition 5.2. This simplifies our analysis since in this case we can expect that the test size is equal to the level α . A statistical test that controls type I error is expected to have empirical test sizes below the 45 degree line in the plots. The lower the line, the more ‘conservative’ the test is. On the other hand, a test does not control type I error if the empirical sizes are above the 45 degree line. In this case the true null hypothesis is rejected too often.

The irregular set-ups (b) and (c) show the advantage of the proposed testing method because the empirical test sizes of the LR test do not hold the nominal levels, that is, the empirical sizes are above the 45 degree line. In comparison, our testing method holds level for all choices of the computational budget parameter N . However, there is an interesting difference in the behaviour of the empirical test sizes between the regular set-up (a) and set-ups (b) and (c) where parameters are close to an irregular point. In the regular set-up, a higher computational budget parameter N tends to yield empirical sizes closer to the nominal level α while in the irregular set-ups this is the other way round. The behaviour coincides with the theoretical guarantees discussed in Section 2 and Remark 3.3, suggesting that the bootstrap approximation holds under mixed degeneracy only when choosing the computational budget N appropriately, in this case $N = \mathcal{O}(n)$.

Besides size, we also study the empirical power of our test. For $\Sigma \in \mathcal{M}(T)$, we consider local alternatives

$$\tilde{\Sigma} = \Sigma + \gamma\gamma^T \frac{b}{\sqrt{n}}, \quad (5.2)$$

where $\gamma = (0, \dots, 0, 1, 1) \in \mathbb{R}^l$ is fixed and $b > 0$ varies. According to Proposition 3.5, our test is consistent against alternatives outside of a neighbourhood of size $\sqrt{\log(p)/n}$ up to an unknown constant. Thus, it is natural to consider alternatives as in equation (5.2) that also depend on the sample size through \sqrt{n} . Moreover, all entries of the alternative covariance matrix $\tilde{\Sigma}$ are equal to the entries of Σ except for four entries, which implies that only few constraints in Proposition 5.2 are violated. However, since our test statistic is the maximum of an incomplete U -statistic, we expect that our test is consistent even for such sparse signals.





Figure 4. Empirical sizes vs. nominal levels for testing tetrads based on 500 experiments. The computational budget parameter N is varied as indicated and empirical sizes of the LR test are also shown. Data are generated from set-ups (b) and (c) with $(l, n) = (15, 500)$.

For the regular set-up (a), the values of b are plotted against the empirical power in Figure 5. This time we considered all constraints in Proposition 5.2, equalities and inequalities. The power increases with b for all choices of the computational budget N . A higher budget empirically yields




Figure 5. Empirical power for different local alternatives based on 500 experiments. The computational budget parameter N is varied as indicated and empirical power of the LR test is also shown. Local alternatives are generated as described in the text for set-up (a) with $(l, n) = (15, 500)$ and level $\alpha = 0.05$.




Figure 6. Mammalian gene expression phylogenies for cerebellum.

better power which is reasonable due to more precise estimation of the individual polynomials. For large computational budgets the power on the local alternatives is comparable to the LR test. However, as discussed before, one should be sceptical about the bootstrap approximation for large computational budgets N when the underlying true parameter is close to singular. The above experiments suggest that the choice $N = 2n$ is a good trade-off between guarding against irregularities and statistical efficiency. A similar plot for the singular set-ups can be found in [Appendix D, online supplementary material](#). Moreover, we provide another simulation study in [Appendix E, online supplementary material](#), where we apply our methodology to test minors in the two-factor analysis model.

5.3 Application to gene expression data

In this section, we analyze data from [Brawand et al. \(2011\)](#), where they obtained gene expression levels from sequencing of polyadenylated RNA from the cerebellum across different species such as gorillas, opossums, or chickens. We focus on data of 20 individuals from different species, where we have 5,636 gene expression levels for each individual. Based on the assumption that species evolved from a common ancestor, a phylogenetic tree describes the evolutionary history among a collection of species ([Semple & Steel, 2003](#)). In the work of [Brawand et al. \(2011\)](#), the authors reconstruct the phylogenetic tree in [Figure 6](#) by neighbour-joining based on pair-wise

distance matrices. If we log-transform the original gene expression data, it is approximately normal distributed and it becomes a natural question to ask whether the data follows a Gaussian latent-tree model where the underlying tree structure is the same.

In our analysis to check tree compatibility we remove rows containing a zero such that we are left with $N = 4,615$ observations. Since the tree has 20 leaves, we have to test $p = 14250$ constraints that define the null hypothesis; recall Proposition 5.2. We compute p -values using the LR test and our test strategy with computational budget $N = 2n$. Both tests indicate overwhelming evidence against the Gaussian latent-tree model since we obtain a p -value of zero in both cases. To check whether the assumption of a Gaussian latent-tree model might hold locally, we also run the same analysis on all substructures of the tree induced by 6 neighbouring leaves. This implies that $p = 110$. Again, we only observe p -values equal to zero, indicating that the Gaussian latent-tree assumption is also too strong for the considered subsets of variables. We summarize that our approach is able to reach the same conclusion as the LR test, but our strategy is optimization-free and safe with respect to possible irregularities.

6 Discussion

We proposed a general methodology to simultaneously test many constraints on a statistical parameter. The test statistic is given by a studentized maximum of an incomplete U -statistic and critical values are approximated by Gaussian multiplier bootstrap. Our method is applicable to all hypotheses that can be defined by estimable equality and inequality constraints. If the constraints are polynomial, we presented a general procedure to construct kernels. In the suggested methodology, there is no need to maximize a possibly multi-modal function as one would do for a LR test. Moreover, our method allows for many constraints, that is, the number of constraints can be much larger than the sample size. If the computational budget N for constructing the incomplete U -statistic is appropriately chosen, typically of the same order as the sample size n , then the test asymptotically controls type I error even if the true parameter is irregular or close to irregular.

The latter fact is due to non-asymptotic Berry–Esseen-type error bounds on the high-dimensional Gaussian and bootstrap approximation of incomplete U -statistics. Compared to previous work, we have shown that the approximation is also valid under mixed degeneracy if we choose a suitable computational budget N . This yields control of type I error as well as consistency of our testing methodology even in irregular settings.

In practice, the requirement $N = \mathcal{O}(n)$ allows some flexibility in choosing the computational budget N . As illustrated in our simulations, a higher computational budget parameter yields more efficient estimates of the constraints and therefore a more powerful test. However, the computations become more involved for higher order kernels and one should be sceptical about the theoretical guarantees of the bootstrap approximation if N is large and the underlying true parameter is close to an irregular point. Thus, there is a trade-off between efficiency and guarding against irregularities. In our simulations we found that choosing the computational budget parameter as $N = 2n$ is a reasonable choice.

If the inequality constraints defining the null hypothesis are strict, then there is a decrease in power. One may improve power by using a two-step approach for testing inequalities that was initially introduced in Romano et al. (2014) and then further developed by Bai et al. (2021) for the high-dimensional set-up. In the first step, a confidence region for the constraints of interest is constructed and in the second step, this set is used to provide information about which constraints are negative. In future work, it would be of interest to improve our strategy using the mentioned two-step approach.

Acknowledgments

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 883818), and from the German Federal Ministry of Education and Research and the Bavarian State Ministry for Science and the Arts. Nils Sturma acknowledges support by the Munich Data Science Institute (MDSI) at Technical University of Munich via the Linde/MDSI PhD Fellowship programme.

Conflict of interests: None declared.

Data availability

We have published an implementation of our tests for Gaussian latent-tree models in the R package TestGGM, available at <https://github.com/NilsSturma/TestGGM>.

Supplementary material

Supplementary material is available online at *Journal of the Royal Statistical Society: Series B*.

References

- Bai Y., Santos A., & Shaikh A. M. (2021). A two-step method for testing many moment inequalities. *Journal of Business & Economic Statistics*, 40(3), 1–11. <https://doi.org/10.1080/07350015.2021.1897016>
- Bollen K. A., & Ting K.-F. (2000). A tetrad test for causal indicators. *Psychological Methods*, 5(1), 3–22. <https://doi.org/10.1037/1082-989X.5.1.3>
- Brawand D., Soumillon M., Necsulea A., Julien P., Csárdi G., Harrigan P., Weier M., Liechti A., Aximu-Petri A., Kircher M., Albert F. W., Zeller U., Khaitovich P., Grützner F., Bergmann S., Nielsen R., Pääbo S., & Kaessmann H. (2011). The evolution of gene expression levels in mammalian organs. *Nature*, 478(7369), 343–348. <https://doi.org/10.1038/nature10532>
- Chen B., Kumor D., & Bareinboim E. (2017). Identification and model testing in linear structural equation models using auxiliary variables. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, ICML'17 (pp. 757–766). JMLR.
- Chen B., Tian J., & Pearl J. (2014). Testable implications of linear structural equation models. In *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI)* (Vol. 28). AAAI Press.
- Chen X. (2018). Gaussian and bootstrap approximations for high-dimensional U-statistics and their applications. *Annals of Statistics*, 46(2), 642–678. <https://doi.org/10.1214/17-AOS1563>
- Chen X., & Kato K. (2019). Randomized incomplete U-statistics in high dimensions. *Annals of Statistics*, 47(6), 3127–3156. <https://doi.org/10.1214/18-AOS1773>
- Chen X., & Kato K. (2020). Jackknife multiplier bootstrap: Finite sample approximations to the U-process supremum with applications. *Probability Theory and Related Fields*, 176(3–4), 1097–1163. <https://doi.org/10.1007/s00440-019-00936-y>
- Chernozhukov V., Chetverikov D., & Kato K. (2013). Gaussian approximations and multiplier bootstrap for maxima of sums of high-dimensional random vectors. *The Annals of Statistics*, 41(6), 2786–2819. <https://doi.org/10.1214/13-AOS1161>
- Chernozhukov V., Chetverikov D., & Kato K. (2017a). Central limit theorems and bootstrap in high dimensions. *The Annals of Probability*, 45(4), 2309–2352. <https://doi.org/10.1214/16-AOP1113>
- Chernozhukov V., Chetverikov D., & Kato K. (2017b). ‘Detailed proof of Nazarov’s inequality’, arXiv:1711.10696, preprint: not peer reviewed.
- Chernozhukov V., Chetverikov D., & Kato K. (2019). Inference on causal and structural parameters using many moment inequalities. *The Review of Economic Studies*, 86(5), 1867–1900. <https://doi.org/10.1093/restud/rdy065>
- Chernozhukov V., Chetverikov D., Kato K., & Koike Y. (2022). Improved central limit theorem and bootstrap approximations in high dimensions. *Annals of Statistics*, 50(5), 2562–2586. <https://doi.org/10.1214/22-AOS2193>
- Chernozhukov V., Chetverikov D., & Koike Y. (2023a). Nearly optimal central limit theorem and bootstrap approximations in high dimensions. *The Annals of Applied Probability*, 33(3), 2374–2425. <https://doi.org/10.1214/22-AAP1870>
- Chernozhukov V., Newey W. K., & Santos A. (2023b). Constrained conditional moment restriction models. *Econometrica*, 91(2), 709–736. <https://doi.org/10.3982/ECTA13830>
- Chetverikov D., Santos A., & Shaikh A. M. (2018). The econometrics of shape restrictions. *Annual Review of Economics*, 10(1), 31–63. <https://doi.org/10.1146/economics.2018.10.issue-1>
- Claassen T., & Heskes T. (2012). A Bayesian approach to constraint based causal inference. In *Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence (UAI)*, UAI'12 (pp. 207–216). AUAI Press.
- Cox D. A., Little J., & O’Shea D. (2015). *Ideals, varieties, and algorithms: An introduction to computational algebraic geometry and commutative algebra*. Undergraduate Texts in Mathematics (4th ed.). Springer.
- Csardi G., & Nepusz T. (2006). The igraph software package for complex network research. *InterJournal, Complex Systems*, 1695(5), 1–9. <https://doi.org/10.5281/zenodo.7682609>
- de la Peña V. H., & Giné E. (1999). *Decoupling: From dependence to independence, randomly stopped processes. U-statistics and processes. Martingales and beyond*. Probability and its Applications. Springer-Verlag.
- Dempster A. P., Laird N. M., & Rubin D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1), 1–22. (With discussion). <https://doi.org/10.1111/j.2517-6161.1977.tb01600.x>

- Drton M. (2009). Likelihood ratio tests and singularities. *The Annals of Statistics*, 37(2), 979–1012. <https://doi.org/10.1214/07-AOS571>
- Drton M., Lin S., Weihs L., & Zwiernik P. (2017). Marginal likelihood and model selection for Gaussian latent tree and forest models. *Bernoulli*, 23(2), 1202–1232. <https://doi.org/10.3150/15-BEJ775>
- Drton M., Massam H., & Olkin I. (2008). Moments of minors of Wishart matrices. *The Annals of Statistics*, 36(5), 2261–2283. <https://doi.org/10.1214/07-AOS522>
- Drton M., Sturmfels B., & Sullivant S. (2007). Algebraic factor analysis: Tetrads, pentads and beyond. *Probability Theory and Related Fields*, 138(3–4), 463–493. <https://doi.org/10.1007/s00440-006-0033-2>
- Drton M., & Xiao H. (2016). Wald tests of singular hypotheses. *Bernoulli*, 22(1), 38–59. <https://doi.org/10.3150/14-BEJ620>
- Dufour J.-M., Renault E., & Zinde-Walsh V. (2013). ‘Wald tests when restrictions are locally singular’, arXiv, arXiv:1312.0569, preprint: not peer reviewed.
- Fang X., & Koike Y. (2021). High-dimensional central limit theorems by Stein’s method. *The Annals of Applied Probability*, 31(4), 1660–1686. <https://doi.org/10.1214/20-AAP1629>
- Friedman N., Ninio M., Pe’er I., & Pupko T. (2002). A structural EM algorithm for phylogenetic inference. *Journal of Computational Biology*, 9(2), 331–353. <https://doi.org/10.1089/10665270252935494>
- Gaffke N., Heiligers B., & Offinger R. (2002). On the asymptotic null-distribution of the Wald statistic at singular parameter points. *Statistics & Risk Modeling*, 20(4), 379–398. <https://doi.org/10.1524/strm.2002.20.14.379>
- Gaffke N., Steyer R., & von Davier A. A. (1999). On the asymptotic null-distribution of the Wald statistic at singular parameter points. *Statistics & Risk Modeling*, 17(4), 339–358. <https://doi.org/10.1524/strm.1999.17.4.339>
- Götze F., Sambale H., & Sinulis A. (2021). Concentration inequalities for polynomials in α -sub-exponential random variables. *Electronic Journal of Probability*, 26(48), 1–22. <https://doi.org/10.1214/21-EJP606>
- Hipp J. R., & Bollen K. A. (2003). Model fit in structural equation models with censored, ordinal, and dichotomous variables: Testing vanishing tetrads. *Sociological Methodology*, 33(1), 267–305. <https://doi.org/10.1111/j.0081-1750.2003.00133.x>
- Junker B. H., & Schreiber F. (2011). *Analysis of biological networks*. John Wiley & Sons.
- Kuchibhotla A. K., & Chakraborty A. (2022). Moving beyond sub-Gaussianity in high-dimensional statistics: Applications in covariance estimation and linear regression. *Information and Inference: A Journal of the IMA*, 11(4), 1389–1456. <https://doi.org/10.1093/imaia/iaac012>
- Leung D., & Drton M. (2018). Algebraic tests of general Gaussian latent tree models. In *Advances in Neural Information Processing Systems 31 (NeurIPS 2018)* (Vol. 31, pp. 6301–6310). Neural Information Processing Systems Foundation.
- Leung D., Drton M., & Hara H. (2016). Identifiability of directed Gaussian graphical models with one latent source. *Electronic Journal of Statistics*, 10(1), 394–422. <https://doi.org/10.1214/16-EJS1111>
- Leung D., & Sturma N. (2024). ‘Singularity-agnostic incomplete U-statistics for testing polynomial constraints in Gaussian covariance matrices’, arXiv, arXiv:2401.02112, preprint: not peer reviewed.
- Lopes M. E. (2022). Central limit theorem and bootstrap approximation in high dimensions: Near $1/\sqrt{n}$ rates via implicit smoothing. *The Annals of Statistics*, 50(5), 2492–2513. <https://doi.org/10.1214/22-AOS2184>
- Mourad R., Sinoquet C., Zhang N. L., Liu T., & Leray P. (2013). A survey on latent tree models and applications. *Journal of Artificial Intelligence Research*, 47, 157–203. <https://doi.org/10.1613/jair.3879>
- Okamoto M. (1973). Distinctness of the eigenvalues of a quadratic form in a multivariate sample. *The Annals of Statistics*, 1(4), 763–765. <https://doi.org/10.1214/aos/1176342472>
- Pearl J. (2009). *Causality: Models, reasoning, and inference* (2nd ed.). Cambridge University Press.
- Pearl J., & Verma T. S. (1995). A theory of inferred causation. In D. Prawitz, B. Skyrms, & D. Westerståhl (Eds.), *Logic, Methodology and Philosophy of Science IX*. Studies in Logic and the Foundations of Mathematics (Vol. 134, pp. 789–811). Elsevier.
- R Core Team (2020). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing.
- Romano J. P., Shaikh A. M., & Wolf M. (2014). A practical two-step method for testing moment inequalities. *Econometrica*, 82(5), 1979–2002. <https://doi.org/10.3982/ECTA11011>
- Semple C., & Steel M. (2003). *Phylogenetics*. Oxford Lecture Series in Mathematics and its Applications (Vol. 24). Oxford University Press.
- Shiers N., Zwiernik P., Aston J. A. D., & Smith J. (2016). The correlation space of Gaussian latent tree models and model selection without fitting. *Biometrika*, 103(3), 531–545. <https://doi.org/10.1093/biomet/asw032>
- Silva R., Scheines R., Glymour C., & Spirtes P. (2006). Learning the structure of linear latent variable models. *Journal of Machine Learning Research*, 7, 191–246. <http://jmlr.org/papers/v7/silva06a.html>
- Song Y., Chen X., & Kato K. (2019). Approximating high-dimensional infinite-order U-statistics: Statistical and computational guarantees. *Electronic Journal of Statistics*, 13(2), 4794–4848. <https://doi.org/10.1214/19-EJS1643>

- Spearman C. (1904). "General intelligence," objectively determined and measured. *The American Journal of Psychology*, 15(2), 201–292. <https://doi.org/10.2307/1412107>
- Spirites P., Glymour C., & Scheines R. (2000). *Causation, prediction, and search*. Adaptive Computation and Machine Learning (2nd ed.). MIT Press. (With additional material by David Heckerman, Christopher Meek, Gregory F. Cooper and Thomas Richardson, A Bradford Book).
- Steyer R. (2005). Analyzing individual and average causal effects via structural equation models. *Methodology*, 1(1), 39–54. <https://doi.org/10.1027/1614-1881.1.1.39>
- Strieder D., Freidling T., Haffner S., & Drton M. (2021). Confidence in causal discovery with linear causal models. In C. de Campos & M. H. Maathuis (Eds.). *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence of Proceedings of Machine Learning Research* (Vol. 161, pp. 1217–1226). PMLR.
- Sullivant S., Talaska K., & Draisma J. (2010). Trek separation for Gaussian graphical models. *The Annals of Statistics*, 38(3), 1665–1685. <https://doi.org/10.1214/09-AOS760>
- Sung W.-K. (2009). *Algorithms in bioinformatics: A practical introduction*. Chapman & Hall/CRC.
- Van der Vaart A. W. (1998). *Asymptotic statistics*. Cambridge Series in Statistical and Probabilistic Mathematics (Vol. 3). Cambridge University Press.
- Vershynin R. (2018). *High-dimensional probability: An introduction with applications in data science*. Cambridge Series in Statistical and Probabilistic Mathematics (Vol. 47). Cambridge University Press. (With a foreword by Sara van de Geer).
- Wishart J. (1928). Sampling errors in the theory of two factors. *British Journal of Psychology*, 19(2), 180–187. <https://doi.org/10.1111/j.2044-8295.1928.tb00508.x>
- Zwiernik P. (2016). *Semialgebraic statistics and latent tree models*. Monographs on Statistics and Applied Probability (Vol. 146). Chapman & Hall/CRC.