

---

# InfoAtlas: A Foundation Model for Zero-Shot Statistical Dependence Estimate

---

Zhengyang Hu<sup>1\*</sup> Yanzhi Chen<sup>2,3\*</sup> Hanxiang Ren<sup>4</sup> Qunsong Zeng<sup>1</sup> Youyi Zheng<sup>4</sup> Adrian Weller<sup>2,5</sup>  
Kaibin Huang<sup>1</sup> Yanchao Yang<sup>1+</sup>

## Abstract

Measuring statistical dependency between high-dimensional random variables is a fundamental task in data science and machine learning. Neural mutual information (MI) estimators offer a promising avenue, but they typically require costly iterative optimization for each new dataset, making them impractical for real-time applications. We present *InfoAtlas*, a foundation model-like architecture that eliminates this bottleneck by directly inferring MI in a single forward pass. Pretrained on large-scale synthetic data with rich dependence patterns, *InfoAtlas* learns to identify diverse dependence structures and predict MI directly from the dataset. Comprehensive experiments demonstrate that *InfoAtlas* matches state-of-the-art neural estimators in accuracy while achieving 100× speedup, can flexibly handle varying dimensions and sample sizes through a single unified model, and generalizes effectively to complex, real-world scenarios. By reformulating MI estimation as an inference task, *InfoAtlas* establishes a foundation for real-time dependency analysis. Project page: [InfoAtlas-project](#)

## 1. Introduction

Understanding statistical dependencies between variables is fundamental to data science and machine learning. Quantifying how variables influence each other uncovers hidden structures and causal mechanisms that drive complex systems. Applications span a wide range of domains: in healthcare, identifying dependencies between lifestyle factors and disease risks enables personalized prevention strategies (Du et al., 2024); in autonomous driving, modeling dependen-

cies between sensor signals and road conditions improves safety and decision making (Maanpää et al., 2025); in biology, assessing the dependence between protein sequences reveals insights for understanding their functional relationship (Gowri et al., 2024); and in robotics, maximizing statistical dependence between observational states is shown useful for policy discovery (Zhou & Yang, 2024).

Mutual information (MI) (Shannon, 1948) has long served as a principled measure for dependency, uniquely capturing complex nonlinear relationships for multivariate variables in interpretable units of bits. Its generality has made it a core tool in data analysis, generative modeling and representation learning (Chen et al., 2016; Oord et al., 2018; Chen et al., 2022). However, computing MI from empirical samples is notoriously difficult: closed-form solutions exist only for certain distributions, and neural estimators (Belghazi et al., 2018; Choi & Lee, 2020; Franzese et al., 2023; Tschannen et al., 2019; Chen et al., 2020; Tsai et al., 2020) require costly gradient-based optimization for every dataset, making them impractical for real-time or large-scale applications.

In this work, we introduce *InfoAtlas*, a foundation model-style architecture for fast and accurate estimation of statistical dependence between *multivariate* random variables. *InfoAtlas* predicts the strength of dependence in a single forward pass—an ability reminiscent of foundation models (Hollmann et al., 2025; Comanici et al., 2025). This ability is acquired through large-scale pretraining on massive synthetic datasets that capture a wide range of dependence structures and data patterns, enabling *InfoAtlas* to directly infer statistical relationships without per-dataset optimization. Crucially, *InfoAtlas* preserves full differentiability, facilitating seamless integration into larger AI pipelines. Extensive experiments demonstrate that *InfoAtlas* generalizes effectively from synthetic benchmarks to complex real-world data, accurately capturing a broad spectrum of dependencies, being a versatile tool for rapid understanding of variable relationships. Our main contributions are:

- We introduce *InfoAtlas*, the first pretrained architecture for zero-shot estimation of mutual information between *multivariate* variables. *InfoAtlas* achieves accuracy on par with state-of-the-art neural methods without any gradient-based optimization, and flexibly handles variables of varying dimensionalities and sample sizes with a *single* model.

---

\* Co-first Author <sup>1</sup>The University of Hong Kong <sup>2</sup>University of Cambridge <sup>3</sup>Microsoft <sup>4</sup>Zhejiang University <sup>5</sup>Alan Turing Institute.  
+ Corresponding Author

Email to: Zhengyang Hu <u3010250@connect.hku.hk>, Yanzhi Chen <yc514@cam.ac.uk>, Yanchao Yang <yanchao@hku.hk>.

*Proceedings of the 43<sup>rd</sup> International Conference on Machine Learning*, Seoul, South Korea. PMLR 306, 2026. Copyright 2026 by the author(s).

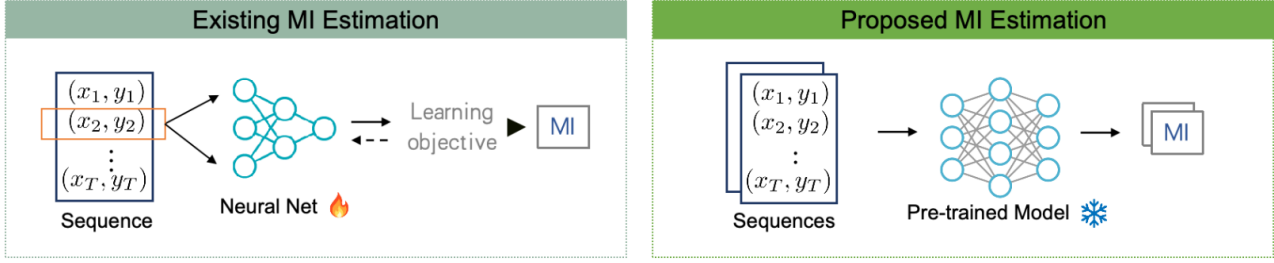


Figure 1. **Conceptual comparison: prior methods vs our method.** Existing neural MI estimators (left) requires iterative gradient-based optimization to train a neural network for each new dataset. In contrast, we use a *pre-trained* architecture to directly generate MI estimates in a single forward pass (right), eliminating per-dataset training and achieving speedup while maintaining comparable accuracy.

- We propose an attentive dual-path hypernetwork-based architecture, which is pretrained on large-scale synthetic datasets covering diverse dependency structures. This design enables *InfoAtlas* to predict dependency strength in a single inference step, and generalizes effectively to unseen real-world scenarios without task-specific finetuning.
- We comprehensively evaluate *InfoAtlas* on both synthetic benchmarks and real-world tasks, including independence testing, CLIP embedding analysis (Radford et al., 2021), motion trajectory modeling and robotics manipulation. Results demonstrate its robust performance and accurate perception of a wide spectrum of dependencies.

## 2. Problem Statement

In this work, we consider the problem of quantifying statistical dependence between two *multivariate* random variables  $\mathbf{x} \in \mathbb{R}^{d_x}$  and  $\mathbf{y} \in \mathbb{R}^{d_y}$ , with  $d_x \geq 1$  and  $d_y \geq 1$ .

**Measuring dependence via mutual information.** Mutual information (MI) offers a principle measure for quantifying statistical dependence between multivariate variables. Unlike linear correlation coefficients that capture only linear relationships, MI effectively captures both linear and non-linear correlations. Formally, MI is defined as the Kullback-Leibler (KL) divergence between the joint distribution  $p_{\mathbf{x}, \mathbf{y}}$  and the product of marginals  $p_{\mathbf{x}} \otimes p_{\mathbf{y}}$  (Kullback, 1997):

$$\begin{aligned} \mathbb{I}(\mathbf{x}, \mathbf{y}) &= \text{KL}(p_{\mathbf{x}, \mathbf{y}} \| p_{\mathbf{x}} \otimes p_{\mathbf{y}}) \\ &= \int_{\mathbf{y}} \int_{\mathcal{X}} p_{\mathbf{x}, \mathbf{y}}(\mathbf{x}, \mathbf{y}) \log \left( \frac{p_{\mathbf{x}, \mathbf{y}}(\mathbf{x}, \mathbf{y})}{p_{\mathbf{x}}(\mathbf{x})p_{\mathbf{y}}(\mathbf{y})} \right) d\mathbf{x}d\mathbf{y}. \end{aligned} \quad (1)$$

Strong correlation manifests as significant divergence between  $p(\mathbf{x}, \mathbf{y})$  and  $p(\mathbf{x})p(\mathbf{y})$ , yielding large MI, while uncorrelated variables satisfy  $p(\mathbf{x}, \mathbf{y}) \approx p(\mathbf{x})p(\mathbf{y})$ , resulting in MI near zero.

While MI offers a principled dependence measure, it rarely admits closed-form solutions except for certain known distributions (Czyż et al., 2023a;b). Thus, practical applications require estimation from finite samples  $\mathcal{D} = \{\mathbf{x}^i, \mathbf{y}^i\}_{i=1}^n$

drawn from  $p_{\mathbf{x}, \mathbf{y}}$ . Recent advances have produced powerful neural estimators (Belghazi et al., 2018; Duong & Nguyen, 2023; Franzese et al., 2023; Tsai et al., 2020; Poole et al., 2019; Song & Ermon, 2019; Letizia et al., 2024; Tsur et al., 2023a), with the most prominent one leveraging the Donsker-Varadhan (DV) representation (Donsker & Varadhan, 1983):

$$\mathbb{I}(\mathbf{x}, \mathbf{y}) := \sup_{\theta} \mathbb{E}_{p_{\mathbf{x}, \mathbf{y}}}[\theta] - \log(\mathbb{E}_{p_{\mathbf{x}} \otimes p_{\mathbf{y}}}[e^{\theta}]), \quad (2)$$

where  $\theta : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$  is a critic function. Mutual Information Neural Estimation (MINE) (Belghazi et al., 2018) parametrizes  $\theta$  as a neural network and approximates the supremum through gradient-based optimization. Besides MINE, there also exist a wide range of neural estimators based on different bounds and learning objectives; see §6.

**Challenge of real-time MI estimation.** Despite differences in theory and algorithm, all existing neural estimators share a critical computational bottleneck in practice: they require training a network  $\theta$  from scratch for each incoming dataset  $\mathcal{D} = \{\mathbf{x}^i, \mathbf{y}^i\}_{i=1}^n$  via gradient descent:

$$\theta^{t+1} \leftarrow \theta^t - \eta \nabla_{\theta^t} \mathcal{L}(\theta^t), \quad t = 1, \dots, T \quad (3)$$

where  $\mathcal{L}(\theta)$  is an estimator-specific objective (e.g., the negative DV bound for MINE). Achieving accurate MI estimates typically requires thousands of gradient steps, resulting in a  $\mathcal{O}(T)$  computational complexity. This prohibitive cost limits real-time applications such as high-frequency financial correlation monitoring or large-scale genomic screening. The recent InfoNet (Hu et al., 2024) aimed to address this inefficiency by pretraining a network to directly output optimal critic value via lookup tables, eliminating inference-time optimization. However, InfoNet is fundamentally limited to univariate inputs, and extending it to  $d$ -dimensional variables would require storing  $\mathcal{O}(e^d)$  values in its lookup table, which quickly becomes intractable even for  $d = 8$ , and it can not process data with varying data dimensionality. These limitations motivate our fundamentally different approach for real-time measurement of statistical dependence, where a unified model is developed to directly process multivariate data with varying dimensionalities and sample sizes.

### 3. Method

We present *InfoAtlas*, a pretrained architecture to address the above challenge of real-time correlation estimation between multivariate random variables  $\mathbf{x} \in \mathbb{R}^{d_x}$  and  $\mathbf{y} \in \mathbb{R}^{d_y}$ . Unlike existing neural estimators that require iterative optimization, *InfoAtlas* directly outputs mutual information (MI) in a single forward pass. This capability is enabled by two key innovations: (i) a dual-path attentive hypernetwork-based architecture, which directly generates distribution-specific critic parameters from observed samples with varying sizes and dimensionality; (ii) a comprehensive pre-training strategy using diverse synthetic distributions, which ensures generalization across different application domains.

#### 3.1. Direct Optimal Critic Generation

Our key innovation is to reformulate MI estimation from a test-time optimization problem into a direct inference task with the aid of a hypernetwork. Specifically, given a dataset  $\mathcal{D} = \{(\mathbf{x}^i, \mathbf{y}^i)\}_{i=1}^n$  drawn from an unknown joint distribution, *InfoAtlas* employs an attention-based hypernetwork  $\mathcal{H} : \mathcal{D} \mapsto \Theta$  that directly outputs the complete parameter set  $\theta^*$  of the optimal critic network in the Donsker-Varadhan representation (Eq. 2) via a single feedforward pass<sup>1</sup>:

$$\theta^* = \mathcal{H}(\mathcal{D}) = \mathcal{H}(\{(\mathbf{x}^i, \mathbf{y}^i)\}_{i=1}^n) \quad (4)$$

An empirical MI estimation is then obtained via

$$\hat{\mathbb{I}}_{\theta}(\mathbf{x}, \mathbf{y}) = \frac{1}{n} \sum_{i=1}^n \theta(\mathbf{x}^i, \mathbf{y}^i) - \log \left( \frac{1}{n} \sum_{j=1}^n e^{\theta(\mathbf{x}^j, \mathbf{y}^{\pi(j)})} \right), \quad (5)$$

where  $\{(\mathbf{x}^j, \mathbf{y}^{\pi(j)})\}_{j=1}^n$  denotes the marginal pairs with  $\pi$  a random permutation of indices  $\{1, \dots, n\}$ . This eliminates the iterative gradient updates required by neural MI estimators while avoiding the exponential value storage of InfoNet’s lookup table approach. This architectural shift fundamentally changes the computational complexity from  $\mathcal{O}(T)$  gradient steps, where  $T$  is the number of optimization iterations, to  $\mathcal{O}(1)$  feedforward propagation.

The hypernetwork  $\mathcal{H}$  leverages attention (Vaswani et al., 2017) and consists of the following key modules:

**The joint distribution path** processes  $n$  paired samples  $\{(\mathbf{x}^i, \mathbf{y}^i)\}_{i=1}^n$  to extract correlation patterns inherent in  $p(\mathbf{x}, \mathbf{y})$ . Each sample pair is treated as a token in a sequence, enabling permutation-invariant processing through attention mechanisms. Specifically, a learnable query vector  $\mathbf{q}_{\text{joint}} \in \mathbb{R}^{d_{\text{model}}}$  initiates cross-attention computation, where the concatenated samples  $[\mathbf{x}^i; \mathbf{y}^i]$  serve simultaneously as

keys and values. This mechanism computes attention weights  $\alpha_i = \text{softmax}(\mathbf{q}_{\text{joint}}^T \mathbf{W}_K [\mathbf{x}^i; \mathbf{y}^i] / \sqrt{d_{\text{model}}})$ , producing an aggregated representation that emphasizes sample pairs exhibiting strong correlations. The aggregated features subsequently pass through 16 self-attention layers ultimately producing a comprehensive encoding  $\mathbf{h}_{\text{joint}} \in \mathbb{R}^{d_{\text{hidden}}}$  that characterizes the joint distribution’s correlation structure.

**The marginal distribution path** processes samples from the product of marginals  $p(\mathbf{x})p(\mathbf{y})$  by breaking the pairing relationship. Specifically, the samples  $\{\mathbf{x}^i\}_{i=1}^n$  and  $\{\mathbf{y}^j\}_{j=1}^n$  are passed through separate projection networks  $f_{\mathbf{x}} : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_{\text{proj}}}$  and  $f_{\mathbf{y}} : \mathbb{R}^{d_y} \rightarrow \mathbb{R}^{d_{\text{proj}}}$ , implemented as Multi-Layer Perceptrons (MLPs) to obtain higher-dimensional representations that facilitate correlation detection. The architecture employs bidirectional cross-attention with two sets of learnable query vectors:  $\mathbf{q}_{\mathbf{x} \rightarrow \mathbf{y}}$  attends from projected  $\mathbf{x}$  representations to projected  $\mathbf{y}$  representations (as keys and values), while  $\mathbf{q}_{\mathbf{y} \rightarrow \mathbf{x}}$  performs the reverse attention. The outputs from both directions are summed element-wise, capturing symmetric independence patterns that should appear when variables lack correlation. This combined representation is then processed by 8 self-attention layers, resulting in encoding  $\mathbf{h}_{\text{marginal}} \in \mathbb{R}^{d_{\text{hidden}}}$  that provides a baseline representation against which the correlation strength can be measured.

**The integration and generation module** fuses information from both distributional paths through a cross-attention mechanism that allows the joint distribution features to be modulated by marginal distribution patterns. In particular, we compute cross-attention where  $\mathbf{h}_{\text{marginal}}$  serves as the query and  $\mathbf{h}_{\text{joint}}$  provides both keys and values, producing a fused representation  $\mathbf{h}_{\text{fused}} = \text{CrossAttention}(\mathbf{h}_{\text{marginal}}, \mathbf{h}_{\text{joint}}, \mathbf{h}_{\text{joint}})$ . This asymmetric fusion ensures that correlation patterns identified in the joint path are evaluated against the independence baseline from the marginal path. The fused features are then processed by a parameter generation MLP serving as a nonlinear mapping from distributional features to critic network parameters. The MLP outputs a flattened vector  $\theta \in \mathbb{R}^{|\Theta|}$  containing all parameters for a critic network, where  $|\Theta| = \sum_{l=1}^L (d_l \times d_{l-1} + d_l)$  accounts for both weight matrices and bias terms across all layers.

**Noise padding module** further addresses the challenge of varying input dimensions through a unified data preprocessing strategy that maintains MI while enabling consistent model architecture. For inputs with dimensions  $d < D$ , we pad variables with independent Gaussian noise  $\mathcal{N}(0, \mathbf{I})$  to reach  $D$  dimensions. This padding preserves mutual information exactly since  $\mathbb{I}(\mathbf{x}, \mathbf{y}) = \mathbb{I}([\mathbf{x}; \mathbf{n}_x], [\mathbf{y}; \mathbf{n}_y])$  when the noise vectors  $\mathbf{n}_x$  and  $\mathbf{n}_y$  are independent of each other and independent of both  $\mathbf{x}$  and  $\mathbf{y}$ , as justified in Proposition A.3.

<sup>1</sup>While primarily focusing on the DV representation, which we find useful for achieving good performance when paired with massive pre-training, our method is fully compatible with other variational estimators (Song & Ermon, 2019; Letizia et al., 2024).

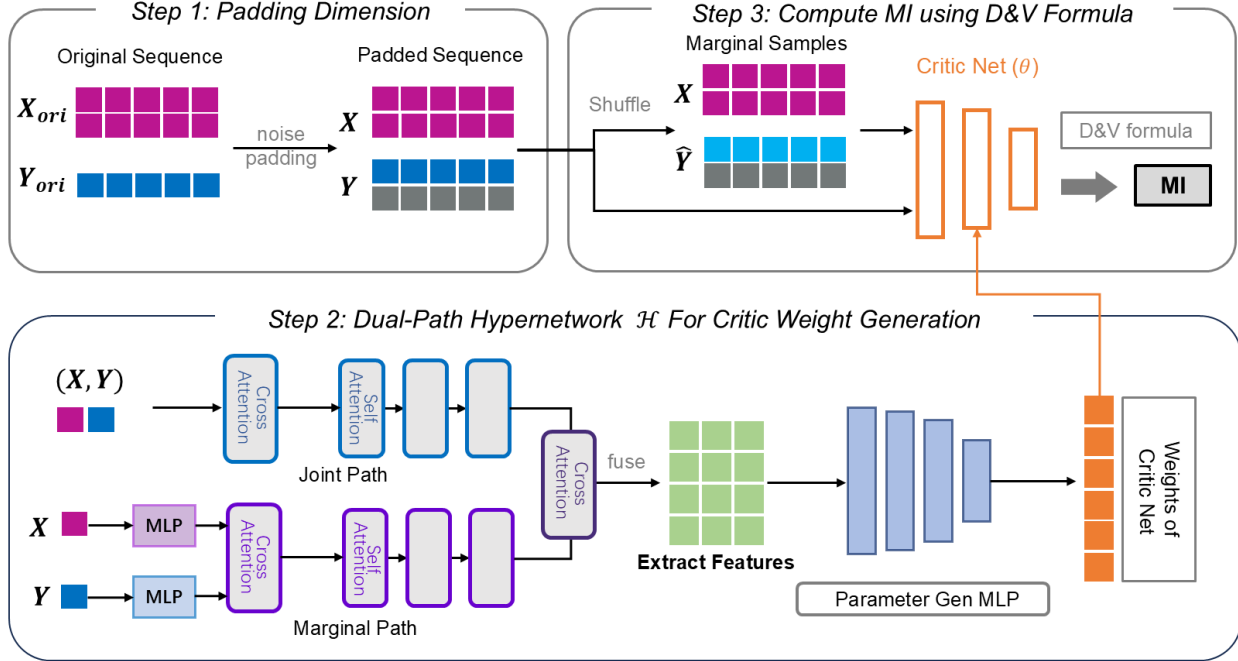


Figure 2. **The InfoAtlas estimation pipeline.** Step 1: We pad input dimensions with noise to ensure all variables share the same dimensionality, while allowing flexible sample sizes. Step 2: A dual-path hypernetwork  $\mathcal{H}$ —with joint and marginal branches—extracts features in alignment with the D-V formulation (Eq. 2). Cross-attention integrates these features, and a parameter-generation MLP is then used to produce the critic parameters. Step 3: The empirical D-V formula (Eq. 5) is applied to joint and marginal samples, with marginals obtained by index permutation, to estimate MI. This pipeline enables single-pass estimation without gradient-based optimization.

### 3.2. Large-Scale Synthetic Pre-training

We pretrain *InfoAtlas* using a comprehensive spectrum of pretraining data (the ‘atlas’). For this purpose, we construct a meta-distribution  $p(\mathcal{D})$  over datasets by systematically generating synthetic data  $\mathcal{D}$  that span diverse statistical properties, drawing from the principle that a model pretrained on sufficiently diverse synthetic data can generalize effectively to real-world unseen data.

**Diversity-driven synthetic distribution generation** Our diversity-aware data generation procedure consists of two complementary steps targeting the diversity of *dependence structure* and *marginal patterns* respectively:

*Dependence diversity via random copula mixture.* We ensure diversity in correlation structure by sampling from a diverse mixtures of copulas with varying dependence properties. Specifically, let  $c_i$  be a copula chosen from a pre-defined pool  $\mathcal{C}$ . We generate samples  $\mathbf{x}, \mathbf{y}$  according to:

$$\mathbf{x}, \mathbf{y} \sim \sum_{i=1}^K \pi_i c_i, \quad c_i \in \mathcal{C} \quad (6)$$

where the parameters of each copula  $c_i$  and the mixture coefficients  $\pi_i$  are randomly initialized. In this work, we employ both Gaussian copulas with rich correlation ma-

trix and Student’s  $t$ -copulas with varying tail dependencies (see Appendix A.2). According to recent vector copula theory (Chen et al., 2025), such copula mixture is a consistent estimator for the dependence structure between  $\mathbf{x}$  and  $\mathbf{y}$  given sufficiently large  $K$ . In this work, we use  $K = 60$  mixtures, substantially beyond  $K = 32$  used in (Chen et al., 2025) for accurate approximation of dependence structure.

*Marginal diversity via random flow transformation.* To complement the correlation diversity, we enhance marginal pattern diversity through flow-based models (Papamakarios et al., 2021; Dinh et al., 2016). Specifically, we apply two flow models  $f_X : \mathbb{R}^{d_X} \rightarrow \mathbb{R}^{d_X}$ ,  $f_Y : \mathbb{R}^{d_Y} \rightarrow \mathbb{R}^{d_Y}$  with randomly initialized parameters to transform data in each training batch:

$$\mathbf{x} \leftarrow f_X(\mathbf{x}), \quad \mathbf{y} \leftarrow f_Y(\mathbf{y}). \quad (7)$$

These invertible transformations preserve mutual information while introducing complex marginal patterns, as  $\mathbb{I}(\mathbf{x}, \mathbf{y}) = \mathbb{I}(f_X(\mathbf{x}), f_Y(\mathbf{y}))$  for any bijective function. Additionally, we apply a differentiable copula transformation using the sofrank function (Blondel et al., 2020), which maps each marginal distribution to approximately uniform distribution  $[0, 1]$ . This normalization enables the model to focus on learning correlation patterns rather than adapting to spurious features irrelevant to the true dependence structure.

**Overall pre-training objective** With comprehensive data generation, the parameters of the hypernetwork  $\mathcal{H}$  are optimized through a meta-learning objective that maximizes the expected accuracy of MI estimation in the distribution of training datasets. Formally, we minimize:

$$\mathcal{L}(\mathcal{H}) = -\mathbb{E}_{\mathcal{D} \sim p(\mathcal{D})} \left[ \hat{\mathbb{I}}_{\mathcal{H}(\mathcal{D})}(\mathbf{x}_{\mathcal{D}}, \mathbf{y}_{\mathcal{D}}) \right], \quad (8)$$

where  $p(\mathcal{D})$  represents the meta-distribution over datasets induced by our synthetic generation process,  $\mathcal{H}(\mathcal{D})$  outputs the critic parameters for dataset  $\mathcal{D}$ , and  $\hat{\mathbb{I}}_{\theta}(\mathbf{x}_{\mathcal{D}}, \mathbf{y}_{\mathcal{D}})$  is the empirical MI estimate using critic parameters  $\theta$  as in Eq. 5.

Proposition A.1 establishes that under mild conditions, the above learning objective yields a *consistent* estimate to the ground truth MI for all  $\mathcal{D}$  such that  $p(\mathcal{D}) > 0$ , thereby guaranteeing convergence to the optimal critic  $\theta^*$ .

We highlight two key differences between the above pre-training pipeline and conventional MI estimator training. First, unlike standard approaches, our synthetic pre-training can generate *unlimited* training datasets, yielding a theoretically infinite sample size per each dataset and a large batch size. Second, our pre-training amortizes learning across dataset through a centralized hypernetwork  $\mathcal{H}$ , enabling knowledge acquired from one dataset to effectively transfer to others. Together, we alleviate known failure modes of conventional neural MI estimation, such as high estimation variances and biases caused by insufficient samples per dataset (McAllester & Stratos, 2020; Song & Ermon, 2019).

#### 4. Scaling to High Dimensions via Slicing

*InfoAtlas* natively supports multivariate inputs up to a pre-defined dimensionality  $D$ . To scale to higher-dimensional data with  $d > D$ , we leverage sliced mutual information (sliced MI), which estimates high-dimensional statistical dependence by aggregating information across multiple low-dimensional projections, referred to as “slices”.

**Slices of high-dimensional dependence** Our intuition is as follows. In a geographical atlas, the information of a specific region may be characterized by different aspects, such as topology, climate or population. Similarly, in a genetic atlas, different slices may reveal gene activity across tissues, cell types, or spatial locations. Each individual view is necessarily partial, but they together complementarily provide a rich characterization of the underlying object.

We apply the same principle to high-dimensional dependence estimation. Rather than estimating MI directly in the original high-dimensional space, we examine many low-dimensional projections of the variables. Each projection captures one specific “view” of the original dependence structure, and aggregating over many such views yields an informative summary of the overall statistical dependence.

Formally, the ( $k$ -)sliced MI under slicing dimensionality  $k$  is defined as (Goldfeld et al., 2022):

$$\begin{aligned} \mathbb{S}\mathbb{I}_k(\mathbf{x}; \mathbf{y}) &= \mathbb{E}_{\mathbf{P}, \mathbf{P}'} \left[ \mathbb{I}(\mathbf{P}^\top \mathbf{x}; \mathbf{P}'^\top \mathbf{y}) \right], \\ &\approx \frac{1}{S} \sum_{j=1}^S \mathbb{I}(\mathbf{P}_j^\top \mathbf{x}; \mathbf{P}'_j^\top \mathbf{y}) \end{aligned} \quad (9)$$

where  $\mathbf{P} \in \text{St}(d_x, k)$  and  $\mathbf{P}' \in \text{St}(d_y, k)$  are random orthonormal projection matrices sampled from the Stiefel manifolds. Note that single-sided slicing could also be used:  $\mathbb{S}\mathbb{I}'_k(\mathbf{x}; \mathbf{y}) = \mathbb{E}_{\mathbf{P}} \left[ \mathbb{I}(\mathbf{P}^\top \mathbf{x}; \mathbf{y}) \right]$ .

Sliced MI has been shown to be an effective measure for quantifying high-dimensional dependence, particularly under moderately large  $k$  and  $S$ . It inherits several key properties of MI, including: (a)  $\mathbb{I} = 0 \Leftrightarrow \mathbb{S}\mathbb{I} = 0$  (Goldfeld & Greenewald, 2021), and (b) maximizing  $\mathbb{I}$  is equivalent to maximizing  $\mathbb{S}\mathbb{I}$  under single-sided slicing (Chen et al., 2023). These properties make sliced MI a theoretically grounded and practically useful alternative to full MI in applications such as independence testing (Tsur et al., 2023b; Goldfeld et al., 2022) and representation learning (Chen et al., 2023; Zhou & Yang, 2024). In many applications, exact MI values is often not of interests; the relative strength of dependence as captured by sliced MI is already highly informative.

**Batched inference across slices** *InfoAtlas* is particularly well suited for sliced MI estimation due to its ability to process multiple slices *in parallel*. Existing neural MI estimators typically require a separate optimization procedure for each projection direction, leading to a time complexity of  $O(ST)$ , where  $T$  denotes the optimization cost, e.g., the number of gradient steps. In contrast, the transformer-based architecture of *InfoAtlas* allows multiple projected datasets to be packaged into a single batch and processed jointly through one feedforward pass, analogous to how large language models process multiple sequences simultaneously:

$$\{\theta_1^*, \dots, \theta_S^*\} = \mathcal{H}(\{\mathcal{D}_j\}_{j=1}^S) \quad (10)$$

where  $\theta_j^*$  denotes the optimal critic predicted by *InfoAtlas* for the  $j$ -th slicing direction and  $\mathcal{D}_j = \{\mathbf{P}_j \mathbf{x}^i, \mathbf{P}'_j \mathbf{y}^i\}_{i=1}^n$  denotes the projected dataset for that direction. This way, we reduce time complexity from  $O(ST)$  to  $O(1)$ , facilitating highly efficient computation of sliced MI.

We note that while theoretically and empirically grounded, sliced MI is not a drop-in replacement for full MI: slicing may miss certain dependence structures, and finite- $S$  averaging can miss rare but informative projection directions. We discuss the failure cases of slicing, together with the effects of  $k, S$  in Appendix A.6 and A.7. Nevertheless, as our slicing-based method strikes a favorable accuracy–efficiency trade-off across a broad range of settings, including high-dimensional real-world tasks, we consider it a practical and scalable approach for statistical dependence measurement.

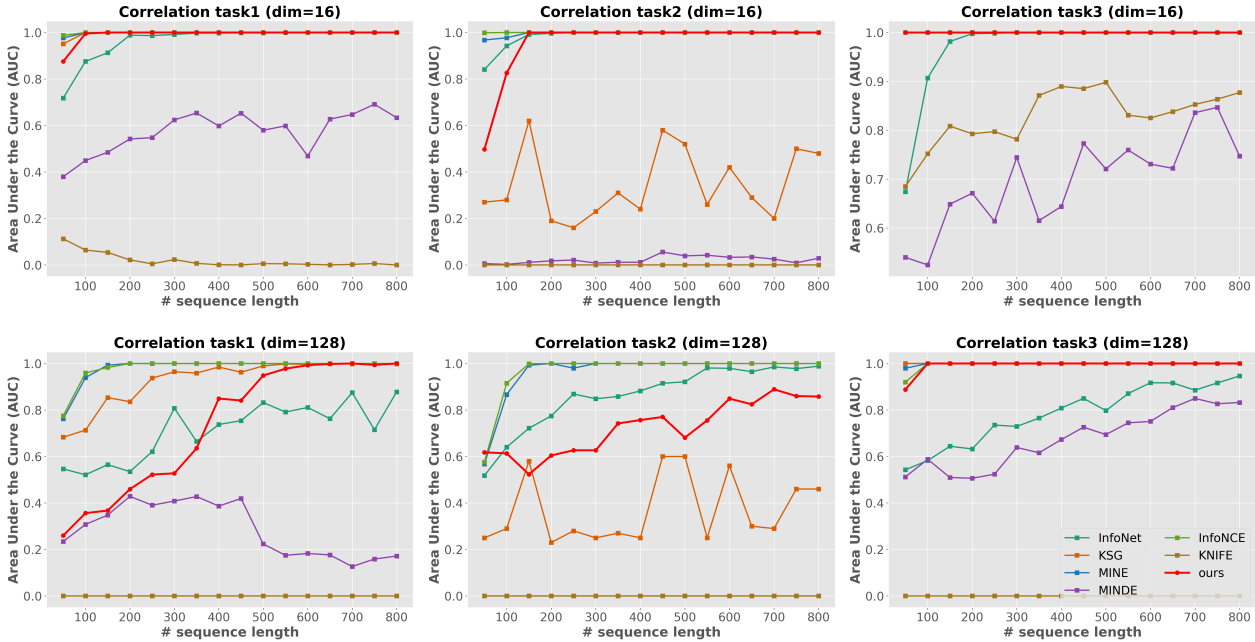


Figure 3. Independence testing under three types of data correlations. Each curve depicts the area under the curve (AUC) of the receiver operating characteristic (ROC) with respect to sequence length  $n$ . Seven MI estimators are compared: *InfoAtlas*, InfoNet, KSG, MINE, MINDE, InfoNCE and KNIFE. *InfoAtlas* uses 5-sliced MI with 32 slices, while InfoNet adopts 1-sliced MI with 128 slices.

## 5. Experiments

### 5.1. Setups

**Slicing.** *InfoAtlas* is pretrained for data with dimensionality up to  $D = 20$ . For high-dimensional inputs where  $d > 20$ , we employ  $k$ -sliced mutual information (Goldfeld et al., 2022), which projects the data onto  $k$ -dimensional random subspaces and averages the MI estimates across projections. Specifically, we compute  $\hat{\mathbb{S}}\mathbb{I}_k(\mathbf{x}, \mathbf{y}) = \frac{1}{S} \sum_{i=1}^S \hat{\mathbb{I}}(\mathbf{P}_i \mathbf{x}, \mathbf{P}'_i \mathbf{y})$  where  $\mathbf{P}_i, \mathbf{P}'_i \in \mathbb{R}^{k \times d}$  are random projection matrices. This approach preserves substantially more correlation structure than 1-dimensional slicing used in previous work (Hu et al., 2024). For other neural methods such as MINDE, we do not employ slicing, as this will involve training  $S$  networks for  $S$  slicing directions, being computationally prohibitive.

**Baselines.** We consider seven MI estimation methods: KSG (Kraskov et al., 2004) and KDE (Silverman, 2018), two classic non-parametric MI estimators that do not require costly iterative optimization; KNIFE (Pichler et al., 2022), which uses KDE with learnable parameters in MI estimate, and neural methods MINE (Belghazi et al., 2018), InfoNCE (Oord et al., 2018), and MINDE (Franzese et al., 2023), all of which require training a network from scratch for each new distribution. We also compare to the pretrained InfoNet (Hu et al., 2024) method whenever appropriate. InfoNet (Hu et al., 2024) is restricted to one dimension, thus requiring using slicing for data dimension beyond  $D = 1$ .

### 5.2. Results

**High-dimensional independence testing.** We first evaluate *InfoAtlas* on its ability to accurately discriminate varying levels of statistical dependency between pairs of random variables in high-dimensional settings. Following the setup in (Goldfeld et al., 2022), we consider various correlation types. For each type, we generate two populations of paired variables: one with no statistical dependence and another with non-trivial dependence. The goal is to evaluate how well the estimated dependency scores separate the two populations. Performance is measured using the area under the precision-recall curve (AUC). All results are the average of 10 independent trials.

As shown in Fig. 3, *InfoAtlas* consistently exhibits strong test power in assessing statistical dependence in high-dimensional settings, particularly when the sample size exceeds 500 (with the exception of correlation type 2 at dimensionality 128). In these settings, *InfoAtlas* achieves performance comparable to gradient-based neural estimators such as MINE, InfoNCE, and MINDE, despite requiring no gradient-based optimization. Compared to InfoNet, *InfoAtlas* attains comparable or superior performance in most cases, especially at moderate-to-large sample sizes (e.g.,  $n \geq 500$ ).

Overall, these results show that *InfoAtlas* can reliably quantify high-dimensional dependence in a zero-shot manner, suggesting its potential for efficient dependence analysis.

Table 1. **Sanity check with BMI benchmark.** We compare *InfoAtlas* with other MI estimators on 8 representative tasks from (Czyż et al., 2023a) with known ground truth MI. Each estimate represents the average over 10 random seeds with  $N = 5000$  samples per task. Task notation indicates distribution type (Mn=Multivariate normal, St=Student- $t$ , Asinh=Arc sinh, Uniform=correlated uniform, Hc=Half cube) and corresponding parameters, with the first two digits indicating dimensionality of  $\mathbf{x}$  and  $\mathbf{y}$  respectively. Methods are color-coded: neural-based methods in green and non-neural methods in blue. **Bold** indicates closest to ground truth, while underlined values show second-best estimates. The rightmost column shows execution time in seconds to compare computational efficiency.

Method*	Tasks								Time (s)
	Mn-dense 5-5-0.5	Spiral 3-3-2-2.0	Asinh@St 5-5-2	St 3-3-3	Uniform 3-3-2-2.0	Hc@Mn 5-5-2	Additive 1-1-0.1	Bimodal 1-1-0.75	
<i>GT</i>	0.59	1.02	0.45	0.18	1.02	1.02	1.71	0.41	–
KSG	0.54	0.75	0.25	0.07	0.79	0.58	1.61	<b>0.41</b>	<u>0.13</u>
KDE	1.59	2.87	2.43	2.36	1.17	2.23	2.94	1.23	2.04
MINE	<b>0.60</b>	<b>1.00</b>	0.53	<u>0.21</u>	<b>1.03</b>	<u>1.06</u>	<b>1.63</b>	0.39	25.9
MINE-5s	<b>0.60</b>	0.90	0.33	0.15	0.93	<u>1.06</u>	1.61	0.38	4.92
MINDE	<b>0.58</b>	0.92	<b>0.43</b>	0.36	0.89	<b>1.01</b>	1.42	0.50	34.2
InfoNCE	<u>0.56</u>	<u>0.98</u>	<u>0.49</u>	<b>0.18</b>	<u>0.97</u>	<b>1.03</b>	<u>1.62</u>	<u>0.40</u>	67.6
KNIFE	0.93	0.10	0.66	0.50	0.07	0.92	0.05	0.65	48.4
<i>InfoAtlas</i>	<b>0.60</b>	0.89	<u>0.41</u>	<u>0.21</u>	0.93	0.96	1.46	0.39	<b>0.09</b>

\*We exclude InfoNet on this task, as InfoNet cannot output exact MI for data beyond 1D.

**Sanity check on benchmark with known MI.** We next consider the benchmarks proposed in (Czyż et al., 2023a), where we select 8 representative tasks with analytically derived ground-truth MI. The test distributions exhibit diverse statistical patterns, ranging from complex Spiral transformation to heavy-tailed Student- $t$  distributions. We generate 5,000 samples for each task.

As presented in Table 1, the MI values predicted by *InfoAtlas* closely align with the ground-truth MI across all tasks. In all tasks, *InfoAtlas* achieves comparable accuracy to the best neural baselines, while being approximately  $300\times$  faster. Compared to KSG and KDE, which requires no neural network training, *InfoAtlas* offers a clear advantage in accuracy.

**CLIP-based image-text embedding analysis** The CLIP model (Radford et al., 2021) encodes images and text into a shared feature space, enabling robust cross-modal understanding by measuring similarity. Here, we assess the correlation between images and their corresponding text annotations by estimating the MI between their latent representations encoded by the pre-trained CLIP model.

We utilize the COCO Captions dataset (Chen et al., 2015), selecting 33,000 image-caption pairs and encoding them into 512-dimensional feature vectors using CLIP. By systematically introducing Gaussian noises to the data, we create conditions where statistical dependence naturally decreases. Our objective is to evaluate whether different estimators can effectively detect these changes with high sensitivity – a spirit similar to the self-consistency test (Song & Ermon, 2019). For each noise level, we conduct 20 experiments, and we report both the mean and the standard deviation.

Our results in Figure 4 demonstrate that *InfoAtlas* achieves a strong performance in detecting noise fluctuations, yielding clearly separated error bounds across different noise levels and hence high sensitivity w.r.t the dependence strengths, while being substantially more efficient than alternative approaches. InfoNet (Hu et al., 2024) is the only method comparable in efficiency, but its accuracy is much worse than our method due to its reliance on 1-slicing, which discards a large amount of information despite using more slicing directions.

**Real-world motion trajectory modeling** To assess *InfoAtlas*’s generalization ability to accurately capture complex real-world relationships, we utilize the PointOdyssey dataset (Zheng et al., 2023), which contains multi-dimensional ground-truth motion trajectories of points on objects across video frames. In this task, a reference point  $P^*$  is selected, and we estimate the mutual information  $\mathbb{I}(\text{trajectory}(P^*), \text{trajectory}(P))$  between its trajectory and those of all other points  $P$  in the video. Since points on the same object  $O$  typically exhibit stronger spatial correlations in their trajectories than those on different objects, the following relationship is expected (where  $t$  is a threshold):

$$\begin{cases} \mathbb{I}(\text{trajectory}(P^*), \text{trajectory}(P)) > t & \text{if } P^*, P \in O \\ \mathbb{I}(\text{trajectory}(P^*), \text{trajectory}(P)) \leq t & \text{if } P^* \in O, P \notin O \end{cases}$$

An accurate mutual information estimator should correctly reflect this relationship. Figure 5 visualizes the raw MI estimates between the reference point and all other points. *InfoAtlas* successfully identifies points that belong to the same object by flagging a high MI value, demonstrating

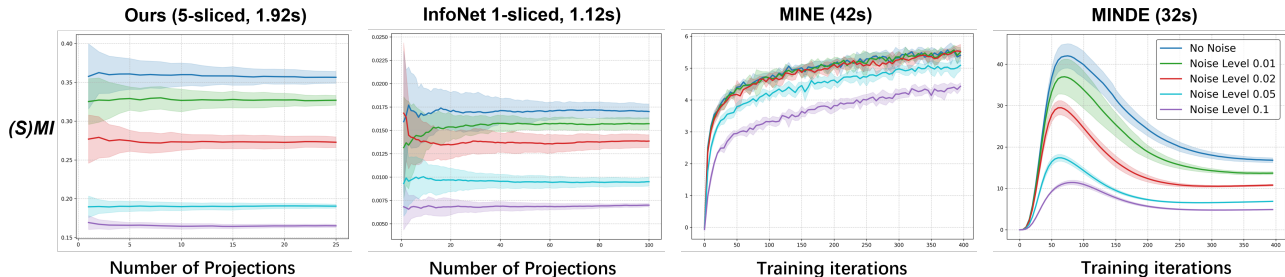


Figure 4. Comparing different methods on 512-dimensional CLIP-encoded image-text representations across five noise levels. The light-colored areas indicate error bounds from 20 repeated experiments. (Left to right) *InfoAtlas* with 5-sliced MI using  $S = 25$  random projections; InfoNet with 1-sliced MI using more projections (up to  $S = 128$ ); MINE and MINDE estimating original MI via gradient-based optimization. *InfoAtlas* demonstrates superior noise level discrimination with clearly separated error bounds, while maintaining significantly faster computation time (noted in parentheses) compared to neural-based alternatives. Slicing are not employed in MINE and MINDE, as this will require training  $S$  different networks for  $S$  slicing directions and does not contribute to efficiency.

its effectiveness in modeling complex spatial dependencies in real-world motion data. Remarkably, *InfoAtlas* delivers results for all point pairs within only a few seconds.

We further evaluate *InfoAtlas* on 12 objects sampled from 6 different videos, where we segment video objects using MI estimators. A visualization of these objects can be found in Fig. 8 and Fig. 9 in the appendix. In Fig. 5, we compare different methods by comparing the estimated pointwise correlations with the ground-truth segmentation masks. On this out-of-distribution dataset, *InfoAtlas* achieves competitive segmentation accuracy while being orders more efficient.

**Robotic manipulation concept discovery** To fully demonstrate our *InfoAtlas*'s strong potential in complex, large-scale applications in real world, we further apply our method to a robotic manipulation concept discovery task. The goal is to identify *key states*—critical moments in a trajectory  $\tau^i = \{s_t^i\}_{t=1}^T$  that carry strong physical significance (e.g., "peg aligned with hole"). Identifying such key states has been shown to significantly improve robotics policy training (Zhou & Yang, 2024). Following recent frameworks (Zhou & Yang, 2024), we extract key states by maximizing the mutual information  $\mathbb{I}(s_t; s_{t-\Delta t})$  between the key state  $s_t$  and its prior states  $s_{t-\Delta t}$  ( $\Delta t > 0$ ):

$$\max \mathbb{I}(s_t; s_{t-\Delta t}), \quad \forall t \in 1, 2, \dots, T$$

where each  $s_t \in \mathbb{R}^{100}$  encodes the states of an environmental observation. Here  $\Delta t = 2$  and  $T = 200$ . We use 5-sliced MI as a surrogate of  $\mathbb{I}(s_t; s_{t-\Delta t})$ , where each sliced MI is computed using *InfoAtlas*. A total number of  $S = 40$  slices are used.

We train manipulation policies using key concepts extracted from different MI estimators and compare the success rates (SR) of the resulting policies. We consider three manipulation tasks from ManiSkill 2 (Gu et al., 2023)—Pick Cube, Stack Cube, and Peg Insertion. As summarized in Table 2,

concepts derived from *InfoAtlas* substantially outperform those from other methods in terms of success rate under equal or even reduced time budgets. These results highlight that our model generalizes effectively to unseen, complex real-world data, serving as a versatile toolbox for dependence assessment in modern machine learning applications.

## 6. Related Work

**Neural mutual information estimators.** A series of powerful, neural network-based methods have been developed for MI estimation. The most prominent among these is the MINE estimator (Belghazi et al., 2018), which builds on the Donsker–Varadhan representation. Other approaches rely on density-ratio estimation (Letizia et al., 2024; Gutmann & Hyvärinen, 2010), direct density modeling (Song & Ermon, 2019), score function estimation (Franzese et al., 2023), or leverage normalizing flows (Duong & Nguyen, 2023; Butakov et al., 2024) and autoencoders (Gowri et al., 2024; Butakov et al., 2023) to construct MI estimates. Despite methodological differences, these methods are primarily designed to improve estimation accuracy, and the computational overhead associated with neural network training is often overlooked in real-world deployment. In contrast, we target the orthogonal dimension of computational efficiency, replacing costly iterative optimization with a lightweight forward pass at inference time.

**Efficient computation of mutual information.** Various approaches have been developed for rapid MI computation, each with different trade-offs. Non-parametric methods (Kraskov et al., 2004; Moon et al., 1995; Silverman, 2018) offer training-free efficiency but typically lack the capacity to capture complex dependencies in high-dimensional data. Copula-based approaches (Keziou & Regnault, 2016; Safaai et al., 2018; Purkayastha & Song, 2024; Zeng et al., 2018) balance efficiency with accuracy by assuming data follows a known copula family (e.g., Gaussian copula), but




Figure 5. Point trajectory mutual information for video object segmentation on PointOdyssey (Zheng et al., 2023). We estimate mutual information  $\mathbb{I}(\text{trajectory}(P^*), \text{trajectory}(P))$  between a reference point trajectory  $P^*$  (marked by  $\star$ ) and every other point trajectory  $P$  across video frames, yielding  $\sim 4 \times 10^3$  MI terms per video. (a,b) The estimated MI is consistently higher for points belonging to the same object as the reference point than for points on other objects. (c) Using trajectory MI  $\mathbb{I}(\text{trajectory}(P^*), \text{trajectory}(P))$  as an affinity for video segmentation, where we report the area under the precision–recall curve (AUC-PR) and the time of different MI estimators.

Table 2. Comparison of policy success rates using key states extracted by different MI estimators across three robotic tasks. MINE-100 denotes training MINE for 100 iterations, while No-MI-Loss removes the MI maximization term when identifying key states. MI is estimated on 100-dimensional variables using a batch size of 100. InfoAtlas uses 25 slices, whereas InfoNet uses 250 slices.

Tasks	Pick Cube		Stack Cube		Peg Insertion		Time (s)
	Seen	Unseen	Seen	Unseen	Seen	Unseen	
No-MI-Loss	66.0	60.0	67.4	<b>41.0</b>	38.6	9.3	–
MINE-100	86.4	81.0	68.0	37.0	55.0	13.5	0.62
MINE-1000	81.2	81.0	61.2	37.0	65.4	17.8	6.01
InfoNet	91.0	76.0	63.0	27.0	46.4	9.8	1.23
InfoAtlas (Ours)	<b>94.2</b>	<b>82.0</b>	<b>68.2</b>	37.0	<b>72.4</b>	<b>18.3</b>	2.17

this assumption limits their applicability to general distributions. The recent InfoNet (Hu et al., 2024) enables fast MI estimation through neural network pretraining – a concept related to our work. However, InfoNet is restricted to scalar inputs due to its lookup table designs and limited pretraining, whereas we support handling of multivariate variables with varying dimensionalities using a single unified model. Concurrent to our work, MIST (Gritsai et al., 2025) also considers pretrained MI estimators for multivariate data, tailored to low-sample regimes. However, it does not support handling varying dimensionalities, and its pretraining relies on a single predefined benchmark (Czyż et al., 2023a), in contrast to our universal dependence generation framework.

**Foundation models for statistical analysis.** Recent advances in large-scale pretrained models have enabled direct statistical analysis on raw data without gradient-based optimization at test time. For instance, LLM-based frameworks (Requeima et al., 2024; Siddiqui et al., 2025) directly leverage large language models to perform classification and regression through direct inference, while (Sun et al., 2026) develops a LLM-based data agent. Beyond off-the-shelf LLMs, the community has also developed pretrained transformers tailored to specific data-analysis tasks, including prediction on small tabular datasets (Hollmann et al., 2025), Bayesian inference (Vetter et al., 2025; Teh et al., 2025), and time-series forecasting (Ansari et al., 2024). Our work is

closely related to this emerging paradigm. However, unlike existing methods primarily focus on *one-way* prediction or forecasting, we study the problem of quantifying *mutual* dependence between two multivariate random variables. This requires modeling interactions between two sets of samples rather than predicting one variable from another, motivating our dual-path attention architecture and diverse synthetic distribution generation strategy for effective generalization.

## 7. Conclusion

We introduce *InfoAtlas*, a foundation model-like architecture for zero-shot estimation of statistical dependence between multivariate variables. Through large-scale pretraining on synthetic distributions spanning diverse dependence structures and marginal patterns, *InfoAtlas* learns to directly infer dependence strengths from raw data in a single forward pass, completely eliminating costly per-dataset optimization. By leveraging slicing techniques, it further scales gracefully and effectively to high-dimensional settings. Extensive evaluations demonstrate that *InfoAtlas* matches state-of-the-art neural methods in accuracy while being orders of magnitude faster, and it generalizes effectively to unseen, real-world scenarios. By reframing MI estimation as one-step inference rather than dataset-specific optimization, *InfoAtlas* enables a paradigm shift toward scalable dependency estimation, with particular promise for large-scale and real-time applications.

## Acknowledgments

This work is supported by the Early Career Scheme of the Research Grants Council (RGC) grant # 27207224, the HKU-100 Award, and the HKU Shanghai Intelligent Computing Research Center (ICRC). Yanzhi Chen acknowledges the support from the Qualcomm Innovation Fellowship.

## Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

## References

- Ansari, A. F., Stella, L., Turkmen, C., Zhang, X., Mercado, P., Shen, H., Shchur, O., Rangapuram, S. S., Arango, S. P., Kapoor, S., et al. Chronos: Learning the language of time series. *arXiv preprint arXiv:2403.07815*, 2024.
- Belghazi, M. I., Baratin, A., Rajeshwar, S., Ozair, S., Bengio, Y., Courville, A., and Hjelm, D. Mutual information neural estimation. In *International conference on machine learning*, pp. 531–540. PMLR, 2018.
- Blondel, M., Teboul, O., Berthet, Q., and Djolonga, J. Fast differentiable sorting and ranking. In *International Conference on Machine Learning*, pp. 950–959. PMLR, 2020.
- Butakov, I., Tolmachev, A., Malanchuk, S., Neopryatnaya, A., Frolov, A., and Andreev, K. Information bottleneck analysis of deep neural networks via lossy compression. *arXiv preprint arXiv:2305.08013*, 2023.
- Butakov, I., Tolmachev, A., Malanchuk, S., Neopryatnaya, A., and Frolov, A. Mutual information estimation via normalizing flows. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- Chen, X., Fang, H., Lin, T.-Y., Vedantam, R., Gupta, S., Dollár, P., and Zitnick, C. L. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015.
- Chen, X., Duan, Y., Houthoofd, R., Schulman, J., Sutskever, I., and Abbeel, P. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. *Advances in neural information processing systems*, 29, 2016.
- Chen, Y., Zhang, D., Gutmann, M., Courville, A., and Zhu, Z. Neural approximate sufficient statistics for implicit models. *arXiv preprint arXiv:2010.10079*, 2020.
- Chen, Y., Li, Y., Weller, A., et al. Scalable infomn learning. *Advances in Neural Information Processing Systems*, 35: 2226–2239, 2022.
- Chen, Y., Gutmann, M. U., and Weller, A. Is learning summary statistics necessary for likelihood-free inference? In *International Conference on Machine Learning*, pp. 4529–4544. PMLR, 2023.
- Chen, Y., Ou, Z., Weller, A., and Gutmann, M. U. Neural mutual information estimation with vector copulas. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.
- Choi, K. and Lee, S. Regularized mutual information neural estimation. 2020.
- Comanici, G., Bieber, E., Schaekermann, M., Pasupat, I., Sachdeva, N., Dhillon, I., Blistein, M., Ram, O., Zhang, D., Rosen, E., et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025.
- Czyż, P., Grabowski, F., Vogt, J. E., Beerenwinkel, N., and Marx, A. Beyond normal: On the evaluation of mutual information estimators. *arXiv preprint arXiv:2306.11078*, 2023a.
- Czyż, P., Grabowski, F., Vogt, J. E., Beerenwinkel, N., and Marx, A. On the properties and estimation of pointwise mutual information profiles. *arXiv preprint arXiv:2310.10240*, 2023b.
- Dinh, L., Sohl-Dickstein, J., and Bengio, S. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.
- Donsker, M. D. and Varadhan, S. S. Asymptotic evaluation of certain markov process expectations for large time. iv. *Communications on pure and applied mathematics*, 36 (2):183–212, 1983.
- Du, Y., de Bock, G. H., Vonk, J. M., Pham, A. T., van der Ende, M. Y., Snieder, H., Smidt, N., Krabbe, P. F., Alizadeh, B. Z., Lunter, G., et al. Lifestyle factors and incident multimorbidity related to chronic disease: a population-based cohort study. *European Journal of Ageing*, 21(1):37, 2024.
- Duong, B. and Nguyen, T. Diffeomorphic information neural estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 7468–7475, 2023.
- Franzese, G., Bounoua, M., and Michiardi, P. Minde: Mutual information neural diffusion estimation. *arXiv preprint arXiv:2310.09031*, 2023.

- Goldfeld, Z. and Greenewald, K. Sliced mutual information: A scalable measure of statistical dependence. *Advances in Neural Information Processing Systems*, 34:17567–17578, 2021.
- Goldfeld, Z., Greenewald, K., Nuradha, T., and Reeves, G.  $k$ -sliced mutual information: A quantitative study of scalability with dimension. *Advances in neural information processing systems*, 35:15982–15995, 2022.
- Gowri, G., Lun, X., Klein, A., and Yin, P. Approximating mutual information of high-dimensional variables using learned representations. *Advances in Neural Information Processing Systems*, 37:132843–132875, 2024.
- Gritsai, G., Richards, M., Méloux, M., Cho, K., and Peyrard, M. Mist: Mutual information via supervised training. *arXiv preprint arXiv:2511.18945*, 2025.
- Gu, J., Xiang, F., Li, X., Ling, Z., Liu, X., Mu, T., Tang, Y., Tao, S., Wei, X., Yao, Y., et al. Maniskill2: A unified benchmark for generalizable manipulation skills. *arXiv preprint arXiv:2302.04659*, 2023.
- Gutmann, M. and Hyvärinen, A. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pp. 297–304. JMLR Workshop and Conference Proceedings, 2010.
- Hollmann, N., Müller, S., Purucker, L., Krishnakumar, A., Körfer, M., Hoo, S. B., Schirrmester, R. T., and Hutter, F. Accurate predictions on small data with a tabular foundation model. *Nature*, 637(8045):319–326, 2025.
- Hu, Z., Kang, S., Zeng, Q., Huang, K., and Yang, Y. Infonet: Neural estimation of mutual information without test-time optimization. *arXiv preprint arXiv:2402.10158*, 2024.
- Jaegle, A., Borgeaud, S., Alayrac, J.-B., Doersch, C., Ionescu, C., Ding, D., Koppula, S., Zoran, D., Brock, A., Shelhamer, E., et al. Perceiver io: A general architecture for structured inputs & outputs. *arXiv preprint arXiv:2107.14795*, 2021.
- Keziou, A. and Regnault, P. Semiparametric estimation of mutual information and related criteria: Optimal test of independence. *IEEE Transactions on Information Theory*, 63(1):57–71, 2016.
- Kraskov, A., Stögbauer, H., and Grassberger, P. Estimating mutual information. *Physical review E*, 69(6):066138, 2004.
- Kullback, S. *Information theory and statistics*. Courier Corporation, 1997.
- Letizia, N. A., Novello, N., and Tonello, A. M. Mutual information estimation via  $f$ -divergence and data derangements. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- Maanpää, J., Pesonen, J., Hytti, H., Melekhov, I., Kannala, J., Manninen, P., Kukko, A., and Hyypä, J. Dense road surface grip map prediction from multimodal image data. In *International Conference on Pattern Recognition*, pp. 387–404. Springer, 2025.
- McAllester, D. and Stratos, K. Formal limitations on the measurement of mutual information. In *International Conference on Artificial Intelligence and Statistics*, pp. 875–884. PMLR, 2020.
- Moon, Y.-I., Rajagopalan, B., and Lall, U. Estimation of mutual information using kernel density estimators. *Physical Review E*, 52(3):2318, 1995.
- Oord, A. v. d., Li, Y., and Vinyals, O. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- Papamakarios, G., Nalisnick, E., Rezende, D. J., Mohamed, S., and Lakshminarayanan, B. Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research*, 22(57):1–64, 2021.
- Pichler, G., Colombo, P. J. A., Boudiaf, M., Koliander, G., and Piantanida, P. A differential entropy estimator for training neural networks. In *International Conference on Machine Learning*, pp. 17691–17715. PMLR, 2022.
- Poole, B., Ozair, S., Van Den Oord, A., Alemi, A., and Tucker, G. On variational bounds of mutual information. In *International Conference on Machine Learning*, pp. 5171–5180. PMLR, 2019.
- Purkayastha, S. and Song, P. X.-K. fastmi: A fast and consistent copula-based nonparametric estimator of mutual information. *Journal of Multivariate Analysis*, 201:105270, 2024.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021.
- Requeima, J., Bronskill, J., Choi, D., Turner, R., and Duvenaud, D. K. Llm processes: Numerical predictive distributions conditioned on natural language. *Advances in Neural Information Processing Systems*, 37:109609–109671, 2024.
- Safaai, H., Onken, A., Harvey, C. D., and Panzeri, S. Information estimation using nonparametric copulas. *Physical Review E*, 98(5):053302, 2018.

- Shannon, C. E. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948.
- Siddiqui, S. A., Chen, Y., Heo, J., Xia, M., and Weller, A. On evaluating llms’ capabilities as functional approximators: A bayesian evaluation framework. In *Proceedings of the 31st International Conference on Computational Linguistics*, pp. 5826–5835, 2025.
- Silverman, B. W. *Density estimation for statistics and data analysis*. Routledge, 2018.
- Song, J. and Ermon, S. Understanding the limitations of variational mutual information estimators. In *International Conference on Learning Representations*, 2019.
- Sun, M., Han, R., Jiang, B., Qi, H., Sun, D., Yuan, Y., and Huang, J. Lambda: A large model based data agent. *Journal of the American Statistical Association*, 121(553): 1–13, 2026.
- Teh, A., Jabbour, M., and Polyanskiy, Y. Solving empirical bayes via transformers. *arXiv preprint arXiv:2502.09844*, 2025.
- Tsai, Y.-H. H., Zhao, H., Yamada, M., Morency, L.-P., and Salakhutdinov, R. R. Neural methods for point-wise dependency estimation. *Advances in Neural Information Processing Systems*, 33:62–72, 2020.
- Tschannen, M., Djolonga, J., Rubenstein, P. K., Gelly, S., and Lucic, M. On mutual information maximization for representation learning. *arXiv preprint arXiv:1907.13625*, 2019.
- Tsur, D., Aharoni, Z., Goldfeld, Z., and Permuter, H. Neural estimation and optimization of directed information over continuous spaces. *IEEE Transactions on Information Theory*, 69(8):4777–4798, 2023a.
- Tsur, D., Goldfeld, Z., and Greenewald, K. Max-sliced mutual information. *Advances in neural information processing systems*, 36:80338–80351, 2023b.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Vetter, J., Gloeckler, M., Gedon, D., and Macke, J. H. Effortless, simulation-efficient bayesian inference using tabular foundation models. *arXiv preprint arXiv:2504.17660*, 2025.
- Zeng, X., Xia, Y., and Tong, H. Jackknife approach to the estimation of mutual information. *Proceedings of the National Academy of Sciences*, 115(40):9956–9961, 2018.
- Zheng, Y., Harley, A. W., Shen, B., Wetzstein, G., and Guibas, L. J. Pointodyssey: A large-scale synthetic dataset for long-term point tracking. *arXiv preprint arXiv:2307.15055*, 2023.
- Zhou, P. and Yang, Y. Maxmi: A maximal mutual information criterion for manipulation concept discovery. In *European Conference on Computer Vision*, pp. 88–105. Springer, 2024.

## A. Technical Appendix

### A.1. Theoretical foundations

**Proposition A.1** (Consistency of the estimator w.r.t sample size  $n$ ). *Let  $p$  be a probability measure over datasets  $\mathcal{D}$ . For each  $\mathcal{D}$ , let  $(\mathbf{x}_{\mathcal{D}}, \mathbf{y}_{\mathcal{D}}) \sim p_{\mathbf{x}_{\mathcal{D}}, \mathbf{y}_{\mathcal{D}}}$  with  $p_{\mathbf{x}_{\mathcal{D}}, \mathbf{y}_{\mathcal{D}}} \ll p_{\mathbf{x}_{\mathcal{D}}} \cdot p_{\mathbf{y}_{\mathcal{D}}}$  and  $I(\mathbf{x}_{\mathcal{D}}; \mathbf{y}_{\mathcal{D}}) < \infty$ . For any admissible critic  $\theta \in \Theta$ , define*

$$\hat{I}_{\theta}(\mathbf{x}_{\mathcal{D}}; \mathbf{y}_{\mathcal{D}}) := \mathbb{E}_{p_{\mathbf{x}_{\mathcal{D}}, \mathbf{y}_{\mathcal{D}}}}[\theta(\mathbf{x}_{\mathcal{D}}, \mathbf{y}_{\mathcal{D}})] - \log \mathbb{E}_{p_{\mathbf{x}_{\mathcal{D}}} \cdot p_{\mathbf{y}_{\mathcal{D}}}}[e^{\theta(\mathbf{x}_{\mathcal{D}}, \mathbf{y}_{\mathcal{D}})}].$$

and let  $\hat{I}_{\theta}^n(\mathbf{x}_{\mathcal{D}}; \mathbf{y}_{\mathcal{D}})$  be its empirical estimate with  $n$  sample.

Assume:

- (i) (**Per- $\mathcal{D}$  attainment**) For  $p$ -a.e.  $\mathcal{D}$ , there exists  $\theta_{\mathcal{D}}^* \in \Theta$  attaining the supremum:  $I(\mathbf{x}_{\mathcal{D}}; \mathbf{y}_{\mathcal{D}}) = \sup_{\theta \in \Theta} \hat{I}_{\theta}(\mathbf{x}_{\mathcal{D}}; \mathbf{y}_{\mathcal{D}})$ .
- (ii) (**Hypernetwork capacity**) The optimal selector  $\mathcal{H} : \mathcal{D} \mapsto \Theta$  is within the class of the hypernetwork.
- (iii) (**Population integrability**) The expectation  $\mathbb{E}_{\mathcal{D}}[I(\mathbf{x}_{\mathcal{D}}, \mathbf{y}_{\mathcal{D}})]$  satisfies  $\mathbb{E}_{\mathcal{D}}[I(\mathbf{x}_{\mathcal{D}}, \mathbf{y}_{\mathcal{D}})] < \infty$

Define  $J(\mathcal{H}) := \mathbb{E}_{\mathcal{D}}[\hat{I}_{\mathcal{H}(\mathcal{D})}(\mathbf{x}_{\mathcal{D}}; \mathbf{y}_{\mathcal{D}})]$  and let  $J^n(\mathcal{H}) := \mathbb{E}_{\mathcal{D}}[\hat{I}_{\mathcal{H}(\mathcal{D})}^n(\mathbf{x}_{\mathcal{D}}; \mathbf{y}_{\mathcal{D}})]$  be its finite sample estimate with  $n$  samples. Let  $\mathcal{H}^* \in \arg \max_{\mathcal{H}} J^n(\mathcal{H})$ . Then the estimator

$$\hat{I}_{\mathcal{H}^*(\mathcal{D})}(\mathbf{x}_{\mathcal{D}}; \mathbf{y}_{\mathcal{D}})$$

is a consistent estimate to  $I(\mathbf{x}_{\mathcal{D}}; \mathbf{y}_{\mathcal{D}})$   $p$ -a.e.  $\mathcal{D}$ .

*Proof.* We begin with the following identity:

$$\sup_{\mathcal{H}} J^n(\mathcal{H}) = \sup_{\mathcal{H}} \mathbb{E}[\hat{I}_{\mathcal{H}(\mathcal{D})}^n(\mathbf{x}_{\mathcal{D}}; \mathbf{y}_{\mathcal{D}})] = \mathbb{E}[\sup_{\theta} \hat{I}_{\theta}^n(\mathbf{x}_{\mathcal{D}}; \mathbf{y}_{\mathcal{D}})]$$

where the second equality comes from the fact that the hypernetwork  $\mathcal{H}$  is a universal selector for  $\mathcal{D} \rightarrow \Theta$ , so that the supremum for each  $\hat{I}_{\mathcal{H}(\mathcal{D})}^n(\mathbf{x}_{\mathcal{D}}; \mathbf{y}_{\mathcal{D}})$  is reachable.

This suggests that for the optimal  $\mathcal{H}^* = \arg \max_{\mathcal{H}} J^n(\mathcal{H})$ , we have

$$\mathcal{H}^*(\mathcal{D}) = \sup_{\theta} \hat{I}_{\theta}^n(\mathbf{x}_{\mathcal{D}}; \mathbf{y}_{\mathcal{D}})$$

According to (Belghazi et al., 2018), the estimator  $\hat{I}(\mathbf{x}_{\mathcal{D}}, \mathbf{y}_{\mathcal{D}}) = \sup_{\theta} \hat{I}_{\theta}^n(\mathbf{x}_{\mathcal{D}}, \mathbf{y}_{\mathcal{D}})$  itself is a consistent estimate of  $I(\mathbf{x}_{\mathcal{D}}, \mathbf{y}_{\mathcal{D}})$ . This suggests that for each  $\mathcal{D}$  and every  $\epsilon > 0$ , there exists  $n(\mathcal{D}) \in \mathbb{N}$ , such that

$$\left| I(\mathbf{x}_{\mathcal{D}}, \mathbf{y}_{\mathcal{D}}) - \sup_{\theta} \hat{I}_{\theta}^n(\mathbf{x}_{\mathcal{D}}, \mathbf{y}_{\mathcal{D}}) \right| \leq \epsilon, \quad \forall n \geq n(\mathcal{D}), \quad a.s.$$

By taking  $n' = \sup_{\mathcal{D}} n(\mathcal{D})$ , substituting  $\sup_{\theta} \hat{I}_{\theta}^n(\mathbf{x}_{\mathcal{D}}; \mathbf{y}_{\mathcal{D}}) = \mathcal{H}^*(\mathcal{D})$ , we have that for every  $\epsilon > 0$ ,

$$\left| I(\mathbf{x}_{\mathcal{D}}, \mathbf{y}_{\mathcal{D}}) - \hat{I}_{\mathcal{H}^*(\mathcal{D})}^n(\mathbf{x}_{\mathcal{D}}, \mathbf{y}_{\mathcal{D}}) \right| \leq \epsilon, \quad \forall \mathcal{D}, \quad \forall n \geq n', \quad a.s.$$

which completes the proof.  $\square$

**Proposition A.2** (Positive Definiteness of Generated Covariance Matrix). *The covariance matrix constructed by Algorithm 2 is positive definite almost surely, with controllable condition number through the rank parameter  $m$ .*

*Proof.* We construct  $\Sigma = \mathbf{W}\mathbf{W}^{\top} + \mathbf{D}$  where  $\mathbf{W} \in \mathbb{R}^{d \times m}$  with  $W_{ij} \sim \mathcal{N}(0, 1)$  and  $\mathbf{D} = \text{diag}(d_1, \dots, d_d)$  with  $d_i \sim \text{Uniform}(0, 1)$ .

For any nonzero  $\mathbf{x} \in \mathbb{R}^d$ :

$$\mathbf{x}^{\top} \Sigma \mathbf{x} = \underbrace{\|\mathbf{W}^{\top} \mathbf{x}\|_2^2}_{\geq 0} + \underbrace{\sum_{i=1}^d d_i x_i^2}_{> 0 \text{ a.s.}} > 0$$

The eigenvalues satisfy  $\lambda_{\min}(\Sigma) \geq \min_i d_i > 0$  and  $\lambda_{\max}(\Sigma) \leq \|\mathbf{W}\|_F^2 + \max_i d_i$ . The expected condition number scales as  $\mathcal{O}(m)$ , allowing control over numerical stability.  $\square$

**Proposition A.3** (Invariance of MI under Noise Padding). *Let  $(X, Y)$  be random variables with  $X \in \mathbb{R}^{d_x}$ ,  $Y \in \mathbb{R}^{d_y}$ . For any independent noise variables  $\epsilon_X \perp \epsilon_Y \perp (X, Y)$  of arbitrary dimensions, defining  $X' = [X, \epsilon_X]$  and  $Y' = [Y, \epsilon_Y]$ :*

$$I(X'; Y') = I(X; Y)$$

This invariance holds for any MI estimator, including the DV representation used in InfoAtlas.

*Proof.* Since  $\epsilon_X \perp \epsilon_Y \perp (X, Y)$ , the joint and marginal densities factor as:

$$p(x', y') = p(x, y) \cdot p(\epsilon_x) \cdot p(\epsilon_y) \quad (11)$$

$$p(x') = p(x) \cdot p(\epsilon_x), \quad p(y') = p(y) \cdot p(\epsilon_y) \quad (12)$$

Therefore, the density ratio is preserved:

$$\frac{p(x', y')}{p(x')p(y')} = \frac{p(x, y)}{p(x)p(y)}$$

For the DV representation specifically:

$$I(X'; Y') = \sup_{\theta'} \mathbb{E}_{p(x', y')}[\theta'] - \log \mathbb{E}_{p(x') \otimes p(y')}[e^{\theta'}] \quad (13)$$

$$= \sup_{\theta} \mathbb{E}_{p(x, y)}[\theta] - \log \mathbb{E}_{p(x) \otimes p(y)}[e^{\theta}] \quad (14)$$

$$= I(X; Y) \quad (15)$$

where the optimal critic  $\theta'^*(x', y') = \theta^*(x, y)$  depends only on the non-noise components.  $\square$

**Corollary A.4** ( $k$ -Sliced MI Invariance and Approximation). *For high-dimensional variables with  $d > d_{\max}$ , the  $k$ -sliced MI with padding satisfies:*

1. **Invariance:** For padded variables  $X', Y'$  and random projections  $\{P_i\}_{i=1}^k$ :

$$I_{k\text{-sliced}}(X'; Y') = \frac{1}{k} \sum_{i=1}^k I(P_i X'; P_i Y') = I_{k\text{-sliced}}(X; Y)$$

2. **Approximation Quality:** Under mild regularity conditions:

$$|I_{k\text{-sliced}}(X; Y) - I(X; Y)| \leq \frac{C}{\sqrt{k}} \cdot \sqrt{\text{Var}_P[I(PX; PY)]}$$

where  $C$  is a universal constant and the variance is over random projections.

*Proof.* We prove the two parts respectively as follows.

**Part 1:** Follows directly from Proposition A.3 applied to each projection.

**Part 2:** By the central limit theorem over independent projections:

$$\sqrt{k}(I_{k\text{-sliced}} - \mathbb{E}_P[I(PX; PY)]) \xrightarrow{d} \mathcal{N}(0, \text{Var}_P[I(PX; PY)])$$

The bias  $|\mathbb{E}_P[I(PX; PY)] - I(X; Y)|$  depends on the projection dimension and decreases as more projections capture the dependency structure.  $\square$

---

**Algorithm 1** Full Training Sequence Generation Pipeline

---

- Require:** Variable dimensions  $d_x, d_y$ , max components  $K_{\max} = 60$ , samples  $N$ , max dim  $d_{\max}$
- 1: Randomly select  $K \in \{1, 2, \dots, K_{\max}\}$  and sample weights  $\{\pi_i\}_{i=1}^K$  s.t.  $\sum_{i=1}^K \pi_i = 1$
  - 2: Set total dimension  $d = d_x + d_y$
  - 3: **for** each component  $i = 1$  to  $K$  **do**
  - 4:   Sample mean  $\mu_i \in \mathbb{R}^d$  with elements from  $\text{Uniform}([-5, 5])$
  - 5:   Select rank  $m \in \{1, 2, \dots, d\}$  and generate  $\mathbf{W} \in \mathbb{R}^{d \times m}$  with  $W_{ij} \sim \mathcal{N}(0, 1)$
  - 6:   Generate diagonal matrix  $\mathbf{D}$  with  $D_{ii} \sim \text{Uniform}(0, 1)$
  - 7:   Compute  $\Sigma'_i = \mathbf{W}\mathbf{W}^T + \mathbf{D}$  and normalize to correlation matrix  $\text{Corr}_i$ :  $\text{Corr}_{i_{jk}} = \frac{\Sigma'_{i_{jk}}}{\sqrt{\Sigma'_{i_{jj}}\Sigma'_{i_{kk}}}}$
  - 8:   Sample variances  $\sigma^2 \in \mathbb{R}^d$  from  $\text{Uniform}([0.01, 10])$
  - 9:   Set  $\Sigma_i = \text{diag}(\sigma) \cdot \text{Corr}_i \cdot \text{diag}(\sigma)$
  - 10: **end for**
  - 11: Define GMM:  $p(\mathbf{z}) = \sum_{i=1}^K \pi_i \mathcal{N}(\mathbf{z} | \mu_i, \Sigma_i)$
  - 12: Sample  $\mathbf{Z} = \{z^1, z^2, \dots, z^N\} \sim p(\mathbf{z})$  and partition each  $z^j$  into  $x^j \in \mathbb{R}^{d_x}$  and  $y^j \in \mathbb{R}^{d_y}$
  - 13: Organize as sequences  $\mathbf{X} = \{x^1, \dots, x^N\}$  and  $\mathbf{Y} = \{y^1, \dots, y^N\}$
  - 14: If needed, pad  $\mathbf{X}$  and  $\mathbf{Y}$  with  $\mathcal{N}(0, 1)$  noise to dimensions  $d_{\max}$
  - 15: **return**  $(\mathbf{X}, \mathbf{Y})$
- 

## A.2. Additional details of synthetic data generation

### Full data generation pipeline

The following outlines the detailed procedure used to generate synthetic data in our experiments.

- *Mixture components number  $K$ .* We uniformly sample  $K$  from the set  $\{1, \dots, 60\}$ .
- *Weights  $\pi_i$ .* We uniformly sample each  $\pi_i$  from the interval  $[0, 1]$ , then set  $\pi_i \leftarrow \pi_i / \sum_{j=1}^K \pi_j$ .
- *Mean vector  $\mu_i$ .* Each element in the mean vector is uniformly sampled from the interval  $[-5, 5]$ .
- *Correlation matrices.* For the correlation matrices in Gaussian copulas and  $t$ -copula, we introduce a novel low-rank factorization method for covariance matrix construction that ensures meaningful inter-dimensional correlations. We construct the covariance matrix as  $\Sigma = \mathbf{W}\mathbf{W}^T + \mathbf{D}$ , where  $\mathbf{W} \in \mathbb{R}^{d \times m}$  with rank  $m \sim \text{Uniform}(\{1, 2, \dots, d\})$  has entries  $W_{ij} \sim \mathcal{N}(0, 1)$ , and  $\mathbf{D} = \text{diag}(d_1, \dots, d_d)$  with  $d_i \sim \text{Uniform}(0, 1)$  ensures positive definiteness (Proposition A.2). This formulation guarantees that the expected absolute correlation between off-diagonal entries scales as  $\mathbb{E}[|\rho_{ij}|] \approx \sqrt{m}/(m + 0.5)$  for  $i \neq j$ , producing stronger correlations when  $m$  is small. By controlling the rank parameter  $m$ , we systematically vary correlation strength from weak (high rank) to strong (low rank), ensuring the hypernetwork encounters the full spectrum of correlation patterns during training.
- *Degree of freedom  $\nu$ .* For student- $t$  copula, we randomly sample the degree of freedom  $\mu$  as  $\nu \sim \text{Uniform}([2, 30])$  to vary tail behavior. This exposes the hypernetwork to both short and heavy-tailed dependences.

### Advantages of the proposed covariance matrix generation mechanism

In this section, we highlight the advantages of our proposed covariance matrix generation mechanism by comparing it to several commonly used alternatives.

We consider three baseline approaches:

- **Full-rank matrix reparameterization**, where the covariance matrix is constructed as  $\mathbf{C} = \mathbf{A}\mathbf{A}^T$ , with  $\mathbf{A} \in \mathbb{R}^{d \times d}$  being a full-rank matrix whose entries are sampled independently from  $\mathcal{N}(0, 1)$ .
- **Cholesky decomposition**, where  $\mathbf{C} = \mathbf{L}\mathbf{L}^T$ , and  $\mathbf{L} \in \mathbb{R}^{d \times d}$  is a lower triangular matrix with positive diagonal elements. The diagonal entries of  $\mathbf{L}$  are sampled from  $\text{Uniform}(0, 1)$ , and the off-diagonal entries from  $\mathcal{N}(0, 1)$ .

**Algorithm 2** Low-Rank Factorization Method for Covariance Matrix Generation

**Require:** Target dimension  $d \in \mathbb{N}^+$ 
**Ensure:** Positive definite covariance matrix  $\Sigma \in \mathbb{R}^{d \times d}$ 

- 1: Sample rank parameter  $m \sim \text{Uniform}(\{1, 2, \dots, d\})$
- 2: Generate factor matrix  $\mathbf{W} \in \mathbb{R}^{d \times m}$  with entries  $W_{ij} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$
- 3: Generate diagonal matrix  $\mathbf{D} = \text{diag}(d_1, \dots, d_d)$  with  $d_i \stackrel{\text{i.i.d.}}{\sim} \text{Uniform}(0, 1)$
- 4: Compute covariance matrix  $\Sigma = \mathbf{W}\mathbf{W}^T + \mathbf{D}$
- 5: **Optional:** Convert to correlation matrix  $\mathbf{R}$  with entries:

$$R_{ij} = \frac{\Sigma_{ij}}{\sqrt{\Sigma_{ii}\Sigma_{jj}}}$$

- 6: **Optional:** Rescale to final covariance  $\Sigma_{\text{final}} = \text{diag}(\sigma)\mathbf{R}\text{diag}(\sigma)$ ,  $\sigma_i \sim \text{Uniform}(0.1, \sqrt{10})$
- 7: **return**  $\Sigma$  (or  $\Sigma_{\text{final}}$  if rescaled)

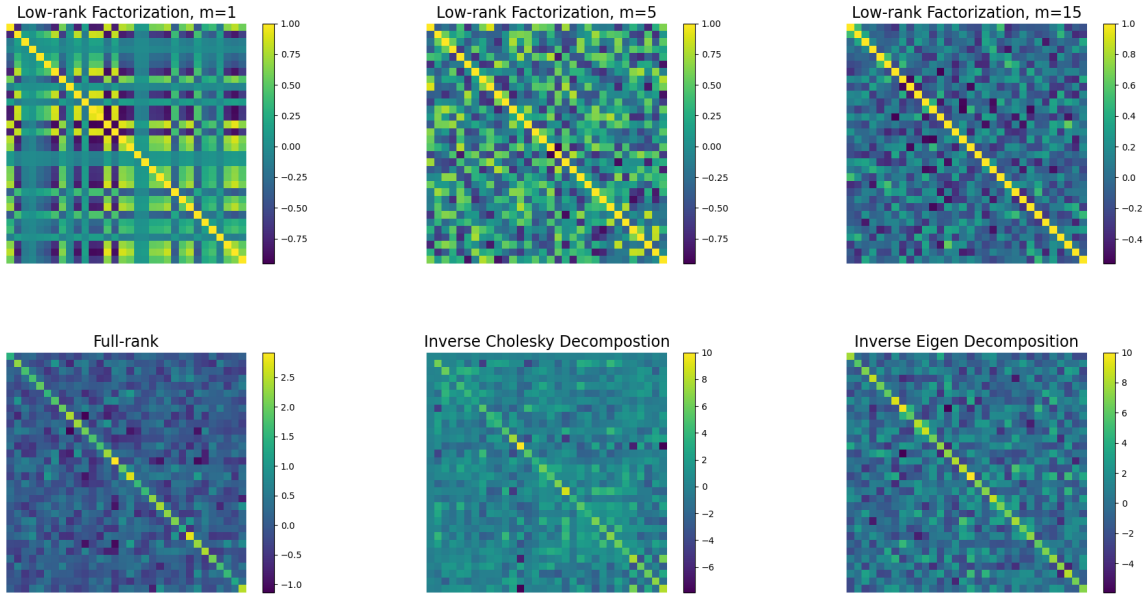


Figure 6. Visualization of correlation matrices generated by various methods. Existing approaches often yield small off-diagonal elements, whereas the low rank factor method adjusts their magnitude by tuning the rank factor  $m$ .

- **Eigenvalue decomposition**, where  $\mathbf{C} = \mathbf{Q}\mathbf{D}\mathbf{Q}^T$ , with  $\mathbf{D} \in \mathbb{R}^{d \times d}$  being a diagonal matrix with positive entries sampled from  $\text{Uniform}[0.1, 10.1)$ , and  $\mathbf{Q} \in \mathbb{R}^{d \times d}$  being an orthogonal matrix obtained via QR decomposition of a random matrix with entries from  $\mathcal{N}(0, 1)$ .

While these methods guarantee positive definiteness, they often produce covariance matrices with relatively small off-diagonal entries compared to the diagonal, resulting in limited diversity in the induced dependence structure; see the lower panel in Fig. 6.

In contrast, our method employs a *low-rank factorization* strategy (see Algorithm 2). By tuning the rank parameter  $m \leq d$ , we can flexibly control the strength of off-diagonal entries, thereby enabling the generation of covariance matrices with highly diverse dependence structures — an important design for ensuring training data diversity. This effect is illustrated in the upper panel of Fig. 6.

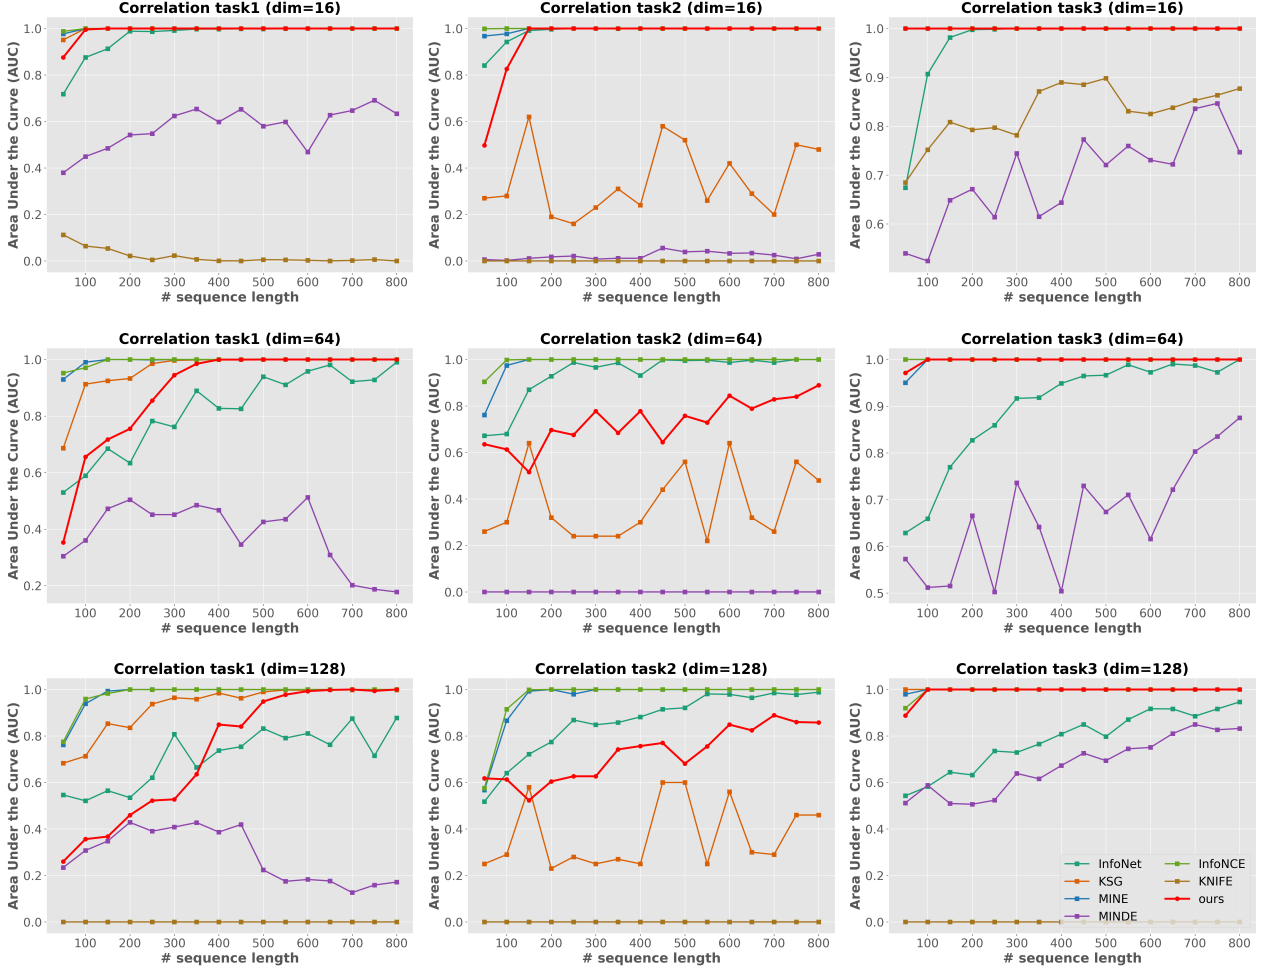


Figure 7. Independence testing across three correlation types and dimensions (16, 64, 128) across seven methods. Each curve plots the ROC-AUC as a function of sequence length  $n$ . The figure demonstrates that performance degrades with increasing dimensionality.

### A.3. Additional experimental details of independent testing experiments

*Test cases details.* Below are three different relationships between  $X$  and  $Y$  in high dimensional independence test in sec. 5.2.

- (a) **One feature (linear):**  $X, Z \sim \mathcal{N}(0, I_d)$  i.i.d. and  $Y = \frac{1}{\sqrt{2}} \left( \frac{1}{\sqrt{d}} (\mathbf{1}^\top X) \mathbf{1} + Z \right)$ , where  $\mathbf{1} := (1, \dots, 1)^\top \in \mathbb{R}^d$ .
- (b) **Two features:**  $X, Z \sim \mathcal{N}(0, I_d)$  i.i.d. and  $Y_i = \frac{1}{\sqrt{2}} \begin{cases} \frac{1}{d} (\mathbf{1}_{[d/2]} 0 \dots 0)^\top X + Z_i, & i \leq \frac{d}{2} \\ \frac{1}{d} (0 \dots 0 \mathbf{1}_{[d/2]})^\top X + Z_i, & i > \frac{d}{2}. \end{cases}$
- (c) **Independent coordinates:**  $X, Z \sim \mathcal{N}(0, I_d)$  i.i.d. and  $Y = \frac{1}{\sqrt{2}}(X + Z)$ .

*Dimensionality trends.* In Fig. 7, we examine the impact of dimensionality on estimation performance by considering three settings with increasing dimensions: 16, 64, and 128. As expected, the test power of our method decreases as dimensionality grows, particularly in small-sample regimes ( $n \leq 400$ ).

### A.4. Full results of validation on out-of-domain motion data

In this section, we provide detailed results of the experiments on motion data. We only consider points that appear throughout the entire video. Fig. 8 and Fig. 9 show the full visualization results of estimated mutual information between one selected point and other points in the videos.

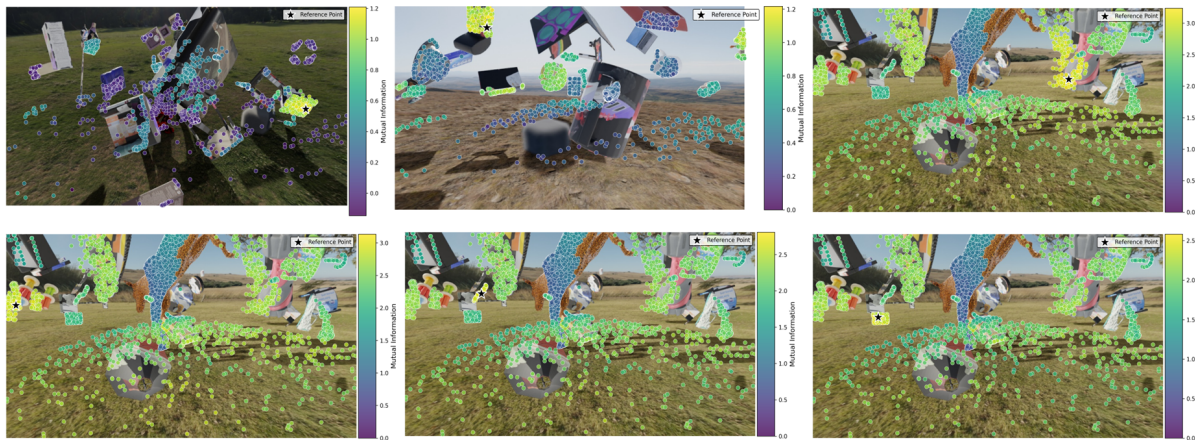


Figure 8. Full visualization results of *InfoAtlas* estimated motion data.

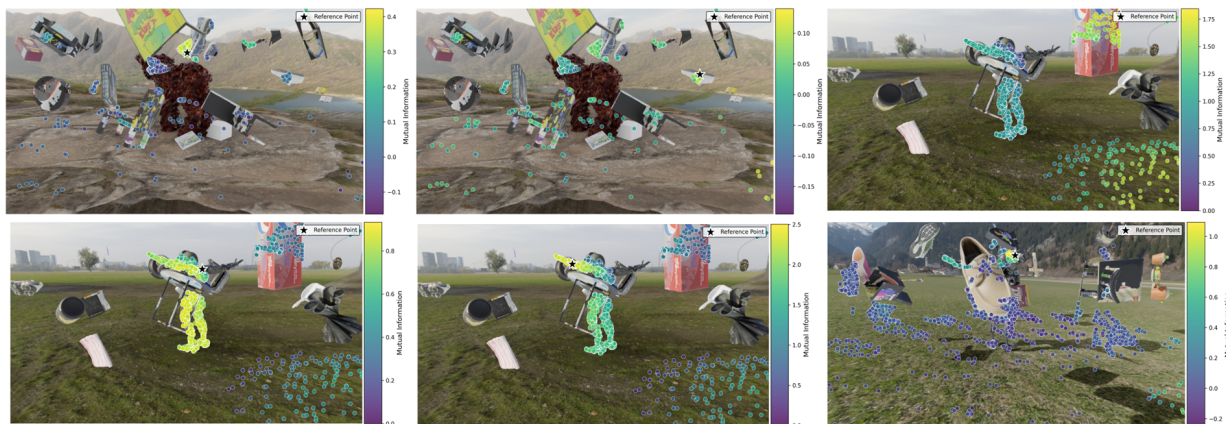


Figure 9. Full visualization results of *InfoAtlas* estimated motion data.

### A.5. Details of model training and architecture

We provide the details of model architecture and training protocol as below.

**Neural architecture details of *InfoAtlas*** For the attention module in *InfoAtlas*, we configure the dimensionality of the key and value to 1536. The Weight-Decoding MLP comprises seven layers, each of which has 8196 hidden units. Both the joint-path and marginal-path in *InfoAtlas* adopt a Perceiver IO-style architecture (Jaegle et al., 2021). Specifically, a fixed set of learnable latent queries first performs cross-attention over the input sequence of  $n$  samples, and the subsequent self-attention layers are applied only within this latent space. This design avoids quadratic self-attention over the raw input sequence and ensures that both memory usage and computational cost scale linearly with the number of input samples  $n$ .

**Optimizer setup** We pre-train *InfoAtlas* using the Adam optimizer with its default settings for 800,000 iterations, which takes approximately two weeks.

**Batch size and distributional diversity** Each training batch contains 256 independently sampled distributions with 5,000 samples per distribution (so the sequence length fed to the attention module is 5,000), providing the hypernetwork with sufficient statistical evidence to accurately estimate the optimal critic parameters for each distribution type.

**Neural network training protocol.** To ensure a fair comparison, all neural estimators MINE, InfoNCE, MINDE are trained for a maximum of 2,000 epochs with a learning rate of  $1 \times 10^{-4}$ , employing early stopping if no improvement is observed within 100 epochs. KNIFE is trained with 200 epochs.

**Computational resource.** We pre-train *InfoAtlas* on a server with 16 NVIDIA H800 GPUs, while all downstream evaluations are conducted using a single H800 GPU and 8-core Intel(R) 8480C CPU.

### A.6. Failure Cases of Sliced Mutual Information

In this section, we discuss how sliced mutual information (SMI) may fail to fully characterize certain statistical dependencies. Before all, we would like to emphasize that *slicing itself does not nullify dependence*: under mild conditions, statistical dependence in the original space remains detectable from one-dimensional projections, i.e.,  $\mathbb{I}(\mathbf{x}; \mathbf{y}) \neq 0 \Leftrightarrow \mathbb{S}\mathbb{I}(\mathbf{x}; \mathbf{y}) \neq 0$ . This property is one of the key reasons why SMI can serve as a useful and scalable measure for high-dimensional dependence assessment. However, SMI could still miss crucial information about dependencies, as demonstrated by the following two failure modes.

**Indistinguishability of different dependence structures.** The first failure mode is that two distinct dependence structures (with different underlying MIs) can induce the same SMI, making their structural differences indistinguishable. Consider two-dimensional jointly Gaussian random variables

$$\mathbf{x} = (X_1, X_2), \quad \mathbf{y} = (Y_1, Y_2),$$

where both  $\mathbf{x}$  and  $\mathbf{y}$  have identity covariance matrices, and their cross-covariance matrix is diagonal. We compare the following two models:

1. *Sparse dependence.* The dependence is concentrated on a single coordinate:

$$\text{corr}(X_1, Y_1) = \rho, \quad \text{corr}(X_2, Y_2) = 0.$$

2. *Dense dependence.* The dependence is evenly distributed across the two coordinates:

$$\text{corr}(X_1, Y_1) = \text{corr}(X_2, Y_2) = \rho'.$$

Although these two models have different dependence structures, they can induce exactly the same SMI. In particular, by setting

$$\rho' = \sqrt{1 - (2e^{-\mathbb{S}\mathbb{I}_{\text{sparse}}(\rho)} - 1)^2},$$

we obtain

$$\mathbb{S}\mathbb{I}_{\text{dense}}(\rho') = \mathbb{S}\mathbb{I}_{\text{sparse}}(\rho).$$

However, their full mutual information values are generally different:

$$\mathbb{I}_{\text{dense}}(\rho') \neq \mathbb{I}_{\text{sparse}}(\rho).$$

For instance, when  $\rho = 1$ , we can derive that for  $\rho' \approx 0.7963$ , the two models have the same SMI. However, in such case, we have

$$\mathbb{I}_{\text{sparse}} = \infty, \quad \mathbb{I}_{\text{dense}} < \infty.$$

This shows that SMI may preserve whether dependence exists while still losing information about how dependence is organized in the original high-dimensional space.

**Dilution of nonlinear dependence under projection.** The second failure mode arises when the main shared information is nonlinear and intrinsically high-dimensional, but becomes ambiguous or difficult to detect after projection. A simple example is the polar construction

$$\mathbf{x} = R(\cos \Theta, \sin \Theta), \quad \mathbf{y} = R(\cos \Phi, \sin \Phi),$$

where  $R \sim p(R)$  is a shared latent radius, while  $\Theta \sim \mathcal{U}[0, 2\pi]$  and  $\Phi \sim \mathcal{U}[0, 2\pi]$  are independent random angles. In the original two-dimensional space,  $\mathbf{x}$  and  $\mathbf{y}$  share information through their common norm:

$$\|\mathbf{x}\| = \|\mathbf{y}\| = R.$$

Thus, knowing the norm of  $\mathbf{x}$  immediately determines the norm of  $\mathbf{y}$ , making the shared information explicit and easy to extract.




Figure 10. Comparison of slice dimension  $k$  and slice number  $S$  on CLIP-generated data (original dimension 1024). Increasing  $k$  recovers more ambient dependence per slice, while increasing  $S$  reduces Monte-Carlo variance but does not remove the bias introduced by low-dimensional projection.

Now consider the one-dimensional projections of  $\mathbf{x}$  and  $\mathbf{y}$  along directions  $u$  and  $v$ . It is straightforward to verify that the projected variables take the form

$$u^\top \mathbf{x} = R \cos(\Theta'), \quad v^\top \mathbf{y} = R \cos(\Phi'),$$

where  $\Theta' \sim \mathcal{U}[0, 2\pi]$  and  $\Phi' \sim \mathcal{U}[0, 2\pi]$  are again independent random angles.

Here, in the projected space, the shared variable  $R$  is multiplied by independent angular noise terms. Consequently, recovering the common radius from a single projected scalar becomes difficult: different values of  $R$  can induce highly overlapping projected distributions due to the random cosine factors, making the shared information in  $R$  much less identifiable. For instance, if  $R \in \{1, 1.01\}$ , then  $R$  is exactly recoverable from either  $\|\mathbf{x}\|$  or  $\|\mathbf{y}\|$  in the original space, but becomes hard to distinguish from the projected samples because the small difference in radius is easily masked by the multiplicative noises. In this sense, the nonlinear dependence encoded by the equality of norms is “diluted” by slicing.

This example shows that SMI can underestimate or obscure dependencies that are clear in the original high-dimensional geometry, but become entangled with nuisance variation after projections.

### A.7. Sensitivity to slicing numbers $S$ and projection dimensions $k$

We sweep  $k$  and  $S$  on CLIP embeddings ( $d = 1024$ ); see Fig. 10. Two trends emerge. For any fixed  $S$ , the estimate grows with  $k$  as larger projections retain more joint structure, with scalar slicing ( $k = 1$ ) remaining saturated as in equation 9. For any fixed  $k$ , increasing  $S$  reduces Monte-Carlo variance but cannot remove the bias of projecting into a  $k$ -dimensional subspace (Corollary A.4). This supports our default  $k = 5$ ,  $S = 25$ , and since *InfoAtlas* amortizes per-slice estimation, the trade-off can be revisited at inference without retraining.

## B. Limitations

One limitation of *InfoAtlas* is that while *InfoAtlas* can effectively handle multivariate data with moderate dimensionalities (e.g. data up to 20 dimensions), it currently requires slicing techniques (Goldfeld & Greenwald, 2021; Goldfeld et al., 2022) to scale to higher dimensions, where extensive pre-training becomes challenging. Despite relying on slicing in high-dimensional setups, *InfoAtlas* is still the first neural method that can directly output MI without iterative optimization for data up to 20D, and it reliably quantifies statistical dependence for data up to 1024D in seconds, as demonstrated by our multiple real-world experiments. Note that for many applications, the exact value of MI is often not the interest; quantifying the orders of statistical dependence is already highly informative.

Another limitation of *InfoAtlas* is that it may fall short in cases with small sample cases (e.g.  $n < 400$ ), as seen in the independent testing experiments. In such case, the transformer fail to extract informative signals from a small population. That said, our method remains reliable for typical sample sizes encountered in reality (e.g.  $n \geq 500$ ), where our approach consistently matches MINE’s accuracy.