
NON-PARAMETRIC HYPOTHESIS TESTS FOR DISTRIBUTIONAL GROUP SYMMETRY

Kenny Chiu

Department of Statistics
The University of British Columbia
kenny.chiu@stat.ubc.ca

Benjamin Bloem-Reddy

Department of Statistics
The University of British Columbia
benbr@stat.ubc.ca

ABSTRACT

Symmetry plays a central role in the sciences, machine learning, and statistics. For situations in which data are known to obey a symmetry, a multitude of methods that exploit symmetry have been developed. Statistical tests for the presence or absence of general group symmetry, however, are largely non-existent. This work formulates non-parametric hypothesis tests, based on a single independent and identically distributed sample, for distributional symmetry under a specified group. We provide a general formulation of tests for symmetry that apply to two broad settings. The first setting tests for the invariance of a marginal or joint distribution under the action of a compact group. Here, an asymptotically unbiased test only requires a computable metric on the space of probability distributions and the ability to sample uniformly random group elements. Building on this, we propose an easy-to-implement conditional Monte Carlo test and prove that it achieves exact p -values with finitely many observations and Monte Carlo samples. The second setting tests for the invariance or equivariance of a conditional distribution under the action of a locally compact group. We show that the test for conditional invariance or equivariance can be formulated as particular tests of conditional independence. We implement these tests from both settings using kernel methods and study them empirically on synthetic data. Finally, we apply them to testing for symmetry in geomagnetic satellite data and in two problems from high-energy particle physics.

1 Introduction

Symmetry has played an important role in statistical problems, from the classical literature on equivariant estimation (Lehmann and Casella, 1998, Ch. 3) and invariant testing (Lehmann and Romano, 2005, Ch. 6), to modern work on the use of transformation groups in statistical estimation (Chen et al., 2020; Huang et al., 2022) and machine learning problems (Cohen et al., 2019). One of the key ideas that emerges from this line of work is that by using models that account for symmetries present in data, one obtains statistical benefits through various forms of optimality (Lehmann and Casella, 1998; Lehmann and Romano, 2005; Eaton and Sudderth, 1999), improved sample efficiency (Chen et al., 2020; Huang et al., 2022), and better out-of-sample generalization (Elesedy and Zaidi, 2021; Elesedy, 2021; Lyle et al., 2020). Such approaches also have a certain appeal that bridges different philosophical positions: decision problems with symmetry are among the only known problems in which frequentist and Bayesian inference coincide, and also present working realizations of other approaches to inference, variously known as fiducial (Fraser, 1961; Hora and Buehler, 1966, 1967), pivotal (Eaton and Sudderth, 1999), or structural (Fraser, 1966, 1968). A pervasive characteristic shared by all of that work is that a specific symmetry group is known or assumed, and the problem is carefully constructed with respect to that group. That work, however, does not address the problem of identifying symmetry from data. Moreover, a symmetry assumption can be difficult to check, and if the assumption is wrong, then the performance of a symmetric model can be much worse than a non-symmetric one.

Separately, symmetry plays a central role in modern science, particularly in the physical sciences where entire theories are constructed around the symmetries that must be obeyed by equations describing the behaviour of physical systems (Gross, 1996). Additionally, detection of new or broken symmetries is playing a role in the search for physics beyond the Standard Model (ATLAS Collaboration, 2017), particularly in data-driven approaches (Karagiorgi et al., 2022; Birman et al., 2022). Recent work in machine learning and physics aims to learn or estimate symmetry groups from data (Krippendorf and Syvaeri, 2020; Zhou et al., 2021; Desai et al., 2021; Dehmamy et al., 2021; Yang et al., 2023) or to detect anomalous symmetry-breaking (Collins et al., 2018; Birman et al., 2022). However, key inferential tools based on hypothesis tests for symmetry are missing. Such tools are crucial if the discovery of symmetry from data is to be a reliable part of the scientific process. For example, they can be used to test for the presence or absence of a particular symmetry in data, with that symmetry specified by hypothesis or by a data-driven method that has learned or estimated a symmetry. In situations with known or assumed symmetry, hypothesis tests for symmetry could also be used as model-checking criteria for models meant to exhibit that symmetry.

The present work formulates non-parametric tests, based on a single independent and identically distributed (i.i.d.) sample, for distributional symmetry under a specified group. We provide abstract formulations of tests that apply to two broad settings. The first setting tests for the invariance of a marginal or joint distribution under the action of a compact group. The test is formulated in such a way that if one has an asymptotically consistent estimator of a metric on the space of probability measures and the ability to sample uniformly random group elements, then it is straightforward to devise an asymptotically unbiased test for invariance. More importantly, we design an easy-to-implement conditional Monte Carlo test that achieves exact p -values with finitely many observations and Monte Carlo samples. The test attains those properties by conditioning on a sufficient statistic induced by the group. The second setting tests for the invariance or equivariance of a conditional distribution under the action of a locally compact group, provided that the group action obeys weak regularity conditions. We show that a test for conditional equivariance can be formulated as a particular test of conditional independence, which inherits the statistical properties of the conditional independence test chosen for implementation. Although conditional independence testing is known to be hard (Shah and Peters, 2020), especially as the dimension of the problem increases, the structure induced by symmetry means that the conditioning variables used in the test are often of much lower dimension than the observations.

In addition to the generic testing methods and the study of their theoretical properties, we provide specific instantiations of the tests using kernel-based methods for non-parametric hypothesis testing. We study these tests empirically on synthetic data, and apply them to geomagnetic satellite data and to two problems in high-energy particle physics. Computer code required to run the experiments can be found on the [GitHub repository](#)¹ for this work.

1.1 Overview

The remainder of this section describes our work and main results informally, refraining for now from addressing technicalities. The mathematical object that encodes symmetry is a group \mathbf{G} . Relevant technical details of groups are given in Section 2. Elements $g \in \mathbf{G}$ act via transformations $x \mapsto gx$ of elements from a sample space $x \in \mathbf{X}$. This action on \mathbf{X} extends to the set $\mathcal{P}(\mathbf{X})$ of probability measures on \mathbf{X} . If P is the distribution of a random element $X \in \mathbf{X}$, then g acts on P via the pushforward, $g_*P(A) := P(g^{-1}A)$, with $A \subseteq \mathbf{X}$ and $g^{-1}A := \{g^{-1}x : x \in A\}$. A key question in many settings is whether the distribution P underlying a set of i.i.d. observations $X_{1:n} := (X_1, \dots, X_n)$ is *invariant* under \mathbf{G} in the sense that

$$g_*P = P, \quad \text{for each } g \in \mathbf{G} .$$

Outside of ill-behaved situations that typically do not arise in practice, this is only possible for a probability measure when \mathbf{G} is compact. Any compact group \mathbf{G} has a unique invariant probability measure that can be thought of as the uniform distribution on \mathbf{G} . We denote this measure by λ .

For a specified group \mathbf{G} , the statistical problem we address is to test the hypotheses

$$H_0 : P \text{ is } \mathbf{G}\text{-invariant} \quad \text{versus} \quad H_1 : P \text{ is not } \mathbf{G}\text{-invariant} .$$

If \mathbf{G} is relatively small and finite, or generated by a small set of elements (say of size m), invariance might be tested with a composite of m two-sample hypothesis tests. For large discrete groups, this approach quickly becomes untenable;

¹<https://github.com/chiukenny/Tests-for-Distributional-Symmetry>

for uncountable groups, it is not possible. Instead, one might formulate hypothesis tests based on other characterizations of distributional invariance. In Proposition 2, we collect a number of known identities that uniquely characterize the invariance of a probability measure and on which hypothesis tests may be based. Perhaps the most well-known of the identities is that $P = P^\circ$ if and only if P is \mathbf{G} -invariant, where P° is obtained by averaging g_*P over \mathbf{G} with respect to the invariant probability measure λ ,

$$P^\circ(A) := \int_{\mathbf{G}} P(g^{-1}A) \lambda(dg). \quad (1)$$

Because both P and P° are probability measures on \mathbf{X} , any metric D on $\mathcal{P}(\mathbf{X})$ can be used in conjunction with the empirical measure and a Monte Carlo estimate of the integral in (1) to define a test statistic of the form

$$T_{n,m}(X_{1:n}) := D\left(\frac{1}{n} \sum_{i=1}^n \delta_{X_i}(\bullet), \frac{1}{nm} \sum_{i=1}^n \sum_{j=1}^m \delta_{G_j X_i}(\bullet)\right), \quad G_j \sim_{\text{iid}} \lambda.$$

This approach is very general and can be used for abstract spaces \mathbf{X} other than \mathbb{R}^d as long as one has access to a metric on $\mathcal{P}(\mathbf{X})$ and the ability to sample random elements of \mathbf{G} acting on \mathbf{X} . An asymptotically consistent estimator of D then yields an asymptotically unbiased test for invariance as $n, m \rightarrow \infty$ (Theorem 1).

Beyond the general-purpose averaging approach, more detailed structure induced by \mathbf{G} is often available. The group action partitions \mathbf{X} into equivalence classes called *orbits* so that x and x' are equivalent if and only if $x = gx'$ for some $g \in \mathbf{G}$. One can choose a representative element $[x]$ of each orbit to define an *orbit selector* $\gamma(x) = [x]$. The orbit selector allows one to decompose probability measures on \mathbf{X} , so that a random variable X has an invariant distribution if and only if it satisfies $X \stackrel{d}{=} G\gamma(X)$, where $G \perp\!\!\!\perp X$ is sampled uniformly from \mathbf{G} . As in the averaging test, one can then conduct a non-parametric test for invariance by comparing the untransformed empirical measure with the empirical measure of the observations $(G_1\gamma(X_1), \dots, G_n\gamma(X_n))$. Such a test is asymptotically unbiased when based on a consistent estimator of a metric on $\mathcal{P}(\mathbf{X})$.

More importantly, $\gamma(X)$ is a special case of a *maximal invariant* statistic, which is an invariant function that takes a different value on each orbit and thus uniquely encodes the orbits. It is known that any maximal invariant is a sufficient statistic for $\mathcal{P}^\circ(\mathbf{X})$, the class of \mathbf{G} -invariant probability distributions. Sufficiency in particular means that for each $P \in \mathcal{P}^\circ(\mathbf{X})$, a sample $X_{1:n} \sim_{\text{iid}} P$ has the same conditional distribution given $\gamma(X)_{1:n}$, and we can generate samples from that conditional distribution as $(G_1^{(b)}\gamma(X_1), \dots, G_n^{(b)}\gamma(X_n))$, with $G_i^{(b)}$ independent of $X_{1:n}$ and sampled i.i.d. from λ . The ability to sample from this shared conditional distribution enables us to design a conditional Monte Carlo test that yields exact p -values with finitely many observations and Monte Carlo samples. Section 3.2 details the test (Algorithm 1) and its statistical properties (Theorem 2). In Section 3.2.1, we describe a method for estimating the conditional power function at the empirical measure of the observed data, \hat{P}_n , which can be combined with standard bootstrap resampling to estimate the power function at P .

If \mathbf{G} acts freely on \mathbf{X} in the sense that $gx = x$ implies g is the identity element of \mathbf{G} , then the orbit selector γ can be “inverted” to obtain the element of \mathbf{G} that sends $[x]$ to x . We call such a function, denoted $\tau: \mathbf{X} \rightarrow \mathbf{G}$, a *representative inversion* because it satisfies $\tau(x)\gamma(x) = \tau(x)[x] = x$. Yet another characterization of \mathbf{G} -invariance is that for $X \sim P$, $\tau(X) \stackrel{d}{=} G$, with $G \perp\!\!\!\perp X$ sampled uniformly from \mathbf{G} . In this case, a metric on the space of probability measures on \mathbf{G} , $\mathcal{P}(\mathbf{G})$, can be used as a test statistic. Theorem 2 is easily adapted to this situation because λ is the unique invariant probability measure on \mathbf{G} , and no appeal to sufficiency is required. If the action of \mathbf{G} is not free, so that $gx = x$ for g in some non-trivial subset $\mathbf{G}_x \subseteq \mathbf{G}$, then τ can be replaced by an appropriate random variable $\tilde{\tau}$ sampled from an *inversion kernel*, $\zeta(x, \bullet)$. The inversion kernel has a number of remarkable properties; the relevant one here is that if $\tilde{\tau} \sim \zeta(x, \bullet)$, then $\tilde{\tau}\gamma(x) = x$ with probability one. From this, it follows that a characterization of \mathbf{G} -invariance is that $\tilde{\tau} \stackrel{d}{=} G$, with $G \sim \lambda$.

In each of the above cases, a non-parametric test for distributional invariance can be constructed by sampling random group transformations and applying them to a sample of data, then using an estimator for a metric on $\mathcal{P}(\mathbf{X})$ (or $\mathcal{P}(\mathbf{G})$, as appropriate) to compare the untransformed sample with the randomly transformed sample. This recipe is generic and can in principle be used with any metric on $\mathcal{P}(\mathbf{X})$ (resp. $\mathcal{P}(\mathbf{G})$). In Section 4, we formulate specific versions of the tests using the kernel maximum mean discrepancy. In Section 7, we present the results of an empirical study of these

tests on synthetic data to validate the theoretical properties obtained in Theorem 2, and apply the tests to two different applications: geomagnetic satellite data and simulated dijet events from the Large Hadron Collider.

1.1.1 Conditional symmetry

In some problems, especially those involving regression, classification, or prediction of a variable $Y \in \mathbf{Y}$ from X , primary interest is in symmetry of the conditional distribution $P_{Y|X}$. The conditional distribution is said to be *equivariant* if for each measurable subset $B \subseteq \mathbf{Y}$,

$$P_{Y|X}(gx, B) = P_{Y|X}(x, g^{-1}B), \quad x \in \mathbf{X}, g \in \mathbf{G}.$$

It is said to be *invariant* if the action of \mathbf{G} on \mathbf{Y} is trivial, so that the above equation holds with $g^{-1}B$ replaced by B . Equivariant conditional distributions arise from the disintegration of jointly invariant probability distributions $P_{X,Y} = P_X \otimes P_{Y|X}$. If \mathbf{G} is compact and the marginal distribution P_X is known to be invariant, then testing for conditional equivariance of $P_{Y|X}$ is equivalent to testing for the joint invariance of $P_{X,Y}$, which could be carried out using the methods described above. However, the marginal distribution of X may not be invariant—in many cases it is known not to be—and the problem cannot be reduced to a test for joint invariance. For example, if \mathbf{G} is non-compact, then P_X cannot be \mathbf{G} -invariant, but $P_{Y|X}$ may be. Instead, in Theorem 3, we obtain a conditional independence property that characterizes equivariance. In particular, we show that $P_{Y|X}$ is equivariant if and only if

$$(\tilde{\tau}, X) \perp\!\!\!\perp \tilde{\tau}^{-1}Y \mid \gamma(X), \quad \text{with } \tilde{\tau} \mid X \sim \zeta(X, \bullet).$$

Here, $\tilde{\tau}^{-1}$ denotes the group inverse of the element $\tilde{\tau} \in \mathbf{G}$. Moreover, $\gamma(X)$ can be replaced by any maximal invariant statistic $M(X)$. This generalizes a related result of [Bloem-Reddy and Teh \(2020\)](#). A consequence of the result is that a test for conditional symmetry (equivariance or invariance) can be formulated as a test for conditional independence. In Section 6, we describe an instantiation of this test using a kernel-based conditional independence test ([Zhang et al., 2011](#)). We apply it to synthetic data in Section 7, as well as to data from three settings in the physical sciences.

1.2 Related work

There is an extensive literature on invariant testing problems; a textbook treatment can be found in [Lehmann and Romano \(2005\)](#). The main idea in invariant testing is as follows. Let Ω represent the set of probability distributions under consideration, with Ω_0 representing the subset that satisfy the null hypothesis and Ω_1 those that do not, so that $\Omega = \Omega_0 \cup \Omega_1$ and $\Omega_0 \cap \Omega_1 = \emptyset$. The testing problem is said to be \mathbf{G} -invariant if each of Ω_0 and Ω_1 are \mathbf{G} -invariant sets. That is, if for every $P \in \Omega_0$, $g_*P \in \Omega_0$ for each $g \in \mathbf{G}$, and likewise for every $P \in \Omega_1$. There is extensive theory for such tests, much of which focuses on the optimality of tests based on maximal invariant statistics. A related line of work studies randomization tests in which random group elements are used for randomization; randomized permutation and sign-flip tests are two of the most common. The literature on such tests for specific finite groups is substantial. Recent work by [Dobriban \(2022\)](#) obtains consistency results for randomization tests using compact groups, under an additive noise decomposition assumption, and includes a review of earlier invariance-based randomization literature.

The testing framework proposed here in Section 3 trivially fits in the invariant testing framework, and the conditional Monte Carlo test proposed in Section 3.2 is closely related to invariance-based randomization techniques, but little follows directly from those connections. Whereas invariant testing and invariance-based randomization deal with accounting for or leveraging symmetry so that the problem of testing *under* symmetry is simplified, the current work addresses the problem of testing *for* symmetry; the two are conceptually somewhat orthogonal, though the mathematical techniques bear some similarities. For example, a special case of our result on the size of the conditional Monte Carlo test (Theorem 2) was obtained for finite groups by [Hemerik and Goeman \(2018\)](#), and our Theorem 2 is perhaps of independent interest for invariance-based randomization tests.

In light of the potential importance of a formal hypothesis testing framework for the presence or absence of distributional symmetry, it is somewhat surprising that with the exception of recent work ([Fraiman et al., 2021](#); [Christie and Aston, 2022, 2023](#)), such a framework and corresponding methods are largely absent from the literature. As we describe below, those recent methods make strong assumptions that limit their applicability. Our methods are broadly applicable; our experiments compare to the existing methods where possible and demonstrate some settings where those methods cannot be used.

The recent work of [Fraiman et al. \(2021\)](#) applied the Cramér–Wold (CW) theorem to formulate non-parametric tests for group invariance. Those tests rely on the group \mathbf{G} being generated by a (small) finite set \mathbf{G}_0 of transformations such that each element of \mathbf{G} can be written as a finite product $g = g_1 \cdots g_m$, where for each j , either $g_j \in \mathbf{G}_0$ or $g_j^{-1} \in \mathbf{G}_0$. When this assumption holds, it can lead to a reduction in the computational complexity of the test. On the other hand, the assumption can only be satisfied by discrete groups, as no uncountable group can be finitely generated. Conversely, both the abstract formulation of our tests and our kernel-based implementations can be applied to any compact group, which includes finite discrete groups. We compare both approaches empirically in Section 7 and find that although the computational complexity of the CW-based tests is favourable, they tend to be less powerful than the kernel-based tests we implemented.

To our knowledge, the test we propose in Sections 5 and 6 constitutes the first test (parametric or non-parametric) for symmetry of a conditional distribution. In recent work, [Christie and Aston \(2023\)](#) proposed two tests for \mathbf{G} -invariance of the conditional expectation $f(x) = \mathbb{E}[Y|X = x]$, $f: \mathbf{X} \rightarrow \mathbb{R}$. Both of those tests require the user to assume that f belongs to some specific class of functions, \mathcal{F} , of bounded variation, and the assumption of an additive noise model, $Y_i = f(X_i) + \varepsilon_i$, for independent mean-zero noise ε_i . One test requires knowledge of the bound $V(x, x') = \sup_{f \in \mathcal{F}} |f(x) - f(x')|$ and a bound on the deviations on the noise variable, $\Pr(|\varepsilon_i - \varepsilon_j| \geq c) \leq p_c$. The other test is less restrictive, instead requiring knowledge of some $\mathcal{V}(x, x')$ satisfying $|f(x) - f(x')| \leq C_f \mathcal{V}(x, x')$. In our experiments in Section 7, these assumptions are too restrictive for the tests to be applicable. We note that the primary aim of [Christie and Aston \(2023\)](#) is to estimate the *maximal* group under which f is invariant, which amounts to conducting a collection of tests over a subgroup lattice of some candidate maximal group. In principle, our tests could be substituted into their procedure, though we do not address that problem here.

Apart from hypothesis testing, researchers in physics and machine learning have developed methods for estimating symmetries from data; see the references in Section 1. Hypothesis tests for symmetry, either as part of the estimation procedure or as validation of the estimated symmetry, have not been developed. To the best of our knowledge, the only exception is ([Birman et al., 2022](#)), which develops a test for anomaly detection, but requires restrictive distributional assumptions and approximations.

2 Background: Groups, actions, and invariant measures

Throughout, \mathbf{X} denotes a topological space and $\mathbf{S}_{\mathbf{X}}$ its Borel σ -algebra, so that $(\mathbf{X}, \mathbf{S}_{\mathbf{X}})$ is a standard Borel measurable space. When there is no chance of confusion, we will refer to \mathbf{X} as a measurable space, and likewise for \mathbf{Y} and \mathbf{M} . Let P be a probability measure defined on \mathbf{X} . For a random variable X taking values in \mathbf{X} , we write $\mathbb{E}_P[X]$ for the expectation of X with respect to P , and $X \sim P$ to denote a random variable sampled from P . For any measurable function f on \mathbf{X} , let f_*P denote the pushforward, or image measure, with $f_*P(A) = P(f^{-1}(A))$ for all measurable sets $A \in \mathbf{S}_{\mathbf{X}}$. We use δ_x to denote the Dirac measure at a point x .

2.1 Groups

Groups are central to the methods developed here, so we review some basic group theory. A group \mathbf{G} is a set with a binary operation \cdot that satisfies the associativity, identity, and inverse axioms. We denote the identity element by id . For notational convenience, we write $g_1 g_2 = g_1 \cdot g_2$ for $g_1, g_2 \in \mathbf{G}$. The group \mathbf{G} is said to be measurable if the group operations $g \mapsto g^{-1}$ and $(g_1, g_2) \mapsto g_1 g_2$ are $\mathbf{S}_{\mathbf{G}}$ -measurable, where $\mathbf{S}_{\mathbf{G}}$ is a σ -algebra of subsets of \mathbf{G} . We assume throughout that \mathbf{G} has a topology that is locally compact, second countable, and Hausdorff (lscH), which makes the group operations continuous. We may then take $\mathbf{S}_{\mathbf{G}}$ as the Borel σ -algebra, making \mathbf{G} a standard Borel space.

For $A \subseteq \mathbf{G}$ and $g \in \mathbf{G}$, we write $gA = \{gh : h \in A\}$ and $Ag = \{hg : h \in A\}$. A measure ν on \mathbf{G} is said to be left-invariant if $\nu(gA) = \nu(A)$ for all $A \in \mathbf{S}_{\mathbf{G}}$, and right-invariant if $\nu(Ag) = \nu(A)$. When \mathbf{G} is lscH, there exist left- and right-invariant σ -finite measures $\lambda_{\mathbf{G}}$ and $\tilde{\lambda}_{\mathbf{G}}$, respectively, that are unique up to scaling ([Folland, 2016](#), Ch. 2.2), known as left- and right-Haar measures. When there is no chance of confusion, we use λ to denote left-Haar measure. If \mathbf{G} is compact, then $\lambda = \tilde{\lambda}$, and the unique normalized Haar measure acts as the uniform probability measure over the group. We use G to denote a random element of \mathbf{G} ; when \mathbf{G} is compact, $G \sim \lambda$ denotes a random group element sampled from λ .

2.2 Group actions

We briefly review the relevant aspects of groups acting on sets and special properties that are used in our work. Many of the mathematical techniques have appeared in various statistical contexts, and a thorough treatment can be found in Eaton (1989); Wijsman (1990); Eaton (2007). Inversion kernels (described below) do not seem to have been used in statistics or machine learning, perhaps owing to their relatively recent appearance in probability (Kallenberg, 2011). However, special cases in which deterministic versions (called representative inversions below) exist have appeared in statistics and machine learning (Bloem-Reddy and Teh, 2020; Winter et al., 2022).

A group \mathbf{G} acts measurably on a set \mathbf{X} if the group action $\Phi: \mathbf{G} \times \mathbf{X} \rightarrow \mathbf{X}$ is measurable relative to $\mathbf{S}_{\mathbf{G}} \otimes \mathbf{S}_{\mathbf{X}}$ and $\mathbf{S}_{\mathbf{X}}$. All actions in this work are assumed to be continuous (and therefore measurable), and for convenience we simply say that \mathbf{G} acts on \mathbf{X} , writing $gx = \Phi(g, x)$ as short-hand. For a set $A \subseteq \mathbf{X}$, the group acts as $gA = \{gx : x \in A\}$. For fixed $x \in \mathbf{X}$, the stabilizer subgroup is $\mathbf{G}_x = \{g \in \mathbf{G} : gx = x\}$. The action is called *free* or *exact* if $gx = x$ implies that $g = \text{id}$, in which case $\mathbf{G}_x = \{\text{id}\}$ for all $x \in \mathbf{X}$. The orbit of $x \in \mathbf{X}$ is the set $O(x) = \{gx : g \in \mathbf{G}\}$. The orbits partition \mathbf{X} into equivalence classes, where two points are equivalent if and only if they belong to the same orbit. If \mathbf{X} has only one orbit, then the action is said to be *transitive*. It is not hard to show that if $hx = x'$ for $x \neq x'$, then $h\mathbf{G}_x h^{-1} = \mathbf{G}_{x'}$. That is, the stabilizer subgroups of the elements of an orbit are all conjugate.

A function f with domain \mathbf{X} is invariant if it is constant on each orbit: $f(gx) = f(x)$, $x \in \mathbf{X}$, $g \in \mathbf{G}$. In general, an invariant function may take the same value on different orbits. A *maximal invariant* is an invariant function $M: \mathbf{X} \rightarrow \mathbf{M}$ that takes a different value on each orbit, so that if $M(x) = M(x')$, then $x = gx'$ for some $g \in \mathbf{G}$. Maximal invariants arise as particularly useful statistics in problems with group symmetry because any invariant function f can be written as $f(x) = k(M(x))$, for some function k . Maximal invariants are typically not unique. However, they are all isomorphic to the canonical projection onto the quotient space, $\pi: \mathbf{X} \rightarrow \mathbf{X}/\mathbf{G}$, $x \mapsto O(x)$. Measurability issues can arise when \mathbf{G} is non-compact; we discuss these below.

Invariance is a special case of a more general property. Suppose \mathbf{G} acts on \mathbf{X} and on another set \mathbf{Y} ; the group action may be different on each. A function $f: \mathbf{X} \rightarrow \mathbf{Y}$ is *\mathbf{G} -equivariant* if $f(gx) = gf(x)$, $x \in \mathbf{X}$, $g \in \mathbf{G}$. These properties extend to measures.

Definition 1. A probability measure P on \mathbf{X} is *\mathbf{G} -invariant* if $P(g^{-1}A) = P(A)$ for all $g \in \mathbf{G}$, $A \in \mathbf{S}_{\mathbf{X}}$.

We write $g_*P(A) = P(g^{-1}A)$ as the pushforward of P under the action of $g \in \mathbf{G}$. In that notation, \mathbf{G} -invariance of P entails $g_*P = P$ for all $g \in \mathbf{G}$.

We say that $P_{X,Y}$ is jointly \mathbf{G} -invariant if it is invariant in the sense of Definition 1 extended to \mathbf{G} acting on $\mathbf{X} \times \mathbf{Y}$. In addition to joint invariance, we may define symmetry in the conditional distribution.

Definition 2. The conditional distribution of Y given X is said to be *\mathbf{G} -equivariant²* if

$$P_{Y|X}(x, B) = P_{Y|X}(gx, gB), \quad x \in \mathbf{X}, B \in \mathbf{S}_{\mathbf{Y}}, g \in \mathbf{G}. \quad (2)$$

If the action of \mathbf{G} on \mathbf{Y} is trivial and $P_{Y|X}$ satisfies (2) so that $P_{Y|X}(gx, B) = P_{Y|X}(x, B)$, then the conditional distribution is said to be \mathbf{G} -invariant.

Both of these definitions also apply to general measures (i.e., probability measures and conditional distributions can be replaced by measures and Markov kernels, respectively).

2.2.1 Representatives and inversions

Our work makes extensive use of special entities that are somewhat non-standard in the recent invariance-based statistics and machine learning literature, so we review them here. We can assign a particular element of each orbit as the *orbit representative*. We write $[x]$ as the representative on the orbit $O(x)$. That is, $[x] = gx$ for some $g \in \mathbf{G}$. The structural properties described below do not depend on which element of the orbit is chosen as the representative. All of the properties are relative to a particular choice, and a different choice would result in the same properties relative to that choice. For a particular choice of representatives, the subset of \mathbf{X} consisting of each orbit's representative is denoted by $[\mathbf{X}]$. Note that $[\mathbf{X}] \cap O(x)$ consists of a single point; namely, $[x]$. A function $\gamma: \mathbf{X} \rightarrow [\mathbf{X}]$ that maps elements of \mathbf{X}

²Some authors refer to (2) as invariance; we use equivariance to avoid confusion with the invariance of marginal and joint distributions, and to be consistent with current usage, especially with respect to equivariant functions.

onto their corresponding orbit representatives in $[\mathbf{X}]$ is called an *orbit selector*. Note that any orbit selector is a maximal invariant by definition. Conversely, a maximal invariant defines a choice of orbit representatives if the value it takes on each orbit is an element of the orbit. If $[\mathbf{X}]$ is a measurable subset of \mathbf{X} and γ is a measurable function relative to $\mathbf{S}_{\mathbf{X}}$ and $\mathbf{S}_{\mathbf{X}} \cap [\mathbf{X}]$, then $[\mathbf{X}]$ is called a *measurable cross-section*.

A function $\tau: \mathbf{X} \rightarrow \mathbf{G}$ is called a *representative inversion* if $\Phi(\tau(x), \gamma(x)) = \tau(x)\gamma(x) = x$ and $\tau(gx) = g\tau(x)$ for all $x \in \mathbf{X}, g \in \mathbf{G}$. The role of τ is to return the element of \mathbf{G} that must be applied to move $[x]$ to x . Conversely, the inverse element, $\tau(x)^{-1}$, moves x to $[x]$. In order for τ to be uniquely defined, the group action must be free. If it is not, an equivariant *inversion probability kernel*, or inversion kernel for short, $\zeta: \mathbf{X} \times \mathbf{S}_{\mathbf{G}} \rightarrow [0, 1]$, can be used in place of τ , so that a sample from $\zeta(X, \bullet)$ will transform $\gamma(X)$ into X with probability one. That is, if $X \sim P$ and $\tilde{\tau} \mid X \sim \zeta(X, \bullet)$, then $X = \tilde{\tau}\gamma(X)$ almost surely. At a high level, one may think of the inversion kernel $\zeta(x, \bullet)$ as the uniform distribution on the left coset $g\mathbf{G}_{\gamma(x)}$, where $g\gamma(x) = x$. In the case of a free action, this simplifies to $\zeta(x, \bullet) = \delta_{\tau(x)}$. In some cases, a representative inversion can still be defined when the action is not free (see Example 1), in which case an equivalent inversion kernel can be defined as $\zeta'(x, B) := \zeta(\gamma(x), \tau(x)^{-1}B)$.

2.2.2 Proper group actions

In the analysis of probabilistic aspects of group actions, measurability issues can arise without regularity conditions. Throughout, we will assume that the group action is *proper*. That is, there exists a strictly positive measurable function $h: \mathbf{X} \rightarrow \mathbb{R}_+$ such that for each $x \in \mathbf{X}$, we have $\int_{\mathbf{G}} h(gx)\lambda(dg) < \infty$ (Kallenberg, 2007).³ This is a standard assumption in statistical applications of group theory (e.g., Eaton, 1989; Wijsman, 1990; McCormack and Hoff, 2023) and is satisfied in many settings of interest. A sufficient condition for proper group action is when \mathbf{G} is compact and acts continuously on \mathbf{X} , which is the setting for our tests for invariance in Section 3. When \mathbf{G} is non-compact, a group action can fail to be proper if \mathbf{G} is “too large” for \mathbf{X} in the sense that the stabilizer subgroups are non-compact. A class of non-compact group actions known to be proper are those of the isometry group of a Riemannian manifold. For the purposes of this work, we rely on the assumption of proper group actions to guarantee the existence of measurable orbit selectors and inversion kernels, which turn out to have extremely useful properties. We gather some of those properties in a proposition, which is a collection of existing results.

To state it, let ν be any bounded measure on $(\mathbf{X}, \mathbf{S}_{\mathbf{X}})$ and let $\bar{\mathbf{S}}_{\mathbf{X}}^{\nu}$ be the completion of $\mathbf{S}_{\mathbf{X}}$ to include all subsets of ν -null sets, and denote by $\bar{\nu}$ the extension of ν to $\bar{\mathbf{S}}_{\mathbf{X}}^{\nu}$ (see, e.g. Çinlar, 2011, Proposition 1.3.10). All statements of $\bar{\nu}$ -measurability in the following proposition are with respect to $\bar{\mathbf{S}}_{\mathbf{X}}^{\nu}$, so that a set $A \subseteq \mathbf{X}$ is $\bar{\nu}$ -measurable if $A \in \bar{\mathbf{S}}_{\mathbf{X}}^{\nu}$. Moreover, a function defined by a particular property is $\bar{\nu}$ -measurable if it is measurable in the usual sense with respect to $\bar{\mathbf{S}}_{\mathbf{X}}^{\nu}$, and if the defining property holds with the possible exception of a $\bar{\nu}$ -null set. Clearly, such a function would also be $\bar{\rho}$ -measurable for any measure $\rho \ll \nu$.

Proposition 1. *Let \mathbf{G} be a lcsch group acting continuously and properly on \mathbf{X} , and ν any bounded measure on \mathbf{X} . Then the following hold:*

1. *The canonical projection $\pi: \mathbf{X} \rightarrow \mathbf{X}/\mathbf{G}$ is a measurable maximal invariant, and any measurable \mathbf{G} -invariant function $f: \mathbf{X} \rightarrow \mathbf{Y}$ can be written as $f = f^* \circ \pi$, for some measurable $f^*: \mathbf{X}/\mathbf{G} \rightarrow \mathbf{Y}$.*
2. *There exists a $\bar{\nu}$ -measurable orbit selector $\gamma: \mathbf{X} \rightarrow [\mathbf{X}]$, which is a maximal invariant statistic, and it induces a $\bar{\nu}$ -measurable cross-section $[\mathbf{X}] = \gamma(\mathbf{X})$.*
3. *For a fixed $\bar{\nu}$ -measurable orbit selector γ , there exists a unique $\bar{\nu}$ -measurable inversion probability kernel $\zeta: \mathbf{X} \times \mathbf{S}_{\mathbf{G}} \rightarrow [0, 1]$ with the following properties:*
 - (a) *ζ is \mathbf{G} -equivariant: For all $g \in \mathbf{G}, x \in \mathbf{X}, B \in \mathbf{S}_{\mathbf{G}}, \zeta(gx, B) = \zeta(x, g^{-1}B)$.*
 - (b) *For each $x \in \mathbf{X}, \zeta(\gamma(x), \bullet)$ is normalized Haar measure on the stabilizer subgroup $\mathbf{G}_{\gamma(x)}$.*
 - (c) *For each $x \in \mathbf{X}$, if $\tilde{\tau} \sim \zeta(x, \bullet)$, then $\tilde{\tau}\gamma(x) = x$ with probability one.*
 - (d) *If there is a $\bar{\nu}$ -measurable representative inversion $\tau: \mathbf{X} \rightarrow \mathbf{G}$ associated with γ such that it satisfies $\tau(x)\gamma(x) = x$ and $\tau(gx) = g\tau(x)$ for each $x \in \mathbf{X}, g \in \mathbf{G}$, then $\zeta'(x, B) = \zeta(\gamma(x), \tau(x)^{-1}B)$ is an*

³This definition of proper group action is a slightly weaker, non-topological version of the definition used in previous work in the statistics literature (e.g., Eaton, 1989; Wijsman, 1990), and only requires the existence of Haar measure. The topological version is as follows: the map $(g, x) \mapsto (gx, x)$ is a proper map, i.e., the inverse image of each compact set in $\mathbf{X} \times \mathbf{X}$ is a compact set in $\mathbf{G} \times \mathbf{X}$. That definition implies the one used here; see Kallenberg (2007) for details.

equivalent inversion kernel. In particular, this holds when the action of \mathbf{G} on \mathbf{X} is free, in which case $\mathbf{G}_{\gamma(x)} = \{id\}$ and the inversion kernel is $\delta_{\tau(x)}$.

The measurability of the canonical projection is a result from functional analysis; see Eaton (1989, Theorem 5.4) for an extended statement and references. Items 2–3c follow directly from results of Kallenberg (2011, 2017) on the existence of universally measurable versions of γ and ζ . Item 3d follows from 3a and 3b.

In the remainder of the paper, we assume that the action of \mathbf{G} on any space is continuous and proper; these conditions will be implicit in statements such as “let \mathbf{G} be a group that acts on \mathbf{X} ”. In particular, measurable orbit selectors and inversion kernels exist under these assumptions.

3 Testing for distributional invariance

In this section, we develop an abstract framework for non-parametric tests for distributional invariance under a specified group. The tests are based on known characterizations of distributional invariance. We briefly review that background here before developing the hypothesis testing framework in Sections 3.1 and 3.2.

We are interested in testing for the \mathbf{G} -invariance of a *probability* measure, which, under the assumption that \mathbf{G} acts properly on \mathbf{X} , requires \mathbf{G} to be compact. For the remainder of this section, we assume that \mathbf{G} is compact; non-compact groups will arise in the treatment of conditional symmetry in Section 5. For a specified compact group \mathbf{G} , given a sample of data $(X_1, \dots, X_n) \sim_{\text{iid}} P$ from an unknown distribution P , decide between:

$$H_0: P \text{ is } \mathbf{G}\text{-invariant} \quad \text{versus} \quad H_1: P \text{ is not } \mathbf{G}\text{-invariant} .$$

Recall that $\mathcal{P}(\mathbf{X})$ denotes the set of probability measures on \mathbf{X} . The set of \mathbf{G} -invariant probability measures is $\mathcal{P}^\circ(\mathbf{X})$, and the non-invariant ones are $\mathcal{P}^\times(\mathbf{X})$, so that $\mathcal{P}^\circ(\mathbf{X}) \cup \mathcal{P}^\times(\mathbf{X}) = \mathcal{P}(\mathbf{X})$ and $\mathcal{P}^\circ(\mathbf{X}) \cap \mathcal{P}^\times(\mathbf{X}) = \emptyset$. For a specified \mathbf{G} , we use H_0 and $\mathcal{P}^\circ(\mathbf{X})$ interchangeably, and similarly for H_1 and $\mathcal{P}^\times(\mathbf{X})$.

Distributional invariance can be characterized in several known ways. One well-known way to obtain an invariant distribution is to average over the group. In particular, when \mathbf{G} is compact, we can define the probability measure obtained by *orbit-averaging* P as

$$P^\circ(A) := \int_{\mathbf{G}} g_* P(A) \lambda(dg) = \int_{\mathbf{G}} P(g^{-1}A) \lambda(dg) , \quad A \in \mathbf{S}_{\mathbf{X}} .$$

We refer to P° as the *orbit-averaged distribution*. The averaging operator yields a useful characterization of invariant probability measures. The following proposition lists additional characterizations that will be useful in developing hypothesis tests for invariance.

Proposition 2. *Let \mathbf{G} be a compact group acting on \mathbf{X} and P a probability measure on \mathbf{X} . Let γ be a measurable orbit selector and ζ a measurable inversion kernel. With $X \sim P$, the following are equivalent:*

- I0. P is \mathbf{G} -invariant.
- I1. $P = P^\circ$.
- I2. If $G \sim \lambda$ with $G \perp\!\!\!\perp X$, then $X \stackrel{d}{=} GX$.
- I3. If $G \sim \lambda$ and $Y \sim \gamma_* P$ with $G \perp\!\!\!\perp Y$, then $X \stackrel{d}{=} GY$. This holds even conditionally on $\gamma(X)$. That is, $(\gamma(X), X, G) \stackrel{d}{=} (\gamma(X), G\gamma(X), G)$, which implies that $X \mid \gamma(X) \stackrel{d}{=} G\gamma(X) \mid \gamma(X)$.
- I4. If $\tilde{\tau} \sim \zeta(X, \bullet)$ and $G \sim \lambda$ with $\tilde{\tau} \perp\!\!\!\perp G \mid \gamma(X)$, then $\tilde{\tau} \stackrel{d}{=} G$. If there exists a representative inversion $\tau(x)$, then this holds with $\tilde{\tau}$ replaced by $\tau(X)H$, where $H \sim \lambda_{\mathbf{G}_{\gamma(X)}}$.

It follows from the invariance of the Haar measure that Properties I0 and I1 imply each other, which is straightforward to verify. Property I2 is just a reformulation of Property I1 in terms of random variables. These properties hold regardless of the existence of a measurable orbit selector and inversion kernel. Proving that Properties I0 and I3 imply each other is only slightly more involved. An accessible proof can be found in Eaton (1989, Theorems 4.3–4.4); see also Kallenberg (2017, Theorem 7.15). Property I4 follows from Property I3 and the identity $x \stackrel{\text{a.s.}}{=} \tilde{\tau}\gamma(x)$.

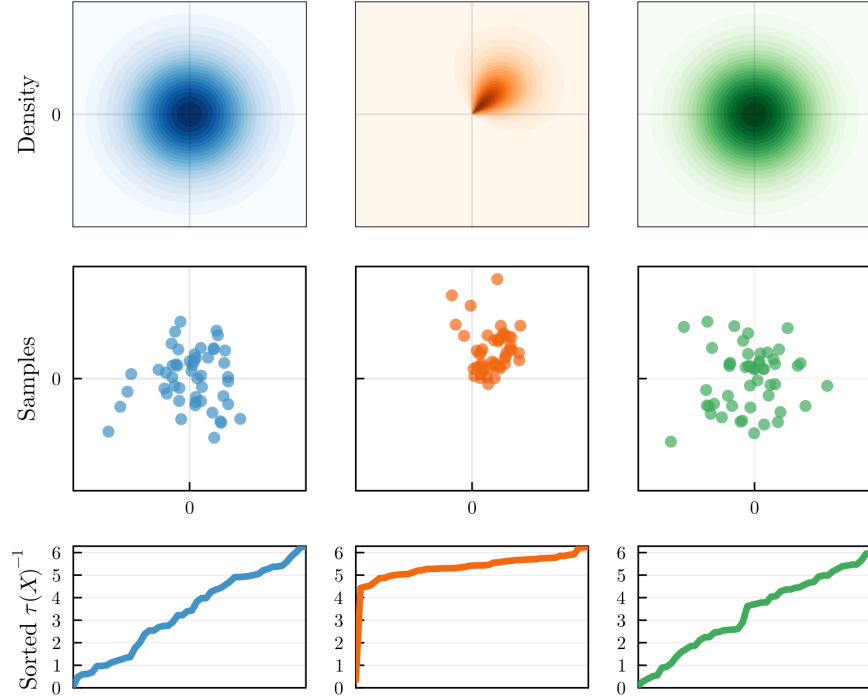


Figure 1: First row: Densities for the 2D multivariate Gaussian $N(0_2, \mathbf{I}_2)$ (blue), Cartesian representation of the distribution $\chi_2 \otimes \text{vonMises}(\pi/4, 4)$ over polar coordinates (orange) and the same distribution averaged over $\text{SO}(2)$ (green). Second row: 50 samples from the respective distributions. Third row: Angles in $[0, 2\pi]$ needed for a counterclockwise rotation of each sample X_i to the point $(\|X_i\|, 0)$, sorted in increasing order.

The following examples illustrate the main ideas.

Example 1. Let $\mathbf{X} = \mathbb{R}^d$, so that X is a random d -dimensional real vector. The isotropic multivariate normal distribution $N(0, \mathbf{I}_d)$ is a well-known example of a distribution that is invariant under the action of $\text{SO}(d)$, the group of d -dimensional rotation matrices, where the action is by matrix-vector multiplication. Properties I1 and I2 in Proposition 2 are straightforward to check. Using the standard formula for an affine transformation of a multivariate normal distribution, if $X \sim N(0, \mathbf{I}_d)$ and $g \in \text{SO}(d)$, then gX has distribution $N(0, g\mathbf{I}_d g^\top) = N(0, \mathbf{I}_d)$. This holds for every g , and therefore it also holds for random G .

The other characterizations are perhaps less well-known. The set of orbit representatives (and the induced cross-section) can be chosen to be the points of an axis, e.g., the axis with unit basis vector $\mathbf{e}_1 = [1, 0, \dots, 0]^\top$. Then for each $x \in \mathbb{R}^d$, $\gamma(x) = \|x\|\mathbf{e}_1$. For $d = 2$, the action is free; for $d > 2$, the stabilizer subgroup $\mathbf{G}_{\gamma(x)}$ is the set of d -dimensional rotations around the axis corresponding to \mathbf{e}_1 . When $X \sim N(0, \mathbf{I}_d)$, $\|X\|$ has a so-called χ_d -distribution (it is the square root of a χ_d^2 -distributed random variable), and $Y \stackrel{d}{=} \|X\|\mathbf{e}_1$ satisfies $X \stackrel{d}{=} GY$, with G a uniform random rotation from $\text{SO}(d)$. The left column of Figure 1 illustrates this for $d = 2$.

In this case, one may construct a representative inversion function corresponding to $\gamma(x) = \|x\|\mathbf{e}_1$ by, for example, rotating $\|x\|\mathbf{e}_1$ to x in the 2D subspace spanned by $x/\|x\|$ and \mathbf{e}_1 . That is, let $\tilde{x} := (x - \langle \mathbf{e}_1, x \rangle \mathbf{e}_1) / \|x - \langle \mathbf{e}_1, x \rangle \mathbf{e}_1\|$, so that $[\mathbf{e}_1, \tilde{x}]$ is a matrix in $\mathbb{R}^{d \times 2}$ whose columns form an orthonormal basis for the 2D subspace spanned by $x/\|x\|$ and \mathbf{e}_1 . Now let θ_x be such that $\cos(\theta_x) = \langle \mathbf{e}_1, x/\|x\| \rangle$, and R_θ the standard 2D rotation matrix of angle θ ,

$$R_\theta = \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix}.$$

Then the d -dimensional rotation defined by

$$\tau(x) = \mathbf{I}_d - \mathbf{e}_1 \mathbf{e}_1^\top - \tilde{x} \tilde{x}^\top + [\mathbf{e}_1, \tilde{x}] R_{\theta_x} [\mathbf{e}_1, \tilde{x}]^\top \quad (3)$$

satisfies $\tau(x)(\|x\|\mathbf{e}_1) = x$. A sample from the corresponding inversion kernel is then generated by taking a uniform random $(d-1)$ -dimensional rotation H and extending it to a d -dimensional rotation H' that fixes the first dimension, so that $\tau(x)H'$ has distribution $\zeta(x, \bullet)$. For $d=2$, $\zeta(\gamma(x), \bullet) = \delta_{\text{id}}$, so Property I4 indicates that $\tau(X) \stackrel{d}{=} G$, with G a uniform random 2D rotation. This is visualized in the bottom-left plot of Figure 1. \diamond

Example 2. As a non-invariant example, consider $\mathbf{X} = \mathbb{R}^d$ and $\mathbf{G} = \text{SO}(d)$ as in Example 1, but now generate $X' = G'Y'$, with $Y' = Z\mathbf{e}_1$, $Z \sim \chi_d$, and G' sampled from the von Mises–Fisher distribution $\text{vMF}(\xi, \kappa)$, a non-invariant distribution on the $(d-1)$ -sphere, which is isomorphic to $\text{SO}(d)$. In this case, although the distribution of Y' is the same as $Y = \|X\|\mathbf{e}_1$ above, rotations toward ξ have higher probability. This is shown in the middle column of Figure 1. Averaging the vMF distribution over $\text{SO}(d)$ results in the uniform distribution on $\text{SO}(d)$, restoring invariance, as shown in the right column of Figure 1. This remains true even when the distribution of G' depends on Y' , such as if G' were sampled from $\text{vMF}(\xi(y'), \kappa(y'))$. \diamond

Example 3. As in the previous examples, let $\mathbf{X} = \mathbb{R}^d$. Now consider $\mathbf{G} = \mathbb{S}_d$, the symmetric group on d elements acting on a vector $x \in \mathbf{X}$ by permutation. The distribution P is said to be *finitely exchangeable* if $gX \stackrel{d}{=} X$ for every permutation $g \in \mathbb{S}_d$. Here, the orbit representative is the vector of order statistics, $X_{(d)}$, which puts the elements of X in increasing order. (We assume for convenience that ties occur with probability zero.) An exchangeable random variable can be generated by first sampling a random order statistic and then applying a uniform random permutation. The representative inversion $\tau(X)$ is the permutation that transforms the order statistics back into X .

As an example, consider again the multivariate normal distribution, $\text{N}(0, \Sigma)$. In order for this distribution to be exchangeable, the covariance matrix Σ must satisfy (Aldous, 1985)

$$\Sigma = (1 - \rho)\sigma^2\mathbf{I}_d + \rho\sigma^2\mathbf{1}_d, \quad \sigma^2 > 0, \quad -\frac{1}{d-1} \leq \rho \leq 1,$$

where $\mathbf{1}_d$ is the $d \times d$ matrix of all ones. If, for example, $\Sigma_{1,d} = \Sigma_{d,1} = \rho^2\sigma^2$ for $|\rho| \neq 1$, then the resulting distribution would not be exchangeable. \diamond

3.1 Non-parametric test statistics

The practical advantages of the equivalent characterizations of \mathbf{G} -invariance will become clear in the following sections: rather than verifying potentially (uncountably) many equalities of the form $P = g_*P$, tests for invariance can be performed via a single comparison between a sample and random transformations of it. The distributional identities in Proposition 2 suggest natural non-parametric test statistics in the form of divergences or metrics on the space of probability distributions on \mathbf{X} , so that any consistent estimator of such a statistic can be used to construct a consistent test of the desired asymptotic level.

To that end, let $D: \mathcal{P}(\mathbf{X}) \times \mathcal{P}(\mathbf{X}) \rightarrow [0, \infty)$ be a metric on $\mathcal{P}(\mathbf{X})$. (The following also holds if D is a divergence or any other continuous function that separates points of $\mathcal{P}(\mathbf{X})$.) The aim will be to recover $D(P, P^\circ)$ in the limit of infinite sample size. Denote by $X_{1:n} := (X_1, \dots, X_n)$ an i.i.d. sample from an unknown distribution P . The samples can be used to estimate P by the empirical measure

$$\hat{P}_n(A) := \frac{1}{n} \sum_{i=1}^n \delta_{X_i}(A), \quad A \in \mathbf{S}_{\mathbf{X}}.$$

Similarly, using i.i.d. samples $G_{i,j} \sim \lambda$, an estimator of P° is the Monte Carlo averaged⁴ empirical measure

$$\hat{P}_{n,m}^\square(A) := \frac{1}{mn} \sum_{i=1}^n \sum_{j=1}^m \delta_{X_i}(G_{i,j}^{-1}A) = \frac{1}{mn} \sum_{i=1}^n \sum_{j=1}^m \delta_{G_{i,j}X_i}(A), \quad A \in \mathbf{S}_{\mathbf{X}}.$$

⁴We assume henceforth that \mathbf{G} is either uncountable or large and discrete, so that enumeration of the group elements is impossible or infeasible. If \mathbf{G} is discrete and small enough to feasibly enumerate, then Monte Carlo averages over \mathbf{G} can be replaced by exact averages.

A natural test statistic is the distance between the two estimates,

$$T_{n,m}(X_{1:n}) := D(\hat{P}_n, \hat{P}_{n,m}^\square), \quad (4)$$

which converges almost surely to $D(P, P^\circ)$ as $n \rightarrow \infty$, for any $m \geq 1$.

In practice, using $m = 1$ amounts to a two-sample test between (X_1, \dots, X_n) and $(G_1 X_1, \dots, G_n X_n)$. For $m > 1$, it can be thought of as an aggregated $(m + 1)$ -sample test, where m samples $(G_{1,j} X_1, \dots, G_{n,j} X_n)$ are known to be i.i.d. and therefore aggregated.

For a sequence of metric-based statistics $(T_{n,m})_{n \geq 1}$ with fixed D and $m \geq 1$, and critical values $(c_n)_{n \geq 1}$, define the corresponding sequence of critical functions, or tests,

$$\phi_{n,m}(X_{1:n}) = \mathbb{1}\{T_{n,m}(X_{1:n}) > c_n\}. \quad (5)$$

The power function of a test based on $\phi_{n,m}$ is

$$\beta_n(P) = \mathbb{E}_{P \otimes \lambda}[\phi_{n,m}(X_{1:n})],$$

where the expectation with respect to $P \otimes \lambda$ is taken over $X_{1:n}$ and the random transformations $G_{i,j} \sim_{\text{iid}} \lambda$.

Recall that $\mathcal{P}^\circ(\mathbf{X})$ is the set of \mathbf{G} -invariant probability measures on \mathbf{X} and $\mathcal{P}^\times(\mathbf{X})$ the set of non-invariant probability measures so that $\mathcal{P}(\mathbf{X}) = \mathcal{P}^\circ(\mathbf{X}) \cup \mathcal{P}^\times(\mathbf{X})$, and that the hypotheses are $H_0: P \in \mathcal{P}^\circ(\mathbf{X})$ versus $H_1: P \in \mathcal{P}^\times(\mathbf{X})$.

Theorem 1. *Fix $m \geq 1$ and a metric or divergence D on $\mathcal{P}(\mathbf{X})$. Let a sequence of tests $(\phi_{n,m})_{n \geq 1}$ (as in (5)) be such that the critical values $(c_n)_{n \geq 1}$ satisfy $\lim_{n \rightarrow \infty} c_n = c \geq 0$. Then $(\phi_{n,m})_{n \geq 1}$ is pointwise asymptotically level α for any $\alpha \in [0, 1]$. That is, for any $c \geq 0$, for any $P \in \mathcal{P}^\circ(\mathbf{X})$,*

$$\limsup_{n \rightarrow \infty} \mathbb{E}_{P \otimes \lambda}[\phi_{n,m}(X_{1:n})] \leq \alpha, \quad \alpha \in [0, 1]. \quad (6)$$

If $c = 0$, then $(\phi_{n,m})_{n \geq 1}$ is also pointwise consistent in power: for any $P \in \mathcal{P}^\times(\mathbf{X})$,

$$\lim_{n \rightarrow \infty} \mathbb{E}_{P \otimes \lambda}[\phi_{n,m}(X_{1:n})] = 1, \quad (7)$$

and therefore the sequence of tests is asymptotically unbiased.

The proof can be found in Appendix A.1. The theorem can be modified in a number of ways. It remains valid if the metric D is replaced by an estimator of the metric, \hat{D} , as long as $\hat{D}(\hat{P}_n, \hat{P}_{n,m}^\square)$ converges in probability to $D(P, P^\circ)$ as $n \rightarrow \infty$. The condition that D distinguishes elements of $\mathcal{P}(\mathbf{X})$ (which is satisfied when D is a metric or divergence) can be relaxed without changing the asymptotic level of the test. All that is required is that $D(P, P) = 0$. However, the power would be reduced if $D(P, P^\circ) = 0$ for $P \neq P^\circ$.

As an alternative, one may compare the distribution of inversions $\tilde{\tau}_i \mid X_i \sim \zeta(X_i, \bullet)$ to Haar measure λ using a metric D on $\mathcal{P}(\mathbf{G})$, in which case an analogous version of Theorem 1 holds by Proposition 2, Property I4.

3.2 Exact conditional Monte Carlo tests of invariance

Theorem 1 shows that tests based on a metric on $\mathcal{P}(\mathbf{X})$ have desirable large-sample properties, but the question of how to set critical values for finite n remains unanswered. In this section, we develop a conditional Monte Carlo method for doing so that results in a test with exact finite-sample size. We do not address the more difficult theoretical question of power in this setting, though our experiments in Section 7 indicate that the methods do perform well in that respect.

Our Monte Carlo procedure is conceptually similar to a resampling approximation of a two-sample permutation test for equality in distribution. In the latter setting, under the null hypothesis that the two samples have the same distribution, a sufficient statistic for H_0 is the empirical measure of the pooled sample; the conditional distribution of the pooled sample given the empirical measure is induced by the uniform distribution over permutations of the pooled sample. Rather than computing the test statistic under every permutation, conditional p -values are estimated by sampling uniformly from the set of permutations. The conditional p -values are valid unconditionally since the conditional p -values are valid for almost every realization of the empirical measure under H_0 .

Similarly, in the case of \mathbf{G} -invariance, one may also condition on a sufficient statistic for $\mathcal{P}^\circ(\mathbf{X})$. Here, any maximal invariant is a sufficient statistic (Farrell, 1962; Dawid, 1985; Bloem-Reddy and Teh, 2020). For example, Property I3 of Proposition 2 indicates that given $\gamma(X)$, which is a maximal invariant, the conditional distribution of X is that induced by λ on the orbit $O(X)$. That is, conditionally on $\gamma(X)$, $X \stackrel{d}{=} G\gamma(X)$ with $G \sim \lambda$. This holds for every invariant probability measure $P \in \mathcal{P}^\circ(\mathbf{X})$.

Thus, the Monte Carlo testing method we propose in Algorithm 1 generates pseudo-samples by first sampling $G_i^{(b)} \sim_{\text{iid}} \lambda$ and then applying them to $\gamma(X)_{1:n}$, so that

$$(G_1^{(b)}\gamma(X_1), \dots, G_n^{(b)}\gamma(X_n)) \sim P(\cdot \mid \gamma(X)_{1:n}), \quad b = 1, \dots, B. \quad (8)$$

Because $P(\cdot \mid \gamma(X)_{1:n})$ is the same for every $P \in \mathcal{P}^\circ(\mathbf{X})$, these samples can be used to estimate conditional quantities that are valid uniformly across the null hypothesis class $\mathcal{P}^\circ(\mathbf{X})$.

If \mathbf{G} is discrete and relatively small, conditional expectations with respect to the right-hand side of (8) can be computed exactly; otherwise, we can use Monte Carlo estimates. In particular, given a sample $X_{1:n}$, we can estimate a conditional p -value by Monte Carlo sampling as in Algorithm 1.⁵

Algorithm 1 Monte Carlo p -value

- 1: **procedure** MCTEST($X_{1:n}, m, B, D$)
- 2: Sample $G_{j,1}, \dots, G_{j,n} \sim_{\text{iid}} \lambda$, for $j = 1, \dots, m$
- 3: Using $(G_{j,1}, \dots, G_{j,n})_{j \leq m}$, compute $T_{n,m}(X_{1:n})$ as in (4)
- 4: **for** b in $1, \dots, B$ **do**
- 5: Sample $G_1^{(b)}, \dots, G_n^{(b)} \sim_{\text{iid}} \lambda$
- 6: Set $X_{1:n}^{(b)} := (G_1^{(b)}X_1, \dots, G_n^{(b)}X_n)$
- 7: (Re)using $(G_{j,1}, \dots, G_{j,n})_{j=1}^m$, compute $T_{n,m}(X_{1:n}^{(b)})$
- 8: **end for**
- 9: **return** p -value p_B computed as

$$p_B := \frac{1 + \sum_{b=1}^B \mathbb{1}\{T_{n,m}(X_{1:n}^{(b)}) \geq T_{n,m}(X_{1:n})\}}{1 + B} \quad (9)$$

10: **end procedure**

As we formalize below, this procedure produces a valid p -value for any $B \geq 1$. The estimate p_B can then be used in a critical function $\mathbb{1}\{p_B \leq \alpha\}$ and the resulting test has level α . A special case of the following result, for finite \mathbf{G} , appeared in Hemerik and Goeman (2018) in the context of invariance-based randomization tests. To state it, for $x \in \mathbb{R}_+$, let $\lfloor x \rfloor$ denote the ‘‘floor’’ function applied to x , i.e., the largest integer that is less than or equal to x . Furthermore, let $X_{1:n}^{(0)} := X_{1:n}$.

Theorem 2. Assume that $\mathbb{E}_{P \otimes \lambda}[\mathbb{1}\{T_{n,m}(X_{1:n}^{(b)}) = T_{n,m}(X_{1:n}^{(b')})\}] = 0$ for $b \neq b'$. For any fixed $B \in \mathbb{N}$, p_B obtained as in Algorithm 1 is a valid p -value in the sense that for any $\alpha \in [0, 1]$, if $P \in \mathcal{P}^\circ(\mathbf{X})$, then for any $(g_{i,j})_{i \leq n, j \leq m} \in \mathbf{G}^{n \times m}$,

$$\mathbb{E}_{P \otimes \lambda}[\mathbb{1}\{p_B \leq \alpha\} \mid (G_{i,j})_{i \leq n, j \leq m} = (g_{i,j})_{i \leq n, j \leq m}] = \frac{\lfloor \alpha(B+1) \rfloor}{B+1} \leq \alpha. \quad (10)$$

The same also holds unconditionally for random $(G_{i,j}^{(b)})_{i \leq n, j \leq m}$ sampled independently of $X_{1:n}$ such that they are exchangeable over the index $b = 1, \dots, B$, which includes using the same random sample $(G_{i,j})_{i \leq n, j \leq m}$ for each b .

The proof can be found in Appendix A.2. As noted by Dufour and Neves (2019), if $\alpha(B+1)$ is an integer, then the inequality in (10) becomes equality, and the critical region for the test, $\{p_B \leq \alpha\}$, has exact size α . We validate Theorem 2 empirically in Section 7, finding that over simulated datasets, the distribution of p_B is approximately uniform under the null hypothesis and highly non-uniform under various alternatives.

⁵Due to the invariance of λ , $GX \stackrel{d}{=} G\gamma(X)$ (even conditioned on $\gamma(X)$), so in practice we can replace $\gamma(X_i)$ in (8) with X_i .

Theorem 2 can be adapted in a few ways depending on the setting. Firstly, if the probability of ties is non-zero (i.e., if $T_{n,m}$ is supported on a discrete set), then a modified version with randomized tie-breaking yields valid p -values, as described by [Dufour \(2006\)](#); [Hemerik and Goeman \(2018\)](#).

Secondly, reusing $(G_{j,1}, \dots, G_{j,n})_{j=1}^m$ is not strictly necessary in that Theorem 2 still holds if a new set $(G_{j,1}^{(b)}, \dots, G_{j,n}^{(b)})_{j=1}^m$ is sampled for each b . However, reusing the group elements in each iteration b amounts to conditioning the p -value on them, as shown in (10). This conditioning reduces computation and potentially reduces estimation variance in the procedure. The trade-offs between the two methods would appear primarily in the power of the test. Furthermore, the p -value remains valid even if $(G_{j,1}, \dots, G_{j,n})_{j=1}^m$ are not sampled from λ . However, the power may suffer because averaging P with respect to a probability measure other than λ does not result in an invariant distribution. We find that reusing $(G_{j,1}, \dots, G_{j,n})_{j=1}^m$ works well in our experiments in Section 7.

Thirdly, instead of reusing $X_{1:n}$ in each sampling iteration, we may combine the procedure with standard bootstrap resampling (sampling $X_{1:n}^{*(b)}$ i.i.d. with replacement from $X_{1:n}$) to obtain $X_{1:n}^{(b)} := (G_1^{(b)} X_1^{*(b)}, \dots, G_n^{(b)} X_n^{*(b)})$. We leave this possibility to future work.

Finally, as with Theorem 1, a version of Theorem 2 holds for a suitably modified version of the Monte Carlo test that uses an observed sample of representative inversions, $(\tilde{\tau}_i)_{i=1}^n$, where $\tilde{\tau}_i \mid X_i \sim \zeta(X_i, \bullet)$. In that case, X_i is replaced in the sampling procedure by $\tilde{\tau}_i$, and the null hypothesis sample iterates $(G_1^{(b)} \tilde{\tau}_1, \dots, G_n^{(b)} \tilde{\tau}_n)$ would be compared to (G_1, \dots, G_n) sampled i.i.d. from λ , which is the unique invariant probability measure on \mathbf{G} .

3.2.1 Power estimates

Given a sample $X_{1:n}$ and a fixed $\alpha \in [0, 1]$, we may obtain an estimate of the power function at \hat{P}_n of the test based on $\mathbb{1}\{p_B \leq \alpha\}$. Conditioned on $X_{1:n}$ and $G_{1:m,1:n} := (G_{j,1}, \dots, G_{j,n})_{j=1}^m$, the conditional power function at \hat{P}_n is

$$\begin{aligned} \beta_{n,m}(\hat{P}_n, G_{1:m,1:n}) &= \mathbb{E}_\lambda[\mathbb{1}\{p_B \leq \alpha\} \mid X_{1:n}, G_{1:m,1:n}] \\ &= \mathbb{E}_\lambda \left[\sum_{b=1}^B \mathbb{1} \left\{ T_{n,m}(X_{1:n}^{(b)}) \geq T_{n,m}(X_{1:n}) \right\} \leq \alpha(B+1) - 1 \mid X_{1:n}, G_{1:m,1:n} \right], \end{aligned}$$

where the expectation is taken over the conditional Monte Carlo samples. Because those samples are conditionally i.i.d., if $p_0 := \mathbb{E}_\lambda[\mathbb{1}\{T_{n,m}(X_{1:n}^{(b)}) \leq T_{n,m}(X_{1:n})\} \mid X_{1:n}, G_{1:m,1:n}]$, then

$$\beta_{n,m}(\hat{P}_n, G_{1:m,1:n}) = \sum_{\ell=0}^{\lfloor \alpha(B+1)-1 \rfloor} \binom{B}{\ell} p_0^\ell (1-p_0)^{B-\ell}. \quad (11)$$

Estimates of p_0 can be obtained from Algorithm 1 as $\hat{p}_0 = (p_B^{(B+1)-1})/B$.

Alternatively, a new set of transformations $G_{1:m,1:n}^{(b)}$ can be sampled for each b . The resulting Monte Carlo samples would still be conditionally i.i.d., so the resulting estimate of p_0 obtained in the same way as described above would be conditioned only on $X_{1:n}$. An unconditional estimate of the power could be obtained using the usual bootstrap resampling methods for $X_{1:n}$, as shown in Algorithm 2. We demonstrate this estimation procedure in some of our experiments in Section 7.

Algorithm 2 Power estimate

- 1: **procedure** POWERESTIMATE($X_{1:n}, m, C, B, D$)
 - 2: **for** c in $1, \dots, C$ **do**
 - 3: Sample $X_{1:n}^{(c)}$ i.i.d. with replacement from \hat{P}_n
 - 4: Obtain $p_B^{(c)}$ from procedure MCTEST (Algorithm 1), either reusing $G_{1:m,1:n}$ or resampling for each b
 - 5: Set $\hat{p}_0^{(c)} = \frac{p_B^{(c)(B+1)-1}}{B}$ and compute $\beta_{n,m}^{(c)}$ as in (11)
 - 6: **end for**
 - 7: Set $\hat{\beta}_{n,m} := \frac{1}{C} \sum_{c=1}^C \beta_{n,m}^{(c)}$.
 - 8: **end procedure**
-

4 Kernel hypothesis tests for invariance

As a demonstration of the abstract framework developed in Section 3, we develop the details of a particular instantiation using the maximum mean discrepancy (MMD) as the metric on $\mathcal{P}(\mathbf{X})$.

Let \mathcal{H} be a reproducing kernel Hilbert space (RKHS) of functions $f: \mathbf{X} \rightarrow \mathbb{R}$ on which the evaluation functional $\varphi_x(f) = f(x)$ is continuous for all $x \in \mathbf{X}$. For any $x \in \mathbf{X}$, the Riesz representation theorem says there exists a unique element $k_x \in \mathcal{H}$ such that $f(x) = \langle f, k_x \rangle_{\mathcal{H}}$ for all $f \in \mathcal{H}$, where $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ is the inner product on \mathcal{H} . The reproducing kernel $k: \mathbf{X} \times \mathbf{X} \rightarrow \mathbb{R}$ of \mathcal{H} is a symmetric positive definite kernel such that $k_x(\cdot) = k(x, \cdot)$ and $k(x, x') = \langle k_x, k_{x'} \rangle_{\mathcal{H}}$. The element k_x can be viewed as a map $k_x: \mathbf{X} \rightarrow \mathcal{H}$ that takes x into a potentially infinite-dimensional feature space. Kernel evaluations $k(x, x') = \langle k_x, k_{x'} \rangle_{\mathcal{H}}$ can then be interpreted as computing (possibly implicit) inner products on the mapped feature space. See [Christmann and Steinwart \(2008\)](#) for a thorough treatment.

The *kernel mean embedding* (KME) of a distribution P on \mathbf{X} is defined as $\mu_P := \int_{\mathbf{X}} k_x P(dx)$ and is the unique element of \mathcal{H} such that $\mathbb{E}_P[f(X)] = \langle f, \mu_P \rangle_{\mathcal{H}}$, for all $f \in \mathcal{H}$ ([Muandet et al., 2017](#)). It follows that $\langle \mu_{P_1}, \mu_{P_2} \rangle = \int k(x, x') P_1(dx) P_2(dx')$. We assume that the kernel k is *characteristic* so that the map $P \mapsto \mu_P$ from $\mathcal{P}(\mathbf{X})$ into \mathcal{H} is injective, which leads to a unique embedding for each probability measure P ([Sriperumbudur et al., 2010](#)). If a non-characteristic kernel were used instead, the test statistic developed below would be unable to separate some distinct elements of $\mathcal{P}(\mathbf{X})$; as discussed in Section 3.1, this would not affect the level of the test, but would reduce the power under non-invariant alternatives P with $\mu_{P \circ} = \mu_P$.

Kernel-based hypothesis tests compare two distributions through their KMEs under a metric defined on functions in \mathcal{H} , and can have an advantage over classical tests in that the same testing framework can be used for any type of data (e.g., vectors, matrices, images, etc.) as long as a kernel is available. [Gretton et al. \(2012\)](#) introduced a two-sample kernel hypothesis test based on the MMD. The MMD is a particular integral probability metric defined as

$$\text{MMD}(P_1, P_2) := \sup_{\|f\|_{\mathcal{H}} \leq 1} \mathbb{E}_{P_1}[f(X)] - \mathbb{E}_{P_2}[f(X)] = \|\mu_{P_1} - \mu_{P_2}\|_{\mathcal{H}}.$$

We focus our attention on the *squared* MMD, which is more practical to estimate due to the identity

$$\begin{aligned} \text{MMD}^2(P_1, P_2) &= \langle \mu_{P_1}, \mu_{P_1} \rangle_{\mathcal{H}} + \langle \mu_{P_2}, \mu_{P_2} \rangle_{\mathcal{H}} - 2\langle \mu_{P_1}, \mu_{P_2} \rangle_{\mathcal{H}} \\ &= \int k(x, x') P_1(dx) P_1(dx') + \int k(x, x') P_2(dx) P_2(dx') - 2 \int k(x, x') P_1(dx) P_2(dx'). \end{aligned}$$

This can be estimated from two samples of data, $X_{1:n_1} \sim_{\text{iid}} P_1$ and $Y_{1:n_2} \sim_{\text{iid}} P_2$, using the U-statistic

$$\widehat{\text{MMD}}^2(\hat{P}_{1,n_1}, \hat{P}_{2,n_2}) = \frac{1}{n_1(n_1 - 1)} \sum_{i \neq j} k(X_i, X_j) + \frac{1}{n_2(n_2 - 1)} \sum_{i \neq j} k(Y_i, Y_j) - \frac{2}{n_1 n_2} \sum_{i,j} k(X_i, Y_j).$$

The U-statistic is just an unbiased version of the empirical estimator, which is a V-statistic,

$$\widehat{\text{MMD}}_V^2(\hat{P}_{1,n_1}, \hat{P}_{2,n_2}) = \frac{1}{n_1^2} \sum_{i=1}^{n_1} \sum_{j=1}^{n_1} k(X_i, X_j) + \frac{1}{n_2^2} \sum_{i=1}^{n_2} \sum_{j=1}^{n_2} k(Y_i, Y_j) - \frac{2}{n_1 n_2} \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} k(X_i, Y_j).$$

For convenience, we refer to the MMD^2 as just the MMD, and similarly for related estimators.

In a standard two-sample test based on the MMD, the rejection region is estimated using a standard bootstrap procedure that involves pooling the two samples and repeatedly subsampling two pseudo-samples from the pool. A test of level α rejects H_0 if the observed MMD estimate is larger than $(1 - \alpha) \times 100\%$ of the bootstrapped values. The two-sample MMD test is summarized in Algorithm 3.

Computing $\widehat{\text{MMD}}$ has a computational cost of $\mathcal{O}((n_1 + n_2)^2)$, and the need to estimate the reference distribution via resampling means that the computation does not scale well with sample size. The kernel literature has proposed cheaper approximations to the MMD, such as those based on random Fourier features and Nyström approximation (e.g., [Rahimi and Recht, 2007](#); [Raj et al., 2017](#); [Chatalic et al., 2022](#)). While these methods do not alleviate the need for the bootstrap, they make the computation closer to linear complexity, which consequently makes the test more feasible in the large sample regime. In the following sections, we discuss how the two-sample MMD test can be repurposed as a test for invariance. We only discuss the standard MMD formulation, but all MMD-based tests can also be reformulated as

Algorithm 3 Two-sample MMD test

- 1: **procedure** 2SMMDTEST($X_{1:n_1}, Y_{1:n_2}, \alpha$)
 - 2: compute $\widehat{\text{MMD}}(X_{1:n_1}, Y_{1:n_2})$
 - 3: **for** b in $1, \dots, B$ **do**
 - 4: sample $X_{1:n_1}^{(b)}, Y_{1:n_2}^{(b)}$ from $X_{1:n_1} \cup Y_{1:n_2}$ with replacement
 - 5: compute $\widehat{\text{MMD}}(X_{1:n_1}^{(b)}, Y_{1:n_2}^{(b)})$
 - 6: **end for**
 - 7: compute p -value p_B as in Equation (9) with $\widehat{\text{MMD}}$ in place of $T_{n,m}$
 - 8: **return** 1 if $p_B \leq \alpha$ and 0 otherwise
 - 9: **end procedure**
-

approximate tests based on random Fourier features (when the kernel is shift-invariant) or Nyström approximation. In our experiments in Section 7, we investigate the properties of a MMD test and its Nyström-approximated version. We provide more details about the Nyström test statistic in Appendix B.1.

4.1 MMD test for invariance based on orbit-averaging

We first consider a test for \mathbf{G} -invariance based on comparing P and P° under the MMD. The quantity of interest is

$$\begin{aligned} \text{MMD}(P, P^\circ) &= \langle \mu_P, \mu_P \rangle_{\mathcal{H}} + \langle \mu_{P^\circ}, \mu_{P^\circ} \rangle_{\mathcal{H}} - 2\langle \mu_P, \mu_{P^\circ} \rangle_{\mathcal{H}} \\ &= \langle \mu_P, \mu_P \rangle_{\mathcal{H}} + \int_{\mathbf{X} \times \mathbf{X}} \int_{\mathbf{G} \times \mathbf{G}} k(gx, hx') \lambda(dg) \lambda(dh) P(dx) P(dx') \\ &\quad - 2 \int_{\mathbf{X} \times \mathbf{X}} \int_{\mathbf{G}} k(x, gx') \lambda(dg) P(dx) P(dx') , \end{aligned}$$

which, given data $X_{1:n} \sim_{\text{iid}} P$ and sampled group actions $G_{1:m}, H_{1:m} \sim_{\text{iid}} \lambda$, is estimated by the test statistic

$$\widehat{\text{MMD}}(\hat{P}_n, \hat{P}_{n,m}^\square) = \frac{1}{n(n-1)} \sum_{i \neq j} \left(k(X_i, X_j) + \frac{1}{m^2} \sum_{\ell=1}^m \sum_{r=1}^m k(G_\ell X_i, H_r X_j) - \frac{2}{m} \sum_{\ell=1}^m k(X_i, G_\ell X_j) \right) .$$

To obtain a p -value in the test based on this statistic, we use the Monte Carlo sampling procedure described in Section 3.2. If the kernel k is (almost) equivariant in the sense that

$$\int_{\mathbf{G}} k(gx, x') \lambda(dg) = \int_{\mathbf{G}} k(x, gx') \lambda(dg) , \quad (12)$$

then a computationally more efficient estimator for the test statistic is possible. We provide more details about the equivariant kernel setting in Appendix B.2.

4.1.1 Baseline: two-sample MMD tests for invariance

Under H_0 , $g_i X_i \stackrel{d}{=} X_i$ for each $g_i \in \mathbf{G}$, $i = 1, \dots, n$. Therefore, a standard two-sample test for equality in distribution (such as that in Algorithm 3) can be applied to the samples $X_{1:n}$ and $Y_{1:n} := (g_1 X_1, \dots, g_n X_n)$. We can randomize the g_i 's and still have a test of the correct level. We use this test, which we refer to as the *transformation two-sample test*, as a sensible baseline since it is a valid test but does not take full advantage of the group structure via the sufficiency argument behind Theorem 2.

4.2 MMD test for invariance based on representative inversions

The tests for \mathbf{G} -invariance described in Sections 4.1 and 4.1.1 can be modified to be based on representative inversions by replacing $X_{1:n}$ with their respective representative inversions $\tau(X)_{1:n}$. The test based on representative inversions can make use of a different resampling procedure where the new samples are directly sampled from λ .

4.3 Other non-parametric tests for invariance

While we focus on kernel-based approaches in this work, any test that involves estimating a metric defined on a space of probability distributions can be used. Other possibilities include, for example, tests based on the Wasserstein distance or other integral probability metrics, or, if density evaluations are available, f -divergences such as the Kullback–Leibler divergence. Among well-known integral probability metrics, the MMD has favorable computational and statistical rates (Sriperumbudur et al., 2012).

Alternatively, any property that uniquely characterizes a distribution can also be used to test for \mathbf{G} -invariance. Recent work by Fraiman et al. (2021) employs the Cramér–Wold (CW) theorem, which states that any two distributions on \mathbb{R}^d are the same if and only if the distributions of their 1D projections via the linear kernel $x \mapsto x^\top t$ are the same, for all projections t . To our knowledge, the resulting CW test is the only existing test for distributional group invariance.

The CW test procedure is as follows. For i.i.d. random variables $Z_{1:n}$ supported on \mathbb{R} , let $\widehat{F}_{Z_{1:n}}$ denote the empirical cumulative distribution function of a random variable Z . The procedure proposed by Fraiman et al. (2021) requires that \mathbf{G} be finitely generated by a subset of group elements of size L . Given data $X_{1:n}$ on \mathbb{R}^d , the group generators $(g_\ell)_{\ell=1}^L$ and J random unit vectors $t_j \in \mathbb{R}^d$ are used to compute the worst-case Kolmogorov–Smirnov statistic,

$$T_{\text{CW}}(X_{1:n}) = \max_{\substack{\ell \in 1:L \\ j \in 1:J}} \sup_{u \in \mathbb{R}} \left| \widehat{F}_{(t_j^\top X)_{1:n}}(u) - \widehat{F}_{(t_j^\top (g_\ell X)_{1:n})}(u) \right|.$$

The p -value is estimated by standard bootstrap resampling from $X_{1:n}$.⁶ It is straightforward to extend the CW test to more general groups by sampling $G_\ell \sim_{\text{id}} \lambda$ and applying the methods of Section 3 to obtain a valid test. We compare our tests for invariance to the extended CW test in the experiments in Section 7.

5 Equivariance: conditional symmetry

Although tests for marginal or joint distributional invariance are useful in a number of settings, many problems require a test for symmetry of a *conditional* distribution. These commonly arise in prediction settings where, for example, one must predict Y from X , and it is of interest to test whether a change $X \mapsto gX$ corresponds to a change $Y \mapsto gY$. These so-called equivariant prediction problems have been a subject of intense study in machine learning (see Bronstein et al., 2021, for a review).

Suppose our data is of the form $(X, Y)_{1:n}$, with each (X, Y) -pair sampled i.i.d. from some distribution $P_{X,Y}$ defined over the product space $\mathbf{X} \times \mathbf{Y}$. We write the disintegration of $P_{X,Y}$ as $P_{X,Y} = P_X \otimes P_{Y|X}$, where $P_{Y|X}$ denotes a regular conditional probability of Y given X (i.e., represented by a Markov probability kernel from \mathbf{X} to \mathbf{Y}). Furthermore, suppose there is a group \mathbf{G} that acts on both \mathbf{X} and \mathbf{Y} ; the group action may be different on each. Recall that we say that $P_{X,Y}$ is jointly \mathbf{G} -invariant if it is invariant in the sense of Definition 1 extended to \mathbf{G} acting on $\mathbf{X} \times \mathbf{Y}$.

We refer to \mathbf{G} -equivariance and \mathbf{G} -invariance of the conditional distribution (as in Definition 2) collectively as *conditional symmetry*. Subsequently, we describe ideas generally in terms of equivariance, which specializes to invariance when \mathbf{G} acts trivially on \mathbf{Y} . In contrast to Section 3, the symmetry here can be viewed as symmetry of a *function on \mathbf{X}* , rather than of a probability measure on $(\mathbf{X}, \mathbf{S}_\mathbf{X})$. To see this, let $f: \mathbf{X} \rightarrow \mathbf{Y}$ be \mathbf{G} -equivariant, so that $f(gx) = gf(x)$, for all $x \in \mathbf{X}$, $g \in \mathbf{G}$. The conditional distribution of $Y := f(X)$ can be represented by the Dirac kernel $\delta_{f(x)}(B)$ in this case. It is not hard to show that this kernel is equivariant in the sense of (2) if and only if f is an equivariant function. In the more general case, conditional symmetry describes how the conditional distribution transforms when x is transformed to gx . Indeed, an equivalent definition to (2) is

$$P_{Y|X}(gx, B) = P_{Y|X}(x, g^{-1}B), \quad x \in \mathbf{X}, B \in \mathbf{S}_\mathbf{Y}, g \in \mathbf{G}.$$

In this way, an equivariant conditional distribution “transmits” the transformation of one variable to the other. In contrast to the invariance of a marginal (or joint) distribution, which must preserve the total probability mass, the transformations of a conditional distribution have no such restrictions and conditional symmetry is well-defined even when \mathbf{G} is non-compact.

⁶There is also a version that does not rely on bootstrapping but requires the sample to be split and the use of a Bonferroni correction, which likely reduces the power.

Equivariant Markov kernels occur naturally in the disintegration of jointly invariant measures. Under quite general conditions, a measure is jointly invariant if and only if it disintegrates into an invariant marginal measure and an equivariant Markov kernel. A measure-theoretic proof of this can be found in [Kallenberg \(2017, Theorem 7.6\)](#). [Bloem-Reddy and Teh \(2020\)](#) proved a version of this result for probability measures and compact groups using different methods that require the existence of a measurable representative inversion τ , and which results in a detailed description of the probabilistic structure of the disintegration. Those authors found that if $P_{X,Y}$ is *jointly* \mathbf{G} -invariant, then

$$X \perp\!\!\!\perp \tau(X)^{-1}Y \mid M(X) . \quad (13)$$

In many applications, we are interested in testing for conditional equivariance of $Y \mid X$, even when the marginal distribution of X is not invariant. As we prove below, the conditional independence in (13) holds even when $P_{X,Y}$ is not jointly invariant. In fact, equivariance of $P_{Y|X}$ is both sufficient and necessary. Moreover, even when a unique inversion τ does not exist—i.e., when the action of \mathbf{G} on \mathbf{X} is not free—an analogous conditional independence relation applies: $P_{Y|X}$ is \mathbf{G} -equivariant if and only if

$$(\tilde{\tau}, X) \perp\!\!\!\perp \tilde{\tau}^{-1}Y \mid M(X) , \quad \text{with } \tilde{\tau} \mid X \sim \zeta(X, \bullet) . \quad (14)$$

Theorem 3. *Let \mathbf{G} be a lcscH group acting on each of \mathbf{X} and \mathbf{Y} , with the action on \mathbf{X} proper, so that a measurable inversion kernel ζ exists. Then $P_{Y|X}$ is conditionally \mathbf{G} -equivariant if and only if (14) holds. If there exists a measurable inversion function $\tau: \mathbf{X} \rightarrow \mathbf{G}$, then the characterization condition (14) reduces to (13). If the action of \mathbf{G} on \mathbf{Y} is trivial, then (14) reduces to $X \perp\!\!\!\perp Y \mid M(X)$.*

The proof can be found in Appendix A.3. Based on Theorem 3, a test for \mathbf{G} -equivariance of a conditional distribution can be performed as a test of the conditional independence in (14). If the action of \mathbf{G} on \mathbf{X} is transitive, then there is a single orbit and the test simplifies to a test for (unconditional) independence.

Any maximal invariant can be used as the conditioning variable. A particularly convenient maximal invariant for testing purposes is the orbit selector $\gamma(X)$. Although all maximal invariants are isomorphic to one another, there may be computational and statistical benefits to using a maximal invariant with an explicit representation in a low-dimensional space. For example, in the case of $\text{SO}(d)$, one may use $M(X) = \|X\|$.

Example 4. For random vectors X and Y in \mathbb{R}^d , suppose that $Y \mid X \sim \mathcal{N}(X, \mathbf{I}_d)$. It is straightforward to check that $P_{Y|X}$ is $\text{SO}(d)$ -equivariant: If $\varepsilon \sim \mathcal{N}(0, \mathbf{I}_d)$ so that $Y = X + \varepsilon$, then for any $g \in \text{SO}(d)$,

$$gY \mid X = gX + g\varepsilon \stackrel{d}{=} gX + \varepsilon = Y \mid gX .$$

Using the representative inversion (3) described in Example 1, we can test for equivariance with the conditional independence test (13). We do so in Section 7.1. \diamond

Example 5. The Lorentz group, or indefinite orthogonal group $\text{O}(1, 3)$, is a fundamental symmetry group in physics, where it is the group of linear isometries of Minkowski spacetime, on which the theory of special relativity is commonly constructed. It preserves the quadratic form

$$Q(p) = E^2 - p_x^2 - p_y^2 - p_z^2 ,$$

where $p = (E, p_x, p_y, p_z)$ is a vector in \mathbb{R}^4 representing a particle's momentum in Minkowski spacetime, known as the *four-momentum*. Mathematically, the Lorentz group is a non-compact non-Abelian real Lie group that is not connected. The Standard Model of physics is invariant under the Lorentz group, and violations of $\text{O}(1, 3)$ -invariance could indicate phenomena beyond the Standard Model. An example application is quark tagging ([Kasieczka et al., 2019](#); [Bogatskiy et al., 2020](#)). High-energy quarks produced in particle collisions, such as those at the Large Hadron Collider, quickly decay through a cascading process of gluon emission into collimated sprays of stable hadrons, which are subatomic particles that can be detected and measured ([Salam, 2010](#)). This spray is known as a *jet*. Identifying, or tagging, which species of quark gave rise to a jet is a crucial task in collider physics, and an active area of research ([Larkoski et al., 2020](#)).

According to the Standard Model, the inertial frame of the parent quark may differ from the lab frame by a Lorentz group transformation, and the task of quark-tagging based on the four-momenta of constituent particles should be invariant to those transformations. That symmetry was incorporated into a neural network architecture designed for

quark-tagging (Bogatskiy et al., 2020). Alternatively, one may wish to test for symmetry in a dataset of jet-quark pairs. In the language of conditional symmetry, given a collection of four-momenta of ℓ decay particles ($\mathbf{X} = \mathbb{R}^{\ell \times 4}$), the conditional probability that they decayed from a particular particle should be $O(1, 3)$ -invariant. In our experiments in Section 7.4, interest is in whether or not the particles decayed from a top quark, so $\mathbf{Y} = \{0, 1\}$. Since $Q(p)$ is a maximal invariant, it can be used as the conditioning variable. Moreover, since conditional invariance is being tested (the action of $O(1, 3)$ on \mathbf{Y} is trivial), the conditional independence test based on $X \perp\!\!\!\perp Y \mid M(X)$ can be used. \diamond

6 Kernel conditional independence test for equivariance

Theorem 3 establishes that a test for \mathbf{G} -equivariance of a conditional distribution can be formulated as a test for conditional independence of the form (14). Any non-parametric conditional independence test could be used for this purpose; see Li and Fan (2020) for a recent review of the literature. We use the kernel conditional independence (KCI) test (Zhang et al., 2011) in our experiments in Section 7.

In all of our experiments, either a representative inversion τ exists or the action of \mathbf{G} on \mathbf{Y} is trivial, so for simplicity we formulate the statistic for the simpler conditional independence test (13). Modification to accommodate random $\tilde{\tau}$ is straightforward, by replacing $\tau^{-1}(X_i)$ in the expression below with a sampled $\tilde{\tau}_i \sim \zeta(X_i, \bullet)$ and extending the kernel \mathbf{K}_{XM} as defined below on $\mathbf{X} \times \mathbf{M}$ to a kernel on $\mathbf{G} \times \mathbf{X} \times \mathbf{M}$.

The test statistic in the KCI test for equivariance is constructed as follows. Let k_X, k_Y and k_M be kernels on \mathbf{X}, \mathbf{Y} , and \mathbf{M} , respectively. Given data $(X, Y)_{1:n}$, define the kernel matrices $\mathbf{K}_Y, \mathbf{K}_M$ and \mathbf{K}_{XM} as

$$[\mathbf{K}_Y]_{ij} = k_Y(\tau(X_i)^{-1}Y_i, \tau(X_j)^{-1}Y_j), \quad [\mathbf{K}_M]_{ij} = k_M(M(X_i), M(X_j)), \quad [\mathbf{K}_{XM}]_{ij} = k_X(X_i, X_j) [\mathbf{K}_M]_{ij}.$$

Let $\bar{\mathbf{K}}_Y = \mathbf{H}\mathbf{K}_Y\mathbf{H}$ denote the centralized kernel matrix, where $\mathbf{H} = \mathbf{I}_n - n^{-1}\mathbf{1}_n$, and similarly for $\bar{\mathbf{K}}_M$ and $\bar{\mathbf{K}}_{XM}$. For fixed $\varepsilon > 0$, define the matrices $\mathbf{R}_M = \varepsilon(\bar{\mathbf{K}}_M + \varepsilon\mathbf{I}_n)^{-1}$, $\bar{\mathbf{K}}_{XM|M} = \mathbf{R}_M\bar{\mathbf{K}}_{XM}\mathbf{R}_M$, and $\bar{\mathbf{K}}_{Y|M} = \mathbf{R}_M\bar{\mathbf{K}}_Y\mathbf{R}_M$. Then the test statistic is given by

$$T_{\text{KCI}}(X_{1:n}, Y_{1:n}) = \frac{1}{n} \text{Tr}(\bar{\mathbf{K}}_{XM|M} \bar{\mathbf{K}}_{Y|M}).$$

The distribution of this test statistic under H_0 can be approximated by samples $T^{(1)}, \dots, T^{(B)}$ drawn through a simulation procedure described by Zhang et al. (2011), and the test rejects H_0 at level α if

$$\frac{1}{B} \sum_{b=1}^B \mathbb{1} \left\{ T_{\text{KCI}}(X_{1:n}, Y_{1:n}) \leq T^{(b)} \right\} \leq \alpha.$$

See Zhang et al. (2011) for more details and explanations.

As a point of comparison, we also implement a test for equivariance based on the conditional permutation (CP) test (Berrett et al., 2019) where kernel conditional density estimation (KCDE, De Gooijer and Zerom 2003) is used in the sampling procedure. We leave the details of this particular test to Appendix C.1.

7 Experiments

We evaluate our proposed tests on several synthetic and real data examples. In all of our experiments, we sample n data points from a distribution or dataset and perform a test for a specified symmetry. We repeat this procedure over $N = 1000$ simulations for each test and record the proportion of simulations in which the test rejected. The proportion of rejections is an estimate of the level of the test when the data distribution has the specified symmetry and otherwise an estimate of the power when the distribution does not. With $N = 1000$ simulations, estimates are precise up to approximately ± 0.016 . We set the desired significance level to be $\alpha = 0.05$ in all of our experiments. We use $m = 2$ sampled group actions except where otherwise specified. We default to Gaussian radial basis function kernels for data that are continuous and use the median distance heuristic (Garreau et al., 2017) for the kernel bandwidth unless otherwise specified. The median distance is computed from a ‘‘training’’ set of n data points randomly split from the ‘‘test’’ set used to estimate the rejection rate. The median distance is recomputed in every simulation.

	$\mathbf{G} = \text{SO}(4)$						$\mathbf{G} = \mathbb{S}_{10}$					
	Invariance			Equivariance			Invariance			Equivariance		
	H_0	H_1	$\hat{\beta}_{n,m}$	H_0	H_1		H_0^+	H_0^-	H_1	$\hat{\beta}_{n,m}$	H_0	H_1
2SMMD	0.041	0.870	0.778				0.012	0.052	0.987	1.000		
MMD	0.050	0.984	0.961				0.047	0.053	1.000	1.000		
NMMD	0.051	0.896	0.743				0.054	0.044	0.122	0.175		
CW	0.068	0.935	0.988				0.069	0.072	0.872	0.993		
KCI				0.041	0.921						0.075	0.997
CP				0.095	1.000						0.124	0.753

Table 1: Test rejection rates over $N = 1000$ simulations.

Let $\lceil x \rceil$ denote the ‘‘ceiling’’ function. The tests that we evaluate for invariance include: the transformation two-sample test based on the standard two-sample MMD (2SMMD, Section 4.1.1), the MMD test (MMD, Section 4.1), the MMD test with Nyström approximation using $J = \lceil \sqrt{n} \rceil$ subsamples (NMMD, Appendix B.1), and the CW test with $J = \lceil \sqrt{n} \rceil$ random projections (CW, Section 4.3). Where applicable, we use the sampling procedure described in Algorithm 1 with $B = 200$. We test for conditional symmetries using the KCI test (KCI, Section 6), and the CP test with KCDE (CP, Appendix C.1, $S = 50$ steps) using the multiple correlation coefficient (Abdi, 2007) as the test statistic. In most conditional symmetry experiments, we find that using the median heuristic for KCI and CP does not lead to meaningful results. In these cases, we select the kernel bandwidths by performing a grid search for each kernel. For each combination of bandwidths, we estimate the size and power of the test over 100 simulations involving training data. We choose the combination that leads to a rejection rate of at most 0.1 on data generated under H_0 and that maximizes rejection rate on data generated under H_1 . If no combination has rejection rate less than 0.1 under H_0 , we then use the combination that leads to the lowest rejection rate. Further details about the grid search for each experiment, along with other experiment details, can be found in Appendix C.

All experiments were implemented in the Julia programming language (version 1.6.1) and run with a single thread on a high-performance computing allocation with 4 CPU cores and 16GB of RAM.

7.1 Synthetic examples

We use \mathbf{N} to denote both univariate and multivariate Gaussian distributions. For multivariate distributions, the dimension d is given in the context of each experiment. We use 0_d (1_d) to denote a vector of zeros (ones) of length d .

7.1.1 Rotation

We first consider symmetries of random vectors $X_i \in \mathbb{R}^d$ with respect to the group $\mathbf{G} = \text{SO}(d)$ of d -dimensional rotations about the origin. We conduct the following experiments for $d = 4$:

- For tests for invariance, we estimate the test size and power by the average rejection rates over $N = 1000$ simulations, each of $n = 200$ i.i.d. samples. Samples were generated from $\mathbf{N}(0_d, \mathbf{I}_d)$ and from $\mathbf{N}(0.4\mathbf{e}_1, \mathbf{I}_d)$ for the null and alternative hypotheses, respectively.
- For tests for equivariance, we first generate samples $X_{1:n}$ of size $n = 50$ from $\mathbf{N}(0_d, \Sigma_d)$, where Σ_d is sampled from $\text{Wishart}(\mathbf{I}_d, d)$. For $i \in \{1, \dots, n\}$, we then generate Y_i given X_i from $\mathbf{N}(X_i, \mathbf{I}_d)$ to simulate the size of the test (rejection rate under the null), and from $\mathbf{N}(|X_i|, \mathbf{I}_d)$ to estimate the rejection rate under an alternative, where $|X_i|$ denotes the vector obtained by taking element-wise absolute values of X_i .

The estimated test sizes and powers are shown in Table 1. For all $\text{SO}(4)$ -symmetries, our tests reject at a rate of approximately $\alpha = 0.05$ for data generated under H_0 and reject at some noticeably higher rate for data generated under H_1 . We also compute power estimates $\hat{\beta}_{n,m}$ for a single dataset using the procedure described in Algorithm 2. The power estimates relatively align with the ‘‘gold-standard’’ estimates based on simulations. Finally, we validate Theorem 2 by checking the distribution of the estimated p -value p_B across the simulations. Histograms of those are shown in Figure 2. The resulting distributions were tested for uniformity using the Kolmogorov–Smirnov test; each plot of Figure 2 also displays the p -value of that test. (The 2SMMD p -values were estimated by standard bootstrap

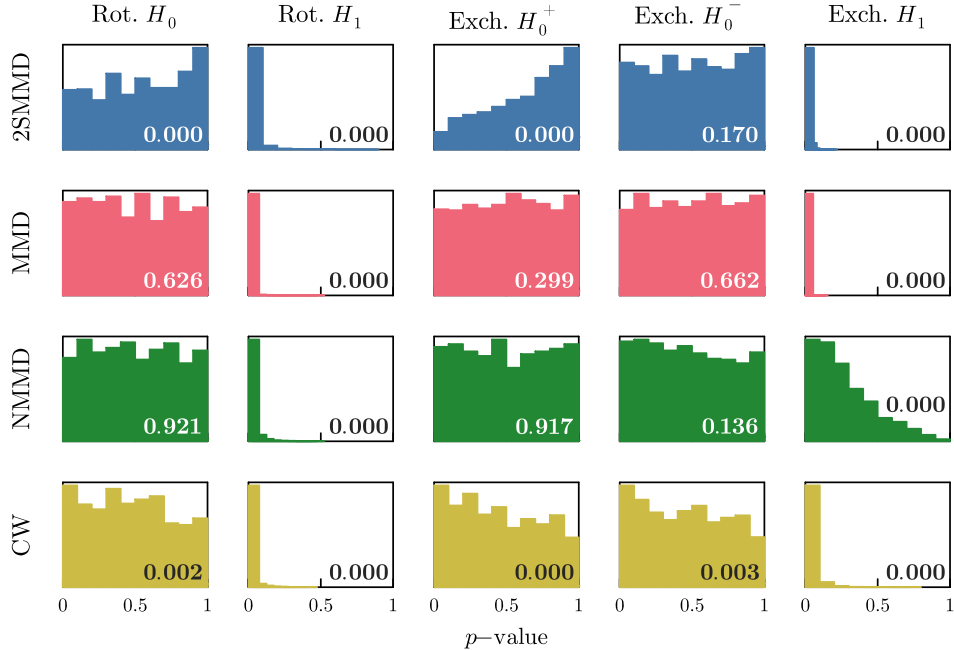


Figure 2: Histograms showing the p -value distributions obtained over $N = 1000$ simulations across different tests and data generation settings. For 2SMMD, the p -value is estimated by standard bootstrapping; for the others, it is estimated by the conditional Monte Carlo method in Algorithm 1. In each plot, the p -value of a Kolmogorov–Smirnov test for uniformity of the distribution is shown in the bottom-right corner.

resampling.) The MMD-based tests using the conditional Monte Carlo method from Algorithm 1 indeed appear to be sampling uniform p -values when H_0 is true.

We use the $SO(d)$ -invariance experiment above to investigate how the properties of these tests change with increasing dimensions d , sample sizes n , and sampled group actions m . When a variable is fixed while varying the others, we set $d = 4$, $n = 200$ and $m = 2$. Figure 3 shows the estimated levels, powers and average computation (wall) times in seconds for $d \in \{5, 10, 15, 20\}$, $n \in \{50, 100, 200, 400\}$, and $m \in \{1, 2, 3, 4, 5\}$. We find that among these tests for invariance, MMD achieves the best statistical performance as the dimension increases at the cost of increasing computation time. MMD also appears to be relatively more efficient in terms of power for small sample sizes. For smaller values of d , our randomized version of CW and the 2SMMD perform nearly as well, particularly for larger sample sizes, with less computation time. We also observe that sampling more than two group actions only leads to a marginal increase in power for most tests but results in an increase in computation time, most notably for MMD.

Separately, we also investigate a MMD test for $SO(3)$ -invariance based on representative inversions. Given an observation $X_i \in \mathbb{R}^3$, we take $\tau(X_i)$ to be the rotation matrix that maps the vector $\|X_i\|_2 \mathbf{e}_1$ to X_i , obtained as in Example 1. A random inversion is then calculated as $\tilde{\tau}_i = \tau(X_i)G_i$, where $G_i \sim \lambda_{\mathbf{G}_{\mathbf{e}_1}}$ is a uniform random sample from the stabilizer subgroup of \mathbf{e}_1 . We sample X_i from $N(0_3, \mathbf{I}_3)$ to estimate the rejection rate under H_0 , and from $N(0_3, \Sigma_3)$ for H_1 , where $\Sigma_3 \sim \text{Wishart}(\mathbf{I}_3, 3)$. We use a characteristic kernel on $SO(3)$ (Fukumizu et al., 2008),

$$k(R_1, R_2) = \frac{\pi\theta(\pi - \theta)}{8 \sin(\theta)}, \quad \text{with } \cos(\theta) = \frac{1}{2} \text{Tr}(R_2^{-1}R_1), \quad \text{for } 0 \leq \theta \leq \pi.$$

The rejection rates in the H_0 setting are 0.014, 0.038, and 0.058 for 2SMMD, MMD, and NMMD, respectively; the rejection rates in the H_1 setting are 0.929, 0.976, and 0.222. Histograms of the p -value distributions are shown in Appendix C.2.

7.1.2 Exchangeability

We next consider symmetries with respect to the group $\mathbf{G} = \mathbb{S}_d$ of permutations acting on \mathbb{R}^d . We conduct the following experiments for $d = 10$:

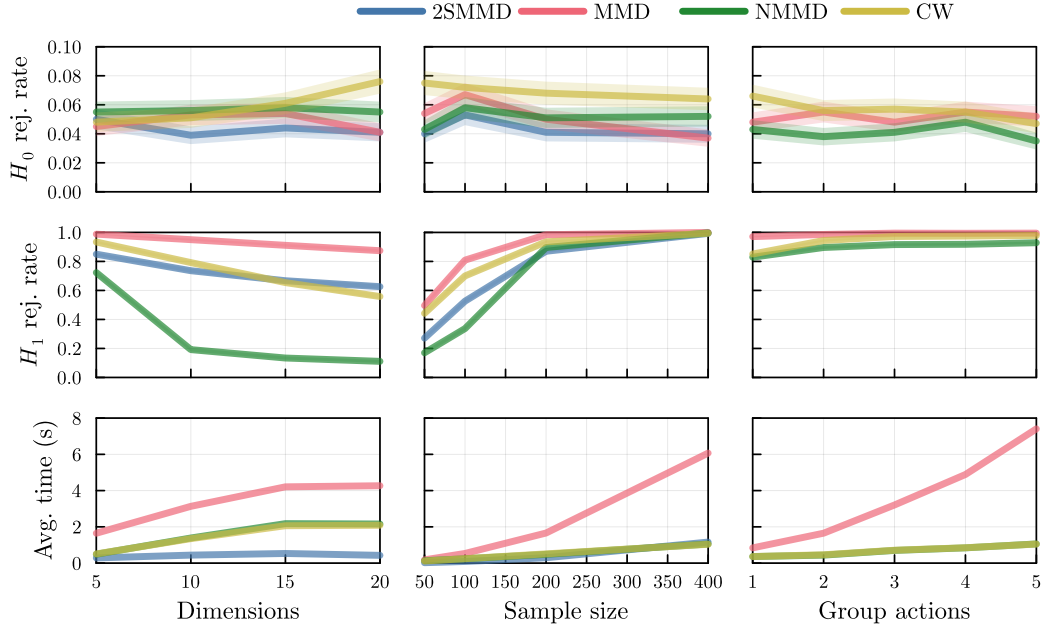


Figure 3: First and second row: Test rejection rates and standard deviations over $N = 1000$ simulations with one of dimensions d , sample size n , or number of group actions m increasing and the others fixed ($d = 4$, $n = 200$, $m = 2$). Third row: Average computation time (in seconds) for a single execution of the test.

- For tests for invariance, in each simulation iteration, we generate i.i.d. samples of size $n = 200$ from three distributions: $N(0_d, \Sigma_d^+)$, where Σ_d^+ is the $d \times d$ matrix with 1 on the diagonal and $1/d$ on the off-diagonals; $N(0_d, \Sigma_d^-)$, where Σ_d^- is similar to Σ_d^+ but with $-1/(d-1)$ on the off-diagonals; and $N(0_d, \Sigma_d)$, where Σ_d is randomly sampled from $\text{Wishart}(I_d, d)$. The first two distributions are \mathbb{S}_d -invariant (Aldous, 1985, pp. 7-8), while the third distribution is almost surely not. We refer to the three settings as H_0^+ , H_0^- , and H_1 , respectively.
- For tests for equivariance, we generate samples $X_{1:n}$ of size $n = 100$ from $N(0_d, \Sigma_d)$, where Σ_d is sampled from $\text{Wishart}(I_d, d)$. For $i \in \{1, \dots, n\}$, we then generate Y_i from $N(X_i, I_d)$ for H_0 and from $N(X_i^\top \mathbf{e}_1 \mathbf{1}_d, I_d)$ for H_1 . For these tests, we select the kernel bandwidths via the grid search procedure described in Section 7, over the grid $\{10^{-6}, \dots, 10^2\}$ for all kernels.

The results of the above tests are shown in Table 1. As in the case of testing $\text{SO}(4)$ -symmetries, most of our tests for \mathbb{S}_{10} -symmetries have rejection rates relatively close to $\alpha = 0.05$ when H_0 is true and otherwise higher rates when H_1 is true. Power estimates are again largely consistent with the simulated rejection rates. Finally, the estimated p -values shown in Figure 2 validate Theorem 2 in this setting as well. In Appendix C.3, we find that increasing the number of random projections in NMMD to $J = 25$ boosts the test power to that comparable with CW.

7.2 SWARM geomagnetic satellite data

The first data application that we consider features the geomagnetic field data collected by the European Space Agency as part of the SWARM constellation mission that was first launched in late 2013 (Olsen et al., 2013). The source dataset includes geomagnetic field measurements (in units of nT) recorded in 1-second intervals from three satellites (labelled A, B, C) that orbit the earth several times a day. The dataset also includes the latitude, longitude, and the radius from the center of the earth at which the measurements were collected. A magnetic dipole model, which is invariant to rotations about the dipole axis (which intersects Earth at the geomagnetic poles), can be used to approximate the geomagnetic field. However, significant deviations from the model—and symmetry—can occur due to variations in mineral composition of Earth’s crust, the effects of solar wind, and other causes.

Following Christie and Aston (2023), we consider only the data collected by satellite A on February 25th, 2023. The data consists of 86,400 measurements collected along satellite orbit trajectories around Earth. We randomly partition




Figure 4: Two-dimensional projections of $n = 220$ 3D Cartesian data points sampled from the SWARM dataset. The label $X_a : X_b$ indicates the projection onto the plane spanning dimensions a and b . The magnetic field strength is represented by the colour; the darker the colour, the greater the measured value.

the 86,400 data points into a training set and a test set of equal size. We transform the latitude, longitude, and radius positions into Cartesian coordinates, and apply a rescaling so that the maximum norm of any observed coordinate vector is 1. We also standardize the geomagnetic field measurements. We take X_i to be the Cartesian coordinates and Y_i to be the standardized field intensity at X_i . A single sample of $n = 220$ data points is visualized in Figure 4. From the visualization, the magnetic field intensity appears to be nearly invariant with respect to rotations around the axis through the poles (represented as the X_3 -axis), though some deviations from invariance are evident.

We test for conditional $\text{SO}(2)$ -invariance of Y_i given X_i with respect to each of the three axes in \mathbb{R}^3 using samples of size $n = 220$ that are sampled (without replacement) from the test set. The rejection rates are 0.464, 0.273, and 0.127, where the last rate corresponds to the rotations about the geographic north pole. We also examine conditional invariance with respect to rotations about the geomagnetic north pole,⁷ for which KCI rejects at a rate of 0.223. Figure 5 shows the distributions of the KCI p -values along with the p -value 0.075 reported by [Christie and Aston \(2023\)](#) for their own test for invariance. For this experiment, we were unable to tune CP to produce meaningful results.

We also test for the conditional invariance of $P_{Y_i|X_i}$ by testing for marginal invariance in X_i and joint invariance in (X_i, Y_i) . Figure 4 shows that the satellite orbited the earth approximately 15 times, and the orbit trajectories are spaced approximately 24 degrees apart, intersecting at the poles. The results are shown in Appendix C.4. We find that the tests generally do not reject marginal invariance of X_i with respect to both discrete 24-degree rotations and continuous rotations about the geographic north pole. However, MMD and CW both reject joint invariance with respect to discrete/continuous rotations in X_i and non-transformations in Y_i at a significantly higher rate. Altogether, the tests indicate that $P_{Y_i|X_i}$ is not conditionally invariant with respect to rotations around any of the tested axes.

Note that by treating the (X_i, Y_i) observations as i.i.d., we are implicitly assuming that $X_{1:n}$ are i.i.d. and that $Y_{1:n}$ are conditionally independent given $X_{1:n}$. The former assumption is justified by the sampling process. The conditional independence of $Y_{1:n}$ is more difficult to justify, but deviations from that should be confined to relatively small local regions. If one is unwilling to make the assumption, the data could be treated as a matrix X of Cartesian coordinates and a vector Y of corresponding geomagnetic field values. A test for conditional invariance is possible by testing based on the maximal invariant $M(X) = X^\top X$.

7.3 Large Hadron Collider dijet events

The second application that we examine is based on the Large Hadron Collider (LHC) Olympics 2020 dataset ([Kasieczka et al., 2021](#)). The LHC dataset consists of 1.1 million simulated dijet events generated by PYTHIA ([Bierlich et al., 2022](#)), a widely-used Monte Carlo generator for high-energy physics processes. A dijet event is two jets of particles that are produced by the collision of subatomic particles. The transverse momentum, polar angle ϕ and pseudorapidity η for up to 200 jet constituents were recorded for each jet. The Cartesian momentum of a particle in the transverse plane is represented by the pair

$$p_x = p_T \cos(\phi), \quad p_y = p_T \sin(\phi).$$

⁷The geomagnetic poles correspond to the axis through Earth for which a dipole approximation (which is $\text{SO}(2)$ -invariant) of the geomagnetic field is best-fitting ([World Data Center for Geomagnetism, Kyoto University](#)).




Figure 5: Histograms showing the KCI p -value distributions across $N = 1000$ simulations for testing conditional invariance with respect to $SO(2)$ -rotations about the three axes and the geomagnetic north pole for the SWARM data. The orange line is the p -value (0.075) reported by [Christie and Aston \(2023\)](#) for their test for invariance.

The leading constituent in a jet is the particle with the largest transverse momentum in any direction. In our analyses, we focus on the joint distribution of the two constituents with the largest transverse momenta in each event (after [Desai et al., 2021](#)). A single observation is therefore a 4D vector $X = (p_{1x}, p_{1y}, p_{2x}, p_{2y})$, where p_1 and p_2 correspond to the momenta of the two leading particles, respectively. As in the previous experiment, we randomly split the dataset into a training set and a test set of equal size. We draw samples of size $n = 100$ in all of the following experiments. Histograms of p -value distributions obtained from the tests can be found in Appendix C.5.

7.3.1 Joint invariance

By conservation of angular momentum, the distribution of the Cartesian momenta of the two leading particles across jet events should be invariant to simultaneous rotations by the same angle, i.e., with respect to the subgroup $\mathbf{G}_0 = \{(g_1, g_2) \in SO(2) \times SO(2) : g_1 = g_2\}$. We conduct tests for invariance with respect to this subgroup, as well as with respect to the full $\mathbf{G}_1 = SO(2) \times SO(2)$ group, and to $\mathbf{G}_2 = SO(4)$. The results are shown in Table 2. We see that 2SMMD, MMD, and CW are able to identify \mathbf{G}_0 -invariance and correctly reject \mathbf{G}_1 - and \mathbf{G}_2 -invariance at a higher rate. In Appendix C.5, we find that increasing the number of random projections from 10 to 15 significantly improves the power of NMMD.

7.3.2 Conditional equivariance

By taking $X_i = (p_{1x}, p_{1y})$ and $Y_i = (p_{2x}, p_{2y})$, invariance of the 4D vector with respect to the subgroup \mathbf{G}_0 can also be viewed as Y_i being conditionally equivariant with respect to $SO(2)$ given X_i . We perform a test for $SO(2)$ -equivariance. We obtain rejection rates 0.0 for KCI and 0.051 for CP in this setting. We also perform a test for conditional $SO(2)$ -invariance, which KCI correctly rejects with rate 1.0 and CP with rate 0.997.

7.4 Top quark tagging

We consider a second particle physics application based on the Top Quark Tagging Reference dataset ([Kasieczka et al., 2019](#)), which also consists of jet events simulated by PYTHIA. The original dataset was constructed for the task of classifying jet events as having decayed from a top quark and consists of a training, validation, and test set. We only use the test set, which contains 404,000 simulated jet events. The four-momenta $p = (E, p_x, p_y, p_z)$ of up to 200 jet constituents are recorded for each event. Each event is also labelled as 1 or 0, representing that the jet decayed from a top quark or did not, respectively. As in Example 5, according to the Standard Model, when predicting whether a jet is the decay of a top quark based on the four-momenta of jet constituents, the distribution of the label should be conditionally invariant with respect to the Lorentz group $O(1, 3)$, which consists of spatial rotations and relativistic boosts. According to Theorem 3, conditional invariance is equivalent to $X \perp\!\!\!\perp Y \mid M(X)$ in this scenario.

For convenience, we take the data to be the four-momenta $X_i = (p_1, p_2)$ of the two leading constituents in each jet (as in [Yang et al., 2023](#)) and the top quark label $Y_i \in \{0, 1\}$. We split the data into a training and test set. We perform a test for conditional invariance of Y_i given X_i with respect to the Lorentz group based on samples of size $n = 200$. In our tests, we use the 2D maximal invariant $M(X_i) = (Q(p_{1i}), Q(p_{2i}))$. For the kernel on $\mathbf{Y} = \{0, 1\}$, we use the kernel $k_Y(x, y) = \mathbb{1}(x = y)$. KCI rejects conditional invariance at a rate of 0.029, which is consistent with the theory of the Standard Model. To verify that KCI is identifying symmetry in a meaningful way, we simulate new labels Y'_i given X_i

	$\mathbf{G}_0 = \{\text{paired SO}(2)\text{-rotations}\}$	$\mathbf{G}_1 = \text{SO}(2) \times \text{SO}(2)$	$\mathbf{G}_2 = \text{SO}(4)$
2sMMD	0.035	0.967	0.983
MMD	0.038	1.000	1.000
NMMD	0.058	0.241	0.214
CW	0.052	0.971	0.999

Table 2: Test rejection rates over $N = 1000$ simulations for the LHC data.

using the model

$$Y_i' \mid X_i \sim \text{Bernoulli}(0.9\mathbb{1}\{E \geq 200\} + 0.1\mathbb{1}\{E < 200\}) .$$

With the new labels, KCI rejects conditional invariance with respect to the Lorentz group at a rate of 0.781. Distributions of the KCI p -values can be found in Appendix C.6. We were unable to tune CP to produce meaningful results.

Acknowledgments

This research was supported in part through computational resources and services provided by Advanced Research Computing at the University of British Columbia. KC and BBR gratefully acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC): RGPIN2020-04995, RGPAS-2020-00095, DGECR-2020-00343.

Appendices

A Proofs of main results

This section contains the proofs for Theorems 1 to 3.

A.1 Proof of Theorem 1

Proof. The results follow easily from Proposition 2 and the fact that $D(\hat{P}_n, \hat{P}_{n,m}^\square)$ converges almost surely (with respect to the product measure $P \otimes \lambda$) to $D(P, P^\circ)$ by the strong law of large numbers. In particular, if $P \in \mathcal{P}^\circ(\mathbf{X})$, then $P = P^\circ$ and so $D(P, P^\circ) = 0$. Since D is continuous, it follows that

$$\lim_{n \rightarrow \infty} D(\hat{P}_n, \hat{P}_{n,m}^\square) = 0, \quad P \otimes \lambda\text{-a.s.}, \quad \text{for any } P \in \mathcal{P}^\circ(\mathbf{X}).$$

Therefore, if $c_n \rightarrow c \geq 0$, then

$$\lim_{n \rightarrow \infty} \mathbb{E}_{P \otimes \lambda}[\phi_{n,m}(X_{1:n})] = \lim_{n \rightarrow \infty} \mathbb{E}_{P \otimes \lambda}[\mathbb{1}\{D(\hat{P}_n, \hat{P}_{n,m}^\square) > c_n\}] = 0,$$

from which (6) follows.

On the other hand, if $P \in \mathcal{P}^\times(\mathbf{X})$, then $P \neq P^\circ$ and therefore $D(P, P^\circ) > 0$. If $c_n \rightarrow 0$, then

$$\lim_{n \rightarrow \infty} \mathbb{E}_{P \otimes \lambda}[\phi_{n,m}(X_{1:n})] = \lim_{n \rightarrow \infty} \mathbb{E}_{P \otimes \lambda}[\mathbb{1}\{D(\hat{P}_n, \hat{P}_{n,m}^\square) > c_n\}] = 1,$$

from which (7) follows. □

A.2 Proof of Theorem 2

The proof of Theorem 2 relies on the following result proven by [Dufour \(2006\)](#). For simplicity, we assume that the probability of ties is zero, but that case can be handled with a randomized tie-breaking procedure described by [Dufour \(2006\)](#).

Lemma 1 (Dufour (2006), Proposition 2.2). *Let S_0, S_1, \dots, S_B be an exchangeable sequence of \mathbb{R} -valued random variables such that $\Pr\{S_i = S_j\} = 0$ for $i \neq j, i, j \in \{0, \dots, B\}$. Set*

$$p_B = \frac{1 + \sum_{b=1}^B \mathbb{1}\{S_b \geq S_0\}}{B + 1}.$$

Then for any $\alpha \in [0, 1]$,

$$\Pr\{p_B \leq \alpha\} = \frac{\lfloor \alpha(B + 1) \rfloor}{B + 1}.$$

Proof of Theorem 2. Due to the sufficiency of $\gamma(X)$ for $\mathcal{P}^\circ(\mathbf{X})$, the samples $X_{1:n}^{(b)} = (G_1^{(b)} X_1, \dots, G_n^{(b)} X_n)$ are conditionally i.i.d. given $\gamma(X)_{1:n}$, with the same conditional distribution as the null conditional distribution. Because of their independence from $X_{1:n}$, conditioning on $(G_{j,1}, \dots, G_{j,n})_{j=1}^m$ does not change that, and therefore $(T_{n,m}(X_{1:n}^{(b)}))_{b=0}^B$ are conditionally i.i.d. given $\gamma(X)_{1:n}$ and $(G_{j,1}, \dots, G_{j,n})_{j=1}^m$, with the same conditional distribution as the null conditional distribution. The sequence $(X_{1:n}^{(b)})_{b=0}^B$ is easily seen to be exchangeable (over the index b) conditioned on $\gamma(X)_{1:n}$ and $(G_{j,1}, \dots, G_{j,n})_{j=1}^m$, and therefore so is $(T_{n,m}(X_{1:n}^{(0)}), \dots, T_{n,m}(X_{1:n}^{(B)}))$. The validity of p_B as a conditional (on $\gamma(X)_{1:n}$) p -value, and (10) in particular, follows from Lemma 1. Since this holds for P -almost every realization of $\gamma(X)_{1:n}$ under each $P \in H_0$, it is also a valid p -value conditioned only on $(G_{j,1}, \dots, G_{j,n})_{j=1}^m$.

The proof remains valid unconditionally if $(G_{j,1}, \dots, G_{j,n})_{j=1}^m$ are sampled independently of $X_{1:n}$, so that $(X_{1:n}^{(b)}, (G_{j,1}, \dots, G_{j,n})_{j=1}^m)_{b=0}^B$ are exchangeable and therefore so is $(T_{n,m}^{(0)}(X_{1:n}^{(0)}), \dots, T_{n,m}^{(B)}(X_{1:n}^{(B)}))$.

The proof also applies unconditionally to random $(G_{i,j}^{(b)})_{i \leq n, j \leq m}$ sampled independently of $X_{1:n}$ in a way such that they are exchangeable over the index $b = 1, \dots, B$, in which case the sequence $(X_{1:n}^{(b)}, (G_{j,1}^{(b)}, \dots, G_{j,n}^{(b)})_{j=1}^m)_{b=0}^B$ is exchangeable and therefore so is $(T_{n,m}^{(0)}(X_{1:n}^{(0)}), \dots, T_{n,m}^{(B)}(X_{1:n}^{(B)}))$. \square

A.3 Proof of Theorem 3

Proof. To simplify notation, let $Q: \mathbf{X} \times \mathbf{S}_Y \rightarrow [0, 1]$ be a regular version (i.e., a Markov probability kernel) of the conditional probability $P_{Y|X}$, and denote the marginal distribution of X by P , so that $P_{X,Y} = P_X \otimes P_{Y|X} = P \otimes Q$. Define the random variable $\tilde{Y} := \tilde{\tau} X$, where $\tilde{\tau} \sim \zeta(X, \bullet)$. The conditional distribution of \tilde{Y} given $(\tilde{\tau}, X)$ is represented by the Markov probability kernel \tilde{Q} so that for any integrable function $f: \mathbf{G} \times \mathbf{X} \times \mathbf{Y} \rightarrow \mathbb{R}$,

$$\int P(dx) \zeta(x, d\tilde{\tau}) \tilde{Q}(\tilde{\tau}, x, d\tilde{y}) f(\tilde{\tau}, x, \tilde{y}) = \int P(dx) \zeta(x, d\tilde{\tau}) Q(x, d\tilde{y}) f(\tilde{\tau}, x, \tilde{\tau}^{-1}y).$$

From this follows the identity $\tilde{Q}(\tilde{\tau}, x, B) = Q(x, \tilde{\tau}B)$.

Now assume that Q is equivariant, so that for each $g \in \mathbf{G}, x \in \mathbf{X}, B \in \mathbf{S}_Y, Q(gx, B) = Q(x, g^{-1}B)$. Then for any $\tilde{\tau} \in \mathbf{G}, x \in \mathbf{X}, g \in \mathbf{G}$ and integrable $f: \mathbf{Y} \rightarrow \mathbb{R}$,

$$\begin{aligned} \int \tilde{Q}(g\tilde{\tau}, gx, d\tilde{y}) f(\tilde{y}) &= \int Q(gx, d\tilde{y}) f((g\tilde{\tau})^{-1}\tilde{y}) \\ &= \int Q(x, d\tilde{y}) f(\tilde{\tau}^{-1}g^{-1}\tilde{y}) \\ &= \int Q(x, d\tilde{y}) f(\tilde{\tau}^{-1}\tilde{y}) \\ &= \int \tilde{Q}(\tilde{\tau}, x, d\tilde{y}) f(\tilde{y}). \end{aligned}$$

This shows that the mapping $(\tilde{\tau}, x) \mapsto \tilde{Q}(\tilde{\tau}, x, \bullet)$ is \mathbf{G} -invariant. Therefore, by Proposition 1, for any measurable maximal invariant $\tilde{M}: \mathbf{G} \times \mathbf{X} \rightarrow \mathbf{M}$, there is a unique Markov probability kernel $\tilde{R}: \mathbf{M} \times \mathbf{S}_Y \rightarrow [0, 1]$ such that

$$\tilde{Q}(\tilde{\tau}, x, B) = \tilde{R}(\tilde{M}(\tilde{\tau}, x), B), \quad \tilde{\tau} \in \mathbf{G}, x \in \mathbf{X}, B \in \mathbf{S}_Y.$$

Because the action of \mathbf{G} on itself is transitive (i.e., there is only one orbit in \mathbf{G}), any maximal invariant M for \mathbf{G} acting on \mathbf{X} is also a maximal invariant for \mathbf{G} acting on $\mathbf{G} \times \mathbf{X}$, and

$$\tilde{Q}(\tilde{\tau}, x, B) = \tilde{R}(M(x), B), \quad \tilde{\tau} \in \mathbf{G}, x \in \mathbf{X}, B \in \mathbf{S}_{\mathbf{Y}}. \quad (15)$$

This is enough to establish the desired conditional independence in (14): For any integrable $f: \mathbf{G} \times \mathbf{X} \times \mathbf{Y} \rightarrow \mathbb{R}$,

$$\int P(dx) \zeta(x, d\tilde{\tau}) \tilde{Q}(\tilde{\tau}, x, d\tilde{y}) f(\tilde{\tau}, x, \tilde{y}) = \int P(dx) \zeta(x, d\tilde{\tau}) \tilde{R}(M(x), d\tilde{y}) f(\tilde{\tau}, x, \tilde{y}).$$

Conversely, assume that $(\tilde{\tau}, X) \perp\!\!\!\perp \tilde{\tau}^{-1}Y \mid M(X)$. Then (15) holds for $P \otimes \zeta$ -almost all $(\tilde{\tau}, x) \in \mathbf{G} \times \mathbf{X}$. In particular, \tilde{Q} is \mathbf{G} -invariant for $P \otimes \zeta$ -almost all $(\tilde{\tau}, x)$. In particular, \tilde{Q} is \mathbf{G} -invariant for $P \otimes \zeta$ -almost all $(\tilde{\tau}, x)$. Recall also that the inversion kernel ζ is \mathbf{G} -equivariant. Therefore, for any integrable $f: \mathbf{G} \times \mathbf{X} \times \mathbf{Y} \rightarrow \mathbb{R}$ and any $g \in \mathbf{G}$,

$$\begin{aligned} \int P(dx) \zeta(x, d\tilde{\tau}) Q(x, dy) f(\tilde{\tau}, x, y) &= \int P(dx) \zeta(x, d\tilde{\tau}) Q(x, dy) f(\tilde{\tau}, x, \tilde{\tau}(\tilde{\tau}^{-1}y)) \\ &= \int P(dx) \zeta(x, d\tilde{\tau}) \tilde{Q}(\tilde{\tau}, x, d\tilde{y}) f(\tilde{\tau}, x, \tilde{\tau}\tilde{y}) \\ &= \int P(dx) \zeta(x, d\tilde{\tau}) \tilde{Q}(g\tilde{\tau}, gx, d\tilde{y}) f(\tilde{\tau}, x, \tilde{\tau}\tilde{y}) \\ &= \int (g_*P)(dx) \zeta(g^{-1}x, d\tilde{\tau}) \tilde{Q}(g\tilde{\tau}, x, d\tilde{y}) f(\tilde{\tau}, g^{-1}x, \tilde{\tau}\tilde{y}) \\ &= \int (g_*P)(dx) \zeta(x, d\tilde{\tau}) \tilde{Q}(\tilde{\tau}, x, d\tilde{y}) f(g^{-1}\tilde{\tau}, g^{-1}x, g^{-1}\tilde{\tau}\tilde{y}) \\ &= \int (g_*P)(dx) \zeta(x, d\tilde{\tau}) Q(x, dy) f(g^{-1}\tilde{\tau}, g^{-1}x, g^{-1}y) \\ &= \int P(dx) \zeta(gx, d\tilde{\tau}) Q(gx, dy) f(g^{-1}\tilde{\tau}, x, g^{-1}y) \\ &= \int P(dx) \zeta(x, d\tilde{\tau}) Q(gx, dy) f(\tilde{\tau}, x, g^{-1}y). \end{aligned}$$

This implies that

$$Q(x, B) = Q(gx, gB), \quad B \in \mathbf{S}_{\mathbf{Y}}, g \in \mathbf{G}, P\text{-a.e. } x \in \mathbf{X}. \quad (16)$$

The subset of \mathbf{X} for which (16) holds is a \mathbf{G} -invariant set (Kallenberg, 2017, Lemma 7.7), and therefore the possible exceptional null set on which Q is not equivariant does not depend on g . If there is such an exceptional null set on which Q is not equivariant, denoted N^\times , define Q' as

$$Q'(x, B) := \begin{cases} Q(x, B) & \text{if } x \notin N^\times \\ \int_{\mathbf{G}} \zeta(x, d\tilde{\tau}) Q(\tilde{\tau}^{-1}x, \tilde{\tau}^{-1}B) & \text{if } x \in N^\times. \end{cases}$$

Since $\zeta(x, \bullet)$ and $Q(x, \bullet)$ are probability kernels, so too is Q' . It is also straightforward to show that Q' is \mathbf{G} -equivariant, so that Q' is another regular version of $P_{Y|X}$ that is \mathbf{G} -equivariant for all $x \in \mathbf{X}$, and equivalent to Q up to the null set N^\times .

If there exists a measurable representative inversion (function) τ , then the same proof holds with the inversion kernel $\zeta(x, \bullet)$ substituted by $\delta_{\tau(x)}$, resulting in the simplified conditional independence statement in (13).

If the action of \mathbf{G} on \mathbf{Y} is trivial, then $\tilde{Y} = Y$. Moreover, $\tilde{\tau} \perp\!\!\!\perp Y \mid X$ by construction, and therefore $(\tilde{\tau}, X) \perp\!\!\!\perp Y \mid M(X)$ is implied by $X \perp\!\!\!\perp Y \mid M(X)$. \square

B Alternative maximum mean discrepancy tests

In this section, we describe two variations of the MMD test for invariance that have cheaper computational costs.

B.1 Nyström approximation MMD test

The Nyström approximation (Raj et al., 2017; Chatalic et al., 2022) can be used to obtain an approximate MMD test based on the *biased* MMD test statistic, which is a V-statistic of the form

$$\begin{aligned} \widehat{\text{MMD}}_{\text{V}}^{\square}(\hat{P}_{1,n_1}, \hat{P}_{2,n_2}) &= \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \left(k(X_i, X_j) + \frac{1}{m^2} \sum_{\ell=1}^m \sum_{r=1}^m k(G_{\ell}X_i, H_rX_j) - \frac{2}{m} \sum_{\ell=1}^m k(X_i, G_{\ell}X_j) \right) \\ &= \frac{1}{n^2} \left(\mathbf{1}_n^{\top} \mathbf{K} \mathbf{1}_n + \frac{1}{m^2} \sum_{\ell=1}^m \sum_{r=1}^m \mathbf{1}_n^{\top} \mathbf{K}_{\ell r}^{(2)} \mathbf{1}_n - \frac{2}{m} \sum_{\ell=1}^m \mathbf{1}_n^{\top} \mathbf{K}_{\ell}^{(1)} \mathbf{1}_n \right), \end{aligned}$$

where the kernel matrices are defined as

$$[\mathbf{K}]_{ij} = k(X_i, X_j), \quad [\mathbf{K}_{\ell r}^{(2)}]_{ij} = k(G_{\ell}X_i, H_rX_j), \quad [\mathbf{K}_{\ell}^{(1)}]_{ij} = k(X_i, G_{\ell}X_j).$$

Nyström approximates the original kernel matrices with matrix products involving J -dimensional random matrices. For $J \ll n$, let \mathbf{t} be J points sampled independently and uniformly with replacement from $\mathbf{x} := X_{1:n}$, and similarly for \mathbf{t}^G from (GX_1, \dots, GX_n) . Applying Nyström approximation to the MMD leads to the test statistic

$$\widehat{\text{MMD}}_{\text{N}}^{\square}(\hat{P}_{1,n_1}, \hat{P}_{2,n_2}) = \psi_{\mathbf{t}}^{\top} \mathbf{K}_{\mathbf{t}, \mathbf{t}} \psi_{\mathbf{t}} + \frac{1}{m^2} \sum_{\ell=1}^m \sum_{r=1}^m \psi_{\mathbf{t}^G_{\ell}}^{\top} \mathbf{K}_{\mathbf{t}^G_{\ell}, \mathbf{t}^G_{r}} \psi_{\mathbf{t}^G_{r}} - \frac{2}{m} \sum_{\ell=1}^m \psi_{\mathbf{t}}^{\top} \mathbf{K}_{\mathbf{t}, \mathbf{t}^G_{\ell}} \psi_{\mathbf{t}^G_{\ell}},$$

where $\mathbf{K}_{\bullet, \bullet}$ denotes the kernel matrix between two sets of points and

$$\psi_{\bullet} = \frac{1}{n} \mathbf{K}_{\bullet, \bullet}^{+} \mathbf{K}_{\bullet, \mathbf{x}} \mathbf{1}_n,$$

with $+$ denoting the Moore-Penrose inverse.

B.2 MMD test with equivariant kernel

If the kernel k is (almost) equivariant in the sense of (12), then the MMD metric in the test for invariance has the form

$$\begin{aligned} &\text{MMD}(P, P^{\circ}) \\ &= \langle \mu_P, \mu_P \rangle_{\mathcal{H}} + \int_{\mathbf{X} \times \mathbf{X}} \int_{\mathbf{G} \times \mathbf{G}} k(gx, hx') \lambda(dg) \lambda(dh) P(dx) P(dx') - 2 \int_{\mathbf{X} \times \mathbf{X}} \int_{\mathbf{G}} k(x, gx') \lambda(dg) P(dx) P(dx') \\ &= \langle \mu_P, \mu_P \rangle_{\mathcal{H}} + \int_{\mathbf{X} \times \mathbf{X}} \int_{\mathbf{G} \times \mathbf{G}} k(x, ghx') \lambda(dg) \lambda(dh) P(dx) P(dx') - 2 \int_{\mathbf{X} \times \mathbf{X}} \int_{\mathbf{G}} k(x, gx') \lambda(dg) P(dx) P(dx') \\ &= \langle \mu_P, \mu_P \rangle_{\mathcal{H}} + \int_{\mathbf{X} \times \mathbf{X}} \int_{\mathbf{G}} k(x, g'x') \lambda(dg') P(dx) P(dx') - 2 \int_{\mathbf{X} \times \mathbf{X}} \int_{\mathbf{G}} k(x, gx') \lambda(dg) P(dx) P(dx') \\ &= \langle \mu_P, \mu_P \rangle_{\mathcal{H}} - \int_{\mathbf{X} \times \mathbf{X}} \int_{\mathbf{G}} k(x, gx') \lambda(dg) P(dx) P(dx'), \end{aligned}$$

where the second equality follows from the equivariance of the kernel. An unbiased estimator for the simplified metric is then

$$\widehat{\text{MMD}}^{\square}(\hat{P}_{1,n_1}, \hat{P}_{2,n_2}) = \frac{1}{n(n-1)} \sum_{i \neq j} \left(k(X_i, X_j) - \frac{1}{m} \sum_{\ell=1}^m k(X_i, G_{\ell}X_j) \right).$$

When the kernel satisfies (12), the KME of the orbit-averaged distribution can also be interpreted as the orbit-averaged KME of the original distribution (Elesedy, 2021, Lemma C.2), i.e., an invariant function. We do not enforce this assumption in this work as our tests for invariance are able to operate without it.

C Additional experiment details

In this section, we provide additional details about the experiments conducted in Section 7.

C.1 Conditional permutation test with kernel conditional density estimation

Let $T_{\text{CP}}: \mathbf{X}^n \times \mathbf{Y}^n \times \mathbf{M}^n \rightarrow \mathbb{R}$ be any statistic. (We use the multiple correlation coefficient of X and Y in our experiments in Section 7.) Let k_Y and k_M be kernels on \mathbf{Y} and \mathbf{M} . Given data $X_{1:n}$ and $Y_{1:n}$, let $Z_{1:n} := (\tau(X)^{-1}Y)_{1:n}$ to simplify notation. Let $Z_{\pi_0(1:n)} := Z_{1:n}$. On iteration s , we sample $\lfloor n/2 \rfloor$ disjoint pairs of indices $(i_1^{(s)}, j_1^{(s)}), \dots, (i_{\lfloor n/2 \rfloor}^{(s)}, j_{\lfloor n/2 \rfloor}^{(s)})$ from $\{1, \dots, n\}$. For each pair $(i_\ell^{(s)}, j_\ell^{(s)})$, we independently perform a swap of the $i_\ell^{(s)}$ -th and $j_\ell^{(s)}$ -th observations with probability $p_\ell^{(s)}$ obtained from the KCDE-estimated conditional density ratio

$$\begin{aligned} \frac{p_\ell^{(s)}}{1 - p_\ell^{(s)}} &= \frac{\hat{f}_{\text{KCDE}} \left(Z_{j_\ell^{(s)}}^{(s-1)} \mid M(X_{i_\ell^{(s)}}) \right) \hat{f}_{\text{KCDE}} \left(Z_{i_\ell^{(s)}}^{(s-1)} \mid M(X_{j_\ell^{(s)}}) \right)}{\hat{f}_{\text{KCDE}} \left(Z_{i_\ell^{(s)}}^{(s-1)} \mid M(X_{i_\ell^{(s)}}) \right) \hat{f}_{\text{KCDE}} \left(Z_{j_\ell^{(s)}}^{(s-1)} \mid M(X_{j_\ell^{(s)}}) \right)} \\ &= \frac{\left\{ \sum_{r=1}^n k_Y \left(Z_{j_\ell^{(s)}}^{(s-1)}, Z_r \right) k_M \left(M(X_{i_\ell^{(s)}}), M(X_r) \right) \right\} \left\{ \sum_{r=1}^n k_Y \left(Z_{i_\ell^{(s)}}^{(s-1)}, Z_r \right) k_M \left(M(X_{j_\ell^{(s)}}), M(X_r) \right) \right\}}{\left\{ \sum_{r=1}^n k_Y \left(Z_{i_\ell^{(s)}}^{(s-1)}, Z_r \right) k_M \left(M(X_{i_\ell^{(s)}}), M(X_r) \right) \right\} \left\{ \sum_{r=1}^n k_Y \left(Z_{j_\ell^{(s)}}^{(s-1)}, Z_r \right) k_M \left(M(X_{j_\ell^{(s)}}), M(X_r) \right) \right\}}. \end{aligned}$$

Note that this density ratio is effectively the ratio of joint densities as the denominators of the conditional density estimators cancel out. Denote by $Z_{\pi_s(1:n)}$ the resulting permutation of $Z_{1:n}$ after all swaps in iteration s have been accepted or rejected. The CP test runs an initial S iterations, after which it then runs B independent sequences initialized at $Z_{\pi_S(1:n)}$, each for another S iterations (Berrett et al., 2019, Algorithm 2). For $b \in \{1, \dots, B\}$, denote the final permutation of each procedure as $Z_{\pi_{2S}^{(b)}(1:n)}$. The p -value of the CP test is then computed as

$$p_{\text{CP}} = \frac{1}{1+B} \left[1 + \sum_{b=1}^B \mathbb{1} \left\{ T_{\text{CP}}(X_{1:n}, Z_{1:n}, M(X)_{1:n}) \leq T_{\text{CP}}(X_{1:n}, Z_{\pi_{2S}^{(b)}(1:n)}, M(X)_{1:n}) \right\} \right].$$

The test rejects H_0 at level α if $p_{\text{CP}} \leq \alpha$. See Berrett et al. (2019) for more details about the CP test.

C.2 Distribution of p -values in SO(3)-invariance experiment

Figure 6 shows the p -value distributions for the tests for SO(3)-invariance conducted in Section 7.1.1.




Figure 6: Histograms showing the p -value distributions obtained over $N = 1000$ simulations in the SO(3)-invariance experiment. The p -value of a Kolmogorov–Smirnov test for uniformity of the distribution is shown in the bottom-right corner of each plot.

C.3 Number of random projections for NMMD and CW in \mathbb{S}_{10} -invariance experiment

Figure 7 shows the rejection rate and average computation time for NMMD and CW as the number of random projections J increases in the \mathbb{S}_{10} -invariance experiment (Section 7.1.2).

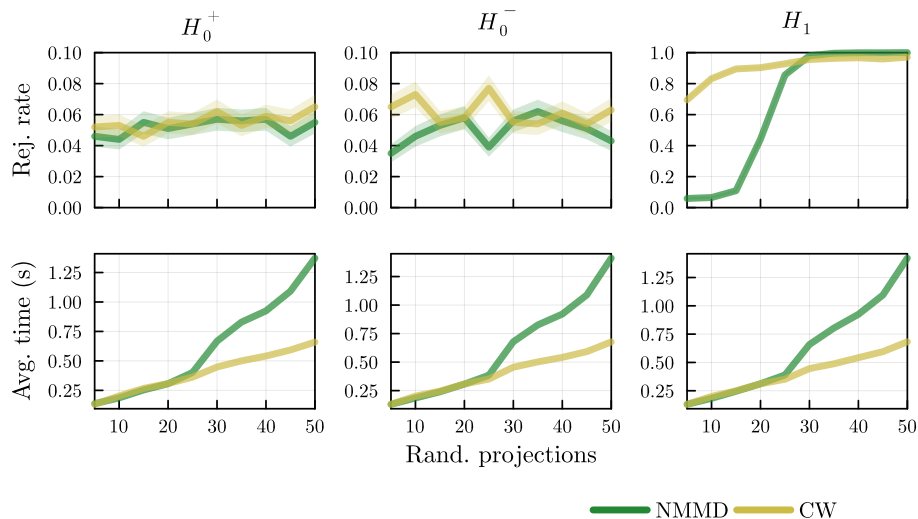


Figure 7: Test for \mathbb{S}_{10} -invariance rejection rates and standard deviations (first row) and average computation time in seconds for a single execution (second row) over $N = 1000$ simulations as the number of random projections increases.

C.4 SWARM experiment

The grids used to train the kernels in KCI were manually tuned through trial and error. The final set of grids used to obtain the results in Section 7.2 were

- k_X : 31 linearly-spaced numbers between $1e^{-3}$ and 3;
- k_Y : 57 linearly-spaced numbers between 1 and 15; and
- $k_{M(X)}$: $\{5e^{-3}, 5.5e^{-3}, 6e^{-3}\}$.

The test results for the marginal invariance and joint invariance experiments are given in Table 3.

	Marginal invariance		Joint invariance	
	Discrete	Continuous	Discrete	Continuous
2SMMD	0.007	0.004	0.006	0.011
MMD	0.047	0.047	0.998	0.999
NMMD	0.060	0.055	0.098	0.094
CW	0.064	0.087	0.903	0.897

Table 3: Test rejection rates over $N = 1000$ simulations for the SWARM data.

Histograms of the p -values for these tests are shown in Figure 8.

C.5 LHC experiment

The grid $\{10^{-2}, 10^{-1}, 0, 10\}$ was used to train the kernels k_X , k_Y , and $k_{M(X)}$ in KCI for the equivariance experiment in Section 7.3. The grid $\{10^{-3}, 10^{-2}, 10^{-1}\}$ was used to train the kernels k_Y and $k_{M(X)}$ in CP.

Figure 9 shows histograms of the p -values obtained from the tests for joint invariance and equivariance in Section 7.3.

Figure 10 shows the rejection rate and average computation time for NMMD and CW as the number of random projections J increases in the LHC joint invariance experiment.

Hypothesis Tests for Distributional Symmetry

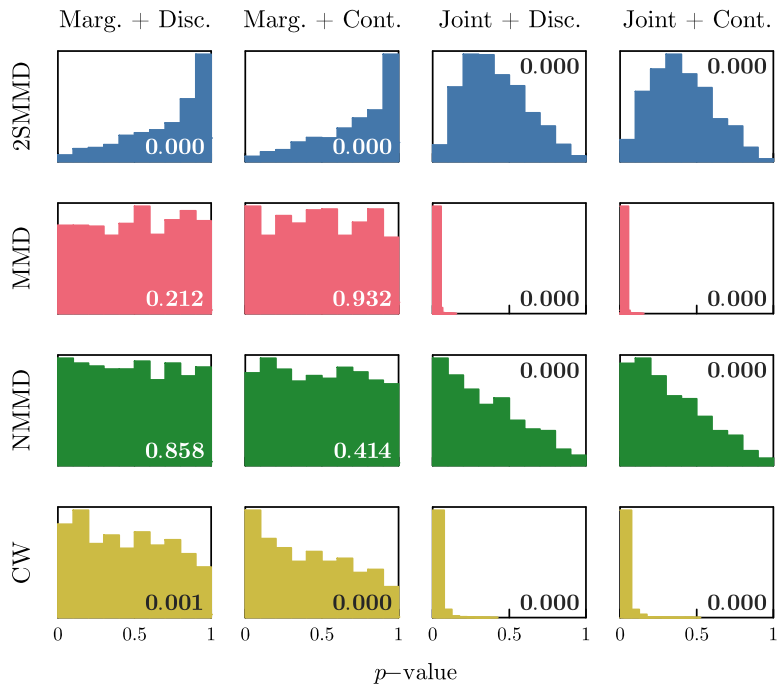


Figure 8: Histograms showing the p -value distributions obtained over $N = 1000$ simulations for tests for X -marginal invariance and (X, Y) -joint invariance with respect to discrete and continuous X -rotations about the geographic north pole. The p -value of a Kolmogorov–Smirnov test for uniformity of the distribution is shown in each plot.

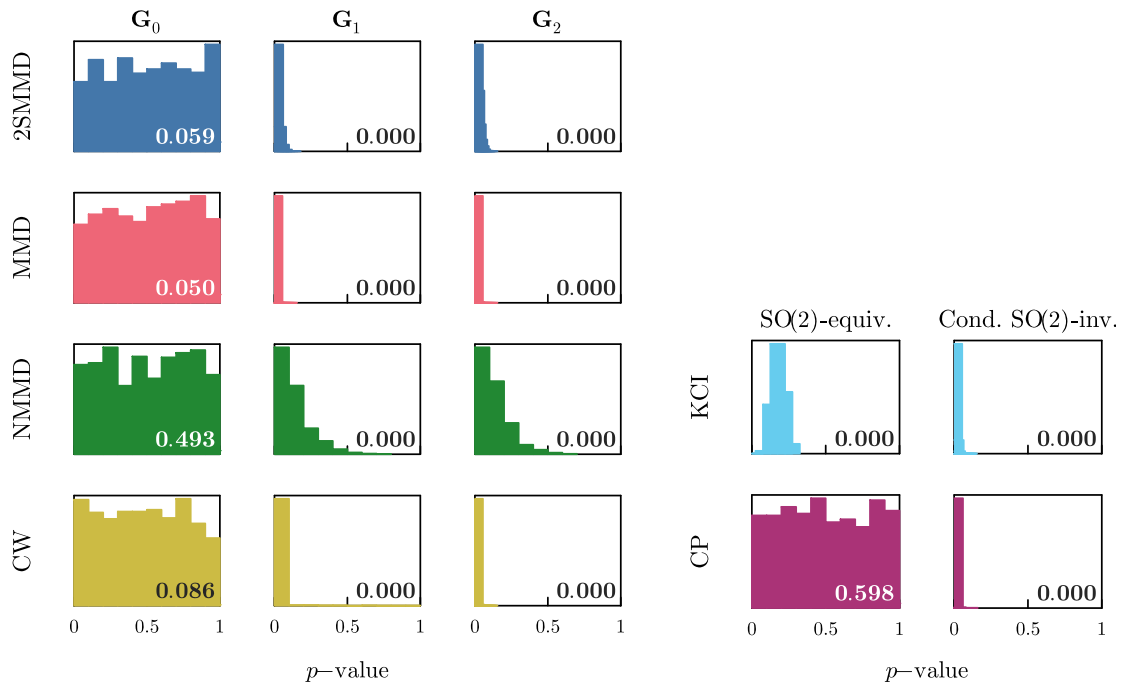


Figure 9: Histograms showing the p -value distributions obtained over $N = 1000$ simulations for tests for joint invariance and equivariance in the LHC experiments. The p -value of a Kolmogorov–Smirnov test for uniformity of the distribution is shown in the bottom-right corner of each plot.

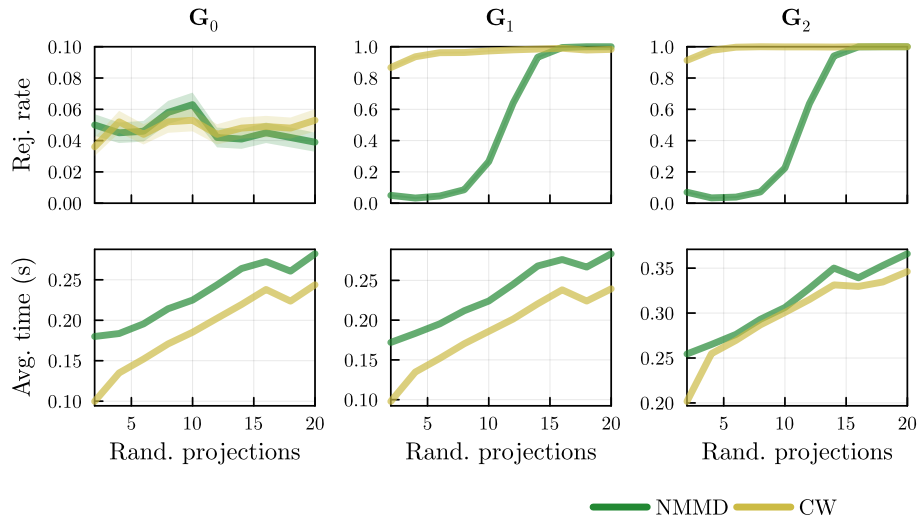


Figure 10: LHC test for joint invariance rejection rates and standard deviations (first row) and average computation time in seconds for a single execution (second row) over $N = 1000$ simulations as the number of random projections increases.

C.6 Top quark experiment

For KCI in the top quark experiment in Section 7.4, the grid $\{5, 7.5, 10, \dots, 50\}$ was used to train the kernel k_X , and the grid $\{5, 7.5, 10, \dots, 100\}$ was used to train the kernel $k_{M(X)}$. The grids were manually selected based on trial and error.

Figure 11 shows the p -value distributions obtained from KCI in the top quark experiment.




Figure 11: Histograms showing the KCI p -value distributions obtained over $N = 1000$ simulations in the top quark experiment.

References

- H. Abdi. Multiple correlation coefficient. *Encyclopedia of Measurement and Statistics*, 648:651, 2007. 19
- D. J. Aldous. Exchangeability and related topics. In P. L. Hennequin, editor, *École d'Été de Probabilités de Saint-Flour XIII - 1983*, number 1117 in Lecture Notes in Mathematics, pages 1–198. Springer, 1985. 10, 21
- ATLAS Collaboration. Searching for new symmetries of nature, April 2017. <https://atlas.cern/updates/briefing/searching-new-symmetries-nature>. Last visited 2023-07-26. 2
- T. B. Berrett, Y. Wang, R. F. Barber, and R. J. Samworth. The conditional permutation test for independence while controlling for confounders. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 82(1):175–197, 10 2019. 18, 28
- C. Bierlich, S. Chakraborty, N. Desai, L. Gellersen, I. Helenius, P. Ilten, L. Lönnblad, S. Mrenna, S. Prestel, C. T. Preuss, T. Sjöstrand, P. Skands, M. Uthheim, and R. Verheyen. A comprehensive guide to the physics and usage of PYTHIA 8.3. *SciPost Physics Codebases*, page 8, 2022. 22
- M. Birman, B. Nachman, R. Sebbah, G. Sela, O. Turetz, and S. Bressler. Data-directed search for new physics based on symmetries of the SM. *The European Physical Journal C*, 82(6):508, 2022. 2, 5
- B. Bloem-Reddy and Y. W. Teh. Probabilistic symmetries and invariant neural networks. *Journal of Machine Learning Research*, 21:90–1, 2020. 4, 6, 12, 17
- A. Bogatskiy, B. Anderson, J. Offermann, M. Roussi, D. Miller, and R. Kondor. Lorentz group equivariant neural network for particle physics. In *Proceedings of the 37th International Conference on Machine Learning*, pages 992–1002. PMLR, 2020. 17, 18
- M. M. Bronstein, J. Bruna, T. Cohen, and P. Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021. 16
- E. Çinlar. *Probability and Stochastics*. Springer New York, 2011. 7
- A. Chatalic, N. Schreuder, L. Rosasco, and A. Rudi. Nyström kernel mean embeddings. In *International Conference on Machine Learning*, pages 3006–3024. PMLR, 2022. 14, 27
- S. Chen, E. Dobriban, and J. H. Lee. A group-theoretic framework for data augmentation. *Journal of Machine Learning Research*, 21(1):9885–9955, 2020. 1
- L. G. Christie and J. A. Aston. Testing for geometric invariance and equivariance. *arXiv preprint arXiv:2205.15280*, 2022. 4
- L. G. Christie and J. A. Aston. Estimating maximal symmetries of regression functions via subgroup lattices. *arXiv preprint arXiv:2303.13616*, 2023. 4, 5, 21, 22, 23
- A. Christmann and I. Steinwart. *Support Vector Machines*. Springer New York, 2008. 14
- T. S. Cohen, M. Geiger, and M. Weiler. A general theory of equivariant CNNs on homogeneous spaces. *Advances in Neural Information Processing Systems*, 32, 2019. 1
- J. Collins, K. Howe, and B. Nachman. Anomaly detection for resonant new physics with machine learning. *Physical Review Letters*, 121:241803, 2018. 2
- A. P. Dawid. Invariance and independence in multivariate distribution theory. *Journal of Multivariate Analysis*, 17(3): 304–315, Dec. 1985. 12
- J. G. De Gooijer and D. Zerom. On conditional density estimation. *Statistica Neerlandica*, 57(2):159–176, 2003. 18
- N. Dehmamy, R. Walters, Y. Liu, D. Wang, and R. Yu. Automatic symmetry discovery with lie algebra convolutional network. *Advances in Neural Information Processing Systems*, 34:2503–2515, 2021. 2
- K. Desai, B. Nachman, and J. Thaler. SymmetryGAN: Symmetry discovery with deep learning. *arXiv preprint arXiv:2112.05722*, 2021. 2, 23
- E. Dobriban. Consistency of invariance-based randomization tests. *The Annals of Statistics*, 50(4), 2022. 4
- J.-M. Dufour. Monte Carlo tests with nuisance parameters: A general approach to finite-sample inference and nonstandard asymptotics. *Journal of Econometrics*, 133(2):443–477, Aug. 2006. 13, 24, 25

- J.-M. Dufour and J. Neves. *Finite-sample inference and nonstandard asymptotics with Monte Carlo tests and R*, volume 41 of *Handbook of Statistics*, chapter 1, pages 3–31. Elsevier, 2019. 12
- M. L. Eaton. *Group Invariance in Applications in Statistics*. Regional Conference Series in Probability and Statistics. Institute of Mathematical Statistics and American Statistical Association, 1989. 6, 7, 8
- M. L. Eaton. Topological groups and invariant measures. In *Multivariate Statistics*, pages 184–232. Institute of Mathematical Statistics, 2007. 6
- M. L. Eaton and W. D. Sudderth. Consistency and strong inconsistency of group-invariant predictive inferences. *Bernoulli*, 5(5):833–854, 1999. 1
- B. Elesedy. Provably strict generalisation benefit for invariance in kernel methods. *Advances in Neural Information Processing Systems*, 34:17273–17283, 2021. 1, 27
- B. Elesedy and S. Zaidi. Provably strict generalisation benefit for equivariant models. In *International Conference on Machine Learning*, pages 2959–2969. PMLR, 2021. 1
- R. H. Farrell. Representation of invariant measures. *Illinois Journal of Mathematics*, 6(3):447–467, 1962. 12
- G. B. Folland. *A Course in Abstract Harmonic Analysis*, volume 29. CRC press, 2016. 5
- R. Fraiman, L. Moreno, and T. Ransford. Application of the Cramér-Wold theorem to testing for invariance under group actions. *arXiv preprint arXiv:2109.01041*, 2021. 4, 5, 16
- D. A. S. Fraser. The fiducial method and invariance. *Biometrika*, 48(3/4):261–280, 1961. 1
- D. A. S. Fraser. Structural probability and a generalization. *Biometrika*, 53(1/2):1–9, 1966. 1
- D. A. S. Fraser. *The Structure of Inference*. Wiley, 1968. 1
- K. Fukumizu, A. Gretton, B. Schölkopf, and B. K. Sriperumbudur. Characteristic kernels on groups and semigroups. *Advances in Neural Information Processing Systems*, 21, 2008. 20
- D. Garreau, W. Jitkrittum, and M. Kanagawa. Large sample analysis of the median heuristic. *arXiv preprint arXiv:1707.07269*, 2017. 18
- A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012. 14
- D. J. Gross. The role of symmetry in fundamental physics. *Proceedings of the National Academy of Sciences*, 93(25):14256–14259, 1996. 2
- J. Hemerik and J. Goeman. Exact testing with random permutations. *TEST*, 27(4):811–825, 2018. 4, 12, 13
- R. B. Hora and R. J. Buehler. Fiducial theory and invariant estimation. *The Annals of Mathematical Statistics*, 37(3):643–656, 1966. 1
- R. B. Hora and R. J. Buehler. Fiducial theory and invariant prediction. *The Annals of Mathematical Statistics*, 38(3):795–801, 1967. 1
- K. H. Huang, P. Orbanz, and M. Austern. Quantifying the effects of data augmentation. *arXiv preprint arXiv:2202.09134*, 2022. 1
- O. Kallenberg. Invariant measures and disintegrations with applications to palm and related kernels. *Probability Theory and Related Fields*, 139(1):285–310, 2007. 7
- O. Kallenberg. Invariant palm and related disintegrations via skew factorization. *Probability Theory and Related Fields*, 149(1):279–301, 2011. 6, 8
- O. Kallenberg. *Random Measures, Theory and Applications*. Springer International, 2017. 8, 17, 26
- G. Karagiorgi, G. Kasieczka, S. Kravitz, B. Nachman, and D. Shih. Machine learning in the search for new fundamental physics. *Nature Reviews Physics*, 4(6):399–412, 2022. 2
- G. Kasieczka, T. Plehn, J. Thompson, and M. Russel. Top quark tagging reference dataset, Mar. 2019. <https://doi.org/10.5281/zenodo.2603256>. 17, 23
- G. Kasieczka, B. Nachman, D. Shih, O. Amram, A. Andreassen, K. Benkendorfer, B. Bortolato, G. Brooijmans, F. Canelli, J. H. Collins, et al. The LHC Olympics 2020 a community challenge for anomaly detection in high energy physics. *Reports on Progress in Physics*, 84(12):124201, 2021. 22

- S. Krippendorf and M. Syvaeri. Detecting symmetries with neural networks. *Machine Learning: Science and Technology*, 2(1):015010, 2020. 2
- A. J. Larkoski, I. Moulton, and B. Nachman. Jet substructure at the Large Hadron Collider: A review of recent advances in theory and machine learning. *Physics Reports*, 841:1–63, 2020. 17
- E. L. Lehmann and G. Casella. *Theory of Point Estimation*. Springer, 2 edition, 1998. 1
- E. L. Lehmann and J. P. Romano. *Testing Statistical Hypotheses*. Springer Texts in Statistics. Springer-Verlag New York, 2005. 1, 4
- C. Li and X. Fan. On nonparametric conditional independence tests for continuous variables. *WIREs Computational Statistics*, 12(3):e1489, 2020. 18
- C. Lyle, M. van der Wilk, M. Kwiatkowska, Y. Gal, and B. Bloem-Reddy. On the benefits of invariance in neural networks. *arXiv preprint arXiv:2005.00178*, 2020. 1
- A. McCormack and P. D. Hoff. Equivariant estimation of Fréchet means. *Biometrika*, page asad014, Feb. 2023. 7
- K. Muandet, K. Fukumizu, B. Sriperumbudur, B. Schölkopf, et al. Kernel mean embedding of distributions: A review and beyond. *Foundations and Trends in Machine Learning*, 10(1-2):1–141, 2017. 14
- N. Olsen, E. Friis-Christensen, R. Floberghagen, P. Alken, C. D. Beggan, A. Chulliat, E. Doornbos, J. T. Da Encarnação, B. Hamilton, G. Hulot, et al. The SWARM satellite constellation application and research facility (SCARF) and SWARM data products. *Earth, Planets and Space*, 65:1189–1200, 2013. 21
- A. Rahimi and B. Recht. Random features for large-scale kernel machines. *Advances in Neural Information Processing Systems*, 20, 2007. 14
- A. Raj, A. Kumar, Y. Mroueh, T. Fletcher, and B. Schölkopf. Local group invariant representations via orbit embeddings. In *Artificial Intelligence and Statistics*, pages 1225–1235. PMLR, 2017. 14, 27
- G. P. Salam. Towards jetography. *The European Physical Journal C*, 67(3):637–686, 2010. 17
- R. D. Shah and J. Peters. The hardness of conditional independence testing and the generalised covariance measure. *The Annals of Statistics*, 48(3), 2020. 2
- B. K. Sriperumbudur, A. Gretton, K. Fukumizu, B. Schölkopf, and G. R. Lanckriet. Hilbert space embeddings and metrics on probability measures. *Journal of Machine Learning Research*, 11:1517–1561, Aug. 2010. 14
- B. K. Sriperumbudur, K. Fukumizu, A. Gretton, B. Schölkopf, and G. R. G. Lanckriet. On the empirical estimation of integral probability metrics. *Electronic Journal of Statistics*, 6, 2012. 16
- R. A. Wijsman. *Invariant measures on groups and their use in statistics*, volume 14 of *Lecture Notes–Monograph Series*. Institute of Mathematical Statistics, 1990. 6, 7
- R. Winter, M. Bertolini, T. Le, F. Noé, and D.-A. Clevert. Unsupervised learning of group invariant and equivariant representations. *Advances in Neural Information Processing Systems*, 35:31942–31956, 2022. 6
- World Data Center for Geomagnetism, Kyoto University. Magnetic north, geomagnetic and magnetic poles. <https://wdc.kugi.kyoto-u.ac.jp/poles/polesexp.html>. Last visited 2023-07-26. 22
- J. Yang, R. Walters, N. Dehmamy, and R. Yu. Generative adversarial symmetry discovery. In *International Conference on Machine Learning*, 2023. 2, 23
- K. Zhang, J. Peters, D. Janzing, and B. Schölkopf. Kernel-based conditional independence test and application in causal discovery. In *Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence, UAI’11*, pages 804–813, 2011. AUAI Press. 4, 18
- A. Zhou, T. Knowles, and C. Finn. Meta-learning symmetries by reparameterization. In *International Conference on Learning Representations*, 2021. 2