



# Graphical test for discrete uniformity and its applications in goodness-of-fit evaluation and multiple sample comparison

Teemu Säilynoja<sup>1</sup> · Paul-Christian Bürkner<sup>2</sup> · Aki Vehtari<sup>1</sup>

Received: 19 March 2021 / Accepted: 16 February 2022  
© The Author(s) 2022

## Abstract

Assessing goodness of fit to a given distribution plays an important role in computational statistics. The *probability integral transformation* (PIT) can be used to convert the question of whether a given sample originates from a reference distribution into a problem of testing for uniformity. We present new simulation- and optimization-based methods to obtain simultaneous confidence bands for the whole *empirical cumulative distribution function* (ECDF) of the PIT values under the assumption of uniformity. Simultaneous confidence bands correspond to such confidence intervals at each point that jointly satisfy a desired coverage. These methods can also be applied in cases where the reference distribution is represented only by a finite sample, which is useful, for example, for simulation-based calibration. The confidence bands provide an intuitive ECDF-based graphical test for uniformity, which also provides useful information on the quality of the discrepancy. We further extend the simulation and optimization methods to determine simultaneous confidence bands for testing whether multiple samples come from the same underlying distribution. This multiple sample comparison test is useful, for example, as a complementary diagnostic in multi-chain Markov chain Monte Carlo (MCMC) convergence diagnostics, where most currently used convergence diagnostics provide a single diagnostic value, but do not usually offer insight into the nature of the deviation. We provide numerical experiments to assess the properties of the tests using both simulated and real-world data and give recommendations on their practical application in computational statistics workflows.

**Keywords** PIT · ECDF · Uniformity test · Simulation-based calibration · MCMC convergence diagnostic

## 1 Introduction

Tests for uniformity play an essential role in computational statistics when estimating goodness of fit to a given distribution (Marhuenda et al. 2005). This is because, even when the distribution of interest is not uniform, there are methods to reduce the problem into testing for uniformity by transforming a sample from the given distribution to a (discrete or continuous) uniform distribution. Common use cases in Bayesian workflow (Gelman et al. 2020) are simulation-based calibration and Markov chain Monte Carlo convergence diagnostic, which we also use as examples in this paper. A graphical test can provide additional insight to

the nature of discrepancy that goes beyond the dichotomy of the uniformity test.

### 1.1 Probability integral transformation

Transforming sampled values to a uniform distribution is usually achieved via the *probability integral transform* (PIT), provided that the distribution of interest has a tractable *cumulative distribution function* (CDF) (D'Agostino and Stephens 1986). Let  $y_1, \dots, y_N \sim g(y)$  be an independent sample from an unknown *continuous* distribution with *probability density function* (PDF)  $g$ . We want to know whether  $g = p$ , where  $p$  is the PDF of a known distribution with a tractable CDF. The PIT of the sampled value  $y_i$  with respect to  $p$  is

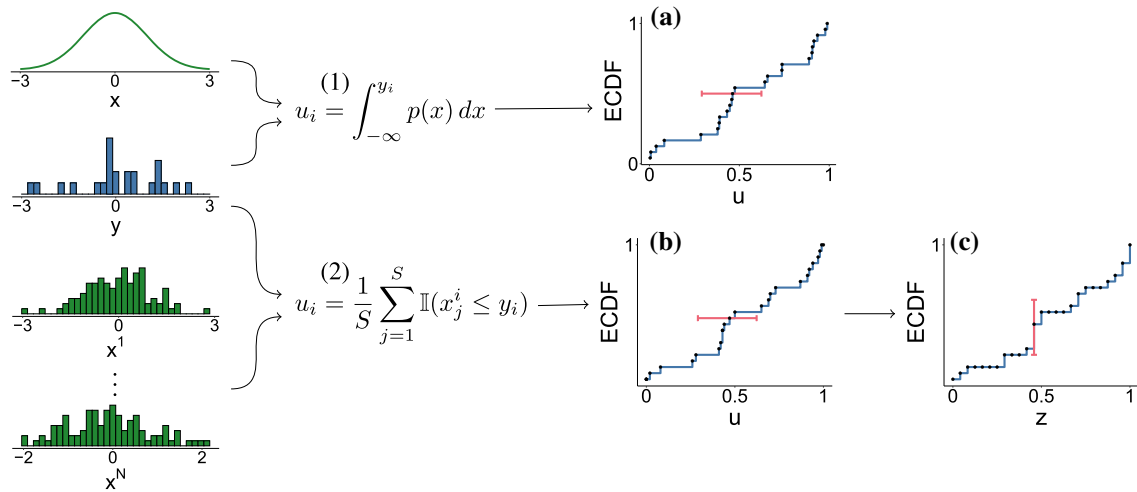
$$u_i = \int_{-\infty}^{y_i} p(x) dx. \quad (1)$$

If  $g = p$ , the transformed values  $u_i$  are continuously, independently, and uniformly distributed on the unit interval

✉ Teemu Säilynoja  
teemu.sailynoja@aalto.fi

<sup>1</sup> Department of Computer Science, Aalto university, Espoo, Finland

<sup>2</sup> Cluster of Excellence SimTech, University of Stuttgart, Stuttgart, Germany



**Fig. 1** Given  $y_1, \dots, y_N \sim g(y)$ , and a distribution  $p(x)$ , the hypothesis  $g = p$  can be assessed in two ways. (1) If the CDF or the PDF of  $p$  has a closed form, the PIT values  $u_i$  are continuous and, if  $g = p$ , uniformly distributed. (a) The pointwise confidence intervals (red bars) for the ordered statistic  $u_{(i)}$  are beta-distributed, and the simultaneous confidence intervals for the ECDF of  $u_i$  are given by Aldor-Noiman et al. (2013). (2) If one can draw an independent sample,  $x_1^i, \dots, x_S^i \sim p(x)$ ,

for each  $y_i$ , the empirical PIT values  $u_i$  are discrete and, given  $g = p$ , uniformly distributed. (b) The pointwise confidence intervals of the discrete ordered statistic  $u_{(i)}$  could be solved from Eq. (3). (c) The pointwise confidence intervals for the values of the ECDF of  $u_i$  at evaluation points  $z_i \in [0, 1]$  are binomially distributed, and the simultaneous confidence intervals are obtained by the method presented in this paper

$[0, 1]$ , reducing the evaluation of the hypothesis into testing for uniformity of the transformed sample  $u_1, \dots, u_N$ . If the integral (1) does not have the closed form, the CDF (and hence the PIT values) can still be computed with sufficient accuracy through numerical integration (e.g. quadrature), if at least the corresponding PDF is tractable.

If neither the CDF nor the PDF have closed form, but a comparison sample of independent values  $x_1^i, \dots, x_S^i \sim p(x)$  can be drawn separately for each  $y_i$ , the hypothesis  $g = p$  can be evaluated through the empirical PIT values

$$u_i = \frac{1}{S} \sum_{j=1}^S \mathbb{I}(x_j^i \leq y_i), \tag{2}$$

where  $\mathbb{I}$  is the indicator function. Now, given  $g = p$ , the transformed values  $u = u_1, \dots, u_N$  are independently distributed according to a discrete uniform distribution with  $S + 1$  values  $(0, 1/S, \dots, (S - 1)/S, 1)$ . Accordingly, we can still apply uniformity tests to assess  $g = p$ , just that this time, we need to test for discrete uniformity.

If either sample has dependencies, like an autocorrelated sample from a Markov chain, the ordered statistics is affected by the dependencies and the empirical PIT values (2) are not distributed uniformly even if  $g = p$  (unless the sample size goes to infinity). For Markov chains, the usual remedy is to thin the chain to obtain an approximately independent sample. This issue is illustrated, and thinning recommendations are provided in Appendix A.

Figure 1 shows an example of the empirical cumulative distribution function (ECDF) for  $u$  obtained through both Eq. (1) and Eq. (2). The figure also shows an example of a pointwise confidence interval for each ECDF. For the continuous integral of Eq. (1), the pointwise confidence interval can be computed from the continuous uniform ordered statistics distribution which is a common beta distribution. For the discrete sum of Eq. (2), the pointwise confidence interval can be computed from the discrete uniform ordered statistics distribution, with the cumulative distribution function of the  $i$ th ordered statistic  $u_{(i)}$  given as

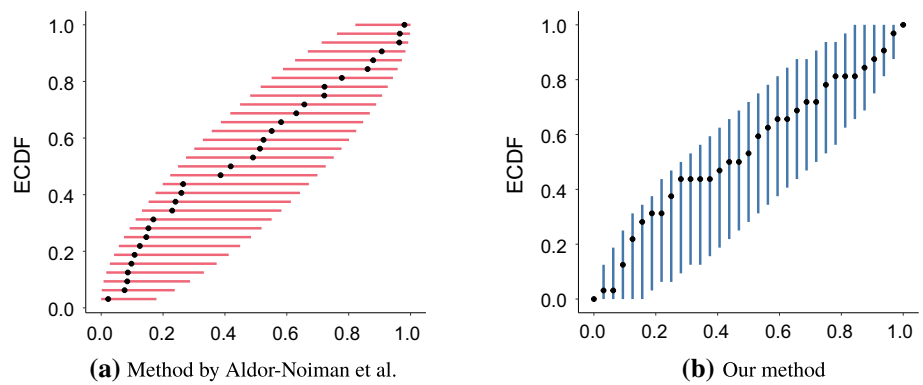
$$F_i(z) = \sum_{k=i}^N \binom{N}{k} z^k (1 - z)^{N-k}, \tag{3}$$

for  $z \in (0, 1/S, \dots, (S - 1)/S, 1)$  (Arnold et al. 2008, Example 3.1). The corresponding pointwise intervals do not have a nice form in general, and, more importantly, the discrete ordered statistics do not exhibit Markovian structure (exploited by our new optimization based approach) if there are possible ties in  $u$  (Arnold et al. 2008, Theorem 3.4.1).

To make the computation of the simultaneous confidence bands more straightforward and efficient, we propose making an additional transformation by computing the ECDF of  $u$  at chosen evaluation points  $z_i$ :

$$F(z_i) = \frac{1}{N} \sum_{j=1}^N \mathbb{I}(u_j \leq z_i). \tag{4}$$

**Fig. 2** Simultaneous confidence bands: **a** Method by Aldor-Noiman et al.: Beta distribution-based 95% simultaneous confidence bands for quantiles are provided for reaching a set of ECDF values (along the x-axis). **b** Our method: For a set of evaluation quantiles, we provide binomial distribution-based 95% simultaneous confidence intervals for the ECDF value (along the y-axis)



We recommend choosing  $z_i$  as the ordered fractional ranks  $\tilde{r}_i$  of  $y_i$ , defined as

$$\tilde{r}_i = \frac{1}{N} \sum_{j=1}^N \mathbb{I}(y_j \leq y_i). \tag{5}$$

The ordered fractional ranks form a uniform partition of the unit interval independent of the distribution of  $y_i$ . Thus, they provide an ECDF that is easier to interpret than the corresponding ECDF based directly on the original sample  $y_i$ . The resulting ECDF is illustrated in Fig. 1c. As we will show, useful properties of this ECDF are that 1) its pointwise confidence intervals can be computed easily from the binomial distribution, with a quantile function already implemented in most widely used environments for statistical computing, and 2) the distribution of the ECDF trajectories is Markovian, which is exploited in Sect. 2.3.

### 1.2 Simultaneous confidence bands

The major challenge that arises when developing a uniformity test based on the ECDF is to obtain simultaneous confidence bands with the desired overall coverage. For this purpose, one needs to take into account the inter-dependency in the ECDF values and adjust the coverage parameter accordingly. (We will discuss this in more detail in Sect. 2.)

When considering whether a given ECDF could present a sample from a uniform distribution, we need to jointly consider all pointwise uncertainties. For a set of evaluation points  $(z_i)_{i=1}^K$ , we provide lower and upper confidence bands  $L_i$  and  $U_i$ , respectively, that jointly satisfy

$$\Pr(L_i \leq F(z_i) \leq U_i \text{ for all } i \in \{1, \dots, K\}) = 1 - \alpha, \tag{6}$$

where  $F(z_i)$  is the ECDF of a sample from either the standard uniform distribution or discrete uniform distribution on the unit interval evaluated at  $z_i \in (0, 1)$  and  $1 - \alpha$  is the desired simultaneous confidence level. In addition to offering a numerical test for uniformity, the simultaneous confidence


bands provide an intuitive graphical representation of possible discrepancies from uniformity.

Aldor-Noiman et al. (2013) presented a simulation-based approach for computing simultaneous confidence band for the ECDF of the transformed sample acquired from Eq. (1) under the assumption of uniformity. In this paper, we present a simulation method inspired by Aldor-Noiman et al. (2013) as well as a new, faster optimization method for computing simultaneous confidence bands under uniformity, when the ECDF is computed from the empirical PIT values using Eqs. (2) and (4). Figure 2 contrasts the simultaneous confidence bands by Aldor-Noiman et al. (2013) against those obtained from our proposed method. Furthermore, we generalize our method and simultaneous confidence bands to test whether multiple samples originate from the same underlying distribution.

### 1.3 Related work

The idea of utilizing the ECDF to test uniformity is not new, but its potential has not yet been realized in full. For example, the well-known Kolmogorov–Smirnov (KS) test, first introduced by Kolmogorov (see e.g. Massey (1951), original article in Italian is Kolmogorov (1933)), is based on evaluating the maximum deviation of the sample ECDF from the theoretical CDF of the distribution to be tested against. Unfortunately, the KS test is relatively insensitive to deviations in the tails of the distribution (Aldor-Noiman et al. 2013), and numerous tests have been proposed to replace the KS test. An extensive comparison of more than thirty tests of uniformity of a single sample is provided by Marhuenda et al. (2005).

Due to its ease of interpretation and familiarity to people even with basic statistical knowledge, a graphical method for assessing uniformity commonly used as part of many statistical workflows is plotting histograms. This can even be turned into a formal test of uniformity with confidence intervals for the individual bins (e.g. Talts et al. 2020). Drawbacks of histograms are that binning discards information, there can be binning artefacts depending on the choice of bin width and



**Fig. 3** Four visualizations depicting the same random sample of 250 values. To assess uniformity of the sample, histograms **a** and **b** show a 95% confidence interval for each bin. Histograms can be sensitive to the number and placement of the bins selected, and the confidence intervals do not take into account possible inter-dependencies between the bin heights. For example, given the same sample, a 20-bin histogram stays

within the confidence interval **a**, but a 50-bin histogram exceeds the confidence interval **b**. The ECDF plot **c** and ECDF difference plot **d** with 95% simultaneous confidence bands for the ECDF both show the sample staying within the given limits with the ECDF difference plot providing a more dynamic range for the visualization

placement, and they ignore the dependency between bins. The proposed ECDF-based method does not require binning or smoothing, provides intuitive visual interpretation, and works for continuous Eq. (1) and discrete Eq. (2) values. An illustration and comparison of histograms with two binning choices and our new method is shown in Fig. 3. The visual range between the simultaneous confidence bands for the ECDF is often narrow when visualizing a sample with a large number of observations. Thus, to achieve a more dynamic range for the visualization, we recommend to show ECDF difference plots, instead, as illustrated in Fig. 3d. The ECDF difference plot is obtained by subtracting the values of the expected theoretical CDF (i.e. the identity function in  $[0, 1]$  in case of standard uniformity) from the observed ECDF values.

## 1.4 Summary of contributions

In this article, we focus on use case examples arising from inference validation and Markov chain Monte Carlo (MCMC) convergence diagnostics as part of a Bayesian workflow (Gelman et al. 2020), but our developed methods are applicable more generally. Our use cases can be divided into two main categories: a single sample test for uniformity and a multiple sample comparison where the hypothesis is tested that the samples are drawn from the same underlying (potentially non-uniform) distribution. We discuss both cases in more detail below.

We offer a graphical test for uniformity by providing simultaneous confidence bands for one or more ECDF trajectories obtained through the empirical probability integral transformation.

As our first contribution, we modify an existing ECDF-based approach proposed by Aldor-Noiman et al. (2013) to take into account the discreteness of the fractional rank-based

PIT values. This forms the basis for our proposed single- and multi-sample tests.

As our second contribution, we provide both a simulation and optimization method to determine the adjustment needed to achieve a desired simultaneous confidence level for the ECDF trajectory given the fractional rank-based PIT values. In addition to presenting a simulation-based adjustment following the method of Aldor-Noiman et al. (2013),

we introduce a new optimization method that is computationally considerably more efficient in determining the needed adjustment, especially when bands with high resolution are desired for a large sample size.

Although our focus is on providing a test with an intuitive graphical representation, we show that our method performs competitively when compared to existing uniformity tests with the state-of-the-art performance. We demonstrate the usefulness of this graphical test in context of simulation-based calibration approach for assessing inference methods (Talts et al. 2020).

Finally, as our third contribution, we generalize the graphical test as well as both the simulation and optimization methods to evaluate the hypothesis that two or more samples are drawn from the same underlying distribution. We demonstrate the usefulness of this graphical test in MCMC convergence diagnostics, where the currently most common graphical tools for assessing convergence are trace plots of the individual sampled chains.

## 1.5 Outline of the paper

In Sect. 2, we first provide a simulation-based method to determine simultaneous confidence bands for the ECDF of a single uniform sample and then present new more efficient optimization-based method.

In Sect. 3, we extend the test to multiple sample comparison and follow a similar structure by offering both a simulation and an optimization-based methods.

We continue in Sect. 4 with simulated and real-world examples illustrating the application of our proposed method and end with a discussion in Sect. 5.

## 2 Simultaneous confidence bands for the empirical cumulative distribution

We propose simulation- and optimization-based approaches to providing the ECDF of a uniform sample with  $1 - \alpha$ -level simultaneous confidence bands that are compatible with empirical PIT values, that is, confidence bands with a type-1-error rate of  $\alpha$ . Our approach is similar to that presented by Aldor-Noiman et al. (2013) with one central distinction illustrated in Fig. 2. The method by Aldor-Noiman et al. (2013) obtains simultaneous confidence bands for the evaluation quantiles with fixed ECDF values based on beta distributions, that is, it obtains confidence bands along the horizontal axis (Fig. 2a). In contrast, our new method provides simultaneous confidence bands for the ECDF values at fixed evaluation quantiles based on binomial distributions, that is, it obtains confidence bands along the vertical axis (Fig. 2b). In the limit, as the sample size approaches infinity, there is no practical difference between the methods. However, when the number of possible unique ranks is small, our proposed method behaves better for smallest and largest ranks and consistently if the ranks are further binned.

### 2.1 Pointwise confidence bands

Determining the pointwise confidence interval for the ECDF value of a sample from the continuous uniform distribution at a given evaluation point  $z_i \in (0, 1)$  is rather straightforward. By definition, given a sample  $u = u_1, \dots, u_N$ , the ECDF value is

$$F(z_i) = \frac{1}{N} \sum_{j=1}^N \mathbb{I}(u_j \leq z_i). \tag{7}$$

As the sampled values,  $u_j \in (0, 1)$ , are expected to be continuously uniformly distributed,  $\Pr(u_j \leq z_i) = z_i$  for each  $j = 1, \dots, N$ . Thus, the values resulting from scaling the ECDF with the sample size  $N$  are binomially distributed as

$$NF(z_i) \sim \text{Bin}(N, z_i). \tag{8}$$

If we instead expect  $u$  to be sampled from a discrete uniform distribution with  $S$  distinct equally spaced values,  $s_j = j/S$ , by choosing the partition points to form a subset

of these category values, we again have  $\Pr(u_j \leq z_i) = z_i$  for  $j = 1, \dots, N$ , and the marginal distribution of the scaled ECDF follows Eq. (8). Therefore, the methods introduced in Sects. 2.2 and 2.3 can be used to determine simultaneous confidence bands for both continuous and discrete uniform samples, allowing for testing uniformity of both the continuous PIT values of Eq. (1) and the discrete empirical PIT values in Eq. (2).

From Equation (8), it is straightforward to determine the  $1 - \alpha$  level pointwise lower and upper confidence bands,  $L_i$  and  $U_i$ , respectively, satisfying for all  $i = 1, \dots, N$  individually

$$\Pr(L_i \leq F(z_i) \leq U_i) = 1 - \alpha. \tag{9}$$

In contrast, determining the simultaneous confidence bands for ECDF trajectories (i.e. sets of ECDF values) is more complicated. In Fig. 4, we illustrate the dependency between ECDF values at distinct evaluation quantiles, together with simultaneous confidence bands computed via either of the new methods described in the following sections. As is illustrated in the figure, ECDF values evaluated at two quantiles close to each other are strongly dependent, while ECDF values evaluated at two quantiles far away from each other are only weakly dependent. In any case, these dependencies need to be taken into account when constructing simultaneous confidence bands.


Another important remark is that, as the marginal distribution of the scaled ECDF is discrete, the simultaneous confidence intervals do not in all cases meet the desired coverage level exactly. Brown et al. (2001) provide a thorough exploration of the effect discreteness plays in the coverage level of various interval estimations for binomial proportion, with listings of what the authors call lucky and unlucky sample lengths. In our experience, even though discreteness plays a role in the coverage level of the pointwise confidence intervals, this effect is reduced to deviation of under  $\pm 1\%$  for  $N \in [50, 2000]$  in the coverage level of the resulting simultaneous confidence bands we introduce next.

### 2.2 Simultaneous confidence bands through simulation

Our goal is to define simultaneous confidence bands for the ECDF of a sample of  $N$  values drawn from the standard uniform distribution so that the interior of the confidence bands contains trajectories induced by that distribution with rate  $1 - \alpha$ , where  $\alpha \in (0, 1)$ .

In this section, we describe a simulation-based method for determining the simultaneous confidence bands for the ECDF trajectory.

We follow steps similar to those introduced by Aldor-Noiman et al. (2013); with the exception that instead of



**Fig. 4** Dependence between the ECDF values of standard uniform samples evaluated at three distinct points. Simultaneous 95% confidence bands for the ECDF computed via either of the new methods introduced in this paper are shown on the left. In the middle, one can see

a stronger dependency between the ECDF values obtained at evaluation points close to each other, whereas on the right the ECDF values are only weakly dependent as the evaluation points far away from each other

determining limits for the Q-Q plot, we now determine the upper and lower limits of the ECDF values at the evaluation points  $z_i$ :

1. Choose a partition  $(z_i)_{i=1}^K$  of the unit interval.
2. Determine coverage parameter  $\gamma$  to account for multiplicity in order to obtain the  $1 - \alpha$ -level simultaneous confidence bands:

$$\Pr(L_i(\gamma) \leq F(z_i) \leq U_i(\gamma) \text{ for all } i \in \{1, \dots, K\}) = 1 - \alpha. \tag{10}$$

In determining these confidence bands, we use the knowledge from that the values of the scaled ECDF at each point  $z_i$  follow a binomial distribution and denote the value of the cumulative binomial distribution function with parameters  $N$  and  $z_i$  at  $k \in \mathbb{N}$  by  $\text{Bin}(k | N, z_i)$  and its inverse by  $\text{Bin}^{-1}(q | N, z_i)$  for quantile  $q \in [0, 1]$ .

To find the desired coverage value  $\gamma$ , we simulate  $M$  draws of size  $N$  from the standard uniform distribution. Let  $F^m$  denote the ECDF of the  $m$ th sample,  $u_1^m, \dots, u_N^m \sim \text{uniform}(0, 1)$ . For each sample, we find the value of  $\gamma$  such that the equal tail quantiles

$$L_i(\gamma) = \frac{1}{N} \text{Bin}^{-1}\left(\frac{\gamma}{2} | N, z_i\right), \tag{11}$$

and

$$U_i(\gamma) = \frac{1}{N} \text{Bin}^{-1}\left(1 - \frac{\gamma}{2} | N, z_i\right) \tag{12}$$

provide the tightest possible lower and upper limits, respectively, to the sample ECDF,  $F^m$ , at each  $z_i$ . This value of  $\gamma$  for the  $m$ th sample is

$$\gamma^m = 2 \min \left\{ \min \left( \text{Bin}(NF^m(z_i) | N, z_i), 1 - \text{Bin}(NF^m(z_i) - 1 | N, z_i) \right) \right\} \tag{13}$$

As now we have for  $\gamma^m$  equally that

$$\gamma^m = \arg \max_{\gamma} \left\{ \frac{\gamma}{2} \leq \text{Bin}(NF^m(z_i) | N, z_i) \leq 1 - \frac{\gamma}{2} \mid \forall i \right\} \tag{14}$$

and as it holds that

$$\begin{aligned} \frac{\gamma}{2} &\leq \text{Bin}(NF^m(z_i) | N, z_i) \leq 1 - \frac{\gamma}{2} \\ &\Rightarrow \text{Bin}^{-1}\left(\frac{\gamma}{2} | N, z_i\right) \\ &\leq NF^m(z_i) \leq \text{Bin}^{-1}\left(1 - \frac{\gamma}{2} | N, z_i\right) \\ &\Rightarrow L_i(\gamma) \leq F^m(z_i) \leq U_i(\gamma), \end{aligned}$$

$\gamma^m$  defines a set of upper and lower limits to the ECDF which is by Eq. (14) the tightest possible pair of limits defining equal tail quantiles for the ECDF at each  $z_i$ . To obtain bands covering a  $1 - \alpha$  fraction of the ECDFs of the simulated samples, we set  $\gamma$  to the  $\alpha$  quantile of the values  $\{\gamma^1, \dots, \gamma^M\}$ . Since  $\gamma^m > 0$  by construction, we also have  $\gamma > 0$ .

The following steps summarize the algorithm for simulating the adjusted coverage parameter  $\gamma$  and determining the  $1 - \gamma$  level simultaneous confidence bands:

1. For  $m = 1, \dots, M$ :
  - (a) Simulate  $u_1^m, \dots, u_N^m \sim \text{uniform}(0, 1)$ .
  - (b) For  $i = 1, \dots, K$ , compute  $F^m(z_i)$ .
  - (c) For  $i = 1, \dots, K$ , compute

$$\text{Bin}(NF^m(z_i) | N, z_i) \text{ and } \text{Bin}(NF^m(z_i) - 1 | N, z_i).$$

(d) Find the minimum probability

$$\gamma^m = 2 \min_i \left\{ \min \left( \text{Bin}(NF^m(z_i) \mid N, z_i), 1 - \text{Bin}(NF^m(z_i) - 1 \mid N, z_i) \right) \right\}.$$

2. Set  $\gamma$  to be the  $100\alpha$  percentile of  $\{\gamma^1, \dots, \gamma^M\}$ .
3. Form the confidence bands

$$[L_i(\gamma), U_i(\gamma)] = \left[ \frac{1}{N} \text{Bin}^{-1} \left( \frac{\gamma}{2} \mid N, z_i \right), \frac{1}{N} \text{Bin}^{-1} \left( 1 - \frac{\gamma}{2} \mid N, z_i \right) \right]$$

for  $i = 1, \dots, K$ .

### 2.3 Simultaneous confidence bands through optimization

We also propose a computationally more efficient optimization-based method for determining the simultaneous confidence bands.

In the following derivation of the optimization method, we denote the interior of the confidence bands for the ECDF at quantile  $z_i$  as  $\tilde{I}_i(\gamma)$ . By denoting  $r_i = NF(z_i)$ , the scaled interior  $I_i(\gamma)$  for  $r_i$  is given by

$$I_i(\gamma) = \left\{ r \in \{0, \dots, N\} \mid \text{Bin}^{-1} \left( \frac{\gamma}{2} \mid N, z_i \right) \leq r \leq \text{Bin}^{-1} \left( 1 - \frac{\gamma}{2} \mid N, z_i \right) \right\}. \tag{15}$$

As is common for discrete statistical tests, we treat the borders between the interior and exterior as being within the confidence bands. Based on  $I_i(\gamma)$ , we can easily obtain  $\tilde{I}_i(\gamma)$  as  $r \in I_i(\gamma)$  is equivalent to  $r/N \in \tilde{I}_i(\gamma)$ .

A scaled ECDF trajectory defined as

$$t_0^K = \left( (z_i)_{i=0}^K, (r_i)_{i=0}^K \right) \tag{16}$$

with  $z_0 = 0$  and  $z_K = 1$  stays within the simultaneous confidence bands completely if and only if  $r_i \in I_i(\gamma)$  for all  $i \in \{0, \dots, K\}$ . If we denote the set of trajectories fulfilling  $r_i \in I_i$  as  $T_i$ , we can write the set of trajectories which are completely within the simultaneous confidence bands as

$$T(\gamma) = \bigcap_{i=0}^K T_i(\gamma). \tag{17}$$

In order for the simultaneous confidence bands to have confidence level  $1 - \alpha$ , we must have

$$\Pr(T(\gamma)) = 1 - \alpha. \tag{18}$$

Due to the pairwise independence of the original draws  $u_i$  (by assumption), the distribution of the ECDF values within a single trajectory is Markovian in the sense that the ECDF value  $F(z_{i+1})$  only depends on the observed value at the previous evaluation point,  $F(z_i)$  and not on the earlier behaviour of the ECDF trajectory.

This implies that, under uniformity of the original distribution, the remaining  $N - NF(z_i) = N - r_i$  samples are uniformly distributed on the interval  $[z_i, 1]$ , and thus the growth of the scaled ECDF from  $r_i$  to  $r_{i+1}$ , between  $z_i$  and  $z_{i+1}$  is binomially distributed with  $N - r_i$  trials and the success probability

$$\tilde{z}_{i+1} = \frac{z_{i+1} - z_i}{1 - z_i}. \tag{19}$$

And so we have

$$\Pr(r_{i+1} \mid r_i) = \text{Bin}(r_{i+1} - r_i \mid N - r_i, \tilde{z}_{i+1}). \tag{20}$$

The probability for  $r_{i+1} = k \in I_{i+1}$  to occur in a scaled ECDF trajectory  $t_0^K$  which stayed within the simultaneous confidence bands until point  $i$ , that is, for which we have

$$t_0^i \in \bigcap_{j=0}^i T_j(\gamma), \tag{21}$$

can thus be written recursively as

$$\Pr \left( r_{i+1} = k \cap \bigcap_{j=0}^i T_j(\gamma) \right) = \sum_{m \in I_i} \Pr \left( r_i = m \cap \bigcap_{n=0}^{i-1} T_n(\gamma) \right) \Pr(r_{i+1} = k \mid r_i = m). \tag{22}$$

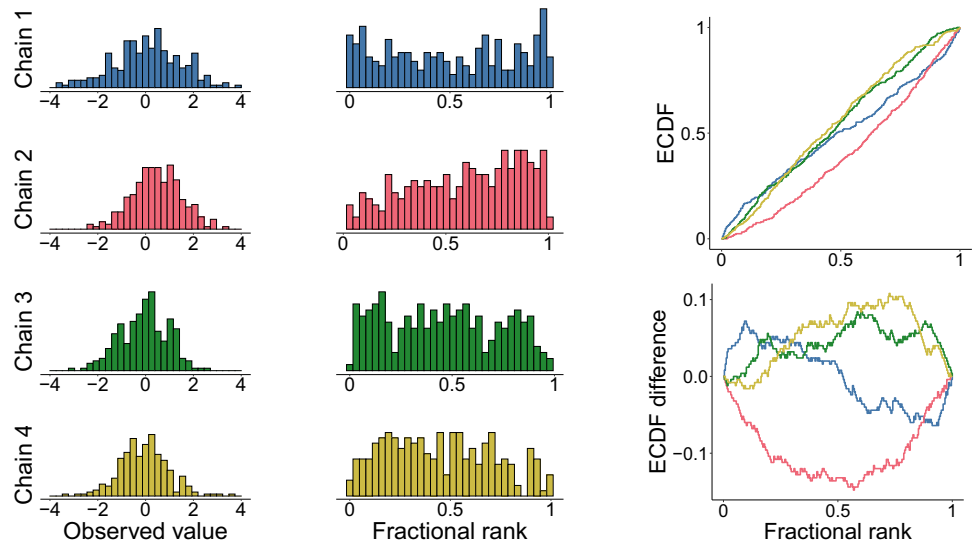
The recursion is initialized at  $z_0 = 0$  with  $\Pr(r_0 = 0) = 1$  so that  $\Pr(T_0(\gamma)) = 1$  for all  $\gamma \in [0, 1]$ . At any point  $i \in \{0, \dots, K\}$ , we can obtain

$$\Pr \left( \bigcap_{j=0}^i T_j(\gamma) \right) = \sum_{m \in I_i} \Pr \left( r_i = m \cap \bigcap_{n=0}^{i-1} T_n(\gamma) \right), \tag{23}$$

which is equal to  $\Pr(T(\gamma))$  when arriving at  $i = K$ . Clearly,  $\Pr(T(\gamma))$  is monotonically decreasing but not continuous in  $\gamma$  due to the discrete nature of the binomial distribution. Thus, Equation (18) will not have an exact solution in general and so we will not be able to meet the simultaneous confidence level  $1 - \alpha$  exactly. We can, however, try to get as close as possible by computing

$$\hat{\gamma} = \arg \min_{\gamma \in [0, \alpha]} |1 - \alpha - \Pr(T(\gamma))| \tag{24}$$

**Fig. 5** We evaluate the hypothesis of the four samples on the left originating from the same underlying distribution by inspecting the distribution of fractional ranks among the joint rank-transformed, Eq. (29), samples presented in the middle. When comparing multiple samples, we would like to take into account the within sample dependency, but also the between sample dependency introduced by the joint transformation. In Sects. 3.2 and 3.3, we extend our methods in order to provide simultaneous confidence bands for the ECDF and the ECDF difference plots shown on the right



with a unidimensional derivative-free optimizer. In our experiments, the optimizer proposed by Brent (1973) (which is implemented, e.g. in the R function `optimize`) converged quickly in all cases to  $\hat{\gamma}$  values implying a simultaneous confidence level very close to the nominal  $1 - \alpha$ .

With a 2015 laptop equipped with a 2.90GHz Intel® Core™ i5-5287U processor, the optimization method reduces the time required to compute the adjustment parameter  $\gamma$  from 10s to 600ms for a sample of length 250 when compared against the time required for 10,000 steps of the simulation method. With  $N = 1000$  this reduction is from 75s to 10s.

Both of the implementations used for this article only use a single computation thread, but would benefit from parallelization, as both methods include independent iterations.

The computation time required can be further reduced by using a grid of pre-computed values as the adjustment parameters, and interpolate for different values of  $N$  in log-log scale.

### 3 Comparison of multiple samples

In this section, we extend the uniformity test of Section 2 to test whether multiple samples originate from the same underlying distribution. In the case of multiple samples sharing the same distribution, the rank statistics of the values within each sample, when ranked jointly across all samples, are uniformly distributed on the interval  $(1, \tilde{N})$ , where  $\tilde{N}$  is the total length of the combined sample (Vehtari et al. 2021). Thus, instead of considering the sampled values directly, we consider the implied jointly rank-transformed values.

Due to this joint rank-transformation, the resulting chains are dependent on each other and the confidence intervals we

construct in the following two sections are used to answer whether all the two or more samples originate from the same underlying distribution. In other words, in the case one or more of the ECDF trajectories leaves the confidence bands, we conclude that at least one of the samples exhibits larger-than-expected deviance from the other samples at hand.

An illustration of the connection between the sampled values, the corresponding fractional rank statistics and the two ECDF plots of these rank statistics are displayed in Fig. 5.

#### 3.1 Pointwise confidence bands

An important distinction to the ECDF case considered in Section 2 is the form of the marginal distribution at quantile  $z_i$  when determining the adjusted coverage parameter  $\gamma$ . As our main application is the comparison of distributions induced by MCMC chains, we speak of the  $L$  different samples as chains and assume all chains to have the same length  $N$ . We define  $r_i$  as the vector (of length  $L$ ) of joint ranks across chains smaller than or equal to the sample size  $s_i = \lfloor z_i N L \rfloor$ . That is, for each of the  $L$  elements  $r_{il}$  of  $r_i$ , we have

$$r_{il} = \left\{ \sum_{j=1}^{s_i} \mathbb{I}_{\{1, \dots, s_i\}} (R(u_{lj} | u)) \right\}, \tag{25}$$

where  $u_{lj}$  is the  $j$ th draw of the  $l$ th chain before transformation,  $R(u_{lj} | u)$  is the rank of  $u_{lj}$  within the vector  $u$  of all draws across all chains, and  $\mathbb{I}$  is the indicator function. Clearly, because of the definition of ranks, we know for all  $i$  that

$$\sum_{l=1}^L r_{il} = s_i, \tag{26}$$



and we define the set of all  $r_i$  satisfying (26) as  $R_i$ . Due to the pairwise independence of the original draws  $u_{lj}$  (by assumption), the marginal distribution of  $r_i$  at quantile  $z_i$  is multivariate hypergeometric

$$r_i \sim \text{MHyp}(\tilde{N}, s_i), \tag{27}$$

where  $\tilde{N} = (N_1, \dots, N_L)$  is the vector chain lengths (i.e. population sizes) and  $N_1 = \dots = N_L = N$  as we assume chains to have equal length. It is well known that, in this case, the marginal distribution of  $r_{il}$ , and thus the distribution defining the pointwise confidence bands, is hypergeometric

$$r_{il} \sim \text{Hyp}(N, N(L - 1), s_i). \tag{28}$$

### 3.2 Simultaneous confidence bands through simulation

In this section, we extend the simulation method presented in Sect. 2.2 to comparison of multiple samples. Our aim is to define simultaneous confidence bands for the ECDFs of multiple, jointly rank-transformed distributions so that the interior of the simultaneous confidence bands jointly contains all trajectories induced by the rank-transformed distributions with rate  $1 - \alpha$ . To this end, we define  $r_i$  and  $s_i$  as in Sect. 3.1 and denote the interior of the simultaneous confidence bands at quantile  $z_i$  as  $\tilde{I}_i(\gamma)$ , with  $\gamma$  being the adjusted coverage parameter to be determined.

We continue the use of fractional ranks in the ECDF plots to provide illustrations independent of the length of the sampled chains. Suppose we have  $L$  chains of length  $N$ . The fractional rank score  $\tilde{r}_{il}$  corresponding to the  $i$ th value of the  $l$ th chain,  $u_{li}$ , is

$$\tilde{r}_{il} = \frac{R(u_{lj} \mid u)}{LN}. \tag{29}$$

Instead of using the adjusted value of  $\gamma$  to obtain the  $1 - \alpha$  level simultaneous confidence bands for a single ECDF trajectory, we adjust  $\gamma$  to account for the dependence between the samples introduced in the transformation into fractional ranks. That is, after choosing the evaluation quantiles  $z_i$ , we adjust  $\gamma$  to find upper and lower simultaneous confidence bands satisfying

$$\Pr(L_i(\gamma) \leq F_l(z_i) \leq U_i(\gamma) \text{ for all } i \in \{1, \dots, K\}, l \in \{1, \dots, L\}) = 1 - \alpha, \tag{30}$$

where  $F_l$  is the ECDF of the fractional rank scores of the  $l$ th chain.

We denote the CDF of the hypergeometric distribution as  $\text{Hyp}$  and its inverse as  $\text{Hyp}^{-1}$ . The algorithm to approximate the adjusted coverage parameter  $\gamma$  when comparing  $L$  samples is as follows:

1. For  $m = 1, \dots, M$ :
  - (a) For  $l = 1, \dots, L$ , simulate  $u_{l1}^m, \dots, u_{lN}^m \sim \text{uniform}(0, 1)$ .
  - (b) For  $j = 1, \dots, N$  and  $l = 1, \dots, L$ , compute  $\tilde{r}_{jl}^m$ .
  - (c) For  $i = 1, \dots, K$  and  $l = 1, \dots, L$ , compute  $F_l^m(z_i)$ .
  - (d) For  $i = 1, \dots, K$  and  $l = 1, \dots, L$ , compute

$$\text{Hyp}(NF_l^m(z_i) \mid N, (L - 1)N, s_i) \text{ and } \text{Hyp}(NF_l^m(z_i) - 1 \mid N, (L - 1)N, s_i),$$

where  $s_i = \lfloor z_i NL \rfloor$ .

- (e) Find the minimum probability

$$\gamma^m = 2 \min_{i,l} \{ \min(\text{Hyp}(NF_l^m(z_i) \mid N, (L - 1)N, s_i), 1 - \text{Hyp}(NF_l^m(z_i) - 1 \mid N, (L - 1)N, s_i)) \}.$$

2. Set  $\gamma$  to be the  $100\alpha$  percentile of  $\{\gamma^1, \dots, \gamma^M\}$ .
3. Form the confidence bands

$$[L_i(\gamma), U_i(\gamma)] = \left[ \text{Hyp}^{-1}\left(\frac{\gamma}{2} \mid N, N(L - 1), s_i\right), \text{Hyp}^{-1}\left(1 - \frac{\gamma}{2} \mid N, N(L - 1), s_i\right) \right],$$

for  $i = 1, \dots, K$ .

### 3.3 Simultaneous confidence bands through optimization

In this section, we extend the optimization method presented in Sect. 2.3 to comparison of multiple samples. With the marginal distribution of  $r_{il}$  being hypergeometric, the rank interior  $I_i(\gamma)$  for  $z_i$  is given by

$$I_i(\gamma) = \left\{ r \in R_i \mid \forall r_l \in r : \text{Hyp}^{-1}\left(\frac{\gamma}{2} \mid N, N(L - 1), s_i\right) \leq r_l \leq \text{Hyp}^{-1}\left(1 - \frac{\gamma}{2} \mid N, N(L - 1), s_i\right) \right\}. \tag{31}$$

We treat the borders between interior and exterior as belonging to the interior. Based on  $I_i(\gamma)$ , we can again easily obtain  $\tilde{I}_i(\gamma)$ , as  $r \in I_i(\gamma)$  is equivalent to  $r/N \in \tilde{I}_i(\gamma)$ .

The remainder of the proof proceeds similar to the one-sample case, except that we replace the binomial distribution with the (multivariate) hypergeometric distribution. A (multivariate) rank ECDF trajectory, defined as

$$t_0^K = \left( (z_i)_{i=0}^K, (r_i)_{i=0}^K \right), \tag{32}$$

where  $z_0 = 0$  and  $z_K = 1$ , stays within the simultaneous confidence bands completely if and only if  $r_i \in I_i(\gamma)$  for all

$i \in \{0, \dots, K\}$ . If we denote the set of trajectories fulfilling  $r_i \in I_i$  as  $T_i$ , we can write the set of trajectories which are completely in the interior of the simultaneous confidence bands as

$$T(\gamma) = \bigcap_{i=0}^K T_i(\gamma). \tag{33}$$

In order for the simultaneous confidence bands to have a confidence level  $1 - \alpha$ , we must satisfy

$$\Pr(T(\gamma)) = 1 - \alpha. \tag{34}$$

Due to the pairwise independence of the original draws  $u_{ij}$  (by assumption), the distribution of the rank ECDF trajectories again exhibits a similar Markovian property as in the single sample case. That is, any ECDF value  $F(z_{i+1})$  beyond a given point  $z_i$  only depends on  $F(z_i)$  but not on the earlier history of the ECDF trajectory. This implies that, under the assumption of all chains coming from the same underlying distribution, the growth  $r_{i+1} - r_i$  of the ECDF from  $z_i$  to  $z_{i+1}$  is multivariate hypergeometric with  $\tilde{N}_i = \tilde{N} - r_i$  and sample size  $\tilde{s}_{i+1} = s_{i+1} - s_i$ . Accordingly, we have

$$\Pr(r_{i+1} \mid r_i) = p_{\text{MHyp}}(r_{i+1} - r_i \mid \tilde{N}_i, \tilde{s}_{i+1}), \tag{35}$$

where  $p_{\text{MHyp}}$  denotes the discrete PDF of the multivariate hypergeometric distribution. The probability for  $r_{i+1} = k \in I_{i+1}$  to occur in a rank ECDF trajectory  $t_0^K$  which stayed in the simultaneous confidence bands until point  $i$ , that is, for which we have

$$t_0^i \in \bigcap_{j=0}^i T_j(\gamma) \tag{36}$$

can thus be written recursively as

$$\Pr\left(r_{i+1} = k \mid \bigcap_{j=0}^i T_j(\gamma)\right) = \sum_{m \in I_i} \Pr\left(r_i = m \mid \bigcap_{n=0}^{i-1} T_n(\gamma)\right) \Pr(r_{i+1} = k \mid r_i = m). \tag{37}$$

The recursion is initialized at  $z_0 = 0$  with  $\Pr(x_0 = (0, \dots, 0)) = 1$  so that  $\Pr(T_0(\gamma)) = 1$  for all  $\gamma \in [0, 1]$ . At any point  $i \in \{0, \dots, K\}$ , we can obtain

$$\Pr\left(\bigcap_{j=0}^i T_j(\gamma)\right) = \sum_{m \in I_i} \Pr\left(r_i = m \mid \bigcap_{n=0}^{i-1} T_n(\gamma)\right), \tag{38}$$

which is equal to  $\Pr(T(\gamma))$  when arriving at  $i = K$ . Clearly,  $\Pr(T(\gamma))$  is monotonically decreasing but not continuous in

$\gamma$  due to the discrete nature of the (multivariate) hypergeometric distribution. We can compute

$$\hat{\gamma} = \arg \min_{\gamma \in [0, \alpha]} |1 - \alpha - \Pr(T(\gamma))| \tag{39}$$

using a unidimensional derivative-free optimizer. In our experiments, the optimizer proposed by Brent (1973) converged in all cases to  $\hat{\gamma}$  values implying a simultaneous confidence level very close to the nominal  $1 - \alpha$ .


Unfortunately, evaluating Eq. (37) suffers from combinatorial explosion as the  $R_i$  are  $L$ -dimensional sets constraint only by Equation (26) and as  $\Pr(r_{i+1} = k \mid r_i = m)$  has to be computed for all combinations of elements  $k \in I_{i+1}$  and  $m \in I_{i+1}$  at each point  $i$ . Several measures can be taken to reduce the complexity of the computation. First, the ranks of one of the  $L$  chains are redundant as they follow deterministically from Equation (26) based on the ranks of the other  $L - 1$  chains. This implies in particular that the 2-chain case has the same computational complexity as the one-sample case as only one of the two chains needs to be evaluated. Second, due to a priori symmetry of the chains, we can, without loss of generality, assume at the first non-zero quantile  $z_1$  that the elements  $r_{1l}$  of  $r_1$  are ordered such that  $r_{11} \leq r_{12} \leq \dots \leq r_{1L}$ . This reduces the number of trajectories to be evaluated by a factor of  $L(L + 1)/2$ . Still even with these measures in place, computation will scale badly with  $L$ , and the simulation-based method, which scales almost linearly, or grid-based interpolation from pre-computed values is faster for larger number of chains.

## 4 Numerical experiments and power analysis

In this section, we provide insights into how the plots produced by our proposed methods should be interpreted. In each of the following cases, we link together the histogram, ECDF plot, and the ECDF difference plot. The code for the experiments and plots is available at <https://github.com/TeemuSailyoja/simultaneous-confidence-bands>.

### 4.1 Uniformity of a Single Sample

We begin by providing two examples connecting the shape of the histogram of the transformed sample to the characteristics of the corresponding ECDF and ECDF difference plots with basic discrepancies between the sample and the comparison distribution. After this, we illustrate an application of our method as part of a workflow to detect issues in model implementation or the computation of the posterior distribution. Lastly we provide power analysis comparing the performance of our proposed method to existing state of the art tests for uniformity.



**Fig. 6** Effect of the sample mean. The histogram, ECDF plot, and ECDF difference plot of the empirical PIT values of  $y = y_1, \dots, y_{250} \sim \text{normal}(0.25, 1)$  with respect to  $x^i = x_1^i, \dots, x_{250}^i \sim \text{normal}(0, 1)$  for  $i = 1, \dots, N$ . The larger-than-expected mean of the sample is some-

what visible as a slant to the right in the histogram, whereas the ECDF difference plot displays a clear U-shape. In the histogram 95% confidence intervals are provided for each of the 50 bins and the ECDF plots show the 95% simultaneous confidence bands

With the exception of the power analysis tests in 4.1.4 where the samples are drawn directly from a continuous uniform distribution, the samples in the following examples are transformed to the unit interval from their respective sampling distributions through empirical PIT and are tested against the hypothesis of discrete uniformity.

**4.1.1 Effect of difference in sample mean**

To observe the typical characteristics of a sample with a mean different than that of the comparison distribution, we draw  $y = y_1, \dots, y_N \sim \text{normal}(0.25, 1)$  and  $N$  independent comparison samples  $x^i = x_1^i, \dots, x_N^i \sim \text{normal}(0, 1)$  with  $N = 250$ . We then test for  $y$  being standard normal distributed by transforming the sampled values to the unit interval through empirical PIT. Figure 6a shows the histogram of the transformed sample exhibiting a higher-than-expected mean. As seen in the figure, a shift in the sample mean leads to the histogram being slanted towards the direction of the shift.

The ECDF plot in Fig. 6b shows this shift through the ECDF of the PIT values remaining under the theoretical CDF, which is also seen in the ECDF difference plot in Fig. 6c.

If the sample in question would instead have a mean lower than expected, the histogram would be slanted to the left and the behaviour of the resulting ECDF plot and ECDF difference plot would be reversed. That is, the ECDF plot would stay above the theoretical CDF as a higher-than-expected density is covered at low fractional ranks and the ECDF difference plot would, respectively, show a  $\cap$ -shape above the zero level.

**4.1.2 Effect of difference in sample variance**

Next, we investigate an example where the sample has a higher-than-expected variance. To this end, we draw  $y =$

$y_1, \dots, y_N \sim \text{normal}(0, 1.25)$  and for each  $y_i$  a standard uniform comparison sample  $x^i = x_1^i, \dots, x_N^i \sim \text{normal}(0, 1)$  with  $N = 250$ . Figure 7a shows the histogram of the empirical PIT values. In general, a larger-than-expected variance leads to a U-shaped histogram and one can indeed see some of the histogram bins breaching the 95% confidence bounds.

In the ECDF plot shown in Fig. 7b, the larger-than-expected variance leads to faster-than-expected growth near the edges and slower-than-expected growth in the middle.

The shape is more clearly seen in the ECDF difference plot in Fig. 7c depicting the difference between the ECDF and the theoretical CDF.

If the sample would instead present a variance lower than expected, the histogram would be  $\cap$ -shaped and the behaviour of the resulting ECDF plot and ECDF difference plot would be reversed.


In the ECDF plot, this is shown as faster increase near the middle.

In general, the ECDF difference plot is decreasing when a smaller-than-expected density of samples is covered, and correspondingly increases when covering a higher-than-expected density.

**4.1.3 Simulation-based calibration: eight schools**

The eight schools (Gelman et al. 2013) is a classic hierarchical model example. The training course effects  $\theta_j$  in eight schools are modelled using an hierarchical varying intercept model.

If the model is constructed with the centred parameterization, the posterior distribution exhibits a funnel shape contracting to a region of high curvature near the population mean  $\mu$  when sampled with small values of the population standard deviation  $\tau$ . This property makes exploring the distribution of  $\tau$  difficult for many MCMC methods. The



**Fig. 7** Effect of the sample variance. The histogram, ECDF plot, and ECDF difference plot of the empirical PIT values of  $y = y_1, \dots, y_{250} \sim \text{normal}(0, 1.25)$  with respect to  $x^i = x_1^i, \dots, x_{250}^i \sim \text{normal}(0, 1)$  for  $i = 1, \dots, N$ . The larger-than-expected variance of the sample is visible as a U-shape in the histogram, whereas the ECDF difference

plot displays rapid growth near the ends of the interval as a larger-than-expected number of values is covered. In the histogram, 95% confidence intervals are provided for each of the 50 bins and the ECDF plots show the 95% simultaneous confidence bands

centred parameterization  $(\mu, \sigma, \mu_0, \tau)$  of the problem is as follows:

$$\mu_j \sim \text{normal}(\mu_0, \tau) \tag{40}$$

$$y_j \sim \text{normal}(\mu_j, \sigma_j). \tag{41}$$

Cook et al. (2006) proposed a simulation-based calibration (SBC) method for validating Bayesian inference software. The idea is based on the fact we can factor the joint distribution of data  $y$  and parameters  $\theta$  in two ways

$$\pi(y, \theta) = \pi(y|\theta)\pi(\theta) = \pi(\theta|y)\pi(y). \tag{42}$$

By considering  $\theta'$  and  $\theta''$  the joint distribution is

$$\pi(y, \theta', \theta'') = \pi(y)\pi(\theta'|y)\pi(\theta''|y), \tag{43}$$

and it is easy to see that  $\theta'$  and  $\theta''$  have the same distribution conditionally on  $y$ . If write the joint distribution in an alternative way

$$\pi(y, \theta', \theta'') = \pi(\theta')\pi(y|\theta')\pi(\theta''|y), \tag{44}$$

$\theta'$  and  $\theta''$  still have the same distribution conditionally on  $y$ . We can sample from the joint distribution  $\pi(y, \theta', \theta'')$  by first sampling from  $\pi(\theta')$  and  $\pi(y|\theta')$ , which is usually easy for generative models. The last step is to sample from the conditional  $\pi(\theta''|y)$ , which is usually not trivial and instead, for example, a Markov chain Monte Carlo algorithm is used. We can validate the algorithm and its implementation used to sample from  $\pi(\theta''|y)$  by checking that the samples obtained have the same distribution as  $\theta'$  (conditionally on  $y$ ).


Cook et al. (2006) operationalize the approach by drawing  $\theta'_i$  from  $\pi(\theta')$ , generating data  $y_i \sim \pi(y_i|\theta'_i)$  and then using

the algorithm to be validated to draw a sample  $\theta''_1, \dots, \theta''_S \sim \pi(\theta''|y_i)$ . If the algorithm and its implementation are correct, then  $\theta'_i, \theta''_1, \dots, \theta''_S$  conditional on  $y_i$  are draws from the same distribution. Cook et al. (2006) propose to compute empirical PIT valued for  $\theta'_i$  that they show to be uniformly distributed given  $S \rightarrow \infty$ . The process is repeated for  $i = 1, \dots, N$  and  $N$  empirical PIT values are used for testing. Cook et al. (2006) propose to use  $\chi^2$ -test for the inverse of the normal CDF of the empirical PIT values. However, with finite  $S$  this approach does not correctly take into account the discreteness or the effect of correlated sample from Markov chain Gelman (2017).

By thinning  $\theta''_1, \dots, \theta''_S$  to be approximately independent, the uniformity of empirical PIT values can be tested with the approach presented in this paper. See Appendix A for more on thinning.

Figure 8 shows the histogram and ECDF plots of 500 prior draws of the population standard deviation  $\tau$ , each ranked based on a thinned posterior sample of 150 draws obtained from a chain of 3000 draws. The graphical test rejects the hypothesis of the prior draws being uniform, moreover the ECDF plots show that the prior draws of the parameter  $\tau$  ranked in relation to the posterior samples obtained from the centred parameterization of the eight schools model are skewed to small ranks. This suggests that the MCMC is not sampling correctly from the target distribution (which in this case is known to be caused by inability to reach the narrow funnel part of the posterior).

In Sect. 4.2.4, we will return to the eight schools model by providing further analysis on the convergence of individual chains in the centred parameterization case and illustrating how our method can be used to detect these convergence issues.



**Fig. 8** SBC of the parameter  $\tau$  in the centred parameterization eight schools model. The empirical PIT values of the prior draws of  $\tau$ , when compared to the corresponding posterior samples, show a strong under-

representation of large PIT values in both the histogram and the ECDF difference plot. This indicates that the MCMC is not sampling correctly from the target distribution

#### 4.1.4 Power analysis

As our primary focus is on providing a graphical uniformity test, which gives the user useful information regarding the possible deviations from uniformity, we want to also ensure that the overall performance of our test is, if not the best, competitive with tests aimed at accurately detecting specific deviations from uniformity. To this end, we compare the sensitivity of our method with existing tests for uniformity, by considering the rejection rate of samples drawn from uniform distribution and then transformed according to the following three transformation families Marhuenda et al. (2005) use in their article comparing various tests for uniformity:

$$f_{A,k}(x) = 1 - (1 - x)^k, \quad 0 \leq x \leq 1, \tag{45}$$

$$f_{B,k}(x) = \begin{cases} 2^{k-1}x^k & \text{if } 0 \leq x \leq 0.5 \\ 1 - 2^{k-1}(1 - x)^k & \text{if } 0.5 < x \leq 1 \end{cases} \tag{46}$$

$$f_{C,k}(x) = \begin{cases} 0.5 - 2^{k-1}(0.5 - x)^k & \text{if } 0 \leq x \leq 0.5 \\ 0.5 + 2^{k-1}(x - 0.5)^k & \text{if } 0.5 < x \leq 1. \end{cases} \tag{47}$$

As Marhuenda et al. (2005) offer an extensive comparison of tests, we limit our comparison to the test specifically recommended to target each of the transformation families in addition to the widely known Kolmogorov–Smirnov test. For each of the test statistics, a critical value is calculated and samples exceeding that value are rejected.

For transformation family A, the recommended test is the mean distance of the  $i$ th value of the ordered sample  $u_{(i)}$  from the expected value  $i/(N + 1)$ :

$$T_1 = \sum_{i=1}^N \frac{|u_{(i)} - i/(N + 1)|}{N}. \tag{48}$$

For family B, the smooth goodness-of-fit test,  $N_h$ , introduced by Neyman (1937) is recommended with the dimension  $h$

chosen according to the method recommended by Ledwina (1994) resulting in the test statistic  $N_S$ , which also has the best overall performance across the transformation families. The test recommended for transformation family C is the statistic recommended by Watson (1961),

$$U^2 = W^2 - i(\bar{u} - 0.5)^2, \tag{49}$$

where  $\bar{u}$  is the mean of the  $u_i$  and  $W^2$  is the Cramér–von Mises statistic,

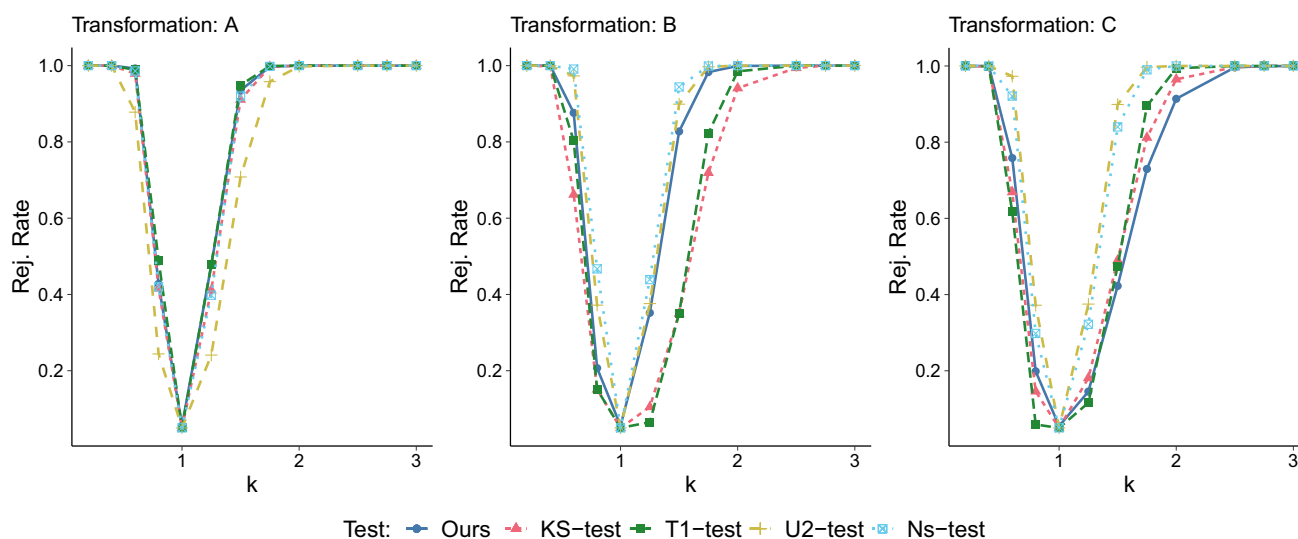
$$W^2 = \sum_{i=1}^N \left\{ u_{(i)} - \frac{2i - 1}{2N} \right\}^2 + \frac{1}{12N}. \tag{50}$$

The rejection rates of these tests and our test through simultaneous confidence bands are shown in Fig. 9 for families A, B, and C with sample size  $N = 100$  and  $k$  varying between 0.20 and 3.00. For each value of  $k$ , the rejection rate among 100,000 samples was computed. As seen from these results, the proposed ECDF simultaneous confidence band method performs in a manner similar to the recommended tests with the exception of family C, where our method exhibits a lower rejection rate compared to some of the other tests.

#### 4.2 Comparing multiple samples

When testing if two or more samples are produced from the same underlying distribution, we can compare the ranks of each sample relative to the sample obtained by combining all the samples in the comparison. As mentioned in Sect. 3, we need to adjust the confidence bands to take into account the dependency of the ranks of the values of one sample on the values in other samples in the comparison.

When using the multiple sample test for MCMC convergence diagnostics, we recommend first using existing numerical convergence statistics, such as the  $\hat{R}$  by Vehtari



**Fig. 9** When compared to existing tests for uniformity, the rejection rate of the graphical test with simultaneous confidence bands performs well in all the three families of deviations introduced in (45). A slightly

lower rejection rate can be observed in family C for  $k \in (1, 3)$ , which corresponds to samples biased towards the centre of the unit interval

et al. (2021) or the  $R^*$  by Lambert and Vehtari (2021) which can assess the convergence of all model parameters jointly and can indicate which parameters have possible convergence issues. In the case that these statistics indicate possible issues, further insight into the nature of these deviations can be obtained with the ECDF plots of fractional ranks.

#### 4.2.1 Effect of difference in means and variances

We first compare two cases of MCMC sampling with four chains containing 250 independent draws, which is enough to reliably estimate the variances and autocorrelations required for  $\hat{R}$  and effective sample size (ESS) as long as the rank-normalized ESS of the sample exceeds 400 (Vehtari et al. 2021), which is the case as the draws are independent. In each case, chains 2 to 4 were sampled from a normal(0, 1) distribution. In the first case, chain 1 is sampled with a larger mean than the other chains, normal(0.5, 1). In the second case, chain 1 is sampled with a larger variance, normal(0, 1.5).

Rank plots for the first case with one chain having a larger mean are shown in Fig. 10a–d. Even though the difference in the sampling distribution of chain 1 can be seen in the histograms with 50 bins, this effect is more clearly represented in the ECDF difference plot in Fig. 10f where chain 1 shows the shape familiar from 4.1.1 and chains 2 to 4 show a reverse shape, indicating similar behaviour between these three chains.

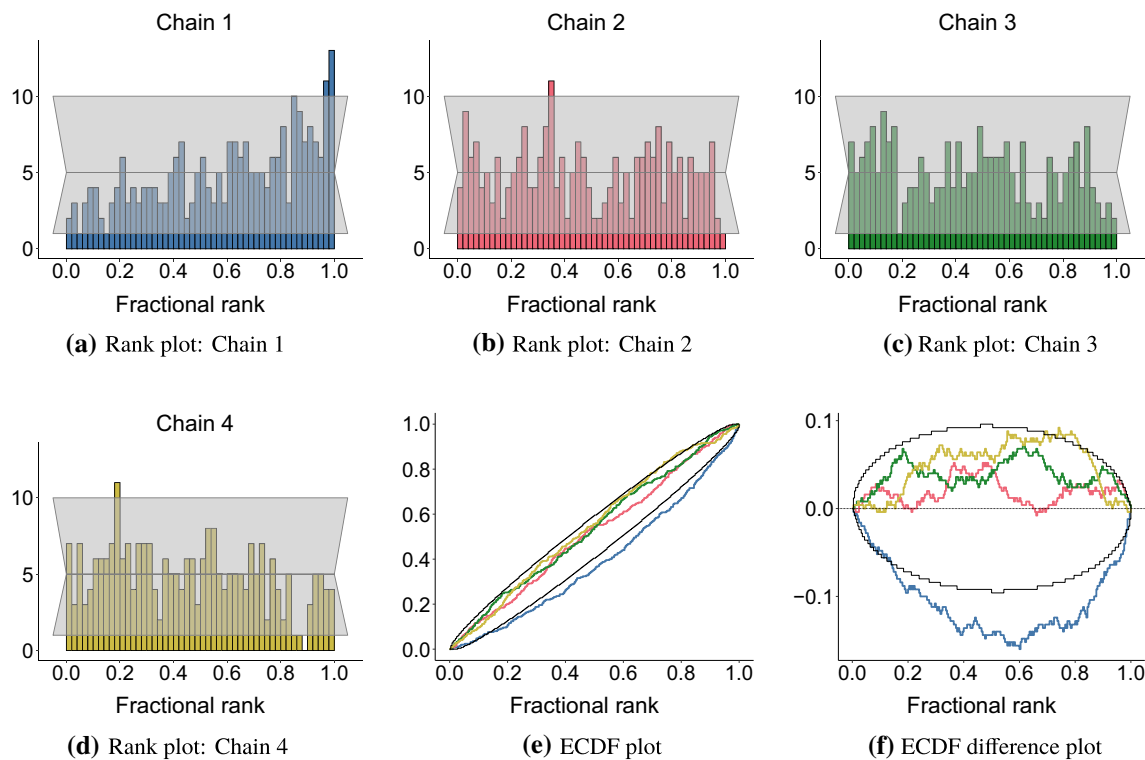
Similar remarks regarding the behaviour of the chains can be made from the ECDF plot in Fig. 10e, but the more dynamic range of the ECDF difference plot in Fig. 10f makes the difference in the behaviour of the chains clearer.

In the second case, where chain 1 is sampled with a higher variance, we can see a U-shape in the rank plot of chain 1 in Fig. 11a, but the behaviour stands out more clearly in the ECDF difference plot in Fig. 11f.

When compared to commonly used convergence diagnostics not offering graphical insight into the nature of the possible underlying problems, both the classical  $\hat{R}$  diagnostic by Gelman and Rubin (1992) and the improved  $\hat{R}$  diagnostic proposed by Vehtari et al. (2021) indicate convergence issues as they give an estimated  $\hat{R}$  values of 1.05 and 1.04, respectively, to both the mean and variance related examples above. Vehtari et al. (2021) suggest that  $\hat{R} > 1.01$  is an indication of potential convergence issues or too short chains.

#### 4.2.2 Test performance under common deviations

To evaluate the performance of the multiple sample comparison test under a set of common deviations, one of the samples was transformed according to the three transformation families defined in equation (45). In the analysis 2, 4, and 8 chains of length 100 were simulated from  $U(0, 1)$  after which one of the chains was transformed according to the transformations  $f_{A,k}$ ,  $f_{B,k}$ , and  $f_{C,k}$ . The rejection rates of the multiple sample comparison test when varying the power,  $k$ , of the transformation were estimated from 10,000 simulations and are recorded in Fig. 12. The observed test performance is independent of the number of chains used in the sample comparison. When compared to the rejection rates observed in the single sample power analysis in 4.1.4, the rejection rates show that the test sensitivity depends in a similar way on the transformation.



**Fig. 10** Effect of differences in sample mean. While chains 2 to 4 are drawn from the standard normal distribution, chain 1 is drawn from normal(0.5, 1), which can be seen as a bias towards large fractional ranks in the rank plot of chain 1 and as a slightly lowered frequency of

large fractional ranks in chains 2-4. In the ECDF plot and the ECDF difference plot, the ECDF of chain 1 obtains values considerably lower than expected resulting in a clear U-shape in the ECDF difference plot

### 4.2.3 Chains with autocorrelation

As samples generated by MCMC processes are typically autocorrelated, it is essential to analyse the performance of the sample comparison test under autocorrelated samples. In Fig. 13, rejection rates of simulated multiple sample test 2, 4, and 8 chains produced by autoregressive models of order 1 (i.e. AR(1) models) with varying AR-parameter values are presented. Each rejection rate is computed as the mean of 100,000 simulations. As seen in the figure, the higher the autocorrelation in the samples is and the more chains are sampled, the more likely the test is to reject the hypothesis that the samples are drawn from the same underlying distribution. Thus, before using the graphical illustration or the corresponding test, the chains should be thinned to have negligible autocorrelation. The same holds for other common uniformity tests as well, as they rely on the assumption of pairwise independence of draws.

### 4.2.4 Detecting model sampling issues: eight schools

We return to the eight schools model used to demonstrate SBC in Sect. 4.1.3. The issues detected with SBC earlier are

apparent when multiple sample comparison is used to inspect the rank distribution between the four individual chains, each containing 1000 posterior draws after a warm-up period of 1000 steps. Even when sampled with more conservative settings of the sampler, we see from Fig. 14 that the chains are not properly exploring the posterior and thus the realized rank transformed chains have clearly different ECDFs.

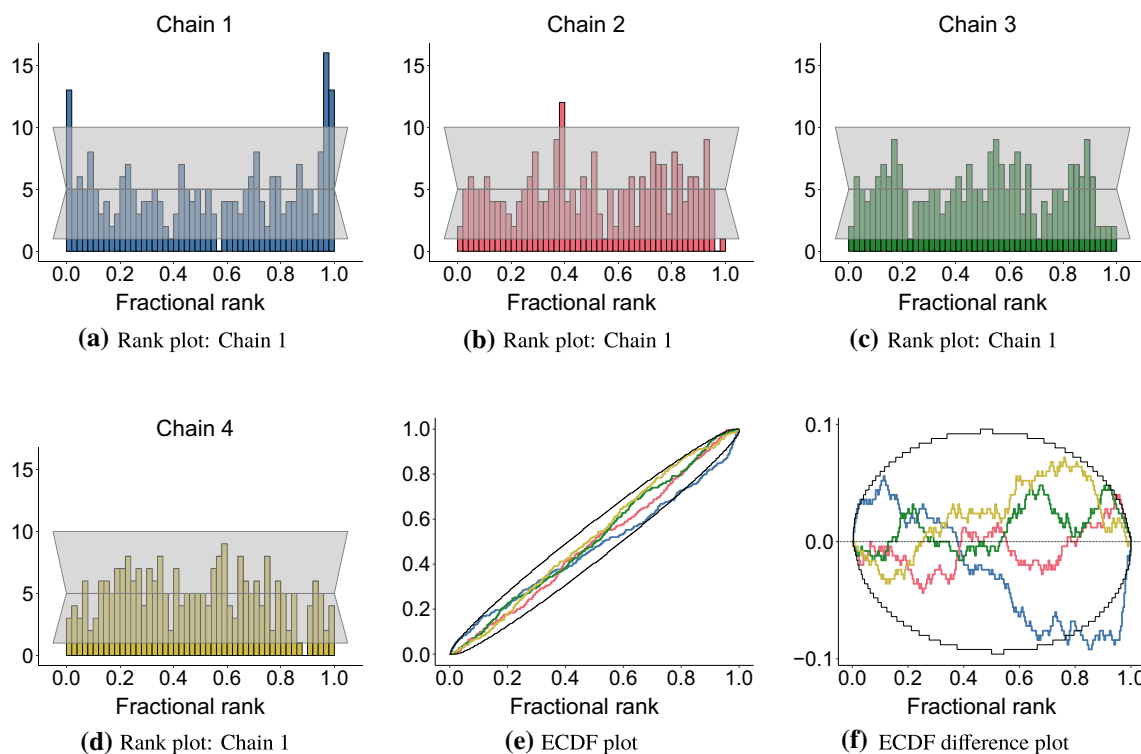
While the classical  $\hat{R}$  is estimated at 1, the improved  $\hat{R}$  diagnostic gives a value of 1.02 indicating possible convergence issues. One should also note that the sampling efficiency for  $\tau$  in the model is very low, as both the bulk-ESS and the tail-ESS by Vehtari et al. (2021) are under 150 for the combined sample.

As recommended in Sect. 22.7 of the Stan User’s Guide (Stan Development Team. xxx), these observed sampling issues of a hierarchical model with weak likelihood contribution can often be avoided by using the non-centred parameterization  $(\tilde{\theta}, \mu, \tau, \sigma)$  of the model:

$$\tilde{\theta}_j \sim \text{normal}(0, 1) \tag{51}$$

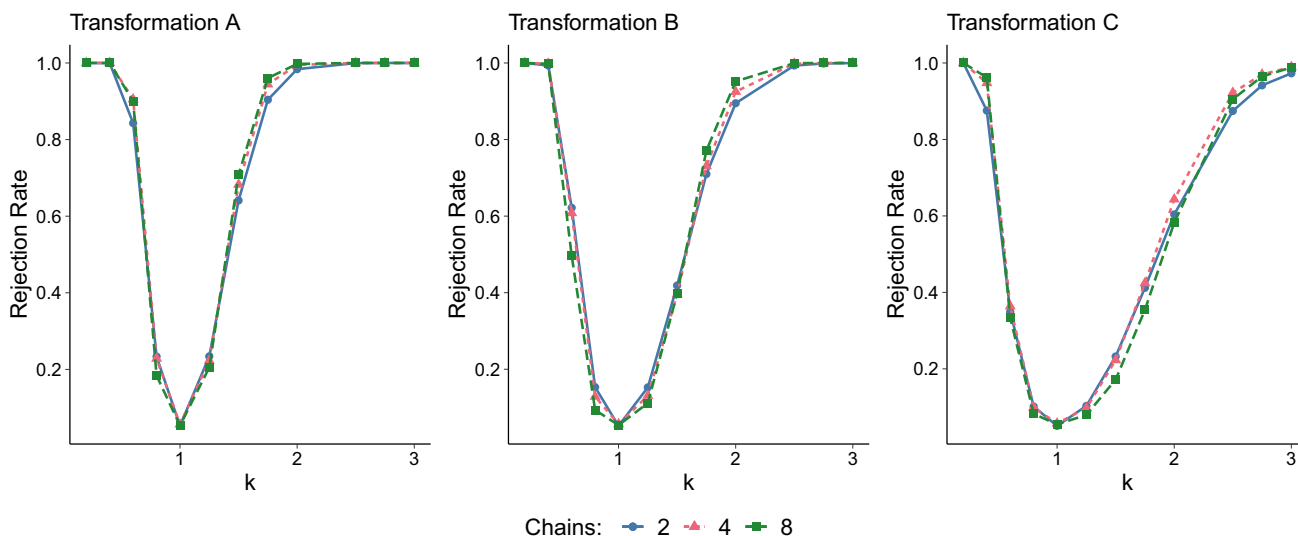
$$\theta_j = \mu + \tau \tilde{\theta}_j \tag{52}$$

$$y_j \sim \text{normal}(\theta_j, \sigma_j) \tag{53}$$



**Fig. 11** Effect of differences in sample variances. While chains 2–4 are drawn from the standard normal distribution, chain 1 is drawn from  $\text{normal}(0, 1.5)$ , which can be seen as a U-shape in the rank plot of chain 1 and as a low frequency of small and large fractional ranks in the rank plots of chains 2–4. In the ECDF plot and the ECDF difference plot, the


ECDF of chain 1 grows fast near the ends of the unit interval, where a higher-than-expected density of fractional ranks is covered. In the middle of the interval, the ECDF difference plot of chain 1 is decreasing, whereas chains 2–4 are increasing



**Fig. 12** An inspection of the rejection rates of the sample comparison test when one of the chains is transformed according to one of the three transformation families in (45), the test performance shows independen-

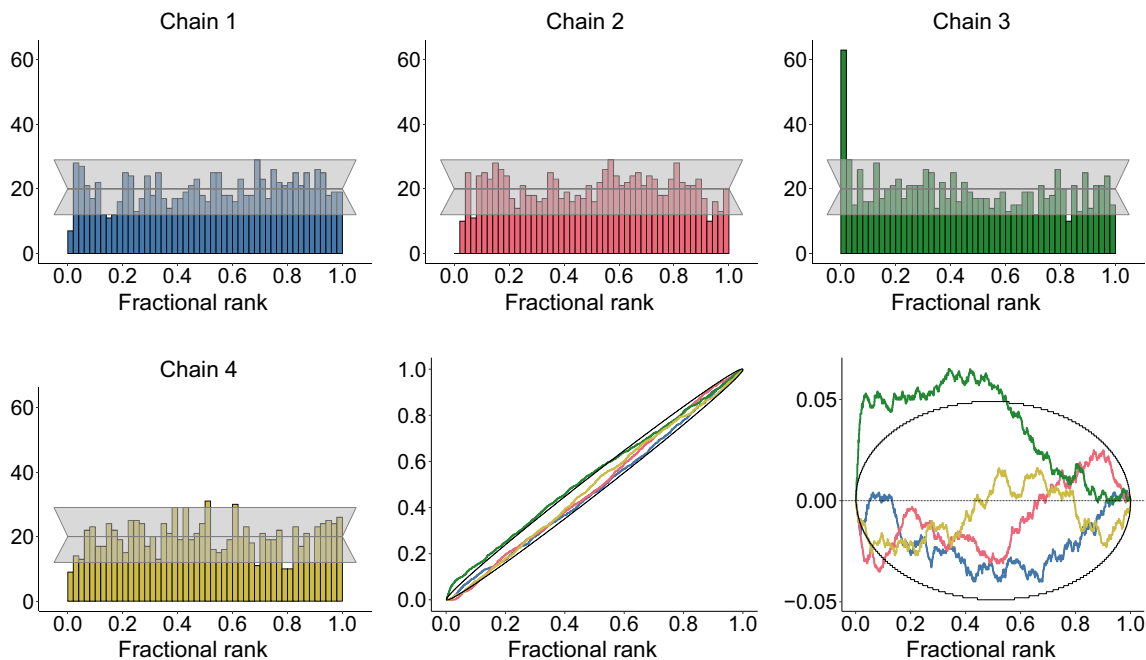
dence of the number of chains and demonstrates similar dependency on the extremity of the transformation





**Fig. 13** Test rejection rate when comparing chains with autocorrelation. On the left as a function of the AR-parameter value and on the right as a function of the ratio between the bulk effective sample size, as defined by Vehtari et al. (2021), and the total sample size. The nomi-

nal rejection rate 0.05 is shown with a vertical line in both plots. As the test expresses sensitivity to autocorrelation, we recommend thinning the samples, in order to reduce autocorrelation, before using the sample comparison test



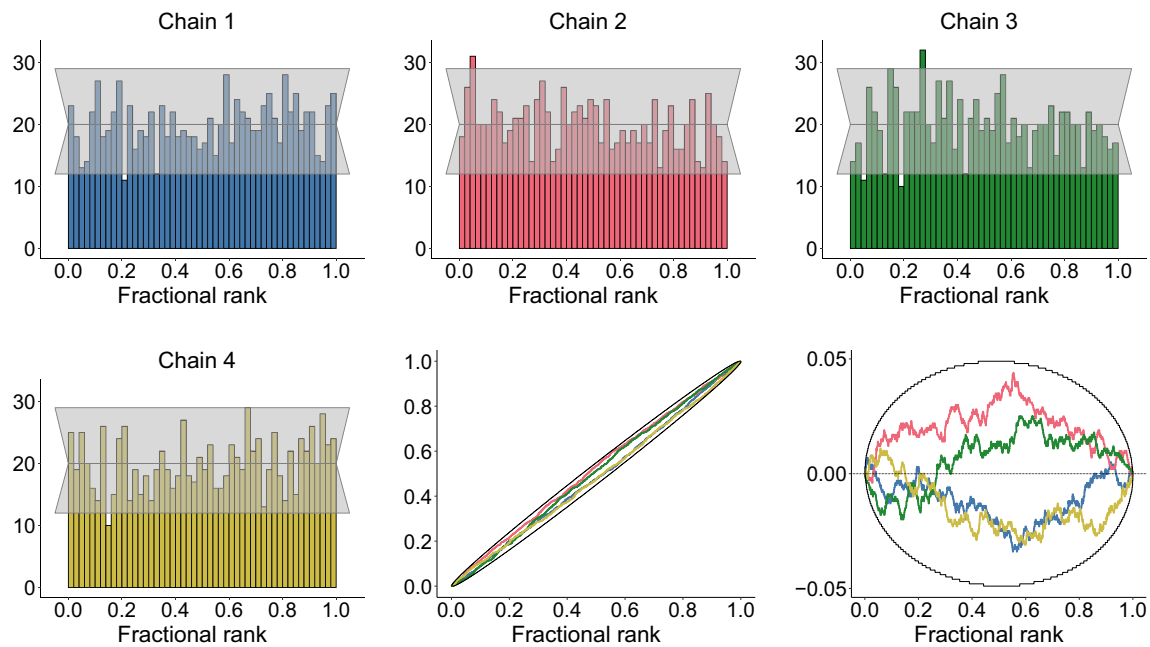
**Fig. 14** Detecting model sampling issues: centred parameter eight schools model. When inspecting the sampling of parameter  $\tau$ , both the rank plot of chain 3 and the ECDF difference plot in the centred

parameter model indicate a convergence issue with chain 3 including small values of the parameter at a rate considerably higher than the rest of the sampled chains

In the above parameterization, the treatment effect  $\theta_j$  is derived deterministically from the other parameter values and instead  $\tilde{\theta}_j$  is sampled. To keep the models comparable, we use the same conservative sampling options for the non-centred model although this is not required to obtain well mixing chains. In Fig. 15, we see an improvement in the sampling compared to the centred parameterization, as the sample ranks are distributed approximately uniformly

among the four chains implying that the chains are mixing well.

Now, both of the  $\hat{R}$  diagnostics agree on convergence with the graphical test, yielding values close to 1.00, while also the sampling efficiency issues detected in the centred parameterization model have disappeared giving samples with bulk-ESS and tail-ESS reaching 2200 and 1600, respectively.



**Fig. 15** Detecting model sampling issues: non-centred parameter eight schools model. When inspecting the sampling of parameter  $\tau$ , even when the 95% confidence bands of the rank plots of chains 2 and 3 are exceeded by one bin each, the ECDF plot and the ECDF difference

plot of the non-centred parameterization eight schools model indicate no mixing issues as the ECDF of the fractional ranks of each chain stay between the 95% simultaneous confidence bands

## 5 Discussion

By providing a graphical test for uniformity and comparison of samples, we offer an accessible tool to be used in many parts of practical statistical workflow.

For assessing the uniformity of a single sample, we recommend the optimization-based adjustment method, as it is efficient even for large sample sizes. For comparing multiple samples, the simulation-based method is likely to be computationally more efficient than the optimization-based method. To speed up the computations, we recommend pre-computing adjusted  $\gamma$  values for a set of sample size and number of samples (chains) and then interpolate (in log-log space) the adjustment as needed.

In the examples, we used empirical PIT with SBC, where the uniformity is expected by construction if the inference algorithm works correctly. PIT has also been used to compare predictive distributions. Specifically, in the LOO-PIT approach, PIT has been used to compare leave-one-out (LOO) cross-validation predictive distributions to the observations (e.g. Gneiting et al. 2007; Czado et al. 2009). Although the graphical LOO-PIT test is useful for visualization of model–data discrepancy, exact uniformity of LOO-PIT values can be expected only asymptotically given the true model. For example, if the data comes from a normal distribution and is modelled with a normal distribution with unknown mean and scale, the posterior predictive distribu-


tion is a Student's  $t$  distribution that approaches normal only asymptotically. Thus use of graphical LOO-PIT tests needs further research.

We have assumed that distributions  $g$  and  $p$  are continuous and only the fractional rank statistics  $u_i$  from Eq. (2) are discrete. Our proposed methods do not work directly if  $g$  and  $p$  are discrete, as values obtained through PIT are no longer uniform. Also, in the multiple sample comparison case, the rank statistics are no longer mutually distinct as ties are possible. The potential approach to handling discrete  $g$  and  $p$  is to use randomized or non-randomized modifications of PIT values for discrete distributions, as discussed by Czado et al. (2009). However, developing proven and efficient algorithms for this purpose requires further work, which is left for future research.

**Acknowledgements** We thank the Academy of Finland (grant 298742), the Finnish Center for Artificial Intelligence, and the Technology Industries of Finland Centennial Foundation (grant 70007503; Artificial Intelligence for Research and Development) for partial support of this research. We also acknowledge the computational resources provided by the Aalto Science-IT project.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence,

**Fig. 16** The difference between the expectation of the first 100 ordered statistics of a sample of size 1000. The expectations are computed from 1000 simulations. One can see that both the AR process with strong positive autocorrelation,  $\phi = 0.95$ , and the process with negative autocorrelation,  $\phi = -0.95$ , produce samples with expected ordered statistics that are biased towards the centre of the distribution. When thinned according to the tail-ESS, resulting in thinning by 18 and 7, respectively, the samples align well with the expected ordered statistics of the target distribution



unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

### Appendix A: Autocorrelated samples

In this appendix, we highlight the effect autocorrelated draws have when they are used to estimate the extreme rank statistics of the target distribution. Accounting for autocorrelation is important when inspecting the distribution of order statistics, including the PIT values in Sect. 2 or the between chain fractional ranks in Sect. 3.

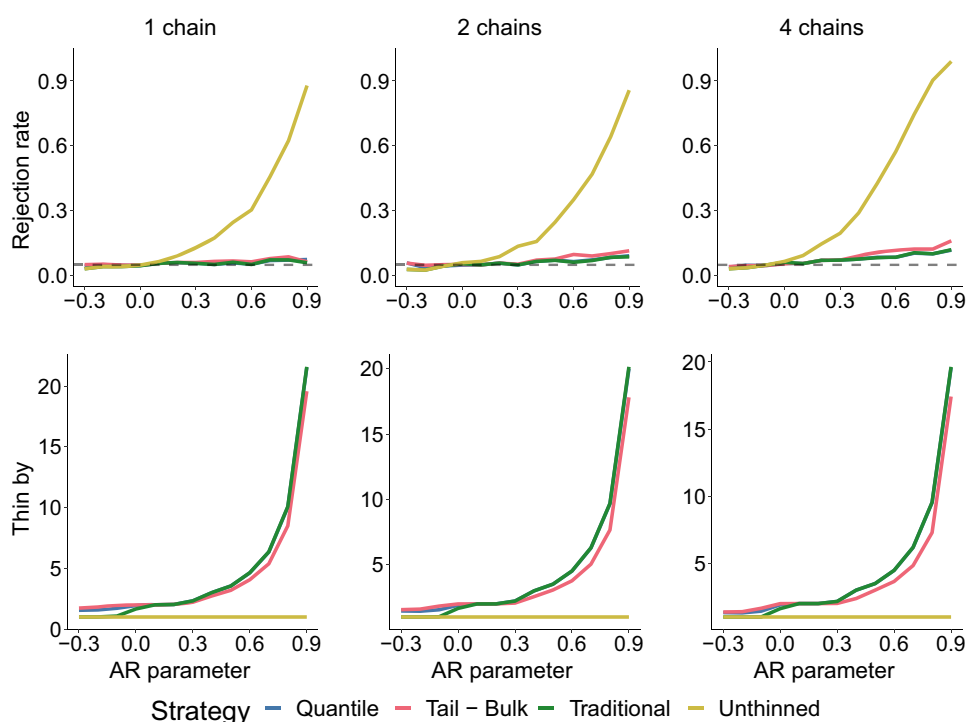
Given finite variance, central limit theorem holds also for correlated samples and many useful expectations can be estimated with desired accuracy by increasing the sample size. However, the bias in extreme ordered statistics can be non-negligible. This manifests in the expected value of the smallest and largest order statistics of an autocorrelated sample being less extreme than expected. This phenomenon is demonstrated with AR(1) processes in Fig. 16, which shows expected values of the 100 smallest ordered statistics computed from a sample of length 1000. The bias is smaller with less extreme ordered statistics, and, for example, estimates of  $p(x < -1.5)$  or 10% quantile in this case are likely to have

negligible bias. In the uniformity test, extreme PIT estimates can have non-negligible bias, increasing the probability that ECDF steps out of the simultaneous confidence band.

A standard approach to reduce sample autocorrelation is to thin the sampled chains by keeping only every  $T$  values in the sample. Below, we compare three thinning strategies. First, the traditional approach, where  $T = S/ESS$ , where ESS is computed for the posterior mean (without rank normalization) (Vehtari et al. 2021). Second, an approach recommended by Talts et al. (2020), where the above ESS is computed for estimating the ECDF,  $P(y < y^*)$  where  $y^*$  are empirical quantiles of the sample  $y$ . The authors recommend using 19 quantiles (0.05, 0.1, . . . , 0.95) and thinning the sample based on the ESS, which would result in the largest thinning factor. This method is targeted to address differences in sampling efficiency between the distribution quantiles. The third method, we introduce, is calculating the tail-ESS and bulk-ESS as defined by Vehtari et al. (2021), and picking the one resulting into the stricter thinning. This method aims to address possible differences in sampling efficiency between the central 90% quantile and the two 5% tail quantiles. R package ‘posterior’ was used for all ESS computations.

In Fig. 16, we additionally show the first 100 order statistics of the standard normal distribution compared to AR processes thinned according to the tail-ESS, as our focus here is on the tails of the distribution. In order to arrive at a thinned sample of equal length, an expected Tail-ESS was

**Fig. 17** The behaviour of the three thinning strategies when applied to a sample consisting of 1, 2, or 4 chains created with an AR(1) process. With AR(1) process, the three strategies agree on the recommended strategy and are fairly successful at recovering the desired confidence level, while the rejection rate of the unthinned chains grows as more autocorrelation is introduced. For large values of the AR parameter,  $\phi$ , the thinning amounts are quite high, which would require larger samples to reliable estimate properties of the target distribution



obtained by averaging over 10,000 simulations and the sample length was chosen accordingly to yield thinned samples of length 1000. After the thinning, the order statistics closely match those drawn independently from the standard normal distribution.

We inspected how the three above-mentioned thinning strategies manage to reduce the autocorrelation which, as shown in Sect. 4.2.3, the ECDF-based test is sensitive to. In this experiment, 1, 2, and 4 chains of length 1000 were drawn from the AR(1) process with varying values of the AR parameter  $\phi$ . The results of this experiment are displayed in Fig. 17, and as one can see, all three of the methods produce very similar thinning recommendations, and thus also test results, managing to reduce the rejection rate near the desired 5%.

If after using some default thinning approach, there are still many extreme PIT estimates, it is possible that there is still substantial autocorrelation in the sample and more careful investigation of the remaining autocorrelation is warranted. There is certainly a trade-off between the computation time and how accurately the behaviour of extreme tails need to be examined. Often the major issues can be seen with less accurate computation, and natural workflow can include iterative refinement of the diagnostic accuracy.

Although thinning may be needed for uniformity test as part of SBC or posterior predictive checks (PPC), when estimating quantities of interest that are not related to extreme tails, better efficiency is obtained by using all the posterior draws.

## References

- Aldor-Noiman, S., Brown, L.D., Buja, A., Rolke, W., Stine, R.A.: The power to see: a new graphical test of normality. *Am. Stat.* **67**(4), 249–260 (2013). <https://doi.org/10.1080/00031305.2013.847865>
- Arnold, B.C., Balakrishnan, N., Nagaraja, H.N.: *A First Course in Order Statistics (Classics in Applied Mathematics)*. Society for Industrial and Applied Mathematics, USA (2008)
- Brent, R.: An algorithm with guaranteed convergence for finding the minimum of a function of one variable. In: *Algorithms for Minimization Without Derivatives*, pp. 61–80. Prentice-Hall, Englewood Cliffs, NJ (1973)
- Brown, L.D., Cai, T.T., DasGupta, A.: Interval estimation for a binomial proportion. *Stat. Sci.* **16**(2), 101–133 (2001). <https://doi.org/10.1214/ss/1009213286>
- Cook, S.R., Gelman, A., Rubin, D.B.: Validation of software for Bayesian models using posterior quantiles. *J. Comput. Graph. Stat.* **15**(3), 675–692 (2006)
- Czado, C., Gneiting, T., Held, L.: Predictive model assessment for count data. *Biometrics* **65**(4), 1254–1261 (2009). <https://doi.org/10.1111/j.1541-0420.2009.01191.x>
- D’Agostino, R.B., Stephens, M.A.: *Goodness-of-Fit Techniques*. Marcel Dekker Inc, USA (1986)
- Gelman, A.: Correction to Cook, Gelman, and Rubin (2006). *J. Comput. Graph. Stat.* **26**(4), 940–940 (2017). <https://doi.org/10.1080/10618600.2017.1377082>
- Gelman, A., Carlin, J.B., Stern, H.S., Dunson, D.B.: Aki VEHTARI a Donald B. RUBIN, 2013. *Bayesian Data Anal.* (2013). <https://doi.org/10.1201/b16018>
- Gelman, A., Rubin, D.B.: Inference from iterative simulation using multiple sequences. *Stat. Sci.* **7**(4), 457–472 (1992)
- Gelman, A., Vehtari, A., Simpson, D., Margossian, C.C., Carpenter, B., Yao, Y., Kennedy, L., Gabry, J., Bürkner, P.C., Modrák, M. (2020). Bayesian workflow. arXiv preprint [arXiv:2011.01808](https://arxiv.org/abs/2011.01808)
- Gneiting, T., Balabdaoui, F., Raftery, A.E.: Probabilistic forecasts, calibration and sharpness. *J. Royal Stat. Soc. Series B (Stat.*

- Methodol.) **69**(2), 243–268 (2007). <https://doi.org/10.1111/j.1467-9868.2007.00587.x>
- Kolmogorov, A.: Sulla determinazione empirica di una legge di distribuzione. *Inst. Ital. Attuari, Giorn.* **4**, 83–91 (1933)
- Lambert, B., Vehtari, A.:  $R^*$ : A robust MCMC convergence diagnostic with uncertainty using decision tree classifiers. *Bayesian Anal.* <https://doi.org/10.1214/20-BA1252>
- Ledwina, T.: Data-driven version of Neyman's smooth test of fit. *J. Am. Stat. Assoc.* **89**(427), 1000–1005 (1994). <https://doi.org/10.1080/01621459.1994.10476834>
- Marhuenda, Y., Morales, D., Pardo, M.C.: A comparison of uniformity tests. *Statistics* **39**(4), 315–327 (2005). <https://doi.org/10.1080/02331880500178562>
- Massey, F.J., Jr.: The Kolmogorov-Smirnov test for goodness of fit. *J. Am. Stat. Assoc.* **46**(253), 68–78 (1951)
- Neyman, J.: Smooth test for goodness of fit. *Scand. Actuar. J.* **1937**(3–4), 149–199 (1937). <https://doi.org/10.1080/03461238.1937.10404821>
- Stan Development Team. (n.d.). Stan user's guide. Retrieved September 29, 2020, from [https://mc-stan.org/docs/2\\_24/stan-users-guide/index.html](https://mc-stan.org/docs/2_24/stan-users-guide/index.html)
- Talts, S., Betancourt, M., Simpson, D., Vehtari, A., & Gelman, A. (2020). Validating Bayesian inference algorithms with simulation-based calibration. arXiv preprint [arXiv:1804.06788v2](https://arxiv.org/abs/1804.06788v2)
- Vehtari, A., Gelman, A., Simpson, D., Carpenter, B., Bürkner, P.C.: Rank-normalization, folding, and localization: an improved  $\hat{R}$  for assessing convergence of MCMC (with discussion). *Bayesian Anal.* **16**(2), 667–718 (2021). <https://doi.org/10.1214/20-BA1221>
- Watson, G.S.: Goodness-of-fit tests on a circle. *Biometrika* **48**(1–2), 109–114 (1961). <https://doi.org/10.1093/biomet/48.1-2.109>

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.