arXiv:2504.09567v3 [stat.ML] 25 Jul 2025

From Conditional to Unconditional Independence: Testing Conditional Independence via Transport Maps

Chenxuan He[a]*, Yuan Gao[b]*, Liping Zhu[a], and Jian Huang[b,c]

[a]Institute of Statistics and Big Data, Renmin University of China
[b]Department of Data Science and Artificial Intelligence, The Hong Kong Polytechnic University
[c]Department of Applied Mathematics, The Hong Kong Polytechnic University

August 1, 2025

Abstract

Testing conditional independence between two random vectors given a third is a fundamental and challenging problem in statistics, particularly in multivariate nonparametric settings due to the complexity of conditional structures. We propose a novel method for testing conditional independence by transforming it to an unconditional independence test problem. We achieve this by constructing two transport maps that transform conditional independence into unconditional independence, this substantially simplifies the problem. These transport maps are estimated from data using conditional continuous normalizing flow models. Within this framework, we derive a test statistic and prove its asymptotic validity under both the null and alternative hypotheses. A permutation-based procedure is employed to evaluate the significance of the test. We validate the proposed method through extensive simulations and real-data analysis. Our numerical studies demonstrate the practical effectiveness of the proposed method for conditional independence testing.

Keywords: Continuous normalizing flow, Deep neural networks, Distance correlation, Independence, Permutation.

1 Introduction

Let $\mathbf{X} \in \mathbb{R}^{d_{\mathbf{x}}}$, $\mathbf{Y} \in \mathbb{R}^{d_{\mathbf{y}}}$, and $\mathbf{Z} \in \mathbb{R}^{d_{\mathbf{z}}}$ be three random vectors. We are interested in testing whether $\mathbf{X}$ and $\mathbf{Y}$ are conditionally independent given $\mathbf{Z}$ (Dawid, 1979), formally defined as:

$$H_0 : \mathbf{X} \perp\!\!\!\perp \mathbf{Y} \mid \mathbf{Z}, \text{ versus } H_1 : \mathbf{X} \not\perp\!\!\!\perp \mathbf{Y} \mid \mathbf{Z}. \tag{1.1}$$

*Chenxuan He and Yuan Gao are co-first authors.

Conditional independence is a fundamental concept in statistics and machine learning, playing a crucial role in various fields such as causal inference (Spirtes et al., 1993; Pearl, 2009; Spirtes, 2010; Constantinou and Dawid, 2017) and graphical models (Geiger and Pearl, 1993; Koldanov et al., 2017; Fan et al., 2020). Understanding the conditional independence between random vectors allows researchers to simplify complex models, reduce dimensionality, and enhance interpretability.

Numerous methods have been proposed for testing conditional independence. Parametric approaches impose explicit structural assumptions on the conditional distributions $\mathbf{X} \mid \mathbf{Z}$ and $\mathbf{Y} \mid \mathbf{Z}$. For example, Song (2009) adopted a single-index model paired with the Rosenblatt transformation (Rosenblatt, 1952), while Fan et al. (2020) assumed linearity in the conditional distributions. Though such methods achieve root-n consistency under their parametric constraints, their reliance on strict model specifications renders them vulnerable to misspecification bias, limiting their applicability to complex, non-linear dependencies.

Beyond parametric constraints, kernel-based methods employ kernel functions to measure dependence without parametric assumptions. For instance, Zhang et al. (2011) and Huang et al. (2022) derived tests using cross-covariance operator in reproducing kernel Hilbert spaces (RKHS). Wang et al. (2015) introduced conditional distance correlation, where they utilized kernel functions to estimate the empirical form. Cai et al. (2022) formulated conditional independence as mutual independence across transformed variables, estimating dependencies via kernel-based empirical measures. Though these methods avoid strict parametric assumptions, they either incur substantial computational costs from kernel matrix operations or exhibit reduced statistical power against alternatives as the dimension of $\mathbf{Z}$ grows, even at moderate dimensionality (e.g., dimension of $\mathbf{Z}$ larger than 3). Recent methodological advances have introduced modern machine learning techniques for conditional independence testing. For example, Sen et al. (2017) formulated conditional independence test as a classification problem, employing gradient-boosted trees and deep neural networks to perform the classification. Bellot and van der Schaar (2019), Shi et al. (2021), Yang et al. (2024), and Duong and Nguyen (2024) propose utilizing generative model, such as GANs and diffusion models, to learn conditional distributions. They further use these learned distributions to get the residuals. These methods represent modern extensions of the partial correlation concept.

In this work, we propose a framework for conditional independence testing based on transport maps. At the population level, we demonstrate that two invertible maps from the conditional distributions to simple reference distributions allow us to formulate conditional independence test as unconditional independence test, thereby simplifying the formulation. We employ conditional continuous normalizing flows and independence measures to effectively address conditional independence testing. Our contributions are summarized as

follows:

1. *A novel equivalent formulation of conditional independence.* We demonstrate that testing conditional independence between $\mathbf{X}$ and $\mathbf{Y}$ given $\mathbf{Z}$ can be simplified to testing unconditional independence through invertible transformations. By transforming $(\mathbf{X}, \mathbf{Z})$ and $(\mathbf{Y}, \mathbf{Z})$ into new variables (e.g., Gaussian variables) independent of $\mathbf{Z}$, conditional independence becomes equivalent to unconditional independence. We rigorously prove this equivalence in Lemma 1. This formulation significantly reduces the complexity associated with conditional independence structure.

2. *A model-free test statistic.* We introduce a model-free test statistic for conditional independence testing. By using the continuous normalizing flow as a transport map, we can apply popular independence measures, such as distance correlation, which do not depend on specific model assumptions about the data distribution. The asymptotic validity is given in Theorem 3.

3. *Performance across dimensions.* Through extensive simulations, we demonstrate that the proposed test maintains robust performance across various dimensions for variables $\mathbf{X}$, $\mathbf{Y}$, and $\mathbf{Z}$, with each dimension ranged from 1 to 50.

4. *Robustness to independence measures.* After transforming a conditional independence problem into an unconditional independence one, our test remains robust to the selection of dependence measures. We illustrate this by comparing the results of the proposed method using distance correlation (Székely et al., 2007) and improved projection correlation (Zhang and Zhu, 2024) in Supplementary Materials B.4 .

We present an application of the proposed method to the problem of sufficient dimension reduction, specifically in testing whether an empirically estimated data representation satisfies the sufficiency requirement, as discussed in Section 4.

Our approach utilizes a permutation-based method for computing p-values. This method is supported by both simulation studies and real-data analysis, which collectively demonstrate its robustness, empirical validity, and practical utility. Theoretical justification for the permutation-based method is provided in Theorem 4, with numerical evaluations detailed in Section 4.

The remainder of the paper is organized as follows: We provide some preliminaries in Section 2 and introduce the methodology in Section 3. Simulation studies and an application to dimension reduction are presented in Sections 4 and 5. A brief conclusion is provided in Section 6.

Notation. Throughout this paper, uppercase letters such as X, Y, and Z denote random variables, with their realizations represented by the corresponding lowercase letters x, y, and z. Bold uppercase letters such as $\mathbf{X}$, $\mathbf{Y}$, and $\mathbf{Z}$ denote random vectors, while bold lowercase letters such as $\mathbf{x}$, $\mathbf{y}$, and $\mathbf{z}$ denote their realizations as real-valued vectors. The Euclidean norm of a vector $\mathbf{x} \in \mathbb{R}^d$ is denoted by $\|\mathbf{x}\|$. The Gamma function is denoted by $\Gamma(\cdot)$, the d-dimensional identity matrix by $\mathbf{I}_d$, and the indicator function by $\mathbf{1}(\cdot)$. We use $\stackrel{d}{=}$ to denote equal in distribution and $\xrightarrow{p}$ to denote convergence in probability. The support of a random vector is denoted by $\mathrm{supp}(\cdot)$. For two sequences a_n and b_n, the notation $a_n = O(b_n)$ or $a_n \lesssim b_n$ indicates that there exists a constant C_0 such that $|a_n| \leq C_0|b_n|$ as n goes to infinity.

2 Basic Formulation

In this section, we present a basic formulation of our proposed method. The core idea is to employ transport maps to transform a conditional independence test problem into an unconditional one. This transformation is detailed in Section 2.1. We then discuss the implementation of transport maps using continuous normalizing flow models in Section 2.2. Following the application of the transport maps, we introduce independence measures in Section 2.3.

2.1 From Conditional Independence to Independence

To begin with, we establish an equivalence between conditional independence and unconditional independence by invertible transport maps. Recall the formulation for conditional independence testing given in (1.1). Denote the original jointly distribution triplet as $(\mathbf{X}, \mathbf{Y}, \mathbf{Z}) \in \mathbb{R}^{d_\mathbf{x}+d_\mathbf{y}+d_\mathbf{z}}$. Suppose there are two invertible transport maps that push the conditional distributions $\mathbf{X} \mid \mathbf{Z}$ and $\mathbf{Y} \mid \mathbf{Z}$ to two simple reference distributions (e.g., standard normal distributions). We say a function $f(\mathbf{x}, \mathbf{z}) : \mathbb{R}^{d_\mathbf{x}} \times \mathbb{R}^{d_\mathbf{z}} \to \mathbb{R}^{d_\mathbf{x}}$ is invertible with respect to the first argument if, for every fixed $\mathbf{z} \in \mathbb{R}^{d_\mathbf{z}}$, there exists an inverse function $f^{-1} : \mathbb{R}^{d_\mathbf{x}} \times \mathbb{R}^{d_\mathbf{z}} \to \mathbb{R}^{d_\mathbf{x}}$ such that $f^{-1}(f(\mathbf{x}, \mathbf{z}), \mathbf{z}) = f(f^{-1}(\mathbf{x}, \mathbf{z}), \mathbf{z}) = \mathbf{x}$ for all $\mathbf{x} \in \mathbb{R}^{d_\mathbf{x}}$ and $\mathbf{z} \in \mathbb{R}^{d_\mathbf{z}}$. We have the following Lemma 1.

Lemma 1 (Conditional independence to unconditional independence)**.** *For the jointly distributed triplet $(\mathbf{X}, \mathbf{Y}, \mathbf{Z}) \in \mathbb{R}^{d_\mathbf{x}+d_\mathbf{y}+d_\mathbf{z}}$, suppose there exist two invertible transport maps*

(with respect to their first arguments) $\boldsymbol{\xi} = \mathbf{f}(\mathbf{X}, \mathbf{Z})$ *and* $\boldsymbol{\eta} = \mathbf{g}(\mathbf{Y}, \mathbf{Z})$ *such that:*

$$\boldsymbol{\xi} = \mathbf{f}(\mathbf{X}, \mathbf{Z}) \overset{d}{=} \big[\mathbf{f}(\mathbf{X}, \mathbf{Z}) \mid \mathbf{Z} = \mathbf{z}\big], \quad \forall \mathbf{z} \in \mathrm{supp}(\mathbf{Z}),$$
$$\boldsymbol{\eta} = \mathbf{g}(\mathbf{Y}, \mathbf{Z}) \overset{d}{=} \big[\mathbf{g}(\mathbf{Y}, \mathbf{Z}) \mid \mathbf{Z} = \mathbf{z}\big], \quad \forall \mathbf{z} \in \mathrm{supp}(\mathbf{Z}).$$

Then, it holds that $\boldsymbol{\xi} \perp\!\!\!\perp \mathbf{Z}, \quad \boldsymbol{\eta} \perp\!\!\!\perp \mathbf{Z}$. *Furthermore, we have:*

$$\mathbf{X} \perp\!\!\!\perp \mathbf{Y} \mid \mathbf{Z} \iff \boldsymbol{\eta} \perp\!\!\!\perp (\boldsymbol{\xi}, \mathbf{Z}) \iff \boldsymbol{\xi} \perp\!\!\!\perp (\boldsymbol{\eta}, \mathbf{Z}).$$

Remark 1. *Lemma 1 establishes equivalences between conditional and unconditional independence via transport maps. This result aligns with, yet generalizes, earlier frameworks in the literature. For instance, let* $F_1(\mathbf{X} \mid \mathbf{Z})$ *and* $F_2(\mathbf{Y} \mid \mathbf{Z})$ *denote the conditional distributions of* $\mathbf{X}$ *and* $\mathbf{Y}$ *given* $\mathbf{Z}$*, and let* $F_3(\mathbf{Z})$ *represent the marginal distribution of* $\mathbf{Z}$*. Under the assumption that* $(F_1(\mathbf{X} \mid \mathbf{Z}), F_2(\mathbf{Y} \mid \mathbf{Z})) \perp\!\!\!\perp \mathbf{Z}$*, Zhou et al. (2020) demonstrated that conditional independence reduces to the unconditional independence* $F_1(\mathbf{X} \mid \mathbf{Z}) \perp\!\!\!\perp F_2(\mathbf{Y} \mid \mathbf{Z})$*. Similarly, Cai et al. (2022) characterized conditional independence via the mutual independence of* $F_1(\mathbf{X} \mid \mathbf{Z})$*,* $F_2(\mathbf{Y} \mid \mathbf{Z})$*, and* $F_3(\mathbf{Z})$*. While these results provide foundational insights, they are inherently constrained to specific distributional representations. Our work extends this paradigm by utilizing transport maps to create a more flexible equivalence framework, and avoid the assumption of joint independence in formulating the problem.*

By applying Lemma 1, we can convert the conditional independence test into an unconditional independence test. This transformation simplifies the analysis and allows us to leverage classical statistical measures of independence. Detailed proofs of Lemma 1 can be found in the Supplementary Materials C.

Lemma 1 establishes a practical framework for conducting conditional independence testing. The conditional Continuous Normalizing Flow model (CNF, Liu et al. 2023, Albergo and Vanden-Eijnden 2022, Lipman et al. 2022), a state-of-the-art generative model, is able to formulate the transport map and is proven to be well-posed under certain regularity conditions (Gao et al., 2024a). Building on this, the next section presents an implementation of Lemma 1 using conditional CNFs.

2.2 Conditional Continuous Normalizing Flows

In this section, we introduce the conditional CNFs with flow matching (Liu et al., 2023; Albergo and Vanden-Eijnden, 2022; Lipman et al., 2022). Let $\mathbf{X}_0 \sim \mathrm{Normal}(\mathbf{0}, \mathbf{I}_{d_\mathbf{x}})$, $\mathbf{Y}_0 \sim \mathrm{Normal}(\mathbf{0}, \mathbf{I}_{d_\mathbf{y}})$ be two random vectors independent of $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$. We illustrate the conditional CNF for the pair of distributions, the distribution of $\mathbf{X}_0$ and the conditional

distribution of $(\mathbf{X} \mid \mathbf{Z})$; and similarly for $\mathbf{Y}_0$ and $(\mathbf{Y} \mid \mathbf{Z})$. We introduce a specific type of CNFs, namely, the rectified flow (Liu, 2022; Liu et al., 2023). A rectified flow is constructed based on a linear stochastic interpolation between $\mathbf{X}_0$ and $\mathbf{X}$ over time $t \in [0, 1]$:

$$\mathbf{X}_t = (1-t)\mathbf{X}_0 + t\mathbf{X}, \quad \forall t \in [0,1]. \tag{2.1}$$

When $t = 0$, since $\mathbf{X}_0 \perp\!\!\!\perp (\mathbf{X}, \mathbf{Z})$, it follows that $(\mathbf{X}_0 \mid \mathbf{Z}) = \mathbf{X}_0$ in distribution. When $t = 1$, we have $(\mathbf{X} \mid \mathbf{Z}) = (\mathbf{X}_1 \mid \mathbf{Z})$. The process defined in (2.1) indicates that as t goes from 0 to 1, the conditional distribution of $(\mathbf{X}_t \mid \mathbf{Z})$ progressively evolves from the standard normal distribution to the conditional distribution of $(\mathbf{X}_1 \mid \mathbf{Z}), t \in [0, 1]$.

We now present the definition of the velocity field related to the conditional distribution of $(\mathbf{X}_t \mid \mathbf{Z})$.

Definition 1 (Velocity field)**.** *For the time-differentiable random process* $\{(\mathbf{X}_t \mid \mathbf{Z}), t \in [0,1]\}$ *defined in* (2.1)*, its velocity field* $\mathbf{v}_{\mathbf{x}}(t, \mathbf{x}, \mathbf{z})$ *is defined as*

$$\begin{aligned}\mathbf{v}_{\mathbf{x}}(t, \mathbf{x}, \mathbf{z}) =& \mathbb{E}\left\{ \frac{\mathrm{d}\mathbf{X}_t}{\mathrm{d}t} \;\middle|\; \mathbf{X}_t = \mathbf{x}, \mathbf{Z} = \mathbf{z} \right\} \\ =& \mathbb{E}\left(\mathbf{X} - \mathbf{X}_0 \mid \mathbf{X}_t = \mathbf{x}, \mathbf{Z} = \mathbf{z}\right), \quad \forall \mathbf{x} \in \mathrm{supp}(\mathbf{X}_t), \mathbf{z} \in \mathrm{supp}(\mathbf{Z}).\end{aligned}$$

By integrating the velocity field starting from $\mathbf{X}_0$, we derive the forward process:

$$(\mathbf{X}_t^\sharp \mid \mathbf{Z}) = \mathbf{X}_0 + \int_0^t \mathbf{v}_{\mathbf{x}}(s, \mathbf{X}_s^\sharp, \mathbf{Z})\mathrm{d}s, \quad \forall t \in [0,1], \quad \mathbf{X}_0^\sharp = \mathbf{X}_0, \tag{2.2}$$

where we use the notation $(\mathbf{X}_t^\sharp \mid \mathbf{Z})$ to represent the state of the process at time t conditional on $\mathbf{Z}$. We will use this notation repeatedly below. Conversely, by integrating backward from $\mathbf{X}_1 = \mathbf{X}$, we obtain the reverse process:

$$(\mathbf{X}_t^\dagger \mid \mathbf{Z}) = \mathbf{X}_1 + \int_1^t \mathbf{v}_{\mathbf{x}}(s, \mathbf{X}_s^\dagger, \mathbf{Z})\mathrm{d}s, \quad \forall t \in [0,1], \quad (\mathbf{X}_1^\dagger \mid \mathbf{Z}) = (\mathbf{X} \mid \mathbf{Z}). \tag{2.3}$$

The forward process $(\mathbf{X}_t^\sharp \mid \mathbf{Z})$ and the reverse process $(\mathbf{X}_t^\dagger \mid \mathbf{Z})$ have been shown to possess the marginal preserving property, as detailed in Lemma 2 below.

Lemma 2 (Marginal preserving property)**.** *Under several regularity conditions, these two processes* $(\mathbf{X}_t^\sharp \mid \mathbf{Z})$ *and* $(\mathbf{X}_t^\dagger \mid \mathbf{Z})$ *defined in* (2.2) *and* (2.3) *are well-posed and unique. In addition, the following property holds:*

$$(\mathbf{X}_t \mid \mathbf{Z} = \mathbf{z}) \overset{d}{=} (\mathbf{X}_t^\sharp \mid \mathbf{Z} = \mathbf{z}) \overset{d}{=} (\mathbf{X}_t^\dagger \mid \mathbf{Z} = \mathbf{z}), \quad \textit{for any } (\mathbf{z}, t) \in \mathrm{supp}(\mathbf{Z}) \times [0,1],$$

where $(\mathbf{X}_t \mid \mathbf{Z})$ *is defined in* (2.1).

The marginal preserving property is well-established in the literature. For instance, Theorem 5.1 and Corollary 5.2 in Gao et al. (2024a) establish this property under suitable assumptions. Similar results can also be found in Liu (2022), Albergo et al. (2023), and Huang et al. (2024a).

According to Lemma 2, the velocity field $\mathbf{v}_{\mathbf{x}}(t, \mathbf{x}, \mathbf{z})$ enables the transformation of the distribution $\mathbf{X}_0$ into $(\mathbf{X} \mid \mathbf{Z})$, and vice versa. Furthermore, since the ordinary differential equation (ODE) in (2.1) and the integration process defined in (2.3) are deterministic, the transformation is one-to-one and invertible. This property is not shared by the stochastic differential equation-based methods, such as diffusion models (Ho et al., 2020; Song et al., 2020). Consequently, the transport map induced by the rectified flow is invertible and meets the conditions outlined in Lemma 1. Thus, we introduce two rectified flows:

$$\begin{aligned}\mathbf{X}_t =&(1-t)\mathbf{X}_0 + t\mathbf{X}, \quad \forall t \in [0,1],\\ \mathbf{Y}_t =&(1-t)\mathbf{Y}_0 + t\mathbf{Y}, \quad \forall t \in [0,1].\end{aligned}$$

These two processes induce two velocity fields $\mathbf{v}_{\mathbf{x}}(t, \mathbf{x}, \mathbf{z})$ and $\mathbf{v}_{\mathbf{y}}(t, \mathbf{y}, \mathbf{z})$ as given in Definition 1. At the population level, by integrating the velocity fields starting from $\mathbf{X}$ and $\mathbf{Y}$, we obtain:

$$\begin{aligned}(\mathbf{X}_0^\dagger \mid \mathbf{Z}) &= \mathbf{X} + \int_1^0 \mathbf{v}_{\mathbf{x}}(s, \mathbf{X}_s^\dagger, \mathbf{Z})\mathrm{d}s, \quad (\mathbf{X}_1^\dagger \mid \mathbf{Z}) = (\mathbf{X} \mid \mathbf{Z})\\ (\mathbf{Y}_0^\dagger \mid \mathbf{Z}) &= \mathbf{Y} + \int_1^0 \mathbf{v}_{\mathbf{y}}(s, \mathbf{Y}_s^\dagger, \mathbf{Z})\mathrm{d}s, \quad (\mathbf{Y}_1^\dagger \mid \mathbf{Z}) = (\mathbf{Y} \mid \mathbf{Z}).\end{aligned}$$

By Lemma 2, we have $(\mathbf{X}_0^\dagger \mid \mathbf{Z}) \overset{d}{=} \mathbf{X}_0$ and $(\mathbf{Y}_0^\dagger \mid \mathbf{Z}) \overset{d}{=} \mathbf{Y}_0$. We define:

$$\boldsymbol{\xi} := (\mathbf{X}_0^\dagger \mid \mathbf{Z}) \overset{d}{=} \mathbf{X}_0, \tag{2.4}$$

$$\boldsymbol{\eta} := (\mathbf{Y}_0^\dagger \mid \mathbf{Z}) \overset{d}{=} \mathbf{Y}_0, \tag{2.5}$$

where $\boldsymbol{\xi}$ is a function of $\mathbf{X}$ and $\mathbf{Z}$, and $\boldsymbol{\eta}$ is a function of $\mathbf{Y}$ and $\mathbf{Z}$. The pair $(\boldsymbol{\xi}, \boldsymbol{\eta})$ is exactly what we expect in Lemma 1. We illustrate this idea in Example 1 below.

Example 1. *We illustrate Lemmas 1 and 2 through a Gaussian example. Let* $Z \sim$ *Normal*$(0,1)$*, with* $X \mid Z \sim$ *Normal*$(Z,1)$ *and* $Y \mid Z \sim$ *Normal*$(Z,1)$*. By construction,* $X \perp\!\!\!\perp Y \mid Z$*. The linear interpolation between* X_0 *and* X *is* $X_t = (1-t)X_0 + tX, \quad \forall t \in [0,1]$*. Conditional on* Z*, the joint distribution of* $(X, X_0, X_t)^\top$ *is Normal*$(\boldsymbol{\mu}, \Sigma)$*, where*

$\boldsymbol{\mu} = (Z, 0, Z - tZ)^\top$ *and*

$$\Sigma = \begin{pmatrix} 1 & 0 & t \\ 0 & 1 & 1-t \\ t & 1-t & 1-2t-t^2 \end{pmatrix}.$$

From the properties of joint normal distributions, the velocity field is given by

$$v_x(t, x, z) = \mathbb{E}(X - X_0 \mid X_t = x, Z = z) = \frac{2t-1}{1-2t-2t^2}\{x - (1-t)z\} = \frac{\mathrm{d}x}{\mathrm{d}t}.$$

Solving this ODE backward in time with terminal condition $X_1 = x$, we obtain that $X_0^\dagger = z - x$. Consequently, we have $\xi = Z - X \sim$ Normal$(0, 1)$. Analogously, $\eta = Z - Y \sim$ Normal$(0, 1)$. Thus, $\xi \perp\!\!\!\perp (\eta, Z)$ and $\eta \perp\!\!\!\perp (\xi, Z)$ are satisfied.

2.3 Independence Measures

Based on Lemma 1, after transforming the conditional variables $(\mathbf{X} \mid \mathbf{Z})$ and $(\mathbf{Y} \mid \mathbf{Z})$ into $\boldsymbol{\xi}$ and $\boldsymbol{\eta}$, respectively, the conditional independence test problem $\mathbf{X} \perp\!\!\!\perp \mathbf{Y} \mid \mathbf{Z}$ is simplified to testing $\boldsymbol{\xi} \perp\!\!\!\perp (\boldsymbol{\eta}, \mathbf{Z})$ or $\boldsymbol{\eta} \perp\!\!\!\perp (\boldsymbol{\xi}, \mathbf{Z})$.

Several popular independence measures are suitable for use, such as Distance Correlation (DC, Székely et al. 2007), Projection Correlation (PC, Zhu et al. 2017), and Improved Projection Correlation (IPC, Zhang and Zhu 2024). They are all equal to zero if and only if two random vectors are independent.

Given the computational efficiency of DC and IPC, we recommend using these two measures to assess the independence. Our numerical studies suggest that the test results remain robust regardless of the chosen dependence measure. For further details, please refer to Appendix B.4.

In the following, we illustrate the idea of independence measures by distance covariance and distance correlation. An introduction for IPC is given in Appendix A.

Let $\mathbf{U}$ and $\mathbf{V}$ be two arbitary random vectors. The distance covariance between them is defined as (Székely et al., 2007):

$$\mathrm{dcov}^2(\mathbf{U}, \mathbf{V}) = \int \left\| f_{\mathbf{U},\mathbf{V}}(\mathbf{u}, \mathbf{v}) - f_{\mathbf{U}}(\mathbf{u}) f_{\mathbf{V}}(\mathbf{v}) \right\|^2 w(\mathbf{u}, \mathbf{v}) \mathbb{E}_{\mathbf{u}} \mathrm{d}\mathbf{v}, \tag{2.6}$$

where $f_{\mathbf{U},\mathbf{V}}(\cdot, \cdot)$ is the joint characteristic function of $(\mathbf{U}, \mathbf{V})$, $f_{\mathbf{U}}(\cdot)$ and $f_{\mathbf{V}}(\cdot)$ are the marginal characteristic functions of $\mathbf{U}$ and $\mathbf{V}$, and $w(\cdot, \cdot)$ is the weight function defined by $w(\mathbf{u}, \mathbf{v}) = \left\{ c_{d_\mathbf{x}} c_{d_\mathbf{y}} \|\mathbf{u}\|^{1+d_\mathbf{x}} \|\mathbf{v}\|^{1+d_\mathbf{y}} \right\}^{-1}$, with $c_d = \pi^{(1+d)/2} / \Gamma((1+d)/2)$.

Since $\mathbf{U} \perp\!\!\!\perp \mathbf{V} \iff f_{\mathbf{U},\mathbf{V}}(\mathbf{u}, \mathbf{v}) = f_{\mathbf{U}}(\mathbf{u}) f_{\mathbf{V}}(\mathbf{v})$, the distance covariance is nonnegative

and equals to zero if and only if $\mathbf{U}$ and $\mathbf{V}$ are independent. Let $(\mathbf{U}_1, \mathbf{V}_1)$, $(\mathbf{U}_2, \mathbf{V}_2)$, and $(\mathbf{U}_3, \mathbf{V}_3)$ be independent copies of $(\mathbf{U}, \mathbf{V})$. If $\mathbb{E}\|\mathbf{U}\|^2 < \infty$ and $\mathbb{E}\|\mathbf{V}\|^2 < \infty$, the distance covariance defined in (2.6) can be expressed as (Székely et al., 2007)

$$\begin{aligned}\mathrm{dcov}^2(\mathbf{U}, \mathbf{V}) =& \mathbb{E}\left(\|\mathbf{U}_1 - \mathbf{U}_2\|\|\mathbf{V}_1 - \mathbf{V}_2\|\right) + \mathbb{E}\|\mathbf{U}_1 - \mathbf{U}_2\|\mathbb{E}\|\mathbf{V}_1 - \mathbf{V}_2\| \\ & - 2\mathbb{E}\left(\|\mathbf{U}_1 - \mathbf{U}_2\|\|\mathbf{V}_1 - \mathbf{V}_3\|\right). \end{aligned} \tag{2.7}$$

The distance correlation is give by

$$\mathrm{dcorr}^2(\mathbf{U}, \mathbf{V}) = \frac{\mathrm{dcov}^2(\mathbf{U}, \mathbf{V})}{\mathrm{dcov}(\mathbf{U}, \mathbf{U})\,\mathrm{dcov}(\mathbf{V}, \mathbf{V})}.$$

With the ability to transform conditional independence into an unconditional independence framework, along with tools to measure unconditional independence, we are now prepared to introduce the methodology in Section 3.

3 Methdology

In this section, we begin by introducing our test statistic in Section 3.1. Our method employs a sample splitting approach, and the asymptotic properties of the test statistic are discussed in Section 3.2 under appropriate splitting conditions. To enhance the power performance and address the limitations of splitting, we propose using multiple splitting in Section 3.3.

3.1 Test Statistic

For the triplet $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$, we transforme the conditional distributions $(\mathbf{X} \mid \mathbf{Z})$ and $(\mathbf{Y} \mid \mathbf{Z})$ into $\boldsymbol{\xi}$ and $\boldsymbol{\eta}$ by conditional CNFs. The dependence between the transformed variables is then assessed via independence measures. We develop the empirical counterpart of this procedure to perform the conditional independence test.

Consider the observed dataset $\mathcal{D}_n = \{(\mathbf{X}_i, \mathbf{Y}_i, \mathbf{Z}_i)\}_{i=1}^n$. We randomly split the dataset into two disjoint subsets, denoted by two index sets $\mathcal{I}_1$ and $\mathcal{I}_2$, where $|\mathcal{I}_1| = n_1$, $|\mathcal{I}_2| = n_2$, $\mathcal{I}_1 \cap \mathcal{I}_2 = \emptyset$, and $\mathcal{I}_1 \cup \mathcal{I}_2 = \{1, \ldots, n\}$. Denote the training dataset as $\mathcal{D}_{n_1} = \{(\mathbf{X}_i, \mathbf{Y}_i, \mathbf{Z}_i), i \in \mathcal{I}_1\}$ and the test dataset as $\mathcal{D}_{n_2} = \{(\mathbf{X}_i, \mathbf{Y}_i, \mathbf{Z}_i), i \in \mathcal{I}_2\}$.

The proposed test statistic is constructed through four steps. First, we estimate velocity fields via deep neural networks based on training dataset $\mathcal{D}_{n_1}$. Next, we compute the transformed variables $\boldsymbol{\xi}$ and $\boldsymbol{\eta}$ using the maps defined in (2.4) and (2.5). Third, we construct the test statistic T_n as the sample version of the chosen independence measure based on

$\mathcal{D}_{n_2}$. Finally, we compute the p-value using permutation-based methods. We now give detailed explanations of the four steps.

Step 1. Estimation of the velocity fields $\mathbf{v}_{\mathbf{x}}(t, \mathbf{x}, \mathbf{z})$ *and* $\mathbf{v}_{\mathbf{y}}(t, \mathbf{y}, \mathbf{z})$.

In the first step, we estimate the velocity fields using deep neural networks. As established in Definition 1, these velocity fields correspond to conditional expectations, making the least squares estimator a natural choice for their estimation. Operationally, we augment the observed dataset $\mathcal{D}_{n_1}$ with three auxiliary random samples: $\{\mathbf{X}_{0,i}, i \in \mathcal{I}_1\} \sim \text{Normal}(\mathbf{0}, \mathbf{I}_{d_{\mathbf{x}}})$, $\{\mathbf{Y}_{0,i}, i \in \mathcal{I}_1\} \sim \text{Normal}(\mathbf{0}, \mathbf{I}_{d_{\mathbf{y}}})$, and $\{t_i, i \in \mathcal{I}_1\} \sim \text{Uniform}([0,1])$. The velocity field estimators are obtained by minimizing the empirical squared losses:

$$\hat{\mathbf{v}}_{\mathbf{x}}(t, \mathbf{x}, \mathbf{z}) = \underset{\mathbf{v}_{\mathbf{x}} \in \mathcal{F}_{\mathbf{x}}}{\arg\min} \frac{1}{n_1} \sum_{i \in \mathcal{I}_1} \|\mathbf{v}_{\mathbf{x}}(t_i, \mathbf{X}_i, \mathbf{Z}_i) - (\mathbf{X}_i - \mathbf{X}_{0,i})\|_2^2, \tag{3.1}$$

$$\hat{\mathbf{v}}_{\mathbf{y}}(t, \mathbf{y}, \mathbf{z}) = \underset{\mathbf{v}_{\mathbf{y}} \in \mathcal{F}_{\mathbf{y}}}{\arg\min} \frac{1}{n_1} \sum_{i \in \mathcal{I}_1} \|\mathbf{v}_{\mathbf{y}}(t_i, \mathbf{Y}_i, \mathbf{Z}_i) - (\mathbf{Y}_i - \mathbf{Y}_{0,i})\|_2^2, \tag{3.2}$$

where $\mathcal{F}_{\mathbf{x}}$ and $\mathcal{F}_{\mathbf{y}}$ refer to neural network function classes. These estimated velocity fields then serve as the foundation for computing the transformed variables $\boldsymbol{\xi}$ and $\boldsymbol{\eta}$ in the subsequent step.

Step 2. Estimation of $\boldsymbol{\xi}$ *and* $\boldsymbol{\eta}$.

Given the estimated velocity fields, the estimted transport map is denoted as

$$(\hat{\mathbf{X}}_t^\dagger \mid \mathbf{Z}) = \mathbf{X} + \int_1^t \hat{\mathbf{v}}_{\mathbf{x}}(s, \hat{\mathbf{X}}_s^\dagger, \mathbf{Z}) \mathrm{d}s, \quad \forall t \in [0,1], \quad (\hat{\mathbf{X}}_1^\dagger \mid \mathbf{Z}) = (\mathbf{X} \mid \mathbf{Z}) \tag{3.3}$$

$$(\hat{\mathbf{Y}}_t^\dagger \mid \mathbf{Z}) = \mathbf{Y} + \int_1^t \hat{\mathbf{v}}_{\mathbf{y}}(s, \hat{\mathbf{Y}}_s^\dagger, \mathbf{Z}) \mathrm{d}s, \quad \forall t \in [0,1], \quad (\hat{\mathbf{Y}}_1^\dagger \mid \mathbf{Z}) = (\mathbf{Y} \mid \mathbf{Z}). \tag{3.4}$$

Similar to $\boldsymbol{\xi} = (\mathbf{X}_0^\dagger \mid \mathbf{Z})$ and $\boldsymbol{\eta} = (\mathbf{Y}_0^\dagger \mid \mathbf{Z})$ given in (2.4)–(2.5), we define $\hat{\boldsymbol{\xi}} = (\hat{\mathbf{X}}_0^\dagger \mid \mathbf{Z})$ and $\hat{\boldsymbol{\eta}} = (\hat{\mathbf{Y}}_0^\dagger \mid \mathbf{Z})$.

We augment the original dataset to $\mathcal{D}_{n_2} = \{(\mathbf{X}_i, \mathbf{Y}_i, \mathbf{Z}_i, \boldsymbol{\xi}_i, \boldsymbol{\eta}_i), i \in \mathcal{I}_2\}$, where $\boldsymbol{\xi}$ and $\boldsymbol{\eta}$ are unobservable and need to be estimated. We denote the estimated version as $\mathcal{D}'_{n_2} = \{(\mathbf{X}_i, \mathbf{Y}_i, \mathbf{Z}_i, \hat{\boldsymbol{\xi}}_i, \hat{\boldsymbol{\eta}}_i), i \in \mathcal{I}_2\}$. After obtaining $\mathcal{D}'_{n_2}$, we proceed to the next step to construct the test statistic.

Step 3. Construction of the test statistic by $\mathcal{D}'_{n_2}$.

We compute the empirical version of the chosen independence measure based on $\mathcal{D}'_{n_2} = \{(\mathbf{X}_i, \mathbf{Y}_i, \mathbf{Z}_i, \hat{\boldsymbol{\xi}}_i, \hat{\boldsymbol{\eta}}_i), i \in \mathcal{I}_2\}$ to construct the test statistic. We demonstrate this procedure using distance correlation.

Following Lemma 1, we can test one of these two independence hypotheses: $\boldsymbol{\eta} \perp\!\!\!\perp (\boldsymbol{\xi}, \mathbf{Z})$ or $\boldsymbol{\xi} \perp\!\!\!\perp (\boldsymbol{\eta}, \mathbf{Z})$. Let $\mathbf{U}$ and $\mathbf{V}$ represent the paired vectors for testing. For example, to test

the independence $\boldsymbol{\eta} \perp\!\!\!\perp (\boldsymbol{\xi}, \mathbf{Z})$, we define $\hat{\mathbf{U}} = \hat{\boldsymbol{\eta}}$ and $\hat{\mathbf{V}} = (\hat{\boldsymbol{\xi}}^\top, \mathbf{Z}^\top)^\top$ using the estimated transport maps. Let $\mathbf{U} = \boldsymbol{\eta}$ and $\mathbf{V} = (\boldsymbol{\xi}^\top, \mathbf{Z}^\top)^\top$ denote the corresponding true (but unobserved) counterpart.

The sample distance covariance is computed as:

$$\mathrm{dcov}^2_{n_2}(\hat{\mathbf{U}}, \hat{\mathbf{V}}) = S_1 + S_2 - 2S_3,$$

where

$$\begin{aligned}
S_1 =& \frac{1}{n_2^2} \sum_{k,l \in \mathcal{I}_2} \|\hat{\mathbf{U}}_k - \hat{\mathbf{U}}_l\| \|\hat{\mathbf{V}}_k - \hat{\mathbf{V}}_l\|, \\
S_2 =& \frac{1}{n_2^2} \sum_{k,l \in \mathcal{I}_2} \|\hat{\mathbf{U}}_k - \hat{\mathbf{U}}_l\| \frac{1}{n_2^2} \sum_{k,l \in \mathcal{I}_2} \|\hat{\mathbf{V}}_k - \hat{\mathbf{V}}_l\|, \\
S_3 =& \frac{1}{n_2^3} \sum_{k \in \mathcal{I}_2} \sum_{l,m \in \mathcal{I}_2} \|\hat{\mathbf{U}}_k - \hat{\mathbf{U}}_l\| \|\hat{\mathbf{V}}_k - \hat{\mathbf{V}}_m\|.
\end{aligned}$$

The test statistic is defined as the squared distance correlation:

$$T_{n_2} = \mathrm{dcorr}^2_{n_2}(\hat{\mathbf{U}}, \hat{\mathbf{V}}) = \frac{\mathrm{dcov}^2_{n_2}(\hat{\mathbf{U}}, \hat{\mathbf{V}})}{\mathrm{dcov}_{n_2}(\hat{\mathbf{U}}, \hat{\mathbf{U}}) \, \mathrm{dcov}_{n_2}(\hat{\mathbf{V}}, \hat{\mathbf{V}})}, \tag{3.5}$$

where $\mathrm{dcov}_{n_2}(\hat{\mathbf{U}}, \hat{\mathbf{U}})$ and $\mathrm{dcov}_{n_2}(\hat{\mathbf{V}}, \hat{\mathbf{V}})$ are computed analogously to $\mathrm{dcov}_{n_2}(\hat{\mathbf{U}}, \hat{\mathbf{V}})$. We then compute the p-value by the permutation-based method as the following step.

Step 4. Obtaining the p-value by permutation.

To obtain the p-value for T_{n_2}, we first randomly permute $\{(\hat{\mathbf{U}}_i, \hat{\mathbf{V}}_i), i \in \mathcal{I}_2\}$ as $\{(\hat{\mathbf{U}}_i, \hat{\mathbf{V}}_{\pi_b(i)}), i \in \mathcal{I}_2\}$, where π_b is a uniformly random permutation of $\mathcal{I}_2$. A permuted statistic $T_{n_2,b}$ is then computed using (3.5) on $\{(\hat{\mathbf{U}}_i, \hat{\mathbf{V}}_{\pi_b(i)}), i \in \mathcal{I}_2\}$. We replicate the permutation B times and calculate:

$$p_B = \frac{1}{B} \sum_{b=1}^{B} \mathbf{1}(T_{n_2,b} \geq T_{n_2}). \tag{3.6}$$

We summarize the full procedure in Algorithm 1 below.

Algorithm 1 FlowCIT: Conditional independence test based on CNFs.

Require: Dataset $\mathcal{D}_n = \{(\mathbf{X}_i, \mathbf{Y}_i, \mathbf{Z}_i)\}_{i=1}^n$; Randomly split index sets $\mathcal{I}_1$ and $\mathcal{I}_2$; Number of permutations B.

Ensure: Test statistic T_{n_2}; p-value.

1: Generate $\{\mathbf{X}_{0,i}, i \in \mathcal{I}_1\} \sim \text{Normal}(\mathbf{0}, \mathbf{I}_{d_\mathbf{x}})$, $\{\mathbf{Y}_{0,i}, i \in \mathcal{I}_1\} \sim \text{Normal}(\mathbf{0}, \mathbf{I}_{d_\mathbf{y}})$, and $\{t_i, i \in \mathcal{I}_1\} \sim \text{Uniform}([0,1])$.

2: (*Step 1*) Estimate the velocity fields $\mathbf{v}_\mathbf{x}$ and $\mathbf{v}_\mathbf{y}$ by $\{(\mathbf{X}_i, \mathbf{Y}_i, \mathbf{Z}_i), i \in \mathcal{I}_1\}$, $\{\mathbf{X}_{0,i}, i \in \mathcal{I}_1\}$, $\{\mathbf{Y}_{0,i}, i \in \mathcal{I}_1\}$, $\{t_i, i \in \mathcal{I}_1\}$ as given in (3.1) and (3.2).

3: (*Step 2*) Compute $\hat{\boldsymbol{\xi}}$ and $\hat{\boldsymbol{\eta}}$ based on $\hat{\mathbf{v}}_\mathbf{x}$ and $\hat{\mathbf{v}}_\mathbf{y}$ as given in (3.3) and (3.4).

4: (*Step 3*) Compute the test statistic T_{n_2} by an independence measure based on the test dataset (e.g., distance correlation by (3.5)).

5: (*Step 4*) **For** $b = 1, \ldots, B$ **do**

- Generate a uniformly random permutation π_b of $\{1, 2, \ldots, n\}$.
- Denote the permuted dataset as $\mathcal{D}'_{n_2,b}$.
- Compute the permuted test statistic $T_{n_2,b}$ based on $\mathcal{D}'_{n_2,b}$.
- **End for**

6: Calculate the p-value as given in (3.6).

7: **return** T_{n_2}; p-value.

3.2 Asmyptotics of the Test Statistic

As outlined in Algorithm 1, our procedure involves a four-step estimation process: estimating the velocity field, estimating the transformed pairs, calculating the dependence measure of the transformed pairs, and getting the p-value.

In this section, we provide a theoretical analysis of the proposed test. We derive the asymptotic distribution of the test statistic and discuss how to properly split the data in Section 3.2. First, we formalize the convergence requirements for the estimated velocity fields $\hat{\mathbf{v}}_\mathbf{x}$ and $\hat{\mathbf{v}}_\mathbf{y}$.

Assumption 1. *As $n \to \infty$, the L_2-errors of the flow matching estimators satisfy:*

$$\mathbb{E}_{(t,\mathbf{X},\mathbf{Z})} \left\| \hat{\mathbf{v}}_\mathbf{x}(t, \mathbf{X}, \mathbf{Z}) - \mathbf{v}_\mathbf{x}(t, \mathbf{X}, \mathbf{Z}) \right\|^2 = O_p(n_1^{-2\kappa_1}),$$
$$\mathbb{E}_{(t,\mathbf{Y},\mathbf{Z})} \left\| \hat{\mathbf{v}}_\mathbf{y}(t, \mathbf{Y}, \mathbf{Z}) - \mathbf{v}_\mathbf{y}(t, \mathbf{Y}, \mathbf{Z}) \right\|^2 = O_p(n_1^{-2\kappa_2}).$$

The convergence results stated in Assumption 1 align with recent theoretical advances in flow matching (Gao et al., 2024b; Fukumizu et al., 2025; Cai and Li, 2025; Huang et al., 2024a). The parameters κ_1 and κ_2 are rates for nonparametric regression, typically ranges in $(0, 1/2)$.

Based on Assumption 1, it can inferred that the estimated transport map satisfy,

$$\mathbb{E}_{(\mathbf{X},\mathbf{Z})}\left\|\hat{\boldsymbol{\xi}} - \boldsymbol{\xi}\right\|^2 = O_p(n_1^{-2\kappa_1}),$$
$$\mathbb{E}_{(\mathbf{Y},\mathbf{Z})}\left\|\hat{\boldsymbol{\eta}} - \boldsymbol{\eta}\right\|^2 = O_p(n_1^{-2\kappa_2}).$$

We provide a detailed discussion of Assumption 1 in the Supplementary Materials E. Intuitively, if the L_2-risk of the estimated transport maps is asymptotically ignorable, the inference based on independence measure is valid.

At the population level, distance covariance admits a decomposition analogous to Pearson's covariance:

$$\mathrm{dcov}^2(\mathbf{U},\mathbf{V}) = \mathrm{cov}(\|\mathbf{U}_1 - \mathbf{U}_2\|, \|\mathbf{V}_1 - \mathbf{V}_2\|) - 2\,\mathrm{cov}(\|\mathbf{U}_1 - \mathbf{U}_2\|, \|\mathbf{V}_1 - \mathbf{V}_3\|).$$

This structure inherits the *double robustness* property: the product of estimation rates $n_1^{-\kappa_1}$ and $n_1^{-\kappa_2}$ yields faster combined convergence (Chernozhukov et al., 2018; Liu et al., 2021; Shi et al., 2021).

However, unlike classical Pearson-type statistics, under H_0 where $\mathbf{U} \perp\!\!\!\perp \mathbf{V}$, the empirical distance correlation becomes a *degenerate* V-statistic. This degeneracy yeilds an n^{-1} convergence rate instead of the standard $n^{-1/2}$. To address this, we employ a data splitting procedure (Zhang et al., 2023), allocating different orders to n_1 and n_2. Thus, based on Assumption 1, we have the following Theorem 3 regarding the asymptotic behavior of our test statistic T_n from (3.5).

Theorem 3. *Suppose Assumption 1 holds and the condition vector* $\mathbf{Z}$ *has a finite second moment. Let* $n_2 = o\big(n_1^{\kappa_1+\kappa_2}\big)$*. As* $n_1, n_2 \to \infty$*, the following convergence holds:*

1. *Under* H_0*,* $n_2 T_{n_2} \xrightarrow{d} \sum_{j=1}^{\infty} \lambda_j Z_j^2$*.*
2. *Under* H_1*,* $n_2 T_{n_2} \xrightarrow{p} \infty$*.*

Here, Z_j *are independent standard normal random variables, and* $\{\lambda_j\}$ *are nonnegative constants determined by the distribution of* $\mathbf{U}$ *and* $\mathbf{V}$*.*

If $\boldsymbol{\xi}$ and $\boldsymbol{\eta}$ are fully observed without estimation, the convergence of the empirical distance correlation is established in Székely et al. (2007). In our setting, however, the velocity fields are estimated via deep neural networks, and their convergence rates $\kappa_1, \kappa_2 \in (0, 1/2)$ depend on the smoothness and dimensionality of the velocity fields under nonparametric assumptions. As the smoothness index grows to infinity, we have $\kappa_1, \kappa_2 \to 1/2$, which simplifies the rate condition to $n_2 = o(n_1)$. Otherwise, if $\kappa_1 + \kappa_2 > 1/2$, which can be achievable (Chernozhukov et al., 2018; Shi et al., 2021), the requirement becomes $n_2 = O(\sqrt{n_1})$.

In practice, we advocate allocating $n_2 \propto \sqrt{n}$. This choice balances theoretical validity with finite-sample performance, effectively controlling type-I error as demonstrated in extensive simulations.

Since the limiting distribution under H_0 is a mixed χ^2-distribution, which is intractable, we propose using permutation to obtain a p-value as given in (3.6). The following Theorem 4 ensures the validity of this approach.

Theorem 4. *Suppose that the conditions in Theorem 3 are satisfied. As $n_1, n_2, B \to \infty$:*

1. *Under H_0, $p_B \xrightarrow{d} U \sim$ Uniform$(0,1)$.*
2. *Under H_1, $p_B \xrightarrow{p} 0$.*

The permutation test is a classical approach that has been widely used in both statistics and modern machine learning literature (Xu and Zhu, 2022; Cai et al., 2024). Under H_0, the permuted statistic shares the same limiting distribution as the original test statistic. Consequently, the limiting distribution of the p-value is standard uniform, ensuring the validity of the test. Proofs of Theorems 3 and 4 are given in Sections D.1 and D.2, respectively. While sample splitting ensures theoretical validity, it incurs power loss. We mitigate this in Section 3.3 through multiple splitting aggregation.

3.3 Combining Results from Multiple Splitting

In this section, we adopt a multiple splitting strategy analogous to the classical K-fold approach. Specifically, we split the data into $\lfloor n/n_2 \rfloor$ disjoint folds. At each iteration, we use one fold as the test set and the remaining folds for training, and compute a p-value following Algorithm 1. To aggregate $m \leq \lfloor n/n_2 \rfloor$ p-values $p_{B,1}, \ldots, p_{B,m}$, we employ the Cauchy combination method (Liu and Xie, 2020):

$$T = \sum_{i=1}^{m} \frac{1}{m} \tan\{(0.5 - p_{B,i})\pi\}.$$

When $p_{B,i}$ follows a standard uniform distribution, the transformation $\tan\{(0.5-p_{B,i})\pi\}$ follows a standard Cauchy distribution. Under independence of $p_{B,i}$, their average T preserves the standard Cauchy distribution. We then transform T back to yield

$$p_c = 0.5 - (\arctan T)/\pi. \tag{3.7}$$

Since the p-values are obtained from disjoint test sets and the estimation error from the training set is ignorable based on Theorem 3, we have the following Corollary 5.

Corollary 5. *Let $m \leq \lfloor n/n_2 \rfloor$ and the test folds be disjoint. Then, the combined p-value p_c given in (3.7) satisfies:*

1. *Under H_0, $p_c \xrightarrow{d} U \sim Uniform(0,1)$.*
2. *Under H_1, $p_c \xrightarrow{p} 0$.*

4 Simulations

We validate the proposed method through extensive simulations. In Section 4.1, we focus on demonstrating the convergence of the p-values under H_0, as given in Theorem 4 and Corollary 5. We will present the empirical type-I error and statistical power performance, as well as comparisons with other tests, in Section 4.2.

4.1 Convergence of the Test Statistic

Model 1 (Convergence under H_0)**.** *Dimensions $(d_{\mathbf{x}}, d_{\mathbf{y}}, d_{\mathbf{z}})$ and sample size n vary across experiments. Data generation proceeds as follows:*

1. *Generate parameter matrices $\mathbf{B}_1 \in \mathbb{R}^{d_{\mathbf{z}} \times d_{\mathbf{x}}}$ and $\mathbf{B}_2 \in \mathbb{R}^{d_{\mathbf{z}} \times d_{\mathbf{y}}}$, where each element is independently drawn from $Normal(0,1)$.*
2. *Generate error terms $\boldsymbol{\epsilon}_{\mathbf{x}} \sim Normal(\mathbf{0}, \mathbf{I}_{d_{\mathbf{x}}})$ and $\boldsymbol{\epsilon}_{\mathbf{y}} \sim Normal(\mathbf{0}, \mathbf{I}_{d_{\mathbf{y}}})$.*
3. *Generate $\mathbf{Z} \sim Normal(\mathbf{0}, \mathbf{I}_{d_{\mathbf{z}}})$, then compute $\mathbf{X} = \mathbf{Z}\mathbf{B}_1 + \boldsymbol{\epsilon}_{\mathbf{x}}$ and $\mathbf{Y} = \mathbf{Z}\mathbf{B}_2 + \boldsymbol{\epsilon}_{\mathbf{y}}$.*

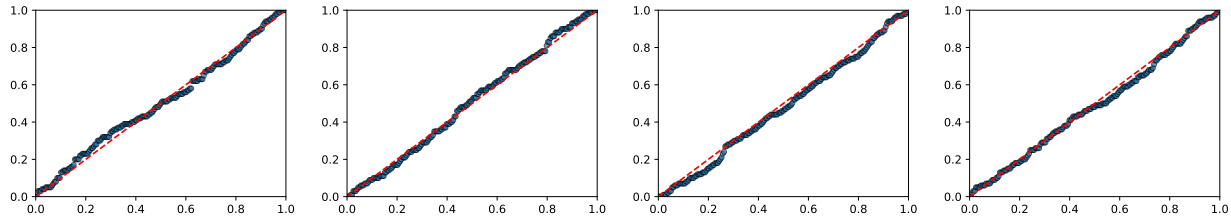

(a) (b) (c) (d)

Figure 1: Finite-sample performance of Theorem 4 under H_0. Q-Q plots compare empirical p-values against the standard uniform distribution for Model 1. (a) $(d_{\mathbf{x}}, d_{\mathbf{y}}, d_{\mathbf{z}}) = (1,1,1)$ and $n = 500$; (b) $(d_{\mathbf{x}}, d_{\mathbf{y}}, d_{\mathbf{z}}) = (1,1,1)$ and $n = 1000$; (c) $(d_{\mathbf{x}}, d_{\mathbf{y}}, d_{\mathbf{z}}) = (5,5,5)$ and $n = 500$; and (d) $(d_{\mathbf{x}}, d_{\mathbf{y}}, d_{\mathbf{z}}) = (5,5,5)$ and $n = 1000$. Results are based on 200 replications.

To perform the test, we split the data with $n_2 = \lfloor 4\sqrt{n} \rfloor$, following the recommendations that $n_2 \propto \sqrt{n}$. We test the hypothesis $\boldsymbol{\eta} \perp\!\!\!\perp (\boldsymbol{\xi}, \mathbf{Z})$ under two configurations: $(d_{\mathbf{x}}, d_{\mathbf{y}}, d_{\mathbf{z}}) = (1,1,1)$ and $(5,5,5)$, with sample sizes $n = 500$ and $n = 1000$.

Figure 1 displays Q-Q plots comparing empirical p-values to the standard uniform distribution. Each subplot corresponds to different dimension and sample size combinations. The close alignment between empirical p-values and the theoretical uniform distribution supports the validity of our asymptotic results.

We further validate the multiple splitting procedure from Section 3.3 by repeating the split $m = 5$ times. The corresponding results, shown in the Supplementary Materials B.1 (Figure S.1). mirror the patterns observed in Figure 1. These findings collectively suggest that Theorem 4 and Corollary 5 hold in finite samples.

4.2 Comparisons with Existing Methods

In this section, we compare our test with the following nonparametric conditional independence test methods: kernel-based conditional independence test (KCI, Zhang et al. 2011), conditional distance correlation test (CDC, Wang et al. 2015), classifier conditional independence test (CCIT, Sen et al. 2017), fast conditional independence test (FCIT, Chalupka et al. 2018), a test proposed by Zhou et al. (2020) (abbreviated as ZLZ), and a test proposed by Cai et al. (2022) (abbreviated as CLZ). The KCI test is implemented by the `R` package `CondIndTests::KCI`, the CDC and FCIT are implemented by the `Python` package `hyppo` (Panda et al., 2024), the CCIT is implemented by its `Python` package `CCIT`, the ZLZ test is conducted based on its original `Matlab` code, and the CLZ test is implemented by its original `R` code. Please refer to the Supplementary Materials B.6 for the details of implementations of these methods.

We consider four different examples in Models S.1, 2, 3, and 4. These models range from linear to nonlinear examples and include univariate, low-dimensional, and moderately high-dimensional scenarios.

As a prelude, we summarize the comparisons of these methods based on our simulation models in Table 1. Additionally, our method effectively controls the type-I error under H_0 and achieves high statistical power performance across all simulation models.

Table 1: Comparisons of nonparametric conditional independence test methods across our simulation settings. A checkmark (✔) indicates that the method is workable under the given setting—effectively controlling the type-I error, demonstrating statistical power against alternatives, and being computationally feasible. A triangle (Δ) indicates that the method is partially workable. Cells are left blank if a method fails to show detectable statistical power under our model settings or is computationally infeasible in practice.

Dimensions		Examples	FlowCIT	KCI	CDC	FCIT	CCIT	ZLZ	CLZ
$\mathbf{Z}$	$\mathbf{X}$ and $\mathbf{Y}$								
Low	Univariate	Model S.1	✔	✔	✔	✔	✔	✔	✔
	Low	Model 2	✔	✔	✔	Δ			
High	Low	Model 3	✔			Δ	Δ		
	High	Model 4	✔			Δ	Δ		

In these models, we use the parameter ψ to control the deviation from H_0: if $\psi = 0$, the model is under H_0; otherwise, a larger $|\psi|$ indicates a greater deviation. We use plots to visualize the empirical type-I errors and statistical powers across various settings.

Since the data points are densely clustered under H_0, detailed empirical type-I error rates are provided in Supplementary Material B.5. In this section, we implement our method using distance correlation to test the independence of $\boldsymbol{\eta}$ and $(\boldsymbol{\xi}, \mathbf{Z})$. We also observe that testing the independence of $\boldsymbol{\xi}$ and $(\boldsymbol{\eta}, \mathbf{Z})$ yields similar performance in Appendix B.3. A comparison between distance correlation and improved projection correlation is presented in Appendix B.4, indicating that their performances are comparable. All results are based on 200 replications and the data is split with $n_2 = \lfloor 4\sqrt{n} \rfloor$.

4.2.1 Low-dimensional Z

We start with examples involving low-dimensional $\mathbf{Z}$. A univariate X and Y example (Model S.1) is deferred to Supplementary Materials B.2. In the following Model 2, we analyze multivariate pairs $\mathbf{X}$ and $\mathbf{Y}$, with all three vectors specified as three-dimensional. The multiple splitting is repeated $m = 5$ times in this section.

Model 2 (Multivariate $\mathbf{X}$, $\mathbf{Y}$, and $\mathbf{Z}$)**.** *Let* $(d_{\mathbf{x}}, d_{\mathbf{y}}, d_{\mathbf{x}}) = (3, 3, 3)$ *and the sample size* $n = 500$. *Each element of* $\mathbf{B}_1$, $\mathbf{B}_2$, $\mathbf{B}_3$, *and* $\mathbf{Z}$ *are generated from Normal*$(0, 1)$. *Vectors* $\mathbf{X}$ *and* $\mathbf{Y}$ *are generated under Settings 2.1–2.4, and we vary* ψ *over the set* $\{0, 0.05, 0.1, 0.15, 0.2\}$. *For the error terms, we generate* $\boldsymbol{\epsilon}_{\mathbf{y}} \sim$ *Normal*$(\mathbf{0}, \mathbf{I}_{d_{\mathbf{y}}})$, *and* $\boldsymbol{\epsilon}_{\mathbf{x}} \sim$ *Normal*$(\mathbf{0}, \mathbf{I}_{d_{\mathbf{x}}})$.

Setting 2.1. $\mathbf{X} = \mathbf{Z}\mathbf{B}_1 + \boldsymbol{\epsilon}_{\mathbf{x}}, \quad \mathbf{Y} = \mathbf{Z}\mathbf{B}_2 + \psi\mathbf{X}\mathbf{B}_3 + \boldsymbol{\epsilon}_{\mathbf{y}}.$

Setting 2.2. $\mathbf{X} = \mathbf{Z}\mathbf{B}_1 + \boldsymbol{\epsilon}_{\mathbf{x}}, \quad \mathbf{Y} = \sin(\mathbf{Z}\mathbf{B}_2) + \psi\mathbf{X}\mathbf{B}_3 + \boldsymbol{\epsilon}_{\mathbf{y}}.$

Setting 2.3. $\mathbf{X} = (\mathbf{Z}\mathbf{B}_1)^2 + \boldsymbol{\epsilon}_\mathbf{x}, \quad \mathbf{Y} = \mathbf{Z}\mathbf{B}_2 + \psi\mathbf{X}\mathbf{B}_3 + \boldsymbol{\epsilon}_\mathbf{y}.$

Setting 2.4. $\mathbf{X} = \mathbf{Z}\mathbf{B}_1 + \boldsymbol{\epsilon}_\mathbf{x}, \quad \mathbf{Y} = \mathbf{Z}\mathbf{B}_2 + |\psi\mathbf{X}\mathbf{B}_3| + \boldsymbol{\epsilon}_\mathbf{y}.$

The results for Model 2 are shown in Figure 2, with the x-axis corresponding to ψ. Since ZLZ is designed for univariate X and Y, we focus on comparing the remaining six methods. Under H_0, all six methods effectively control the type-I error. However, under H_1, CCIT and CLZ demonstrate reduced statistical power. FCIT shows limited statistical power in the linear Setting 2.1 but achieve modest statistical power in the nonlinear Settings 2.2, 2.3, and 2.4, though still under-performing FlowCIT, KCI, and CDC. Our proposed FlowCIT consistently attains the highest or comparable statistical power to KCI and CDC across all settings.

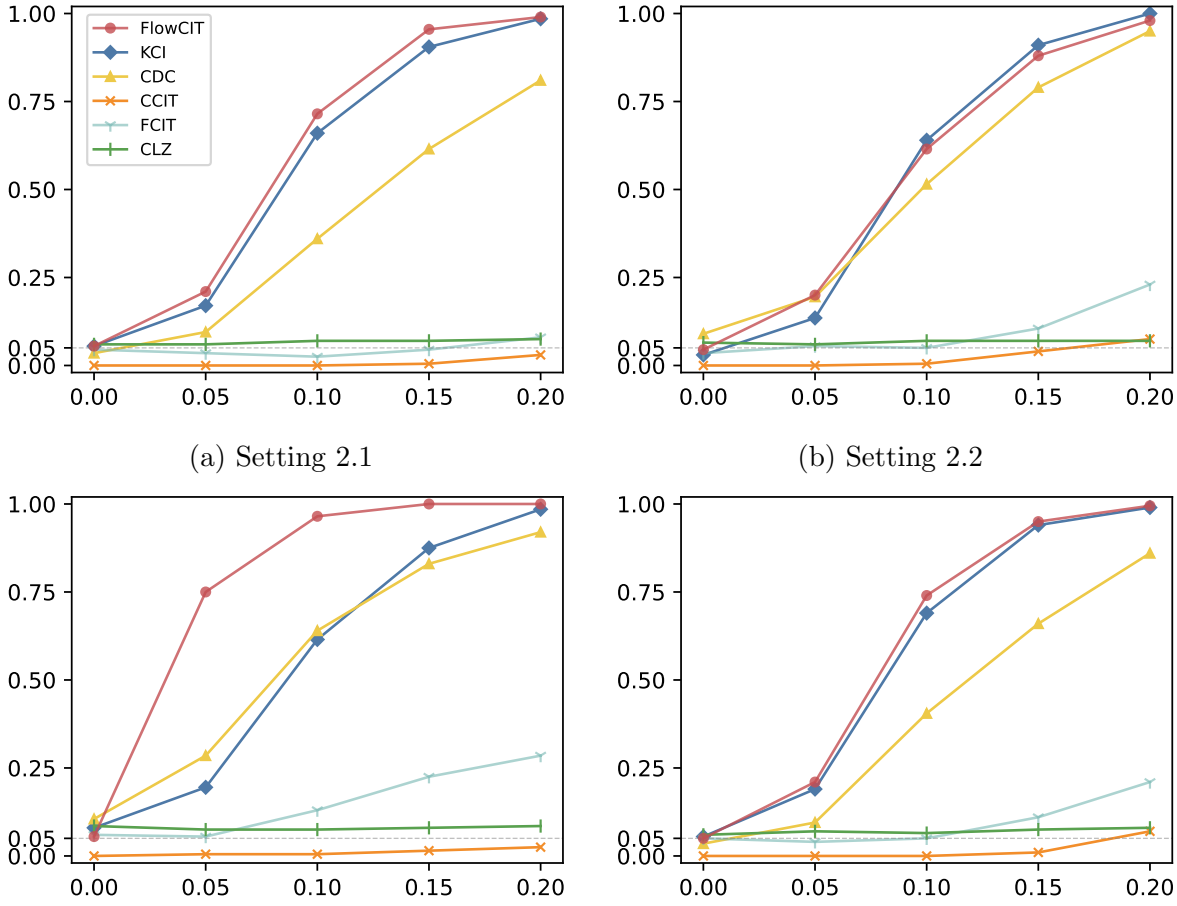

(a) Setting 2.1

(b) Setting 2.2

(c) Setting 2.3

(d) Setting 2.4

Figure 2: Empirical type-I errors and statistical powers of Model 2 with sample size $n = 500$. The dashed gray line below refers to significance level of 0.05. The x-axis refers to the value of ψ.

4.2.2 Moderately high-dimensional Z

Since CLZ and KCI are kernel-based methods that struggle in high-dimensional settings due to computational burdens, we limit our comparison in Models 3–4 to CDC, FCIT, and CCIT. The multiple splitting is repeated $m = 2$ times in this section.

In high-dimensional $\mathbf{Z}$ examples (Models 3–4), we first consider a low-dimensional pair $(\mathbf{X}, \mathbf{Y})$ in Model 3, and then increase the dimensions of $(\mathbf{X}, \mathbf{Y})$ in Model 4. In these models, our test demonstrates higher statistical power compared to other methods. FCIT and CCIT are effective in certain settings but ineffective in others.

Model 3 (Low-dimensional $\mathbf{X}$ and $\mathbf{Y}$, and moderately high-dimensional $\mathbf{Z}$)**.** *Let* $(d_{\mathbf{x}}, d_{\mathbf{y}}, d_{\mathbf{z}}) = (5, 5, 50)$ *and the sample size* $n = 1000$*. Each element of* $\boldsymbol{\epsilon}_{\mathbf{x}}$*,* $\boldsymbol{\epsilon}_{\mathbf{y}}$*, and* $\mathbf{Z}$ *is generated from Normal*$(0, 1)$*. Matrices* $\mathbf{B}_1$*,* $\mathbf{B}_2$*,* $\mathbf{B}_3$*, and vectors* $\mathbf{X}$ *and* $\mathbf{Y}$ *are generated under Settings 3.1–3.4, and the parameter* ψ *is varied over the set* $\{0, 0.1, 0.2, 0.3, 0.4\}$*.*

Setting 3.1 (Row-sparse $\mathbf{B}_1$ and $\mathbf{B}_2$, and dense $\mathbf{B}_3$)*. Each element of* $\mathbf{B}_3$*, as well as the first* $3 \times d_{\mathbf{x}}$ *sub-matrix of* $\mathbf{B}_1$*, and the first* $3 \times d_{\mathbf{y}}$ *sub-matrix of* $\mathbf{B}_2$*, is generated from Normal*$(0, 1)$*. All other elements of* $\mathbf{B}_1$ *and* $\mathbf{B}_2$ *are set to zero. We then generate:* $\mathbf{X} = \mathbf{Z}\mathbf{B}_1 + \boldsymbol{\epsilon}_{\mathbf{x}}, \quad \mathbf{Y} = \mathbf{Z}\mathbf{B}_2 + \psi\mathbf{X}\mathbf{B}_3 + \boldsymbol{\epsilon}_{\mathbf{y}}$.

Setting 3.2 (Sparse $\mathbf{B}_1$, $\mathbf{B}_2$, and $\mathbf{B}_3$)*. Each element of the first* 3×1 *sub-matrices of* $\mathbf{B}_1$*,* $\mathbf{B}_2$*, and* $\mathbf{B}_3$ *is generated from Normal*$(0, 1)$*. All other elements of* $\mathbf{B}_1$*,* $\mathbf{B}_2$*, and* $\mathbf{B}_2$ *are set to zero. We then generate:* $\mathbf{X} = \mathbf{Z}\mathbf{B}_1 + \boldsymbol{\epsilon}_{\mathbf{x}}, \quad \mathbf{Y} = (\mathbf{Z}\mathbf{B}_2)^2 + 4\psi\mathbf{X}\mathbf{B}_3 + \boldsymbol{\epsilon}_{\mathbf{y}}$.

Setting 3.3 (Row-sparse $\mathbf{B}_1$ and $\mathbf{B}_2$, and dense $\mathbf{B}_3$)*. Each element of* $\mathbf{B}_3$*, as well as the first* $2 \times d_{\mathbf{x}}$ *sub-matrix of* $\mathbf{B}_1$*, and the first* $2 \times d_{\mathbf{y}}$ *sub-matrix of* $\mathbf{B}_2$*, is generated from Normal*$(0, 1)$*. All other elements of* $\mathbf{B}_1$*,* $\mathbf{B}_2$*, and* $\mathbf{B}_3$ *are set to zero. We then generate:* $\mathbf{X} = \mathbf{Z}\mathbf{B}_1 + \boldsymbol{\epsilon}_{\mathbf{x}}, \quad \mathbf{Y} = \mathbf{Z}\mathbf{B}_2 + \psi\mathbf{X}\mathbf{B}_3 + \boldsymbol{\epsilon}_{\mathbf{y}}$.

Setting 3.4 (Dense $\mathbf{B}_1$, $\mathbf{B}_2$, and $\mathbf{B}_3$)*. Each element of* $\mathbf{B}_1$*,* $\mathbf{B}_2$*, and* $\mathbf{B}_3$ *is generated from Normal*$(0, 1)$*. We then generate:* $\mathbf{X} = \sin(\mathbf{Z}\mathbf{B}_1) + \boldsymbol{\epsilon}_{\mathbf{x}}, \quad \mathbf{Y} = \mathbf{Z}\mathbf{B}_2 + |5\psi\mathbf{X}\mathbf{B}_3| + \boldsymbol{\epsilon}_{\mathbf{y}}$.

The results for Model 3 are shown in Figure 3. In sparse settings (Settings 3.1–3.3), our method demonstrates robust performance. FCIT and CCIT also achieve statistical power under the alternative hypothesis, though they underperform compared to our method. In the dense setting (Setting 3.4), our method retains its effectiveness, while FCIT and CCIT show no statistical power under the alternative hypothesis.

In the next Model 4, we increase the dimension of $\mathbf{X}$ and $\mathbf{Y}$, considering settings with $(d_{\mathbf{x}}, d_{\mathbf{y}}, d_{\mathbf{z}}) = (50, 50, 50)$.

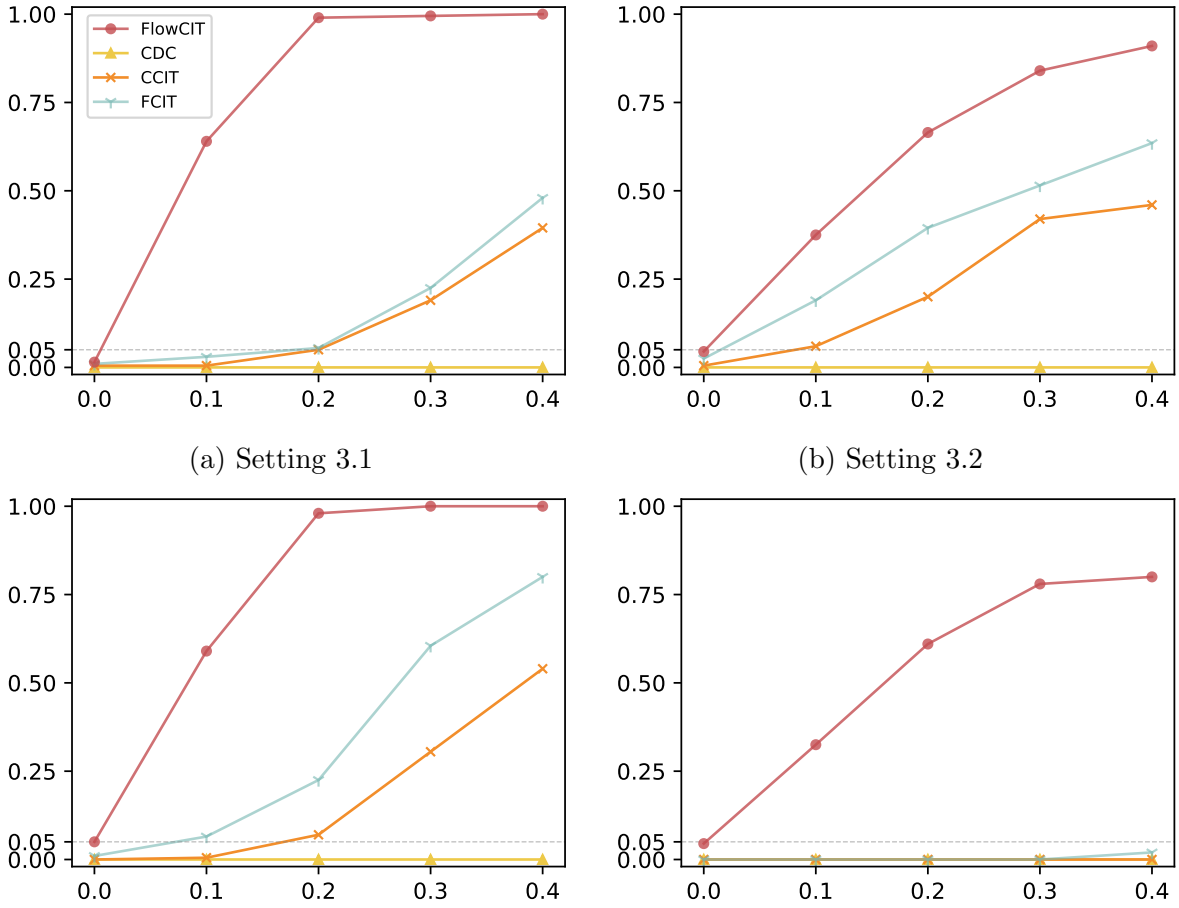

(a) Setting 3.1

(b) Setting 3.2

(c) Setting 3.3

(d) Setting 3.4

Figure 3: Empirical type-I errors and statistical powers of Model 3 with sample size $n = 1000$. The dashed gray line below refers to significance level of 0.05. The x-axis refers to the value of ψ.

Model 4 (Moderately high-dimensional $\mathbf{X}$, $\mathbf{Y}$, and $\mathbf{Z}$)**.** *Let* $(d_{\mathbf{x}}, d_{\mathbf{y}}, d_{\mathbf{z}}) = (50, 50, 50)$ *and the sample size* $n = 1000$*. Each element of* $\boldsymbol{\epsilon}_{\mathbf{x}}$*,* $\boldsymbol{\epsilon}_{\mathbf{y}}$*, and* $\mathbf{Z}$ *is generated from Normal*$(0, 1)$*. Matrices* $\mathbf{B}_1$*,* $\mathbf{B}_2$*,* $\mathbf{B}_3$*, and vectors* $\mathbf{X}$ *and* $\mathbf{Y}$ *are generated under Settings 4.1–4.4, and the parameter* ψ *is varied over the set* $\{0, 0.2, 0.4, 0.6, 0.8\}$*.*

Setting 4.1 (Row-sparse $\mathbf{B}_1$ and $\mathbf{B}_2$, and dense $\mathbf{B}_3$)*. Each element of* $\mathbf{B}_3$ *and the first two rows of* $\mathbf{B}_1$ *and* $\mathbf{B}_2$ *is generated from Normal*$(0, 1)$*. All other elements of* $\mathbf{B}_1$ *and* $\mathbf{B}_2$ *are set to zero. The sparsity percentages of* $\mathbf{B}_1$ *and* $\mathbf{B}_2$ *are* 4%*. We then generate:*

$$\mathbf{X} = \mathbf{Z}\mathbf{B}_1 + \boldsymbol{\epsilon}_{\mathbf{x}}, \quad \mathbf{Y} = \mathbf{Z}\mathbf{B}_2 + \psi\mathbf{X}\mathbf{B}_3 + \boldsymbol{\epsilon}_{\mathbf{y}}.$$

Setting 4.2 (Row-sparse $\mathbf{B}_1$ and $\mathbf{B}_2$, and dense $\mathbf{B}_3$)*. Each element of* $\mathbf{B}_3$ *and the first row*

of $\mathbf{B}_1$ *and* $\mathbf{B}_2$ *is generated from Normal*$(0, 1)$*. The sparsity percentages of* $\mathbf{B}_1$ *and* $\mathbf{B}_2$ *are* 2%*. The other setttings are the same as Setting 4.1.*

Setting 4.3 (Sparse $\mathbf{B}_1$, $\mathbf{B}_2$, and $\mathbf{B}_3$). *Each element of the first* 3×3 *sub-matrices of* $\mathbf{B}_1$*,* $\mathbf{B}_2$*, and* $\mathbf{B}_3$ *is generated from Normal*$(0, 1)$*. All other elements of* $\mathbf{B}_1$*,* $\mathbf{B}_2$*, and* $\mathbf{B}_2$ *are set to zero. The sparsity percentages of* $\mathbf{B}_1$*,* $\mathbf{B}_2$*, and* $\mathbf{B}_3$ *are* 0.18%*,* 0.18%*, and* 0.36%*. We then generate:*

$$\mathbf{X} = \mathbf{Z}\mathbf{B}_1 + \boldsymbol{\epsilon}_{\mathbf{x}}, \quad \mathbf{Y} = \mathbf{Z}\mathbf{B}_2 + |\psi \mathbf{X}\mathbf{B}_3| + \boldsymbol{\epsilon}_{\mathbf{y}}.$$

Setting 4.4 (Sparse $\mathbf{B}_1$, $\mathbf{B}_2$, and $\mathbf{B}_3$). *Each element of matrices* $\mathbf{B}_1$*,* $\mathbf{B}_2$*, and* $\mathbf{B}_3$ *is randomly taken 0 with probability* 90% *and generated from Normal*$(0, 1)$ *with probability* 10%*. The expectation of the sparsity percentages of these matrices are* 10%*. We then generate:*

$$\mathbf{X} = \cos(\mathbf{Z}\mathbf{B}_1) + \boldsymbol{\epsilon}_{\mathbf{x}}, \quad \mathbf{Y} = \sin(\mathbf{Z}\mathbf{B}_2) + \psi \mathbf{X}\mathbf{B}_3 + \boldsymbol{\epsilon}_{\mathbf{y}}.$$

The results for Model 4 are shown in Figure 4. Our method effectively controls the type-I error and demonstrates greater statistical power compared to other methods. FCIT is partially effective in the scenarios with sparse $\mathbf{B}_1$, $\mathbf{B}_2$, and $\mathbf{B}_3$ when the sparsity percentage is below 1% (Setting 4.3). CCIT is partially workable in scenarios with row-sparse $\mathbf{B}_1$ and $\mathbf{B}_2$ and dense $\mathbf{B}_3$ (Setting 4.1 and Setting 4.2). When the sparsity percentage increases to around 10% (Setting 4.4), our proposed FlowCIT is more effective.

5 Application to Dimension Reduction

Conditional independence testing holds broad applicability across statistical methodologies. We demonstrate its practical utility through real-data applications in dimension reduction.

Dimension reduction is a cornerstone of modern statistics and machine learning. Given paired data $(\mathbf{Y}, \mathbf{X})$, the goal is to identify a low-dimensional representation $\mathcal{R}(\mathbf{X})$—linear or nonlinear—that satisfies the conditional independence criterion: $\mathbf{Y} \perp\!\!\!\perp \mathbf{X} \mid \mathcal{R}(\mathbf{X})$. The efficacy of this reduction hinges critically on both the dimensionality and function form of $\mathcal{R}(\mathbf{X})$. To rigorously evaluate whether a candidate $\mathcal{R}(\mathbf{X})$ achieves sufficiency, we can leverage our conditional independence test framework.

We evaluate dimension reduction methods using the wine quality dataset from Cortez et al. (2009), available at `https://archive.ics.uci.edu/dataset/186/wine+quality`. The dataset includes $n = 4898$ samples of white wine, with predictors (fixed acidity, volatile acidity, citric acid, residual sugar, chlorides, free sulfur dioxide, total sulfur dioxide, density, pH, sulfates, alcohol) and a response variable (wine quality rating).

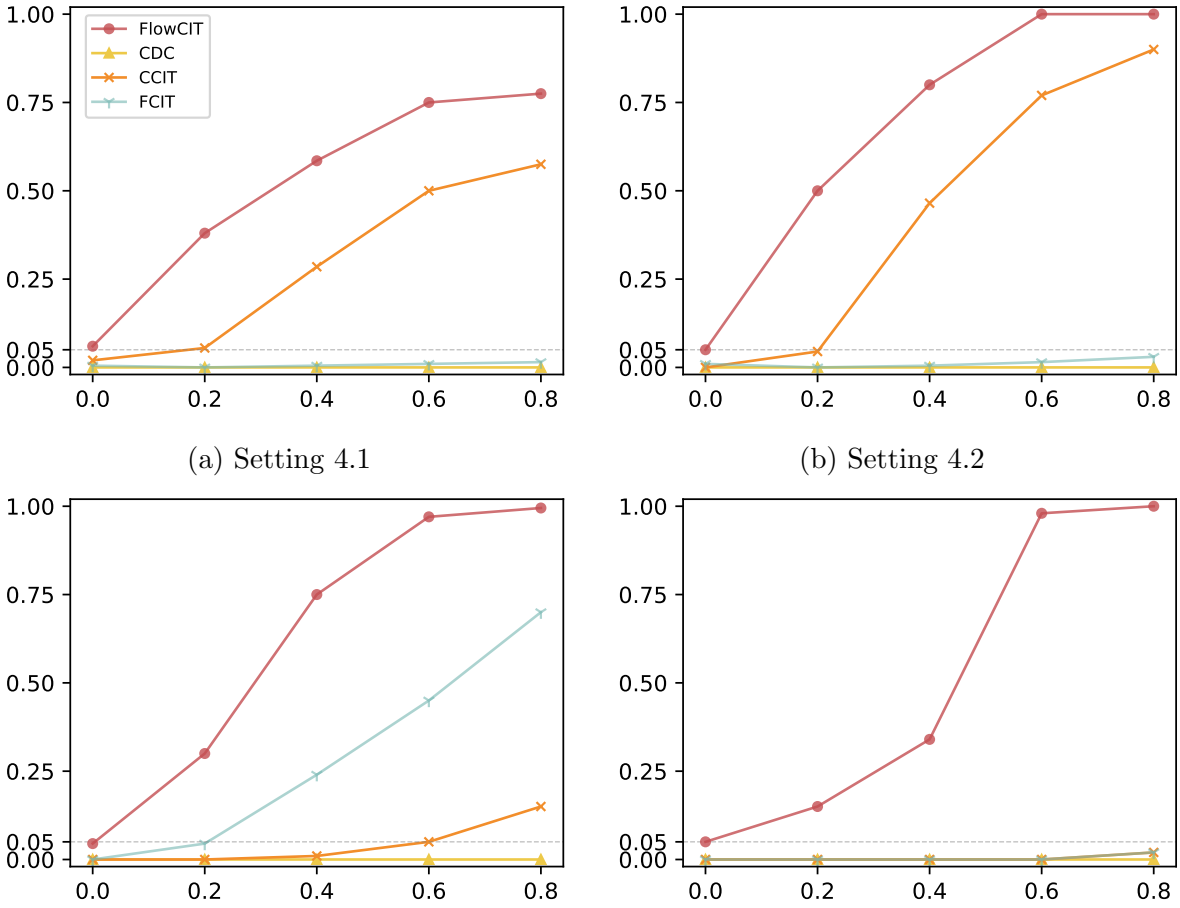

(a) Setting 4.1

(b) Setting 4.2

(c) Setting 4.3

(d) Setting 4.4

Figure 4: Empirical type-I errors and statistical powers of Model 4 with sample size $n = 1000$. The dashed gray line below refers to significance level of 0.05. The x-axis refers to the value of ψ.

Our goal is to test whether a dimension reduction method, represented as $\mathcal{R}(\mathbf{X})$, satisfies the conditional independence $\mathbf{Y} \perp\!\!\!\perp \mathbf{X} \mid \mathcal{R}(\mathbf{X})$. We evaluate four dimension reduction methods:

1. Principal Component Analysis (PCA, Pearson 1901), an unsupervised linear method;
2. Sliced Inverse Regression (SIR, Li 1991), a supervised linear method;
3. Uniform Manifold Approximation and Projection (UMAP, McInnes et al. 2018), an unsupervised nonlinear method;
4. Deep Dimension Reduction (DDR, Huang et al. 2024b), a supervised nonlinear method.

For each method, we test reduced dimensions $d = 1$ and $d = 4$. The data are split

into training (90%) and testing (10%) subsets: $\mathcal{R}(\mathbf{X})$ is learned on the training set, and conditional independence is tested on the held-out set. As a baseline, we set $\mathcal{R}(\mathbf{X}) = \mathbf{X}$, which trivially satisfies $\mathbf{Y} \perp\!\!\!\perp \mathbf{X} \mid \mathbf{X}$.

Table 2: p-values of conditional independence test of $\mathbf{Y} \perp\!\!\!\perp \mathbf{X} \mid \mathcal{R}(\mathbf{X})$ after learning $\mathcal{R}(\mathbf{X})$ via specific methods. Bold: $p \leq 0.05$; Underlined: $0.05 < p \leq 0.1$.

Methods	Supervised	Nonlinear	FlowCIT	KCI	CDC	CCIT	FCIT	CLZ
PCA-1	✗	✗	**0.00**	**0.04**	**0.01**	<u>0.09</u>	**0.00**	**0.00**
PCA-4	✗	✗	**0.00**	**0.00**	**0.01**	0.68	0.27	**0.00**
SIR-1	✔	✗	**0.04**	<u>0.09</u>	**0.02**	0.19	0.94	**0.00**
SIR-4	✔	✗	**0.02**	**0.00**	**0.05**	0.27	0.79	<u>0.07</u>
UMAP-1	✗	✔	**0.00**	<u>0.08</u>	**0.01**	0.29	0.22	**0.00**
UMAP-4	✗	✔	**0.00**	**0.00**	**0.01**	0.37	**0.01**	**0.00**
DDR-1	✔	✔	**0.01**	0.12	0.13	0.32	0.59	0.99
DDR-4	✔	✔	0.30	0.14	0.38	0.53	0.82	0.81
$\mathcal{R}(\mathbf{X}) = \mathbf{X}$			0.71	0.52	0.28	0.68	0.18	1.00

Results are summarized in Table 2. Under the baseline ($\mathcal{R}(\mathbf{X}) = \mathbf{X}$), all tests fail to reject H_0, confirming the validity of the testing framework. For PCA, SIR, and UMAP, four tests (FlowCIT, KCI, CDC, and CLZ) reject H_0 at $\alpha = 0.1$, indicating $\mathbf{Y} \not\perp\!\!\!\perp \mathbf{X} \mid \mathcal{R}(\mathbf{X})$ and suggesting these dimension reduction methods inadequately preserve sufficient information. In contrast, the supervised nonlinear DDR method with $d = 4$ satisfies $\mathbf{Y} \perp\!\!\!\perp \mathbf{X} \mid \mathcal{R}(\mathbf{X})$ across all tests. In addition, DDR with $d = 1$ is rejected by FlowCIT, which suggests that reducing predictors into one dimension is not sufficient.

6 Conclusion

In this work, we propose a novel framework for conditional independence testing that utilizes transport maps. By constructing two transport maps using conditional continuous normalizing flows, we transform the conditional independence problem into an unconditional independence test on the transformed variables. This transformation allows for the application of standard independence measures while preserving statistical validity. Numerical experiments demonstrate the framework's superior performance across a variety of scenarios, including univariate, low-dimensional, and high-dimensional settings, and demonstrate its robustness to the choice of dependence measures.

In future work, we plan to explore the use of alternative transport maps for converting conditional independence problems into unconditional ones. We will also consider applying

the proposed method to other formulations of conditional independence to further enhance the framework's applicability and effectiveness.

APPENDIX

In the Appendix, we include additional discussions, numerical studies, and proofs for the main paper.

A Improved Projection Correlation

We first provide a brief introduction for the improved projection correlation (IPC, Zhang and Zhu 2024) in this section. For independence copies $(\boldsymbol{\xi}_1, \boldsymbol{\eta}_1)$, $(\boldsymbol{\xi}_2, \boldsymbol{\eta}_2)$, and $(\boldsymbol{\xi}_3, \boldsymbol{\eta}_3)$, recall the distance covariance defined in (2.7):

$$\begin{aligned}\mathrm{dcov}^2(\boldsymbol{\xi}, \boldsymbol{\eta}) =& \mathbb{E}\left(\|\boldsymbol{\xi}_1 - \boldsymbol{\xi}_2\|\|\boldsymbol{\eta}_1 - \boldsymbol{\eta}_2\|\right) + \mathbb{E}\|\boldsymbol{\xi}_1 - \boldsymbol{\xi}_2\|\mathbb{E}\|\boldsymbol{\eta}_1 - \boldsymbol{\eta}_2\| \\ & - 2\mathbb{E}\left(\|\boldsymbol{\xi}_1 - \boldsymbol{\xi}_2\|\|\boldsymbol{\eta}_1 - \boldsymbol{\eta}_3\|\right).\end{aligned}$$

Improved projection covariance is defined in a similar manner:

$$\begin{aligned}\mathrm{ipcov}^2(\boldsymbol{\xi}, \boldsymbol{\eta}) =& \mathbb{E}\left\{A(\boldsymbol{\xi}_1, \boldsymbol{\xi}_2)A(\boldsymbol{\eta}_1, \boldsymbol{\eta}_2)\right\} + \mathbb{E}\{A(\boldsymbol{\xi}_1, \boldsymbol{\xi}_2)\}\mathbb{E}\{A(\boldsymbol{\eta}_1, \boldsymbol{\eta}_2)\} \\ & - 2\mathbb{E}\left\{A(\boldsymbol{\xi}_1, \boldsymbol{\xi}_2)A(\boldsymbol{\eta}_1, \boldsymbol{\eta}_3)\right\},\end{aligned}$$

where $A(\mathbf{U}_1, \mathbf{U}_2) = \arccos\{(\sigma_u^2 + \mathbf{U}_1^\top \mathbf{U}_2)(\sigma_u^2 + \mathbf{U}_1^\top \mathbf{U}_1)^{-1/2}(\sigma_u^2 + \mathbf{U}_2^\top \mathbf{U}_2)\}$ is the arc-cosine function, and σ_z^2 is suggested to choose as the median of $\mathbf{U}^\top \mathbf{U}$. The improved projection correlation is defined as

$$\mathrm{ipcorr}^2(\boldsymbol{\xi}, \boldsymbol{\eta}) = \frac{\mathrm{ipcov}^2(\boldsymbol{\xi}, \boldsymbol{\eta})}{\mathrm{ipcov}(\boldsymbol{\xi}, \boldsymbol{\xi})\,\mathrm{ipcov}(\boldsymbol{\eta}, \boldsymbol{\eta})}.$$

In contrast to DC, IPC imposes no moment conditions on the random vectors, rendering it more robust. Furthermore, Zhang and Zhu (2024) demonstrates that under mild regularity conditions, the scaled statistic $n\,\mathrm{ipcorr}^2(\boldsymbol{\xi}, \boldsymbol{\eta})$ converges in distribution to a normal limit, thereby yielding a tractable asymptotic distribution for hypothesis testing.

B Additional Details of Numerical Studies

B.1 Additional Results for the Convergence of the Test Statistic

We provide supplementary evidence for Corollary 5 under H_0 through Figure S.1. These results mirror the patterns observed in Figure 1, demonstrating close alignment with the standard uniform distribution.

In Figure S.1c, a slight deviation from the uniform distribution is observed. This discrepancy arises because the model operates in a relatively high-dimensional space ($d_{\mathbf{x}}+d_{\mathbf{y}}+d_{\mathbf{z}} = 15$) with a limited sample size (n=500), and the Cauchy combination method amplifies such deviations. We further note that the empirical type-I error in this setting is 0.06 under the nominal significance level $\alpha = 0.05$. While this inflation exceeds the nominal level, the deviation remains acceptable in practice. When the sample size increases to $n = 1000$, the empirical p-value aligns closely with the standard uniform distribution.

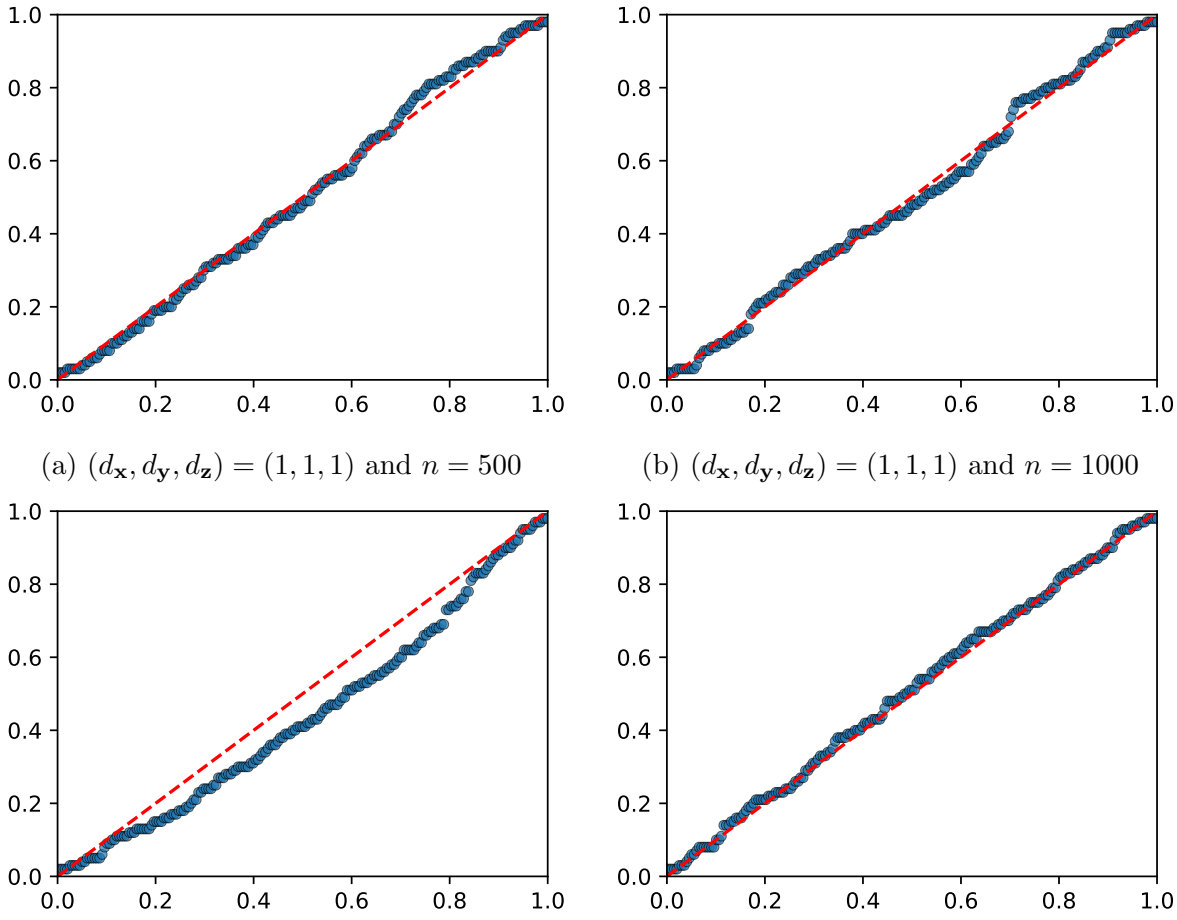

(a) $(d_{\mathbf{x}}, d_{\mathbf{y}}, d_{\mathbf{z}}) = (1, 1, 1)$ and $n = 500$

(b) $(d_{\mathbf{x}}, d_{\mathbf{y}}, d_{\mathbf{z}}) = (1, 1, 1)$ and $n = 1000$

(c) $(d_{\mathbf{x}}, d_{\mathbf{y}}, d_{\mathbf{z}}) = (5, 5, 5)$ and $n = 500$

(d) $(d_{\mathbf{x}}, d_{\mathbf{y}}, d_{\mathbf{z}}) = (5, 5, 5)$ and $n = 1000$

Figure S.1: Finite-sample performance of Corollary 5 under H_0 with $m = 5$ splits. Q-Q plots compare empirical p-values against the standard uniform distribution for Model 1. Results are based on 200 replications.

B.2 A Univariate X and Y Example

In this section, we present an additional univariate X and Y example comparing with several tests. In Model S.1, we consider a simple example where X and Y are univariate and $\mathbf{Z}$ is two-dimensional. Specifically, in Settings S.1.3–S.1.4, we consider heavy-tailed distributions.

Model S.1 (Univariate X and Y, and low-dimensional $\mathbf{Z}$)**.** *Let* $(d_{\mathbf{x}}, d_{\mathbf{y}}, d_{\mathbf{z}}) = (1, 1, 2)$ *and the sample size* $n = 500$*. Each element of* $\mathbf{B}_1 \in \mathbb{R}^{d_{\mathbf{z}} \times d_{\mathbf{x}}}$, $\mathbf{B}_2 \in \mathbb{R}^{d_{\mathbf{z}} \times d_{\mathbf{y}}}$, $\mathbf{B}_3 \in \mathbb{R}^{d_{\mathbf{x}} \times d_{\mathbf{y}}}$, *and* $\mathbf{Z} \in \mathbb{R}^{d_{\mathbf{z}}}$ *are generated from Normal*$(0, 1)$*. Variables* X *and* Y *are generated under Settings S.1.1–S.1.4, and we vary* ψ *over the set* $\{0, 0.1, 0.2, 0.3, 0.4\}$*. For the error terms, we generate* $\epsilon_{\mathbf{y}} \sim$ *Normal*$(0, 1)$*, and* $\epsilon_{\mathbf{x}}$ *is set differently. Let* t_3 *denote the* t*-distribution with 3 degrees of freedom.*

Setting S.1.1. $X = \mathbf{Z}\mathbf{B}_1 + \epsilon_{\mathbf{x}}, \quad Y = \mathbf{Z}\mathbf{B}_2 + \psi X \mathbf{B}_3 + \epsilon_{\mathbf{y}}, \quad \epsilon_{\mathbf{x}} \sim$ *Normal*$(0, 1)$.

Setting S.1.2. $X = \mathbf{Z}\mathbf{B}_1 + \epsilon_{\mathbf{x}}, \quad Y = \mathbf{Z}\mathbf{B}_2 + \exp(\psi X \mathbf{B}_3) + \epsilon_{\mathbf{y}}, \quad \epsilon_{\mathbf{x}} \sim$ *Normal*$(0, 1)$.

Setting S.1.3. $X = \mathbf{Z}\mathbf{B}_1 + \epsilon_{\mathbf{x}}, \quad Y = \mathbf{Z}\mathbf{B}_2 + \psi X \mathbf{B}_3 + \epsilon_{\mathbf{y}}, \quad \epsilon_{\mathbf{x}} \sim t_3$.

Setting S.1.4. $X = \cos(\mathbf{Z}\mathbf{B}_1) + \epsilon_{\mathbf{x}}, \quad Y = \mathbf{Z}\mathbf{B}_2 + \psi X \mathbf{B}_3 + \epsilon_{\mathbf{y}}, \quad \epsilon_{\mathbf{x}} \sim t_3$.

The results of Model S.1 are shown in Figure S.2, where the x-axis represents ψ. When $\psi = 0$, all six methods effectively control the type-I error under H_0. As ψ increases, our proposed FlowCIT consistently achieves the highest statistical power under H_1. The second and the third most powerful methods are KCI and CDC tests, respectively. In contrast, CCIT is the most conservative method.

B.3 Comparison of Two Formulations in Lemma 1

In this section, we refer to testing $\boldsymbol{\eta} \perp\!\!\!\perp (\boldsymbol{\xi}, \mathbf{Z})$ as DC-1, and testing $\boldsymbol{\xi} \perp\!\!\!\perp (\boldsymbol{\eta}, \mathbf{Z})$ as DC-2. We demonstrate in Table S.1 that testing $\boldsymbol{\eta} \perp\!\!\!\perp (\boldsymbol{\xi}, \mathbf{Z})$ or $\boldsymbol{\xi} \perp\!\!\!\perp (\boldsymbol{\eta}, \mathbf{Z})$ yeilds similar performance.

B.4 Robustness and Comparison of Dependence Measures

In Section 2.3, we recommend using distance correlation (DC, Székely et al. 2007) and improved projection correlation (IPC, Zhang and Zhu 2024) to evaluate the transformed variables due to their computational efficiency. For both theoretical analysis and numerical studies of our proposed FlowCIT method, we adopt DC as the independence measure. Moreover, FlowCIT exhibits robustness to the choice of independence measures. In this section, we conduct a comparative performance analysis between DC and IPC.

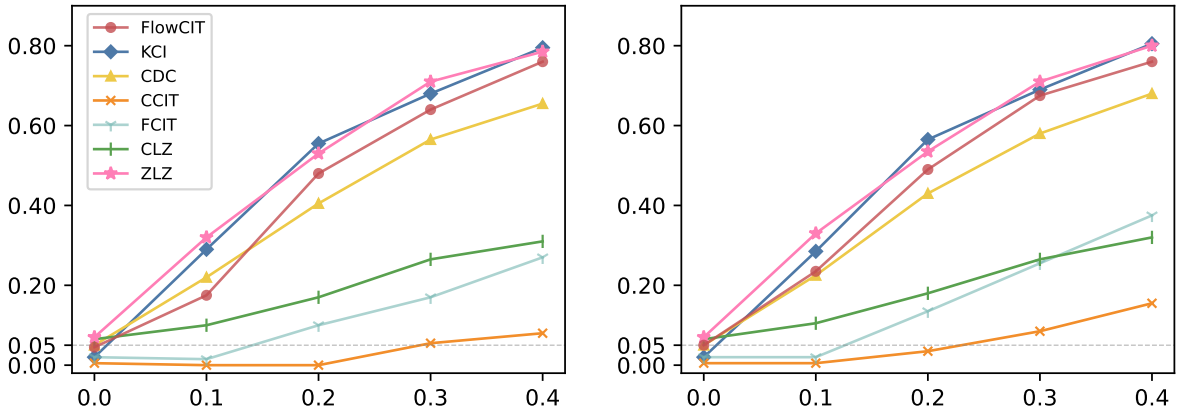

(a) Setting S.1.1

(b) Setting S.1.2

(c) Setting S.1.3

(d) Setting S.1.4

Figure S.2: Empirical type-I errors and statistical powers of Model S.1 with sample size $n = 500$. The dashed gray line below refers to significance level of 0.05. The x-axis refers to the value of ψ.

According to Appendix B.3, testing $\boldsymbol{\eta} \perp\!\!\!\perp (\boldsymbol{\xi}, \mathbf{Z})$ or $\boldsymbol{\xi} \perp\!\!\!\perp (\boldsymbol{\eta}, \mathbf{Z})$ yeilds similar performance. To save space, we only report the results for testing $\boldsymbol{\eta} \perp\!\!\!\perp (\boldsymbol{\xi}, \mathbf{Z})$ when comparing DC and IPC.

B.4.1 Comparison in simulations

We conducted simulation studies aligned with the simulation setups of Models S.1, 2, 3, 4 to evaluate the sensitivity of FlowCIT to different dependence measures. Results demonstrate that FlowCIT exhibits low sensitivity to the choice of measure, reinforcing its robustness across diverse dependency structures.

As summarized in Table S.2, the empirical type-I errors and statistical powers of FlowCIT-DC and FlowCIT-IPC are similar across all tested scenarios. This consistency

Table S.1: Empirical type-I errors and statistical powers under different projection measures used in our FlowCIT. The results suggest that our method is robust to independence measures.

		ψ	0	0.1	0.2	0.3	0.4
Model S.1	Setting S.1.1	DC-1	0.045	0.175	0.480	0.640	0.760
		DC-2	0.050	0.170	0.485	0.665	0.735
	Setting S.1.2	DC-1	0.050	0.235	0.490	0.675	0.760
		DC-2	0.055	0.235	0.500	0.685	0.740
	Setting S.1.3	DC-1	0.050	0.375	0.660	0.800	0.860
		DC-2	0.070	0.380	0.675	0.810	0.875
	Setting S.1.4	DC-1	0.035	0.380	0.650	0.820	0.875
		DC-2	0.050	0.390	0.675	0.805	0.860
		ψ	0	0.05	0.1	0.15	0.2
Model 2	Setting 2.1	DC-1	0.055	0.210	0.715	0.955	0.990
		DC-2	0.065	0.195	0.710	0.935	0.980
	Setting 2.2	DC-1	0.045	0.200	0.615	0.880	0.980
		DC-2	0.065	0.185	0.630	0.885	0.980
	Setting 2.3	DC-1	0.055	0.750	0.965	1.000	1.000
		DC-2	0.065	0.750	0.970	1.000	1.000
	Setting 2.4	DC-1	0.050	0.210	0.740	0.950	0.995
		DC-2	0.065	0.190	0.730	0.940	0.990
		ψ	0	0.1	0.2	0.3	0.4
Model 3	Setting 3.1	DC-1	0.015	0.640	0.990	0.995	1.000
		DC-2	0.025	0.640	0.990	0.995	1.000
	Setting 3.2	DC-1	0.045	0.375	0.665	0.840	0.910
		DC-2	0.065	0.380	0.685	0.825	0.895
	Setting 3.3	DC-1	0.050	0.590	0.980	1.000	1.000
		DC-2	0.035	0.575	0.980	1.000	1.000
	Setting 3.4	DC-1	0.045	0.325	0.610	0.780	0.800
		DC-2	0.060	0.255	0.510	0.630	0.670
		ψ	0	0.2	0.4	0.6	0.8
Model 4	Setting 4.1	DC-1	0.060	0.380	0.585	0.750	0.775
		DC-2	0.070	0.340	0.605	0.745	0.800
	Setting 4.2	DC-1	0.050	0.500	0.800	1.000	1.000
		DC-2	0.060	0.520	0.800	1.000	1.000
	Setting 4.3	DC-1	0.045	0.300	0.750	0.970	0.995
		DC-2	0.055	0.290	0.730	0.950	1.000
	Setting 4.4	DC-1	0.050	0.150	0.340	0.980	1.000
		DC-2	0.055	0.160	0.330	0.950	1.000

highlights that the performance of FlowCIT remains stable irrespective of whether DC or IPC is employed.

B.4.2 Comparison in real-data analysis

As shown in Table S.3, we also conduct a comparison between FlowCIT-DC and FlowCIT-IPC in the real-data analysis, and the resulting p-values are also very close.

B.5 Empirical Type-I Error Control

In the main article, we present our results using plots to effectively illustrate the behavior of each method as ψ varies. Since, under H_0, the data points for these methods are densely clustered, we provide the empirical type-I errors in tabular form in Table S.4. As shown in Table S.4, the evaluated methods achieve effective type-I error control across a majority of configurations. Mild type-I error inflation is observed in specific cases, such as CDC under the nonlinear settings of Model 2 (Settings 2.2–2.3), though these deviations remain marginal and within acceptable thresholds.

B.6 Implementations

Among the compared methods, FlowCIT and CDC require permutation-based method, both configured with $B = 100$ permutations. For CDC and FCIT, all other hyperparameters follow the default configurations specified in the Python package `hyppo` (Panda et al., 2024). The KCI test is implemented using the default settings from the `R` package `CondIndTests::KCI`, while CLZ adheres to the original `R` code. CCIT is executed with the default parameters in the `Python` package `CCIT`. ZLZ is implemented by its original `Matlab` code.

For the FlowCIT, we utilize neural networks to learn the velocity fields. Specifically, we design a ReLU-activated network with two hidden layers: the first layer contains p_1 neurons, and the second layer reduces dimensionality to $p_2 = p_1/2$ neurons. For low-dimensional conditional independence scenarios (Models S.1 and 2), we set $p_1 = 32$. In Model 3, we increase p_1 to 80, and in Model 4), we further raise p_1 to 600 accommodate the higher dimensionality.

C Proof of Lemma 1

Proof. In this section, we present the proof for our central idea: transforming the conditional independence structure into an unconditional structure using transport maps. To

Table S.2: Empirical type-I errors and statistical powers under different projection measures used in our FlowCIT. The results suggest that our method is robust to independence measures.

		ψ	0	0.1	0.2	0.3	0.4
Model S.1	Setting S.1.1	DC	0.045	0.175	0.480	0.640	0.760
		IPC	0.045	0.120	0.325	0.485	0.575
	Setting S.1.2	DC	0.050	0.235	0.490	0.675	0.760
		IPC	0.045	0.125	0.340	0.485	0.610
	Setting S.1.3	DC	0.050	0.375	0.660	0.800	0.860
		IPC	0.060	0.205	0.470	0.640	0.750
	Setting S.1.4	DC	0.035	0.380	0.650	0.820	0.875
		IPC	0.070	0.195	0.505	0.645	0.750
		ψ	0	0.05	0.1	0.15	0.2
Model 2	Setting 2.1	DC	0.055	0.210	0.715	0.955	0.990
		IPC	0.040	0.080	0.405	0.750	0.915
	Setting 2.2	DC	0.045	0.200	0.615	0.880	0.980
		IPC	0.060	0.105	0.315	0.600	0.830
	Setting 2.3	DC	0.055	0.750	0.965	1.000	1.000
		IPC	0.055	0.335	0.655	0.860	0.935
	Setting 2.4	DC	0.050	0.210	0.740	0.950	0.995
		IPC	0.040	0.085	0.450	0.775	0.925
		ψ	0	0.1	0.2	0.3	0.4
Model 3	Setting 3.1	DC	0.015	0.640	0.990	0.995	1.000
		IPC	0.050	0.460	0.965	1.000	1.000
	Setting 3.2	DC	0.045	0.375	0.665	0.840	0.910
		IPC	0.065	0.195	0.370	0.550	0.625
	Setting 3.3	DC	0.050	0.590	0.980	1.000	1.000
		IPC	0.065	0.380	0.940	0.995	1.000
	Setting 3.4	DC	0.045	0.325	0.610	0.780	0.800
		IPC	0.060	0.140	0.170	0.515	0.720
		ψ	0	0.2	0.4	0.6	0.8
Model 4	Setting 4.1	DC	0.060	0.380	0.585	0.750	0.775
		IPC	0.060	0.240	0.450	0.570	0.620
	Setting 4.2	DC	0.050	0.500	0.800	1.000	1.000
		IPC	0.065	0.420	0.685	0.900	1.000
	Setting 4.3	DC	0.045	0.300	0.750	0.970	0.995
		IPC	0.045	0.240	0.450	0.800	0.900
	Setting 4.4	DC	0.050	0.150	0.340	0.980	1.000
		IPC	0.055	0.100	0.210	0.480	0.640

Table S.3: Comparisons of p-values between FlowCIT-DC and FlowCIT-IPC in wine-quality dataset. Bold: $p \leq 0.05$.

Methods	Supervised	Nonlinear	FlowCIT-DC	FlowCIT-IPC
PCA-1	✗	✗	**0.00**	**0.00**
PCA-4	✗	✗	**0.00**	**0.00**
SIR-1	✔	✗	**0.04**	**0.02**
SIR-4	✔	✗	**0.02**	**0.02**
UMAP-1	✗	✔	**0.00**	**0.00**
UMAP-4	✗	✔	**0.00**	**0.00**
DDR-1	✔	✔	**0.01**	**0.01**
DDR-4	✔	✔	0.30	0.48
$\mathcal{R}(\mathbf{X}) = \mathbf{X}$			0.71	0.96

Table S.4: Empirical type-I errors of Models S.1, 2, 3, 4. ZLZ does not work for multivariate $\mathbf{X}$ and $\mathbf{Y}$. KCI and CLZ are not included in the comparisons for Models 3, 4 due to their high computational burden.

		FlowCIT	KCI	CDC	CCIT	FCIT	CLZ	ZLZ
Model S.1	Setting S.1.1	0.050	0.020	0.050	0.005	0.020	0.065	0.070
	Setting S.1.2	0.050	0.020	0.050	0.005	0.020	0.065	0.070
	Setting S.1.3	0.050	0.065	0.040	0.005	0.010	0.070	0.055
	Setting S.1.4	0.035	0.055	0.045	0.000	0.020	0.080	0.040
Model 2	Setting 2.1	0.055	0.055	0.035	0.000	0.045	0.060	
	Setting 2.2	0.045	0.030	0.090	0.000	0.035	0.065	
	Setting 2.3	0.055	0.080	0.105	0.000	0.060	0.085	
	Setting 2.4	0.050	0.055	0.035	0.000	0.050	0.060	
Model 3	Setting 3.1	0.015		0.000	0.005	0.010		
	Setting 3.2	0.065		0.000	0.005	0.025		
	Setting 3.3	0.050		0.000	0.000	0.010		
	Setting 3.4	0.045		0.000	0.000	0.000		
Model 4	Setting 4.1	0.060		0.000	0.020	0.005		
	Setting 4.2	0.060		0.000	0.000	0.010		
	Setting 4.3	0.045		0.000	0.000	0.000		
	Setting 4.4	0.050		0.000	0.000	0.000		

proceed, we first prove that by construction, $\boldsymbol{\xi} \perp\!\!\!\perp \mathbf{Z}$ and $\boldsymbol{\eta} \perp\!\!\!\perp \mathbf{Z}$ in the first step. We then prove the equivalance in the second step.

Step 1. Prove that $\boldsymbol{\xi} \perp\!\!\!\perp \mathbf{Z}$ *and* $\boldsymbol{\eta} \perp\!\!\!\perp \mathbf{Z}$.

We show that $\boldsymbol{\xi} \perp\!\!\!\perp \mathbf{Z}$ in the following, and $\boldsymbol{\eta} \perp\!\!\!\perp \mathbf{Z}$ can be proven similarly.

Firstly, $\boldsymbol{\xi}$ and $\mathbf{Z}$ are independent if their distribution function satisfy:

$$P(\boldsymbol{\xi}, \mathbf{Z}) = P(\boldsymbol{\xi})P(\mathbf{Z}). \tag{S.1}$$

Based on (S.1), the conditional density of $(\boldsymbol{\xi} \mid \mathbf{Z})$ is then given by

$$P(\boldsymbol{\xi} \mid \mathbf{Z}) = \frac{P(\boldsymbol{\xi}, \mathbf{Z})}{P(\mathbf{Z})} = P(\boldsymbol{\xi}).$$

Thus, to prove the independence between $\boldsymbol{\xi}$ and $\mathbf{Z}$, it suffices to show that the law of $(\boldsymbol{\xi} \mid \mathbf{Z} = \mathbf{z})$ coincides with the law of $\boldsymbol{\xi}$ for all $\mathbf{z} \in \text{supp}(\mathbf{Z})$. This is satisfied by:

$$\boldsymbol{\xi} = f(\mathbf{X}, \mathbf{Z}) \overset{d}{=} \big[f(\mathbf{X}, \mathbf{Z}) \mid \mathbf{Z} = \mathbf{z}\big].$$

Step 2. Prove the equivalance.

We are proving that the following three statements are equivalent:

1. $\boldsymbol{\xi} \perp\!\!\!\perp \mathbf{Z}$, $\boldsymbol{\eta} \perp\!\!\!\perp \mathbf{Z}$, and $\boldsymbol{\xi} \perp\!\!\!\perp \boldsymbol{\eta} \mid \mathbf{Z}$.
2. $\boldsymbol{\xi} \perp\!\!\!\perp \mathbf{Z}$, $\boldsymbol{\eta} \perp\!\!\!\perp \mathbf{Z}$, and $\boldsymbol{\xi} \perp\!\!\!\perp (\boldsymbol{\eta}, \mathbf{Z})$.
3. $\boldsymbol{\xi} \perp\!\!\!\perp \mathbf{Z}$, $\boldsymbol{\eta} \perp\!\!\!\perp \mathbf{Z}$, and $\boldsymbol{\eta} \perp\!\!\!\perp (\boldsymbol{\xi}, \mathbf{Z})$.

All these statements above can be inferred by mutual independence between $\boldsymbol{\xi}$, $\boldsymbol{\eta}$, and $\mathbf{Z}$. Thus, it suffices to show that all these statements are equal to mutual independence between these three random vectors, which means that the joint distribution $P(\boldsymbol{\xi}, \boldsymbol{\eta}, \mathbf{Z}) = P(\boldsymbol{\xi})P(\boldsymbol{\eta})P(\mathbf{Z})$.

Based on the first statement, it follows that:

$$P(\boldsymbol{\xi}, \boldsymbol{\eta}, \mathbf{Z}) = P(\boldsymbol{\xi}, \boldsymbol{\eta} \mid \mathbf{Z})P(\mathbf{Z}) = P(\boldsymbol{\xi} \mid \mathbf{Z})P(\boldsymbol{\eta} \mid \mathbf{Z})P(\mathbf{Z}) = P(\boldsymbol{\xi})P(\boldsymbol{\eta})P(\mathbf{Z}).$$

According to the second statement, it can be inferred that:

$$P(\boldsymbol{\xi}, \boldsymbol{\eta}, \mathbf{Z}) = P(\boldsymbol{\xi})P(\boldsymbol{\eta}, \mathbf{Z}) = P(\boldsymbol{\xi})P(\boldsymbol{\eta})P(\mathbf{Z}),$$

and it is similar the third statement. This completes the proof of Lemma 1.

□

D Proofs for the Test Statistic

In this section, we prove the properties of our test statistic T_{n_2} as stated in Theorem 3, with detailed proof provided in Section D.1. The proofs of Theorem 4 and Corollary 5 are

presented in Section D.2.

D.1 Proof of Theorem 3

Proof. Let $\mathbf{U} = (\boldsymbol{\xi}^\top, \mathbf{Z}^\top)^\top$ and $\mathbf{V} = \boldsymbol{\eta}$ denote the true but not observed version while $\hat{\mathbf{U}} = (\hat{\boldsymbol{\xi}}^\top, \mathbf{Z}^\top)^\top$ and $\hat{\mathbf{V}} = \hat{\boldsymbol{\eta}}$ represent the estimated counterpart.

Let $a_{ij} = \|\mathbf{U}_i - \mathbf{U}_j\|$ and $b_{ij} = \|\mathbf{V}_i - \mathbf{V}_j\|$ where $i, j \in \mathcal{I}_2$. Further denote the doubly centered distances:

$$A_{ij} = a_{ij} - \bar{a}_{i\cdot} - \bar{a}_{\cdot j} + \bar{a}_{\cdot\cdot},$$
$$B_{ij} = b_{ij} - \bar{b}_{i\cdot} - \bar{b}_{\cdot j} + \bar{b}_{\cdot\cdot},$$

where $\bar{a}_{i\cdot}$ is the i-th row mean, $\bar{a}_{\cdot j}$ is the j-th column mean, and $\bar{a}_{\cdot\cdot}$ is the grand mean (Székely and Rizzo, 2013).

We first note that $\mathbb{E}(A_{ij} - \hat{A}_{ij})^2 = O_p(n_1^{-2\kappa_1})$, which is followed by that $a_{ij} - \hat{a}_{ij} = \|\mathbf{U}_i - \mathbf{U}_j\| - \|\hat{\mathbf{U}}_i - \hat{\mathbf{U}}_j\| \leq \|\mathbf{U}_i - \hat{\mathbf{U}}_i\| + \|\mathbf{U}_j - \hat{\mathbf{U}}_j\|$, and the condition that $\mathbb{E}_{(\mathbf{X},\mathbf{Z})}\|\hat{\boldsymbol{\xi}} - \boldsymbol{\xi}\|^2 = O_p(n_1^{-2\kappa_1})$.

The empirical distance correlation with true $\mathbf{U}$ and $\mathbf{V}$ can be written as

$$\mathrm{dcov}^2_{n_2}(\mathbf{U}, \mathbf{V}) = \frac{1}{n^2} \sum_{i,j \in \mathcal{I}_2} A_{ij} B_{ij}.$$

For estimated $\hat{\mathbf{U}}$ and $\hat{\mathbf{V}}$, it is denoted by

$$\mathrm{dcov}^2_{n_2}(\hat{\mathbf{U}}, \hat{\mathbf{V}}) = \frac{1}{n^2} \sum_{i,j \in \mathcal{I}_2} \hat{A}_{ij} \hat{B}_{ij}.$$

Consider the difference between these two versions:

$$\begin{aligned}
&\mathrm{dcov}^2_{n_2}(\mathbf{U}, \mathbf{V}) - \mathrm{dcov}^2_{n_2}(\hat{\mathbf{U}}, \hat{\mathbf{V}}) \\
&= \frac{1}{n_2^2} \sum_{i,j \in \mathcal{I}_2} (A_{ij} - \hat{A}_{ij}) B_{ij} + \frac{1}{n_2^2} \sum_{i,j \in \mathcal{I}_2} A_{ij} (B_{ij} - \hat{B}_{ij}) + \frac{1}{n_2^2} \sum_{i,j \in \mathcal{I}_2} (\hat{A}_{ij} - A_{ij})(B_{ij} - \hat{B}_{ij}) \\
&= T_a + T_b + T_c.
\end{aligned}$$

We first bound the term T_a. Since $\hat{\mathbf{U}}_i, i \in \mathcal{I}_2$ is fitted based on training data, it is independent of $\mathbf{V}_i, i \in \mathcal{I}_2$. The expectation of T_a is

$$\mathbb{E}(T_a) = \frac{1}{n_2^2} \sum_{i,j \in \mathcal{I}_2} \mathbb{E}\{(A_{ij} - \hat{A}_{ij}) B_{ij}\} = \frac{1}{n_2^2} \sum_{i,j \in \mathcal{I}_2} \mathbb{E}\{(A_{ij} - \hat{A}_{ij})\} \mathbb{E}(B_{ij}) = 0.$$

Next, the variance of T_a is given by

$$\mathrm{var}(T_a) = \frac{1}{n_2^4}\,\mathrm{var}\Big\{\sum_{i,j\in\mathcal{I}_2}(A_{ij}-\hat{A}_{ij})B_{ij}\Big\} = \frac{1}{n_2^4}\sum_{i,j,k,l\in\mathcal{I}_2}\mathrm{cov}\{(A_{ij}-\hat{A}_{ij})B_{ij},(A_{kl}-\hat{A}_{kl})B_{kl}\}.$$

For each covariance term above, if $\{i,j\}\cap\{k,l\}=\emptyset$, which means they do not share index, we have $(B_{ij},B_{kl}) \perp\!\!\!\perp \{(A_{ij}-\hat{A}_{ij}),(A_{kl}-\hat{A}_{kl})\}$ and $\mathrm{cov}(B_{ij},B_{kl}) = O(n_2^{-2})$ by the property of double-centering, and it follows that:

$$\begin{aligned}
&\mathrm{cov}\{(A_{ij}-\hat{A}_{ij})B_{ij},(A_{kl}-\hat{A}_{kl})B_{kl}\}\\
&=\mathbb{E}\{(A_{ij}-\hat{A}_{ij})B_{ij}\}\mathbb{E}\{(A_{kl}-\hat{A}_{kl})B_{kl}\}-\mathbb{E}\{(A_{ij}-\hat{A}_{ij})B_{ij}(A_{kl}-\hat{A}_{kl})B_{kl}\}\\
&=\mathbb{E}\{(A_{ij}-\hat{A}_{ij})\}\mathbb{E}B_{ij}\mathbb{E}\{(A_{kl}-\hat{A}_{kl})\}\mathbb{E}B_{kl}-\mathbb{E}\{(A_{ij}-\hat{A}_{ij})(A_{kl}-\hat{A}_{kl})\}\mathbb{E}(B_{ij}B_{kl})\\
&=O(n_2^{-2}n_1^{-2\kappa_1}).
\end{aligned}$$

Similarly, if (i,j) and (k,l) share one index, by the property of double-centering it follows that $\mathrm{cov}(B_{ij},B_{kl}) = O(n_2^{-1})$, and we have $\mathrm{cov}\{(A_{ij}-\hat{A}_{ij})B_{ij},(A_{kl}-\hat{A}_{kl})B_{kl}\} = O(n_2^{-1}n_1^{-2\kappa_1})$. Otherwise, if (i,j) and (k,l) share two indices, by the Cauchy-Schwarz inequality, we have

$$\begin{aligned}
&\mathrm{cov}\{(A_{ij}-\hat{A}_{ij})B_{ij},(A_{kl}-\hat{A}_{kl})B_{kl}\}\\
&\le\sqrt{\mathbb{E}(B_{ij}^2)\mathbb{E}\{(A_{ij}-\hat{A}_{ij})^2\}\mathbb{E}(B_{ij}^2)\mathbb{E}\{(A_{ij}-\hat{A}_{ij})^2\}} = O_p(n_1^{-2\kappa_1}).
\end{aligned}$$

The numbers of terms where (i,j) and (k,l) share no index, one index, and two indices are $O(n_2^4)$, $O(n_2^3)$, and $O(n_2^2)$. Therefore,

$$\mathrm{var}(T_a) = \frac{1}{n_2^4}O_p(n_2^2\times n_1^{-2\kappa_1}) = O_p(n_2^{-2}n_1^{-2\kappa_1}).$$

By Chebyshev's inequality, we have $T_a = O_p(n_2^{-1}n_1^{-\kappa_1}\log n_2)$. Similarly, we have $T_b = O_p(n_2^{-1}n_1^{-\kappa_2}\log n_2)$.

For the term T_c, by Cauchy-Schwarz inequality, it follows that

$$T_c \le \sqrt{\frac{1}{n_2^2}\sum_{ij\in\mathcal{I}_2}(\hat{A}_{ij}-A_{ij})^2\frac{1}{n_2^2}\sum_{ij\in\mathcal{I}_2}(\hat{B}_{ij}-B_{ij})^2} = O_p(n_1^{-\kappa_1-\kappa_2}).$$

Since $n_2T_a = O_p(n_1^{-\kappa_1}\log n_2) = o_p(1)$, $n_2T_b = o_p(1)$, and $n_2T_c = O_p(n_2n_1^{-\kappa_1-\kappa_2}) = o_p(1)$, the difference between $n_2\,\mathrm{dcov}_{n_2}^2(\mathbf{U},\mathbf{V})$ and $n_2\,\mathrm{dcov}_{n_2}^2(\hat{\mathbf{U}},\hat{\mathbf{V}})$ is ignorable.

Similarly, the difference between $\mathrm{dcov}_{n_2}^2(\mathbf{U},\mathbf{U})$ and $\mathrm{dcov}_{n_2}^2(\hat{\mathbf{U}},\hat{\mathbf{U}})$, as well as between

$\mathrm{dcov}^2_{n_2}(\mathbf{V},\mathbf{V})$ and $\mathrm{dcov}^2_{n_2}(\hat{\mathbf{V}},\hat{\mathbf{V}})$, are also $o_p(1)$. Thus, by Slutsky's theorem,

$$n_2 T_{n_2} - n_2 \,\mathrm{dcorr}^2_{n_2}(\mathbf{U},\mathbf{V}) = o_p(1). \tag{S.2}$$

Under H_0, $n_2 \,\mathrm{dcorr}^2_{n_2}(\mathbf{U},\mathbf{V}) \xrightarrow{d} \sum_{j=1}^{\infty} \lambda_j Z_j^2$ (Székely et al., 2007). Together with (S.2), this completes the proof for the asymptotic distribution of T_{n_2} under H_0.

Under H_1, it follows that $n_2 \,\mathrm{dcorr}_{n_2}(\mathbf{U},\mathbf{V}) \xrightarrow{p} \infty$. We can similarly show that $\mathrm{dcorr}_{n_2}(\mathbf{U},\mathbf{V}) - T_{n_2} \xrightarrow{p} 0$, and thus $n_2 T_{n_2} \xrightarrow{p} \infty$. □

D.2 Proof of Theorem 4 and Corollary 5

Proof. Let $F(t)$ denote the limiting (continuous) distribution function of the test statistic T_{n_2} under the null hypothesis. For a given dataset $\mathcal{D}_{n_2}$, and let $T_{n_2,b}$ denote the test statistic computed on the b-th permuted sample, for $b = 1, \ldots, B$. Define the empirical distribution function based on the B permuted statistics as

$$\hat{F}_{n_2,B}(t) = \frac{1}{B} \sum_{b=1}^{B} \mathbf{1}\left(n_2 T_{n_2,b} \le t\right),$$

and define the permutation p-value as

$$p_B = 1 - \hat{F}_{n_2,B}(n_2 T_{n_2}),$$

where T_{n_2} is the observed test statistic.

Since $F(t)$ is continuous and, conditional on $\mathcal{D}_{n_2}$, the $T_{n_2,b}$ are independent, it follows from the properties of permutation tests that, as $n_2 \to \infty$, the distribution of $n_2 T_{n_2,b}$ converges to $F(t)$. Therefore, by the Glivenko–Cantelli theorem, as $B, n_2 \to \infty$, $\sup_t \left|\hat{F}_{n_2,B}(t) - F(t)\right| \xrightarrow{p} 0$.

Consequently, we can write

$$p_B = 1 - \hat{F}_{n_2,B}(n_2 T_{n_2}) = 1 - F(n_2 T_{n_2}) + F(n_2 T_{n_2}) - \hat{F}_{n_2,B}(n_2 T_{n_2}).$$

Under the null hypothesis H_0, the distribution of $n_2 T_{n_2}$ converges to $F(t)$, so that $1 - F(n_2 T_{n_2}) \xrightarrow{d} \mathrm{Uniform}(0,1)$. Combined with the uniform convergence of $\hat{F}_{n_2,B}(t)$ to $F(t)$, it follows that $p_B \xrightarrow{d} \mathrm{Uniform}(0,1)$.

Under the alternative hypothesis H_1, since $n_2 T_{n_2} \xrightarrow{p} \infty$, we have $1 - F(n_2 T_{n_2}) \xrightarrow{p} 0$, and thus $p_B \xrightarrow{p} 0$.

For Corollary 5, since the data used in each fold are disjoint, the p-values obtained from multiple splits are asymptotically independent, and each converges in distribution

to Uniform$(0,1)$ under H_0. Therefore, the limiting distribution of the combined p-value remains unchanged.

□

E Proofs and Discussions for the Conditional CNFs

The convergence result of flow-matching stated in Assumption 1 has been established in literature. For instance, interested readers may refer to Theorem 5.12 of Gao et al. (2024b) and Theorem 5.5 of Huang et al. (2024a). Here, we use general rates $n_1^{-\kappa_1}$ and $n_1^{-\kappa_2}$ to denote. They are typically $0 < \kappa_1, \kappa_2 < 1/2$. Based on the following convergence

$$\mathbb{E}_{\mathcal{D}_{n_1}} \mathbb{E}_{(t,\mathbf{X},\mathbf{Z})} \left\| \hat{\mathbf{v}}_{\mathbf{x}}(t,\mathbf{X},\mathbf{Z}) - \mathbf{v}_{\mathbf{x}}(t,\mathbf{X},\mathbf{Z}) \right\|^2 \lesssim n_1^{-\kappa_1}. \tag{S.3}$$

Following Markov inequality, we can directly infer that,

$$\mathbb{E}_{(t,\mathbf{X},\mathbf{Z})} \left\| \hat{\mathbf{v}}_{\mathbf{x}}(t,\mathbf{X},\mathbf{Z}) - \mathbf{v}_{\mathbf{x}}(t,\mathbf{X},\mathbf{Z}) \right\|^2 \lesssim O_p(n_1^{-\kappa_1}).$$

Based on the convergence in (S.3), the convergence of $\hat{\boldsymbol{\xi}}$ and $\hat{\boldsymbol{\eta}}$ is followed by

$$\begin{aligned}
&\mathbb{E}_{(\mathbf{X},\mathbf{Z})} \|\hat{\boldsymbol{\xi}} - \boldsymbol{\xi}\|^2 = \mathbb{E}_{(\mathbf{X},\mathbf{Z})} \| \int_0^1 (\hat{\mathbf{v}}_{\mathbf{x}} - \mathbf{v}_{\mathbf{x}}) \mathrm{d}t \|^2 \\
&\leq \mathbb{E}_{(\mathbf{X},\mathbf{Z})} \int_0^1 \|(\hat{\mathbf{v}}_{\mathbf{x}} - \mathbf{v}_{\mathbf{x}})\|^2 \mathrm{d}t = \mathbb{E}_{(t,\mathbf{X},\mathbf{Z})} \|\hat{\mathbf{v}}_{\mathbf{x}} - \mathbf{v}_{\mathbf{x}}\|^2 = O_p(n_1^{-\kappa_1}),
\end{aligned}$$

which follows from Jensen's inequality.

REFERENCE

Albergo, M. S., Boffi, N. M., and Vanden-Eijnden, E. (2023). Stochastic interpolants: A unifying framework for flows and diffusions. *arXiv preprint arXiv:2303.08797*.

Albergo, M. S. and Vanden-Eijnden, E. (2022). Building normalizing flows with stochastic interpolants. In *The Eleventh International Conference on Learning Representations*.

Bellot, A. and van der Schaar, M. (2019). Conditional independence testing using generative adversarial networks. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.

Cai, C. and Li, G. (2025). Minimax optimality of the probability flow ODE for diffusion models. *arXiv preprint arXiv:2503.09583*.

Cai, Z., Lei, J., and Roeder, K. (2024). Asymptotic distribution-free independence test for high-dimension data. *Journal of the American Statistical Association*, 119(547):1794–1804.

Cai, Z., Li, R., and Zhang, Y. (2022). A distribution free conditional independence test with applications to causal discovery. *Journal of Machine Learning Research*, 23(85):1–41.

Chalupka, K., Perona, P., and Eberhardt, F. (2018). Fast conditional independence test for vector variables with large sample sizes. *arXiv preprint arXiv:1804.02747*.

Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68.

Constantinou, P. and Dawid, A. P. (2017). Extended conditional independence and applications in causal inference. *The Annals of Statistics*, 45(6):2618–2653.

Cortez, P., Cerdeira, A., Almeida, F., Matos, T., and Reis, J. (2009). Modeling wine preferences by data mining from physicochemical properties. *Decision Support Systems*, 47(4):547–553.

Dawid, A. P. (1979). Conditional independence in statistical theory. *Journal of the Royal Statistical Society: Series B (Methodological)*, 41(1):1–15.

Duong, B. and Nguyen, T. (2024). Normalizing flows for conditional independence testing. *Knowledge and Information Systems*, 66(1):357–380.

Fan, J., Feng, Y., and Xia, L. (2020). A projection-based conditional dependence measure with applications to high-dimensional undirected graphical models. *Journal of Econometrics*, 218(1):119–139.

Fukumizu, K., Suzuki, T., Isobe, N., Oko, K., and Koyama, M. (2025). Flow matching achieves almost minimax optimal convergence. In *The Thirteenth International Conference on Learning Representations*.

Gao, Y., Huang, J., and Jiao, Y. (2024a). Gaussian interpolation flows. *Journal of Machine Learning Research*, 25(253):1–52.

Gao, Y., Huang, J., Jiao, Y., and Zheng, S. (2024b). Convergence of continuous normalizing flows for learning probability distributions. *arXiv preprint arXiv:2404.00551*.

Geiger, D. and Pearl, J. (1993). Logical and algorithmic properties of conditional independence and graphical models. *The Annals of Statistics*, 21(4):2001–2021.

Ho, J., Jain, A., and Abbeel, P. (2020). Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851. Curran Associates, Inc.

Huang, D., Huang, J., Li, T., and Shen, G. (2024a). Conditional stochastic interpolation for generative learning. *arXiv preprint arXiv:2312.05579*.

Huang, J., Jiao, Y., Liao, X., Liu, J., and Yu, Z. (2024b). Deep dimension reduction for supervised representation learning. *IEEE Transactions on Information Theory*, 70(5):3583–3598.

Huang, Z., Deb, N., and Sen, B. (2022). Kernel partial correlation coefficient — a measure of conditional dependence. *Journal of Machine Learning Research*, 23(216):1–58.

Koldanov, P., Koldanov, A., Kalyagin, V., and Pardalos, P. (2017). Uniformly most powerful unbiased test for conditional independence in Gaussian graphical model. *Statistics & Probability Letters*, 122:90–95.

Li, K.-C. (1991). Sliced inverse regression for dimension reduction. *Journal of the American Statistical Association*, 86(414):316–327.

Lipman, Y., Chen, R. T. Q., Ben-Hamu, H., Nickel, M., and Le, M. (2022). Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*.

Liu, M., Zhang, Y., and Zhou, D. (2021). Double/debiased machine learning for logistic partially linear model. *The Econometrics Journal*, 24(3):559–588.

Liu, Q. (2022). Rectified flow: A marginal preserving approach to optimal transport. *arXiv preprint arXiv:2209.14577*.

Liu, X., Gong, C., and Liu, Q. (2023). Flow straight and fast: Learning to generate and transfer data with rectified flow. In *The Eleventh International Conference on Learning Representations*.

Liu, Y. and Xie, J. (2020). Cauchy combination test: A powerful test with analytic p-value calculation under arbitrary dependency structures. *Journal of the American Statistical Association*, 115(529):393–402.

McInnes, L., Healy, J., Saul, N., and Großberger, L. (2018). UMAP: Uniform manifold approximation and projection. *Journal of Open Source Software*, 3(29):861.

Panda, S., Palaniappan, S., Xiong, J., Bridgeford, E. W., Mehta, R., Shen, C., and Vogelstein, J. T. (2024). Hyppo: A multivariate hypothesis testing python package. *arXiv preprint arXiv:1907.02088*.

Pearl, J. (2009). *Causality: Models, Reasoning and Inference*. Cambridge University Press, Cambridge, U.K. ; New York, 2nd edition edition.

Pearson, K. (1901). On lines and planes of closest fit to systems of points in space. *Philosophical Magazine*, 2(11):559–572.

Rosenblatt, M. (1952). Remarks on a multivariate transformation. *The Annals of Mathematical Statistics*, 23(3):470–472.

Sen, R., Suresh, A. T., Shanmugam, K., Dimakis, A. G., and Shakkottai, S. (2017). Model-powered conditional independence test. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.

Shi, C., Xu, T., Bergsma, W., and Li, L. (2021). Double generative adversarial networks for conditional independence testing. *Journal of Machine Learning Research*, 22(285):1–32.

Song, J., Meng, C., and Ermon, S. (2020). Denoising diffusion implicit models. In *International Conference on Learning Representations*.

Song, K. (2009). Testing conditional independence via Rosenblatt transforms. *The Annals of Statistics*, 37(6B):4011–4045.

Spirtes, P. (2010). Introduction to causal inference. *Journal of Machine Learning Research*, 11(54):1643–1662.

Spirtes, P., Glymour, C., and Scheines, R. (1993). *Causation, Prediction, and Search*, volume 81 of *Lecture Notes in Statistics*. Springer, New York, NY.

Székely, G. J. and Rizzo, M. L. (2013). The distance correlation t-test of independence in high dimension. *Journal of Multivariate Analysis*, 117:193–213.

Székely, G. J., Rizzo, M. L., and Bakirov, N. K. (2007). Measuring and testing dependence by correlation of distances. *The Annals of Statistics*, 35(6):2769–2794.

Wang, X., Pan, W., Hu, W., Tian, Y., and Zhang, H. (2015). Conditional distance correlation. *Journal of the American Statistical Association*, 110(512):1726–1734.

Xu, K. and Zhu, L. (2022). Nonparametric two-sample tests for equality of distributions using projections. *Scientia Sinica(Mathematica)*, 52(10):1183–1202.

Yang, Y., Li, S., Zhang, Y., Sun, Z., Shu, H., Chen, Z., and Zhang, R. (2024). Conditional diffusion models based conditional independence testing. *arXiv preprint arXiv:2412.11744*.

Zhang, J., Ding, J., and Yang, Y. (2023). Is a classification procedure good enough?—a goodness-of-fit assessment tool for classification learning. *Journal of the American Statistical Association*, 118(542):1115–1125.

Zhang, K., Peters, J., Janzing, D., and Schölkopf, B. (2011). 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, Arlington, Virginia, USA. AUAI Press.

Zhang, Y. and Zhu, L. (2024). Projective independence tests in high dimensions: The curses and the cures. *Biometrika*, 111(3):1013–1027.

Zhou, Y., Liu, J., and Zhu, L. (2020). Test for conditional independence with application to conditional screening. *Journal of Multivariate Analysis*, 175(C).

Zhu, L., Xu, K., Li, R., and Zhong, W. (2017). Projection correlation between two random vectors. *Biometrika*, 104(4):829–843.