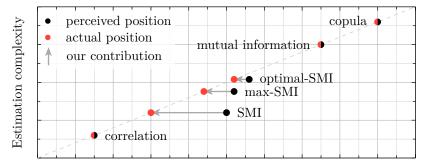
# CURSE OF SLICING: WHY SLICED MUTUAL INFORMATION IS A DECEPTIVE MEASURE OF STATISTICAL DEPENDENCE

## Anonymous authors

Paper under double-blind review

## Abstract

Sliced Mutual Information (SMI) is widely used as a scalable alternative to mutual information for measuring non-linear statistical dependence. Despite its advantages, such as faster convergence, robustness to high dimensionality, and nullification only under statistical independence, we demonstrate that SMI is highly susceptible to data manipulation and exhibits counterintuitive behavior. Through extensive benchmarking and theoretical analysis, we show that SMI saturates easily, fails to detect increases in statistical dependence (even under linear transformations designed to enhance the extraction of information), prioritizes redundancy over informative content, and in some cases, performs worse than simpler dependence measures like the correlation coefficient.



Ability to capture complex statistical dependencies

#### 1 Introduction

Mutual information (MI) is a fundamental and invariant measure of nonlinear statistical dependence between two random vectors, defined as the Kullback-Leibler divergence between the joint distribution and the product of marginals [1]:

$$\mathsf{I}(X;Y) = \mathsf{D}_{\mathsf{KL}} \big( \mathbb{P}_{X,Y} \, \big\| \, \mathbb{P}_X \otimes \mathbb{P}_Y \big).$$

Due to several outstanding properties, such as nullification only under statistical independence, invariance to invertible transformations, and ability to capture non-linear dependencies, MI is used extensively for theoretical analysis of overfitting [2], [3], hypothesis testing [4], feature selection [5], [6], [7], representation learning [8], [9], [10], [11], [12], [13], and studying the mechanisms behind generalization in deep neural networks (DNNs) [14], [15], [16], [17].

In practical scenarios,  $\mathbb{P}_{X,Y}$  and  $\mathbb{P}_X \otimes \mathbb{P}_Y$  are unknown, requiring MI to be estimated from finite samples. Despite all the aforementioned merits, this reliance on empirical estimates leads to the curse of dimensionality: the sample complexity of MI grows exponentially with the number of dimensions [18], [19]. A common strategy to mitigate this issue is to use alternative measures of statistical dependence that are more stable in high dimensions.

However, such measures usually offer only a fraction of MI capabilities. Therefore, it is crucial to maintain a balance between robustness to the curse of dimensionality and the ability to detect complex dependency structures.

To strike this balance, popular techniques often retain MI as a backbone statistical measure but employ dimensionality reduction before estimation. While some studies explore sophisticated nonlinear compression methods [17], [20], others favor more scalable linear projection approaches [21], [22], [23], [24], [25]. Among the latter group, the *Sliced Mutual Information* (SMI) [22], [23] stands out, leveraging random projections to cover all directions uniformly:

$$\mathsf{SI}(X;Y) = \frac{1}{\oint_{\mathbb{S}^{d_x-1}} \mathrm{d}\theta} \frac{1}{\oint_{\mathbb{S}^{d_y-1}} \mathrm{d}\phi} \oint_{\mathbb{S}^{d_x-1}} \oint_{\mathbb{S}^{d_y-1}} \mathsf{I}(\theta^\mathsf{T} X; \phi^\mathsf{T} Y) \, \mathrm{d}\theta \, \mathrm{d}\phi. \tag{1}$$

Uniform slicing allows SMI to maintain some crucial properties of MI (e.g., being zero if and only if X and Y are independent), while remaining completely free from additional optimization problems (e.g., from finding optimal projections, as in [24], [25]). Combined with fast convergence rates, this has established SMI as a scalable alternative to MI. Consequently, it has been widely adopted for studying DNNs [26], [27], [28], [29], [30], deriving generalization bounds [31], independence testing [32] and auditing differential privacy [33]. It was also proposed to use SMI for feature selection [22] and preventing mode collapse in generative models [23].

Despite its popularity, the research community has largely overlooked potential shortcomings of SMI. Some studies prematurely attribute their results to underlying phenomena without rigorously investigating whether they stem from artifacts introduced by random projections. Furthermore, existing works fail to comprehensively address issues related to random slicing, focusing primarily on suboptimality of random projections for information preservation [24], [25].

**Contribution.** In this article, we address this gap by systematically analyzing SMI across diverse settings, demonstrating that it frequently exhibits counterintuitive behavior and fails to accurately capture statistical dependence dynamics. Our key contributions are:

- 1. Saturation and Sensitivity Analysis. Through theoretical analysis and extensive benchmarking, we show that SMI saturates prematurely, even for low-dimensional synthetic problems, and fails to detect significant increases in statistical dependence.
- 2. **Redundancy Bias.** We refute the prevailing assumption that SMI favors linearly extractable information by constructing an explicit example where introducing such structure increases MI and even linear correlation, but decreases SMI. In fact, we show that SMI prioritizes information redundancy over information content. We argue that this bias can lead to catastrophic failures in some applications, e.g. collapses in representation learning.
- 3. Curse of Dimensionality. We revisit the dynamics of SMI for increasing dimensionality and argue that SMI is, in fact, cursed, with the curse of dimensionality manifesting itself not through sample complexity, but via asymptotic decay to zero in high-dimensional regimes due to diminishing redundancy.
- 4. Reestablishing the Trade-off. Finally, we discuss to which extent the aforementioned problems can be solved by using non-uniform/non-random slicing strategies, and how they affect the trade-off between scalability and utility of different measures of statistical dependence.

Our paper is structured as follows. In Section 2, we provide the mathematical background that is necessary for our analysis. In Section 3, we discuss previous findings which are related to the research topic of this work. Section 4 consists of our main theoretical results, with the complete proofs being provided in Section B. In Section 5, we employ synthetic benchmarks to show the disconnection between dynamics of MI and SMI. Section 6 illustrates that tasks related to SMI maximization may yield degenerate solutions, contrary to MI maximization. Finally, we discuss our results in Section 7.

# 2 Preliminaries

Elements of Information Theory. Let  $(\Omega, \mathcal{F}, \mathbb{P})$  be a probability space with sample space  $\Omega$ ,  $\sigma$ -algebra  $\mathcal{F}$ , and probability measure  $\mathbb{P}$  defined on  $\mathcal{F}$ . Consider random vectors  $X:\Omega\to\mathbb{R}^{d_x}$  and  $Y:\Omega\to\mathbb{R}^{d_y}$  with joint distribution  $\mathbb{P}_{X,Y}$  and marginals  $\mathbb{P}_X$  and  $\mathbb{P}_Y$ , respectively. Wherever it is needed, we assume the relevant Radon-Nikodym derivatives exist. For any probability measure  $\mathbb{Q}\ll\mathbb{P}$ , the Kullback-Leibler (KL) divergence is  $\mathrm{D}_{\mathsf{KL}}(\mathbb{Q}\,\|\,\mathbb{P})=\mathbb{E}_{\mathbb{Q}}\left[\log\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right]$ , which is non-negative and vanishes if and only if (iff)  $\mathbb{P}=\mathbb{Q}$ . The mutual information (MI) between X and Y quantifies the divergence between the joint distribution and the product of marginals:

$$\mathsf{I}(X;Y) = \mathbb{E}\log\frac{\operatorname{d}\mathbb{P}_{X,Y}}{\operatorname{d}\mathbb{P}_X\otimes\mathbb{P}_Y} = \mathrm{D}_{\mathsf{KL}}\big(\mathbb{P}_{X,Y} \mathbin{\big\|} \mathbb{P}_X\otimes\mathbb{P}_Y\big).$$

When  $\mathbb{P}_X$  admits a probability density function (PDF) p(X) with respect to (w.r.t.) the Lebesgue measure, the differential entropy is defined as  $h(X) = -\mathbb{E}[\log p(X)]$ , where  $\log(\cdot)$  denotes the natural logarithm. Likewise, the joint entropy h(X,Y) is defined via the joint density p(X,Y), and conditional entropy is  $h(X|Y) = -\mathbb{E}[\log p(X|Y)] = -\mathbb{E}_Y[\mathbb{E}_{X|Y}\log p(X|Y)]$ . Under the existence of PDFs, MI satisfies the identities

$$\mathsf{I}(X;Y) = \mathsf{h}(X) - \mathsf{h}(X \,|\, Y) = \mathsf{h}(Y) - \mathsf{h}(Y \,|\, X) = \mathsf{h}(X) + \mathsf{h}(Y) - \mathsf{h}(X,Y). \tag{2}$$

In this work, we denote by  $\mu_{\mathrm{M}}$  the normalized Haar (uniform) probability measure on a compact manifold M, i.e., the unique bi-invariant measure satisfying  $\mu_{\mathrm{M}}(\mathrm{M})=1.$  Hence, to sample uniformly from specific spaces we write W  $\sim \mu_{\mathrm{O}(d)}, \theta \sim \mu_{\mathbb{S}^{d-1}}, \mathrm{A} \sim \mu_{\mathrm{St}(k,d)},$  indicating draws from the Haar measures on orthogonal group  $\mathrm{O}(d)=\{\mathrm{Q}\in\mathbb{R}^{d\times d}:\mathrm{Q}^{\mathsf{T}}\mathrm{Q}=\mathrm{Q}\mathrm{Q}^{\mathsf{T}}=\mathrm{I}\},$  the unit sphere  $\mathbb{S}^{d-1}=\{X\in\mathbb{R}^d:\|X\|_2=1\},$  and the Stiefel manifold  $\mathrm{St}(k,d)=\{\mathrm{Q}\in\mathbb{R}^{d\times k}:\mathrm{Q}^{\mathsf{T}}\mathrm{Q}=\mathrm{I}\},$  respectively.

**Sliced Mutual Information.** To mitigate the curse of dimensionality, one may average MI over all k-dimensional projections. The k-sliced mutual information (k-SMI) [23] between X and Y is defined as

$$\mathsf{SI}_k(X;Y) = \int_{\mathrm{St}(k,d_x)} \int_{\mathrm{St}(k,d_y)} \mathsf{I}\big(\Theta^\mathsf{T} X;\Phi^\mathsf{T} Y\big) \,\mathrm{d}\mu_{\mathrm{St}(k,d_x)}(\Theta) \,\mathrm{d}\mu_{\mathrm{St}(k,d_y)}(\Phi),$$

Setting k = 1 recovers the standard sliced mutual information (1).

#### 3 Background

Merits of SMI are straightforward and have been investigated thoroughly in [22], [23]. We remind the reader of the two most important of them:

- 1. Scalability enabled by low-dimensional projections.
- 2. Nullification Property (i.e.,  $\mathsf{SI}_k(X;Y) = 0$  iff X and Y are independent), which stems from the projections being random and independent.

In contrast, demerits of SMI are not very obvious and not well-covered in the literature. In this section, we recapitulate and analyze previous works which address the shortcomings of SMI. To facilitate the analysis, we divide them into three main categories.

Suboptimality of random slicing. In [24] and [25], it is argued that a uniform slicing strategy can produce suboptimal projections, impairing SMI's ability to capture dependencies in the presence of noisy or non-informative components. To address this issue, [24] proposed max-sliced MI (mSMI), which selects non-random projectors that maximize the MI between projected representations. This approach is also claimed to improve interpretability and convergence rates.

However, deterministic slicing may overlook dependencies captured by non-optimal components. To mitigate this, [25] extends the max-sliced approach by optimizing SMI over

probability distributions of projectors, with regularization to maintain slice diversity. While the authors emphasize that optimization should occur over *joint* distributions, their motivation primarily addresses the issue of non-optimal marginal distributions of  $\theta$  and  $\phi$ —specifically, the presence of non-informative components in X and Y. We contend that this represents only a partial understanding of the problem, as many SMI artifacts arise from other factors. Needless to say that optimization over probability distributions is also a heavy burden, which does not align with the slicing philosophy.

**Data Processing Inequality violation.** A fundamental property of MI is that it cannot be increased by deterministic processing or, more generally, by Markov kernels. Furthermore, MI is preserved under invertible transformations. This is formalized by the *data processing inequality* (DPI).

**Theorem 3.1.** (Theorem 3.7 in [1]) For a Markov chain  $X \to Y \to Z$ ,  $I(X;Y) \ge I(X;Z)$ . Additionally, if Z = f(Y) where f is measurably invertible, then equality holds.

In contrast to MI, SMI violates the DPI (see Section 3.2 in [22] for an example). While the intuition behind DPI is clear (raw data already contains full information, and processing can only destroy it), the implications of DPI violation are less straightforward.

Existing works suggest that SMI's violation of DPI can reflect a preference for linearly extractable features, framing this as a useful property that aligns with the informal understanding of "practically available" (i.e., easily accessible) information [22], [26], [30]. However, this interpretation can be misleading if the factors behind SMI increases are misidentified. Our analysis reveals that this is indeed the case, as SMI exhibits more inherent biases than previously recognized.

Asymptotics in high-dimensional regime. Convergence analysis suggests that the sample complexity of SMI estimation is far less sensitive to data dimensionality compared to that of MI. In fact, it has been argued that the estimation error may even decrease with dimensionality in some cases (see Remark 4 in [23]). However, an analysis of SMI itself reveals that this behavior may result from the fact that SMI can decrease as dimensionality grows. Specifically, Theorem 3 in [23] provides an asymptotic expression (as  $d \to \infty$ ) for SMI in the case of jointly normal X and Y, which decays hyperbolically with d under some circumstances.

To date, no explanation for this phenomenon has been provided in the literature. We therefore elaborate on this finding by deriving non-asymptotic expressions, along with experimental results for non-Gaussian data, which reveal further nuances behind the decay.

#### 4 Theoretical analysis

We start our analysis with considering a simple example, which (a) admits closed-form expression for SMI and (b) is capable of illustrating severe problems of the quantity in question.

**Lemma 4.1.** Consider the following pair of jointly Gaussian d-dimensional random vectors:

$$(X,Y) \sim \mathcal{N}\bigg(0, \begin{pmatrix} \mathbf{I} & \rho \mathbf{I} \\ \rho \mathbf{I} & \mathbf{I} \end{pmatrix}\bigg), \quad \rho \in (-1;1).$$

In this setup, MI and SMI can be calculated analytically:

$$\mathsf{I}(X;Y) = -\frac{d}{2}\log \bigl(1-\rho^2\bigr), \qquad \mathsf{SI}(X;Y) = \frac{\rho^2}{2d} \ _3F_2\biggl(1,1,\frac{3}{2};\frac{d}{2}+1,2;\rho^2\biggr),$$

where  $_3F_2$  is the generalized hypergeometric function. Additionally, the following limits hold:

$$\begin{split} &\lim_{d\to\infty}\mathsf{I}(X;Y) = +\infty \qquad \lim_{d\to\infty}\mathsf{SI}(X;Y) = 0 \\ &\lim_{\rho^2\to 1}\mathsf{I}(X;Y) = +\infty \qquad \lim_{\rho^2\to 1}\mathsf{SI}(X;Y) = \psi(d-1) - \psi\bigg(\frac{d-1}{2}\bigg) - \log 2 \leq \frac{1}{d-1}, \end{split}$$

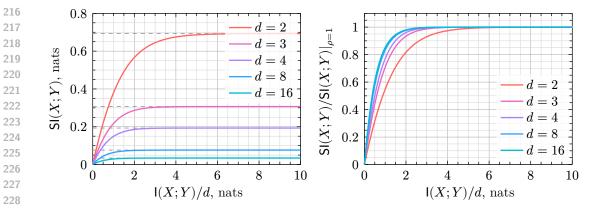


Figure 2: Saturation of SI(X;Y) as function of I(X;Y)/d for the example from Lemma 4.1, non-normalized (left) and normalized (right) versions. Note that the problem becomes more prominent in higher dimensions, both because of lower plateau and faster saturation.

with  $\psi$  being the digamma function.

Note that while MI correctly captures the growing statistical dependence as  $d \to \infty$  (since additional components contribute shared information), SMI drops to zero, exposing a fundamental problem. We interpret this behavior as a distinct manifestation of the **curse** of dimensionality: as d grows, SMI uniformly decays to zero and becomes ineffective for statistical analysis.

The second pair of limits reveals another critical flaw of SMI. When  $\rho^2 \to 1$ , the X-Y relationship becomes deterministic — a property MI reflects successfully. In stark contrast, SMI remains bounded by a dimension-dependent factor that decays hyperbolically. Furthermore, plotting SMI against MI shows this bound is reached prematurely, demonstrating SMI's **rapid saturation** with increasing dependence (Figure 2). In this saturated regime, SMI becomes effectively insensitive to further growth in shared information. Moreover, this renders estimates of SMI for different dimensionalities fundamentally incomparable, as they are theoretically bounded by factors depending on d.

These phenomena can not be explained by suboptimality of individual projections. In fact, each individual projection is optimal, as  $I(\theta^T X; Y)$  does not depend on  $\theta$  in this particular example. The proof of Lemma 4.1 suggests that the problem arises from the majority of pairs of projectors being suboptimal, yielding near-independent  $\theta^T X$  and  $\phi^T Y$  in the most outcomes, even for d=2. Although similar analysis for k-SMI is extremely challenging, we argue that the problems in question prevail even when employing k-rank projectors.

**Proposition 4.2.** Under the setup of Lemma 4.1, k-SMI has the following representation

$$\mathsf{SI}_k(X;Y) = -\frac{1}{2} \int_{[0,1]^k} \sum_{i=1}^k \log \bigl(1 - \rho^2 \lambda_i\bigr) \, p(\pmb{\lambda}) \, \mathrm{d} \pmb{\lambda},$$

where 
$$p(\pmb{\lambda}) \propto \prod_{i < j} \lvert \lambda_j - \lambda_i \rvert \underbrace{\prod_{i=1}^k \left(1 - \lambda_i\right)^{(d-2k-1)/2}}_{(\star)}.$$

**Remark.** 4.3. As d grows,  $(\star)$  asymptotically concentrates  $\lambda_i$  near zero, driving  $\mathsf{Sl}_k$  to zero.

We argue that the limitations we uncovered can be attributed to a strong bias of SMI toward **information redundancy**. That is, SMI favors repetition of information across different axes, and suffers from the curse of dimensionality if X and Y have high entropy. The following proposition and remark present a simple example to clarify this bias.

**Proposition 4.4.** Let X and Y be  $d_x, d_y$ -dimensional random vectors respectively, with  $d_x, d_y < k$ . Let  $A \in \mathbb{R}^{m_x \times d_x}$  and  $B \in \mathbb{R}^{m_y \times d_y}$  be full column rank matrices. Then  $\mathsf{SI}_k(AX; BY) = \mathsf{I}(X; Y)$ .

Corollary 4.5. Consider the following pair of Gaussian d-dimensional random vectors:

$$(X,Y) \sim \mathcal{N}\bigg(0, \begin{pmatrix} \mathbf{J} & \rho \mathbf{J} \\ \rho \mathbf{J} & \mathbf{J} \end{pmatrix}\bigg), \quad \rho \in (-1;1),$$

where  $\mathbf{J} = \mathbf{1} \cdot \mathbf{1}^\mathsf{T}$  with  $\mathbf{1}^\mathsf{T} = (1,...,1)$ . Then  $\mathsf{SI}_k(X;Y) = \mathsf{I}(X;Y) = -\frac{1}{2}\log(1-\rho^2)$ .

**Remark.** 4.6. Applying  $\mathbf{1} \cdot e_1^{\mathsf{T}}$  to the random vectors from Lemma 4.1 individually yields the example from Corollary 4.5. Therefore, this linear transform increases SMI despite decreasing MI.

#### 4.1 Extension to optimal slicing

Although our work primarily focuses on conventional (average) sliced mutual information (SMI), as it is the most widely used variant, we also provide some intuition regarding the limitations of its "optimal" counterparts: max-sliced MI (mSMI) [24] and optimal-sliced MI (oSMI) [25]. Since mSMI is a special case of oSMI without regularization constraints, we restrict our discussion to mSMI, though our reasoning extends to oSMI as well. The k-mSMI is defined as:

$$\overline{\mathsf{SI}}_k(X;Y) = \sup_{\substack{\Theta \in \, \mathrm{St}(d_x,k) \\ \Phi \in \, \mathrm{St}(d_y,k)}} \mathsf{I}\big(\Theta^\mathsf{T} X;\Phi^\mathsf{T} Y\big) \tag{3}$$

To highlight the shortcomings of linear compression, we revisit a Gaussian example. The following proposition demonstrates that even in this simple setting, mSMI captures only a subset of dependencies and can exhibit opposite trends to MI. This occurs, for instance, when dependencies become more evenly distributed across components, which again returns us to the **redundancy bias**.

**Proposition 4.7.** (Proposition 2 in [24]) Let  $(X,Y) \sim \mathcal{N}(\mu,\Sigma)$ , with marginal covariances  $\Sigma_X$ ,  $\Sigma_Y$  and cross-covariance  $\Sigma_{XY}$ . Suppose the matrix  $\Sigma_X^{-2} \Sigma_{XY} \Sigma_Y^{-\frac{1}{2}}$  exists, and let  $\{\rho_i\}_{i=1}^d$  denote its singular values in descending order, where  $d = \min(d_x, d_y)$ . Then

$$\mathsf{I}(X;Y) = -\frac{1}{2} \sum_{i=1}^d \log \bigl(1 - \rho_i^2\bigr), \qquad \overline{\mathsf{SI}}_k(X;Y) = -\frac{1}{2} \sum_{i=1}^k \log \bigl(1 - \rho_i^2\bigr).$$

# 5 Synthetic Experiments

To complement the theoretical analysis from the previous section and address complex, non-Gaussian cases, we conduct an extensive benchmarking of SMI using synthetic tests from [34], based on the works of [17], [35], [36]. This benchmark suite is used to evaluate MI estimators. However, here we do not assess whether SMI estimates converge to ground-truth MI values. SMI is a distinct measure of statistical dependance, and should not be viewed as an approximation of MI. Instead, our analysis focuses on the relationship between the two measures: since MI captures the true degree of statistical dependence, opposing trends in MI and SMI reveal problems with the latter quantity.

For the experiments, we use correlated normal, correlated uniform, smoothed uniform and log-gamma-exponential distributions, for which the ground-truth value of MI is available. To increase the dimensionality, we use independent components with equally distributed per-component MI. For each distribution, we vary both the data dimensionality (d) and the projection dimensionality (k < d). In Section C, we also utilize MI-preserving mappings to transform low-dimensional Gaussian vectors into high-dimensional synthetic images, as described in [17]; the examples of such images are displayed in Figure 3.



Figure 3: Examples of synthetic images from additional experiments in Section C. Note that images are high-dimensional, but admit latent structure, which is similar to real datasets.

To estimate MI between projections, we use the KSG estimator [35] with the number of neighbors fixed at 1. For each configuration, we conduct 10 independent runs with different random seeds to compute means and standard deviations. Our experiments use  $10^4$  samples for (X,Y) and 128 samples for  $(\Theta,\Phi)$ .

To experimentally verify the saturation, we plot SMI against MI normalized by dimensionality d in Figure 4. The plots clearly show that SMI reaches a plateau relatively early for all the featured distributions. The results for the normal distribution also align well with those from Lemma 4.1. We further confirm the saturation of k-SMI for  $k \in \{2,3\}$  experimentally in Section C. Finally, we plot the saturated values against d on a log-log scale, demonstrating that the 1/d trend from Lemma 4.1 also holds for non-Gaussian distributions.

## 6 SMI FOR INFOMAX-LIKE TASKS

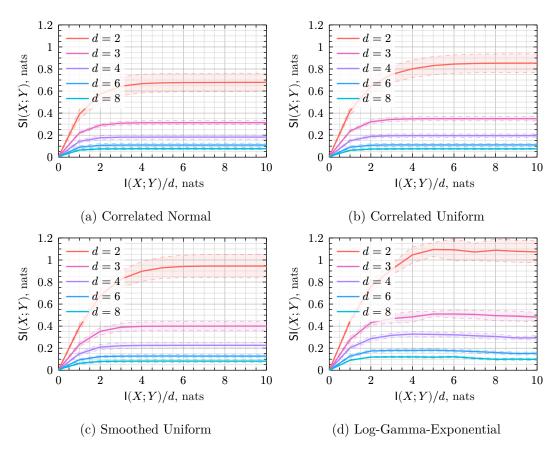


Figure 4: Results of synthetic experiments with different distributions for SMI. We report mean values and standard deviations computed across 10 runs, with  $10^4$  samples used for MI estimation and 128 for averaging across projections.

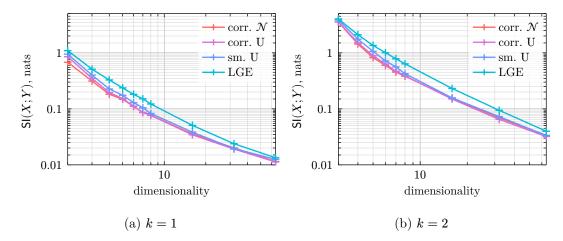


Figure 5: Decaying trends of k-SMI for correlated normal (corr.  $\mathcal{N}$ ), correlated uniform (corr. U), smoothed uniform (sm. U) and log-gamma-exponential (LGE). We plot saturated values of k-SMI against data dimensionality d. Log scale is used to illustrate the 1/d trend predicted in Lemma 4.1.

Since mutual information is interpretable and captures non-linear dependencies, it is widely used as a training objective. Many applications involve maximizing MI (InfoMax) for feature selection [5], [6], [7] and self-supervised representation learning [8], [9], [10], [11], [12], [13]. However, due to the curse of dimensionality, alternative objectives have been proposed, with some works using sliced mutual information maximization for feature extraction [22] and disentanglement in InfoGAN [23].

In this section, we argue that SMI is not a suitable alternative to MI for InfoMax tasks. Since SMI exhibits a strong preference for redundancy, SMI maximization may lead to collapsed (high-redundancy) solutions. To demonstrate this, we examine its performance in learning compressed representations through mutual information maximization (*Deep InfoMax*) [8]. This approach is known to be equivalent to many popular contrastive self-supervised methods [13].

In Deep InfoMax, an encoder network f is trained to maximize a lower bound on I(X; f(X)), where X represents input data and f(X) its compressed representation. This method is theoretically sound, as maximizing MI ensures the most informative embeddings under the latent space dimensionality constraint. For our study, we replace MI with SMI in this framework. This substitution is straightforward since both MI and SMI admit Donsker-Varadhan variational lower bounds [37]:

$$\begin{split} \mathsf{I}(X;Y) &= \sup_{T:\Omega \to \mathbb{R}} \left[ \mathbb{E}_{\mathbb{P}_{X,Y}} \, T(X,Y) - \log \left( \mathbb{E}_{\mathbb{P}_X \otimes \mathbb{P}_Y} \, e^{T(X,Y)} \right) \right], \\ \mathsf{SI}_k(X;Y) &= \sup_{T:\Omega \to \mathbb{R}} \mathbb{E}_{\Theta,\Phi} \left[ \mathbb{E}_{\mathbb{P}_{X,Y}} \, T(\Theta^\mathsf{T} X, \Phi^\mathsf{T} Y, \Theta, \Phi) - \log \left( \mathbb{E}_{\mathbb{P}_X \otimes \mathbb{P}_Y} \, e^{T(\Theta^\mathsf{T} X, \Phi^\mathsf{T} Y, \Theta, \Phi)} \right) \right], \end{split}$$

where T is a critic function, which is also approximated in practice by a neural network. For detailed derivations of these bounds, we refer the reader to [38] (MI) and [22], [23] (SMI).

We strictly follow the experimental protocol from [13]. In particular, we use MNIST hand-written digits dataset [39], employ InfoNCE loss [40] to approximate (4), use convolutional network for f and fully-connected network for T. Latent space dimensionality is fixed at d=2 for visualization purposes. Small Gaussian noise is added to the outlet of the encoder to combat representation collapse [13]. More details are provided in Section D. We focus on this simple setup because our objective is to show that SMI produces degenerate results even in elementary tasks, making more complex configurations unnecessary for this demonstration.

Results are presented in Figure 6. As expected, maximization of SMI immediately leads to collapsed representations, while conventional InfoMax yields embeddings with low redundancy (components are close to  $\mathcal{N}(0, I)$ ). This behavior is consistent across different runs.

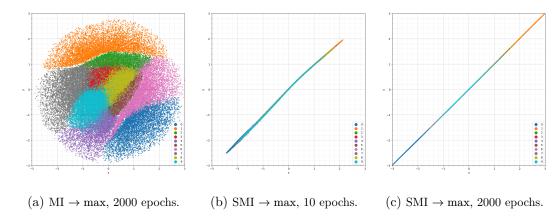


Figure 6: Visualizations of embeddings from the representation learning experiments, with points colored by class. Note that mutual information maximization (left) produces clustered low-redundancy representations, while SMI maximization results in immediate collapse.

# 7 Discussion

**Results.** Sliced mutual information (SMI) has been proposed as a scalable alternative to Shannon's mutual information. While SMI enables efficient computation in high-dimensional settings and satisfies the nullification property, our findings reveal critical deficiencies that undermine its reliability for feature extraction and related tasks.

We demonstrate that SMI saturates rapidly, failing to capture variations in statistical dependence. This makes it difficult to distinguish between intrinsic SMI fluctuations and genuine changes in dependence structure. Furthermore, we invalidate the common hypothesis that SMI favors linear features through a counterexample where even correlation coefficients reflect dependence more faithfully than SMI, which exhibits inverted behavior.

In high dimensions, SMI decays with increasing dimensionality, contrary to MI's monotonic behavior. This is established analytically for Gaussian cases and validated empirically across diverse synthetic experiments. Consequently, SMI variations may reflect redundancy or high-dimensional artifacts without a principled way to disentangle these factors.

**Impact.** Thanks to fast convergence rates and the absence of additional optimization problems, SMI has been widely applied across various fields of statistics and machine learning. Given our findings, it is therefore crucial to recognize how the inherent biases of SMI affect practical applications.

The works [22] and [23] propose using SMI in a Deep InfoMax setting. However, we demonstrate that maximizing SMI can lead to collapsed solutions due to redundancy bias. Meanwhile, [26], [27], [28], [30] study deep neural networks by measuring SMI between intermediate layers. Yet, as our analysis reveals, changes in SMI do not always reflect true shifts in statistical dependence; they may instead result from differences in layer dimensionality, redundancy in intermediate representations, low sensitivity in saturated regimes, or other factors. Finally, [33] suggests using SMI for independence testing in differential privacy tasks. We contend that this approach poses critical issues, as SMI estimates can become statistically indistinguishable from zero in high-dimensional or low-redundancy settings.

Limitations. While we support our claims with both theoretical analysis and experimental evidence, we were able to derive analytical expressions for the Gaussian case only. Nevertheless, we demonstrate that our findings are more than sufficient to expose fundamental limitations of SMI, and to support all the claims we made.

- Ethics statement. This work is not subject to any ethical concerns.
- Reproducibility statement. To ensure reproducibility of our results, we provide complete proofs in Section B and implementation details in Section D. We also provide our code for the experiments in the supplementary material.
  - **LLM usage.** Large Language Models (LLMs) were used only to assist with rephrasing sentences and improving the clarity of the text.

#### References

- [1] Y. Polyanskiy and Y. Wu, Information Theory: From Coding to Learning. Cambridge University Press, 2024. [Online]. Available: https://books.google.ru/books?id=CySoOAEACAAJ
- [2] A. Asadi, E. Abbe, and S. Verdu, "Chaining Mutual Information and Tightening Generalization Bounds," in *Advances in Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds., Curran Associates, Inc., 2018, p. . [Online]. Available: https://proceedings.neurips.cc/paper\_files/paper/2018/file/8d7628dd7a710c8638dbd22d4421ee46-Paper.pdf
- [3] J. Negrea, M. Haghifam, G. K. Dziugaite, A. Khisti, and D. M. Roy, "Information-Theoretic Generalization Bounds for SGLD via Data-Dependent Estimates," in *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Eds., Curran Associates, Inc., 2019, p. . [Online]. Available: https://proceedings.neurips.cc/paper\_files/paper/2019/file/05ae 14d7ae387b93370d142d82220f1b-Paper.pdf
- [4] B. Duong and T. Nguyen, "Conditional Independence Testing via Latent Representation Learning," in 2022 IEEE International Conference on Data Mining (ICDM), Los Alamitos, CA, USA: IEEE Computer Society, Dec. 2022, pp. 121–130. doi: 10.1109/ICDM54844.2022.00022.
- [5] S. Yang and J. Gu, "Feature selection based on mutual information and redundancy-synergy coefficient," *J. Zhejiang Univ. Sci.*, vol. 5, no. 11, pp. 1382–1391, Nov. 2004.
- [6] N. Kwak and C.-H. Choi, "Input feature selection by mutual information based on Parzen window," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 12, pp. 1667–1671, 2002, doi: 10.1109/TPAMI.2002.1114861.
- [7] M. A. Sulaiman and J. Labadin, "Feature selection based on mutual information," in 2015 9th International Conference on IT in Asia (CITA), 2015, pp. 1–6. doi: 10.1109/CITA.2015.7349827.
- [8] R. D. Hjelm *et al.*, "Learning deep representations by mutual information estimation and maximization," in *International Conference on Learning Representations*, 2019. [Online]. Available: https://openreview.net/forum?id=Bklr3j0cKX
- [9] P. Bachman, R. D. Hjelm, and W. Buchwalter, "Learning Representations by Maximizing Mutual Information Across Views," in *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Eds., Curran Associates, Inc., 2019, p. . [Online]. Available: https://proceedings.neurips.cc/paper\_files/paper/2019/file/ddf354219aac374f1d40b7 e760ee5bb7-Paper.pdf
- [10] P. Veličković, W. Fedus, W. L. Hamilton, P. Liò, Y. Bengio, and R. D. Hjelm, "Deep Graph Infomax," in *International Conference on Learning Representations*, 2019. [Online]. Available: https://openreview.net/forum?id=rklz9iAcKQ
- [11] M. Tschannen, J. Djolonga, P. K. Rubenstein, S. Gelly, and M. Lucic, "On Mutual Information Maximization for Representation Learning," in *International Conference on Learning Representations*, 2020. [Online]. Available: https://openreview.net/forum?id=rkxoh24FPH

- 540 [12] X. Yu, "Leveraging Superfluous Information in Contrastive Representation Learning." [Online]. Available: https://arxiv.org/abs/2408.10292
  - [13] I. Butakov, A. Semenenko, A. Tolmachev, A. Gladkov, M. Munkhoeva, and A. Frolov, "Efficient Distribution Matching of Representations via Noise-Injected Deep InfoMax," in *The Thirteenth International Conference on Learning Representations*, 2025. [Online]. Available: https://openreview.net/forum?id=mAmCdASmJ5
  - [14] N. Tishby and N. Zaslavsky, "Deep learning and the information bottleneck principle," 2015 IEEE Information Theory Workshop (ITW), pp. 1–5, 2015.
  - [15] R. Shwartz-Ziv and N. Tishby, "Opening the Black Box of Deep Neural Networks via Information." 2017.
  - [16] Z. Goldfeld et al., "Estimating Information Flow in Deep Neural Networks," in Proceedings of the 36th International Conference on Machine Learning, K. Chaudhuri and R. Salakhutdinov, Eds., in Proceedings of Machine Learning Research, vol. 97. PMLR, 2019, pp. 2299–2308. [Online]. Available: https://proceedings.mlr.press/v97/goldfeld19a.html
  - [17] I. Butakov, A. Tolmachev, S. Malanchuk, A. Neopryatnaya, A. Frolov, and K. Andreev, "Information Bottleneck Analysis of Deep Neural Networks via Lossy Compression," in *The Twelfth International Conference on Learning Representations*, 2024. [Online]. Available: https://openreview.net/forum?id=huGECz8dPp
  - [18] Z. Goldfeld, K. Greenewald, J. Niles-Weed, and Y. Polyanskiy, "Convergence of Smoothed Empirical Measures With Applications to Entropy Estimation," *IEEE Transactions on Information Theory*, vol. 66, no. 7, pp. 4368–4391, 2020, doi: 10.1109/ TIT.2020.2975480.
  - [19] D. McAllester and K. Stratos, "Formal Limitations on the Measurement of Mutual Information," in *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, S. Chiappa and R. Calandra, Eds., in Proceedings of Machine Learning Research, vol. 108. PMLR, 2020, pp. 875–884. [Online]. Available: https://proceedings.mlr.press/v108/mcallester20a.html
  - [20] G. Gowri, X. Lun, A. M. Klein, and P. Yin, "Approximating mutual information of high-dimensional variables using learned representations," in *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. [Online]. Available: https://openreview.net/forum?id=HN05DQxyLl
  - [21] K. H. Greenewald, B. Kingsbury, and Y. Yu, "High-Dimensional Smoothed Entropy Estimation via Dimensionality Reduction," in *IEEE International Symposium on Information Theory, ISIT 2023, Taipei, Taiwan, June 25-30, 2023*, IEEE, 2023, pp. 2613–2618. doi: 10.1109/ISIT54713.2023.10206641.
  - [22] Z. Goldfeld and K. Greenewald, "Sliced Mutual Information: A Scalable Measure of Statistical Dependence," in *Advances in Neural Information Processing Systems*, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds., 2021. [Online]. Available: https://openreview.net/forum?id=27qon5Ut4PSl
  - [23] Z. Goldfeld, K. Greenewald, T. Nuradha, and G. Reeves, "\$k\$-Sliced Mutual Information: A Quantitative Study of Scalability with Dimension," in *Advances in Neural Information Processing Systems*, A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, Eds., 2022. [Online]. Available: <a href="https://openreview.net/forum?id=L-ceBdl2DPb">https://openreview.net/forum?id=L-ceBdl2DPb</a>
  - [24] D. Tsur, Z. Goldfeld, and K. Greenewald, "Max-Sliced Mutual Information," in *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. [Online]. Available: https://openreview.net/forum?id=ce9B2x3zQa
  - [25] A. Fayad and M. Ibrahim, "On Slicing Optimality for Mutual Information," in Thirty-seventh Conference on Neural Information Processing Systems, 2023. [Online]. Available: https://openreview.net/forum?id=JMuKfZx2xU

- [26] S. Wongso, R. Ghosh, and M. Motani, "Understanding Deep Neural Networks Using Sliced Mutual Information," in 2022 IEEE International Symposium on Information Theory (ISIT), 2022, pp. 133–138. doi: 10.1109/ISIT50566.2022.9834357.
  - [27] S. Wongso, R. Ghosh, and M. Motani, "Using Sliced Mutual Information to Study Memorization and Generalization in Deep Neural Networks," in *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, F. Ruiz, J. Dy, and J.-W. van de Meent, Eds., in Proceedings of Machine Learning Research, vol. 206. PMLR, 2023, pp. 11608–11629. [Online]. Available: https://proceedings.mlr.press/v 206/wongso23a.html
  - [28] S. Wongso, R. Ghosh, and M. Motani, "Pointwise Sliced Mutual Information for Neural Network Explainability," in 2023 IEEE International Symposium on Information Theory (ISIT), 2023, pp. 1776–1781. doi: 10.1109/ISIT54713.2023.10207010.
  - [29] J. Dentan, D. Buscaldi, A. Shabou, and S. Vanier, "Predicting and analyzing memorization within fine-tuned Large Language Models." [Online]. Available: https://arxiv.org/abs/2409.18858
  - [30] S. Wongso, R. Ghosh, and M. Motani, "Sliced Information Plane for Analysis of Deep Neural Networks," Jan. 2025, doi: 10.36227/techrxiv.173833980.08812687/v1.
  - [31] K. Nadjahi, K. Greenewald, R. B. Gabrielsson, and J. Solomon, "Slicing Mutual Information Generalization Bounds for Neural Networks," in *ICML 2023 Workshop Neural Compression: From Information Theory to Applications*, 2023. [Online]. Available: https://openreview.net/forum?id=cbLcwK3SZi
  - [32] Z. Hu, S. Kang, Q. Zeng, K. Huang, and Y. Yang, "InfoNet: Neural Estimation of Mutual Information without Test-Time Optimization," in Forty-first International Conference on Machine Learning, 2024. [Online]. Available: https://openreview.net/forum?id=40hCy8n5XH
  - [33] T. Nuradha and Z. Goldfeld, "Pufferfish Privacy: An Information-Theoretic Study," *IEEE Trans. Inf. Theor.*, vol. 69, no. 11, pp. 7336–7356, Nov. 2023, doi: 10.1109/TIT.2023.3296288.
  - [34] I. Butakov *et al.*, "MUTINFO." [Online]. Available: https://github.com/VanessB/mutinfo
  - [35] A. Kraskov, H. Stögbauer, and P. Grassberger, "Estimating mutual information," *Phys. Rev. E*, vol. 69, no. 6, p. 66138, Jun. 2004, doi: 10.1103/PhysRevE.69.066138.
  - [36] F. Czyż Pawełand Grabowski, J. Vogt, N. Beerenwinkel, and A. Marx, "Beyond Normal: On the Evaluation of Mutual Information Estimators," in Advances in Neural Information Processing Systems, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, Eds., Curran Associates, Inc., 2023, pp. 16957–16990. [Online]. Available: https://proceedings.neurips.cc/paper\_files/paper/2023/file/36b 80eae70ff629d667f210e13497edf-Paper-Conference.pdf
  - [37] M. D. Donsker and S. R. Varadhan, "Asymptotic evaluation of certain markov process expectations for large time. IV," *Communications on Pure and Applied Mathematics*, vol. 36, no. 2, pp. 183–212, Mar. 1983, doi: 10.1002/cpa.3160360204.
  - [38] M. I. Belghazi et al., "Mutual Information Neural Estimation," in *Proceedings of the 35th International Conference on Machine Learning*, J. Dy and A. Krause, Eds., in Proceedings of Machine Learning Research, vol. 80. PMLR, 2018, pp. 531–540. [Online]. Available: https://proceedings.mlr.press/v80/belghazi18a.html
  - [39] L. Deng, "The mnist database of handwritten digit images for machine learning research," *IEEE Signal Processing Magazine*, vol. 29, no. 6, pp. 141–142, 2012.
  - [40] A. van den Oord, Y. Li, and O. Vinyals, "Representation Learning with Contrastive Predictive Coding." [Online]. Available: https://arxiv.org/abs/1807.03748

- [41] A. Edelman and B. D. Sutton, "The beta-Jacobi matrix model, the CS decomposition, and generalized singular value problems," *Foundations of Computational Mathematics*, vol. 8, no. 2, pp. 259–285, 2008.
- [42] A. McBride, "Special functions, by George E. Andrews, Richard Askey and Ranjan Roy. Pp. 664.£ 60. 1999. ISBN 0 521 62321 9 (Cambridge University Press.)," The Mathematical Gazette, vol. 83, no. 497, pp. 355–357, 1999.
- [43] N. Elezovic, C. Giordano, and J. Pecaric, "The best bounds in Gautschi's inequality," Math. Inequal. Appl, vol. 3, no. 2, pp. 239–252, 2000.
- [44] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization." 2017.

A SUPPLEMENTARY THEORY

Lemma A.1. (Example 2.4 in [1])  $h(\mathcal{N}(\mu, \Sigma)) = \frac{1}{2} \log((2\pi e)^d \det \Sigma)$ .

Corollary A.2. For  $(X,Y) \sim \mathcal{N}(\mu,\Sigma)$  with non-singular  $\Sigma$ 

$$\begin{split} \mathsf{I}(X;Y) &= \frac{1}{2} \log \det \Sigma_X + \frac{1}{2} \log \det \Sigma_Y - \frac{1}{2} \log \det \Sigma \\ &= -\frac{1}{2} \sum_{i=1}^d \log (1 - \rho_i^2), \end{split}$$

where  $\Sigma_X$ ,  $\Sigma_Y$  are marginal covariances,  $\Sigma_{XY}$  is cross-covariance,  $d = \min(d_x, d_y)$ , and  $\{\rho_i\}_{i=1}^d$  are singular values of  $\Sigma_X^{-\frac{1}{2}}\Sigma_{XY}\Sigma_Y^{-\frac{1}{2}}$ .

Proof of Corollary A.2. Combining Lemma A.1 and (2) yields the first result. Now note that

$$\mathsf{I}(X;Y) = \mathsf{I}\bigg(\Sigma_X^{-\frac{1}{2}}X;\Sigma_Y^{-\frac{1}{2}}Y\bigg) = \mathsf{I}\bigg(\mathbf{U}^\mathsf{T}\Sigma_X^{-\frac{1}{2}}X;\mathbf{V}\Sigma_Y^{-\frac{1}{2}}Y\bigg),$$

where  $U \operatorname{diag}(\rho_i) V^{\mathsf{T}}$  is the SVD of  $\Sigma_X^{-\frac{1}{2}} \Sigma_{XY} \Sigma_Y^{-\frac{1}{2}}$ . However,

$$\bigg(\mathbf{U}^\mathsf{T} \boldsymbol{\Sigma}_X^{-\frac{1}{2}} \boldsymbol{X}, \mathbf{V} \boldsymbol{\Sigma}_Y^{-\frac{1}{2}} \boldsymbol{Y} \bigg) \sim \mathcal{N} \bigg( \boldsymbol{\mu}', \begin{pmatrix} \mathbf{I} & \mathrm{diag}(\rho_i) \\ \mathrm{diag}(\rho_i) & \mathbf{I} \end{pmatrix} \bigg),$$

from which we arrive at the second expression.

**Lemma A.3.** Let  $A \in \mathbb{R}^{n \times m}$  be full column rank matrix and  $\Theta \sim \mu_{St(n,k)}$  Then  $\Theta^{\mathsf{T}}A$  is full-rank with probability one.

Proof of Lemma A.3. Performing QR decomposition of A yields  $\Theta^{\mathsf{T}} A = \Theta^{\mathsf{T}} QR \stackrel{d}{=} \Theta^{\mathsf{T}} {I_m \choose 0} R$ . Since A is full-rank, R is invertible and rank  $\Theta^{\mathsf{T}} A = \operatorname{rank} \Theta^{\mathsf{T}} {I_m \choose 0}$ . Therefore,

$$\mathbb{P}\big\{\Theta^\mathsf{T}\mathbf{A} \text{ is full-rank}\big\} = 1 - \mathbb{P}\bigg\{\Theta^\mathsf{T}\binom{\mathbf{I}_m}{0} \text{ is not full-rank}\bigg\} = 1.$$

**Lemma A.4.** (Theorem 1.5 in [41]) Let W  $\sim \mu_{O(d)}$  and partition

$$\mathbf{W} = \begin{pmatrix} \mathbf{W}_{11} & \mathbf{W}_{12} \\ \mathbf{W}_{21} & \mathbf{W}_{22} \end{pmatrix}$$

with  $W_{11}$  of size k by k. Then the eigenvalues  $\{\lambda_i\}_{i=1}^k$  of  $W_{11}W_{11}^\mathsf{T}$  follow the Jacobi ensemble

$$p(\boldsymbol{\lambda}) \propto \prod_{i < i} |\lambda_i - \lambda_j|^{\beta} \prod_{i=1}^k \lambda_i^{\frac{\beta}{2}(a+1)-1} (1 - \lambda_i)^{\frac{\beta}{2}(b+1)-1}$$

with parameters a = 0, b = d - 2k, and  $\beta = 1$  (over  $\mathbb{R}$ ).

Proof of Lemma A.4. Let  $A_1 \in \mathbb{R}^{k \times d}$  and  $A_2 \in \mathbb{R}^{(d-k) \times d}$  be independent matrices with i.i.d. entries from  $\mathcal{N}(0,1)$ . By stacking  $A_1$  atop  $A_2$  and then performing a block QR decomposition on the resulting Gaussian matrix, the orthogonal invariance of the Gaussian law implies that the two Q-blocks are independent of the upper-triangular factor R, with  $Q_1$  and  $Q_2$  uniformly distributed on O(k) and St(k,d-k), respectively. Finally, computing the SVD of the block rows together with R yields the generalized singular value decomposition (GSVD) of the pair  $(A_1,A_2)$ :

$$\begin{pmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{pmatrix} = \begin{pmatrix} \mathbf{Q}_1 \\ \mathbf{Q}_2 \end{pmatrix} \mathbf{R} = \begin{pmatrix} \mathbf{U}_1 \\ \mathbf{U}_2 \end{pmatrix} \begin{pmatrix} \tilde{\mathbf{C}} \\ 0 \\ -\tilde{\mathbf{S}} \\ 0 \end{pmatrix} \tilde{\mathbf{V}}^\mathsf{T} \mathbf{R},$$

where  $U_1 \in O(k)$ ,  $U_2 \in O(d-k)$ ,  $\tilde{V} \in O(k)$ , and  $C = \text{diag}(c_i)$ ,  $S = \text{diag}(s_i)$  with  $c_i \geq 0$ ,  $s_i \geq 0$ , and  $c_i^2 + s_i^2 = 1$  for all i. The diagonal entries of  $\tilde{C}$  are known as the generalized singular values of the pair  $(A_1, A_2)$ .

For a matrix  $P = \text{diag}(p_1,...,p_k)$  with i.i.d.  $p_i$  sampled uniformly from  $\{-1,1\}$ , we have  $Q_1P \stackrel{\text{d}}{=} W_{11}$ . Let  $W_{11} = \text{UCV}^\mathsf{T}$  be the SVD of  $W_{11}$ , then one has

$$U_1 \begin{pmatrix} \tilde{C} \\ 0 \end{pmatrix} \tilde{V}^\mathsf{T} P \stackrel{d}{=} U C V^\mathsf{T}.$$

Since  $U_1, \tilde{V}$ , and U, V are uniformly distributed and independent of  $\tilde{C}, C$ , we have  $\tilde{C} \stackrel{d}{=} C$  by the invariance of the Haar measure under orthogonal transformations. On the other hand, the generalized singular values  $\tilde{C}$  of a pair  $(A_1, A_2)$  follow the law of the Jacobi ensemble with parameters a = 0, b = d - 2k, and  $\beta = 1$  (Proposition 1.2 in [41]). Therefore, the squared singular values of  $W_{11}$  follow the Jacobi ensemble with the same parameters.

**Corollary A.5.** The squared inner product  $|\theta^{\mathsf{T}}\phi|^2$  between two independent random vectors  $\theta, \phi \sim \mu_{\mathbb{S}^{d-1}}$  follows  $\mathrm{Beta}(\frac{1}{2}, \frac{d-1}{2})$ . Moreover, the shifted inner product  $(1 + \theta^{\mathsf{T}}\phi)/2$  is symmetrically distributed as  $\mathrm{Beta}(\frac{d-1}{2}, \frac{d-1}{2})$ .

Proof of Corollary A.5. Setting Jacobi parameters k=1, a=0, b=d-2 and  $\beta=1$ , the density is proportional to  $x^{-1/2}(1-x)^{(d-3)/2}$  on [0,1], which matches the Beta $\left(\frac{1}{2},\frac{d-1}{2}\right)$  distribution.

Next, observe that  $\theta^{\mathsf{T}} \phi$  has a density proportional to  $(1-t)^{\frac{d-3}{2}}$  for  $t \in [-1,1]$ . Under the change of variables  $\eta \sim \text{Beta}(\frac{d-1}{2},\frac{d-1}{2})$ .

## B Complete proofs

*Proof of Lemma 4.1.* One can acquire  $I(X;Y) = -\frac{d}{2}\log(1-\rho^2)$  from a general expression for MI of two jointly Gaussian random vectors (see Corollary A.2).

Recall that  $(\theta^{\mathsf{T}}X, \phi^{\mathsf{T}}Y)$  is also Gaussian with cross-covariance  $\rho \theta^{\mathsf{T}} \phi$ . Therefore, by Corollary A.2 we have

$$\mathrm{SI}(X;Y) = \mathrm{I}\big(\theta^\mathsf{T} X; \phi^\mathsf{T} Y \,\big|\, \theta, \varphi\big) = -\frac{1}{2}\,\mathbb{E}\big[\log\big(1-\rho^2 \,\,|\theta^\mathsf{T} \phi|^2\big)\big].$$

From Corollary A.5, we note that  $|\theta^{\mathsf{T}}\phi|^2 \sim \mathrm{Beta}(\frac{1}{2}, \frac{d-1}{2})$ , so

$$\begin{split} \mathsf{SI}(X;Y) &= -\frac{1}{2\mathsf{B}\left(\frac{1}{2},\frac{d-1}{2}\right)} \int_0^1 \log(1-\rho^2 x) (1-x)^{\frac{d-3}{2}} x^{-\frac{1}{2}} \, \mathrm{d}x \\ &= \frac{\rho^2}{2} \frac{\Gamma\left(\frac{d}{2}\right)}{\Gamma\left(\frac{1}{2}\right) \Gamma\left(\frac{d-1}{2}\right)} \int_0^1 x^{\frac{1}{2}} (1-x)^{\frac{d-3}{2}} \, {}_2F_1\left(1,1;2;\rho^2 x\right) \, \mathrm{d}x, \end{split}$$

where the last equality follows from the identity  $\log(1-z) = -z \,_2 F_1(1,1;2;z)$  with hypergeometric function  $_2F_1$ . Appling Euler's integral transform ([42], Eq. (2.2.3)) gives

$$\begin{split} \mathsf{SI}(X;Y) &= \frac{\rho^2}{2d} \frac{\Gamma\left(\frac{d}{2}+1\right)}{\Gamma\left(\frac{3}{2}\right) \Gamma\left(\frac{d-1}{2}\right)} \int_0^1 x^{\frac{3}{2}-1} (1-x)^{\left(\frac{d}{2}+1\right)-\frac{3}{2}-1} \, _2F_1\left(1,1;2;\rho^2 x\right) \mathrm{d}x \\ &= \frac{\rho^2}{2d} \, _3F_2\left(1,1,\frac{3}{2};\frac{d}{2}+1,2;\rho^2\right). \end{split}$$

Here  $_3F_2$  denotes the generalized hypergeometric function.

Finally, we calculate the limit of SI(X;Y) as  $\rho^2 \to 1$  using properties of beta-distribution. Denoting  $\eta = (1 + \theta^T \phi)/2 \sim \text{Beta}(\frac{d-1}{2}, \frac{d-1}{2})$  (see Corollary A.5), we get

$$\mathsf{SI}(X;Y) = -\log 2 - \mathbb{E}\log(1-\eta) = -\log 2 - \mathbb{E}\log \eta = \psi(d-1) - \psi\left(\frac{d-1}{2}\right) - \log 2,$$

where  $\psi$  is the digamma function. Using the bounds on digamma function [43]

$$\log\Bigl(x+\frac{1}{2}\Bigr)-\frac{1}{x}\leq \psi(x)\leq \log\bigl(x+e^{\psi(1)}\bigr)-\frac{1}{x},$$

we derive an upper bound on this expression:

$$\psi(d-1) - \psi\left(\frac{d-1}{2}\right) - \log 2 \leq \frac{1}{d-1} + \log\left(1 + \frac{e^{\psi(1)}-1}{d}\right)$$

To simplify the bound, one can note that  $e^{\psi(1)}-1<0,$  as  $\psi(1)<0.$ 

Proof of Proposition 4.2. Let  $Q_X, Q_Y \sim \mu_{St(k,d)}$ . Then  $\left[Q_X^\mathsf{T} X, Q_Y^\mathsf{T} Y\right] \sim \mathcal{N}(0, \Sigma)$ , where  $\Sigma$  is a  $2k \times 2k$  covariance matrix with the following block structure

$$\Sigma = \begin{pmatrix} \mathbf{I}_k & \rho \, \mathbf{Q}_{\mathbf{X}}^\mathsf{T} \mathbf{Q}_{\mathbf{Y}} \\ \rho \, \mathbf{Q}_{\mathbf{Y}}^\mathsf{T} \mathbf{Q}_{\mathbf{X}} & \mathbf{I}_k \end{pmatrix}.$$

Using the formula for the determinant of a block matrix  $\Sigma$  yields

$$\mathsf{SI}_k(X;Y) = -\frac{1}{2}\,\mathbb{E}[\log\det(\Sigma)] = -\frac{1}{2}\,\mathbb{E}\Big[\log\det\Big(\mathbf{I} - \rho^2\big(\mathbf{Q}_{\mathbf{X}}^\mathsf{T}\mathbf{Q}_{\mathbf{Y}}\big)\big(\mathbf{Q}_{\mathbf{X}}^\mathsf{T}\mathbf{Q}_{\mathbf{Y}}\big)^\mathsf{T}\Big)\Big].$$

By the invariance of the Haar measure under left and right multiplication,  $Q_X^TQ_Y \stackrel{d}{=} W_{11}$ , where  $W_{11}$  is a k by k left upper block of the matrix  $W \sim \mu_{O(d)}$ . According to Lemma A.4, the eigenvalues of  $W_{11}W_{11}^T$  follow Jacobi ensemble with parameters a=0, b=d-2k and  $\beta=1$ :

$$p(\lambda) \propto \prod_{i < j} |\lambda_j - \lambda_i| \prod_{i=1}^k (1 - \lambda_i)^{\frac{d-2k-1}{2}}.$$

Thus, we get a general expression for k-SMI

$$\mathsf{SI}_k(X;Y) = -\frac{1}{2} \int_{[0,1]^k} \sum_{i=1}^k \log \bigl(1 - \rho^2 \lambda_i\bigr) p(\lambda) \,\mathrm{d}\lambda.$$

Proof of Proposition 4.4. Using Lemma A.3 and  $d_x, d_y < k$ , we get that  $\Theta^\mathsf{T} A$  and  $\Phi^\mathsf{T} B$  are injective with probability one for independent  $\Theta$ ,  $\Phi$  distributed uniformly on  $\mathrm{St}(d_x, k)$  and  $\mathrm{St}(d_y, k)$ . Therefore, according to Theorem 3.1,  $[\mathsf{I}(\Theta^\mathsf{T} AX; \Phi^\mathsf{T} BY) \mid \Theta, \Phi] = \mathsf{I}(X; Y)$  almost sure. As a result,  $\mathsf{SI}_k(AX; BY) = \mathsf{I}(\Theta^\mathsf{T} AX; \Phi^\mathsf{T} BY \mid \Theta, \Phi) = \mathsf{I}(X; Y)$ .

Proof of Proposition 4.7. Direct corollary of Corollary A.2.

# C Additional experiments

In this section, we conduct supplementary experiments to evaluate SMI under a broader range of setups. We begin by assessing k-SMI on the same set of benchmarks from Section 5. The results for k=1,2,3 are presented in Figure 4, Figure 7, and Figure 8, respectively. Notably, saturation remains consistent even for k=d-1 (i.e., when only one component is discarded).

Next, we examine a setup involving randomized distribution parameters, following the methodology of [34]. Among other adjustments, this includes randomizing per-component mutual information (e.g., assigning interactions unevenly in this experiment). In some cases (e.g., the log-gamma-exponential distribution), this increases linear redundancy, as component pairs with higher mutual information also exhibit higher variance in this particular scenario. Our results are displayed in Figure 9.

Due to numerical constraints, we do not track  $\mathsf{I}(X;Y)/d$  in this particular setup, instead plotting the results against the total mutual information. While this makes saturation slightly less evident, the general trend of SMI decreasing with d remains observable. We also highlight the log-gamma-exponential distribution (Figure 9d), where SMI is less prone to saturation under parameter randomization due to the reasons mentioned earlier.

Finally, using the MI-preserving smooth injective mappings from [34], we reproduce the synthetic datasets used in [17]. These datasets consist of high-dimensional images (see Figure 3) with known ground-truth mutual information. The results presented in Figure 10 again prove our findings.

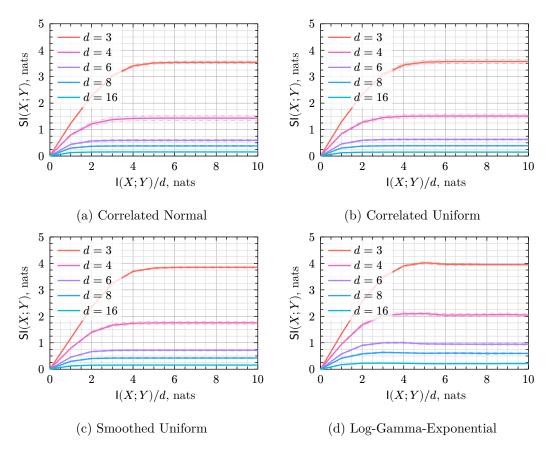


Figure 7: Results of synthetic experiments with different distributions for 2-SMI. We report mean values and standard deviations computed across 10 runs, with  $10^4$  samples used for MI estimation and 128 for averaging across projections.

## D IMPLEMENTATION DETAILS

#### D.1 SYNTHETIC EXPERIMENTS

For the experiments from Section 5, we use implementation of Kraskov-Stoegbauer-Grassberger (KSG) [35] mutual information estimator and random slicing from [34]. The number of neighbors is set to  $k_{\rm NN}=1$  for the KSG estimator. For each configuration, we conduct 10 independent runs with different random seeds to compute means and standard deviations. Our experiments use  $10^4$  samples for (X,Y) and 128 samples for  $(\Theta,\Phi)$ .

For the experiments from Section 5, we use independent components with equally distributed per-component MI. For the supplementary experiments from Figure 9, parameters of each distribution (e.g., covariance matrices) are randomized via the algorithm implemented in [34]. This includes randomization of per-component MI (which is done using a uniform distribution over a (d-1)-dimensional simplex).

For the experiments, we used AMD EPYC 7543 CPU, one core per distribution. Each experiment (fixed k, varying d) took no longer then 3 days to compute.

#### D.2 Representation learning experiments

For experiments on MNIST dataset, we use a simple ConvNet with three convolutional and two fully connected layers. A three-layer fully-connected perceptron serves as a critic network for the InfoNCE loss. We provide the details in Table 1. We use additive Gaussian

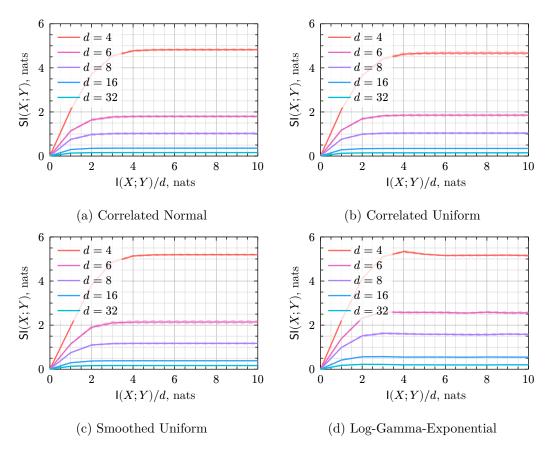


Figure 8: Results of synthetic experiments with different distributions for 3-SMI. We report mean values and standard deviations computed across 10 runs, with  $10^4$  samples used for MI estimation and 128 for averaging across projections.

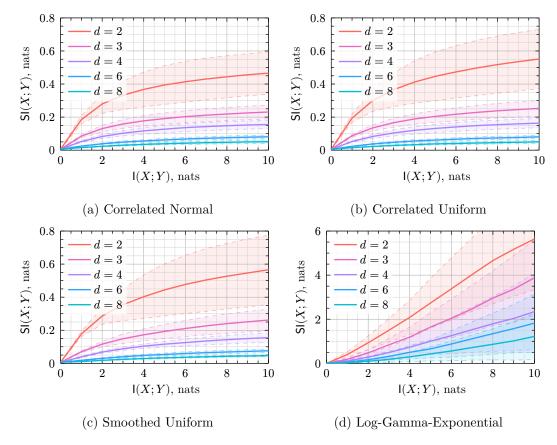


Figure 9: Results of synthetic experiments with different distributions. We report mean values and standard deviations computed across 10 runs, with  $10^4$  samples used for MI estimation and 128 for averaging across projections.

noise with  $\sigma = 0.2$  as an input augmentation. Training hyperparameters are as follows: batch size = 512, 2000 epochs, Adam optimizer [44] with learning rate  $10^{-3}$ .

For the experiments, we used Nvidia A100 GPUs. Each experiment took no longer then 1 day to compute. Table 1: The NN architectures used to conduct the tests on MNIST images in Section 6.

NN	Architecture
ConvNet, $24 \times 24$ images	$ \begin{array}{llll} \times 1: & Conv2d(1,32,ks{=}3),MaxPool2d(2),BatchNorm2d,LeakyReLU(0.01) \\ \times 1: & Conv2d(32,64,ks{=}3),MaxPool2d(2),BatchNorm2d,LeakyReLU(0.01) \\ \times 1: & Conv2d(64,128,ks{=}3),MaxPool2d(2),BatchNorm2d,\\ & LeakyReLU(0.01) \\ \times 1: & Dense(128,128),LeakyReLU(0.01),Dense(128,dim) \end{array} $
Critic NN, pairs of vectors	$\times$ 1: Dense(dim + dim, 256), LeakyReLU(0.01) $\times$ 1: Dense(256, 256), LeakyReLU(0.01), Dense(256, 1)

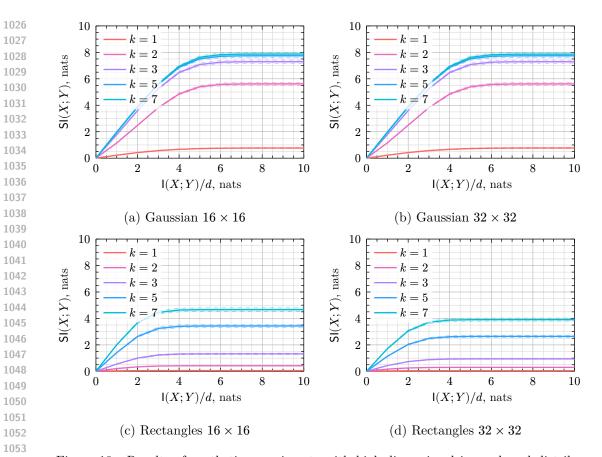


Figure 10: Results of synthetic experiments with high-dimensional image-based distributions for k-SMI. We report mean values and standard deviations computed across 10 runs, with  $10^4$  samples used for MI estimation and 128 for averaging across projections.