Rethinking Disentanglement under Dependent Factors of Variation

Anonymous authors
Paper under double-blind review

Abstract

Representation learning enables the discovery and extraction of underlying factors of variation from data. A representation is typically considered disentangled when it isolates these factors in a way that is interpretable to humans. Existing definitions and metrics for disentanglement often assume that the factors of variation are statistically independent. However, this assumption rarely holds in real-world settings, limiting the applicability of such definitions and metrics in real-world applications. In this work, we propose a novel definition of disentanglement grounded in information theory, which remains valid even when the factors are dependent. We show that this definition is equivalent to requiring the representation to consist of minimal and sufficient variables. Based on this formulation, we introduce a method to quantify the degree of disentanglement that remains effective in the presence of statistical dependencies among factors. Through a series of experiments, we demonstrate that our method reliably measures disentanglement in both independent and dependent settings, where existing approaches fail under the latter.

1 Introduction

In representation learning, data are commonly assumed to be characterized by a set of factors of variation (henceforth, factors) and nuisances. The distinction between them lies in task relevance: factors directly influence the outcome of interest, whereas nuisances do not. For instance, in a dataset of fruit images where the task is to describe the fruits, relevant factors may include fruit type, color, or size, while irrelevant nuisances might be the background color or the shadows cast. A large body of work has emphasized the importance of learning disentangled representations, where distinct dimensions of the representation correspond to different factors of variation (Schmidhuber, 1992; Bengio et al., 2013; Ridgeway, 2016; Lake et al., 2017; Achille & Soatto, 2018a). Such representations are valuable across a range of applications, including: (i) interpretability of learned features and model predictions (Hsu et al., 2017; Worrall et al., 2017; Gilpin et al., 2018; Zhang & Zhu, 2018; Liu et al., 2020; Zhu et al., 2021; Nauta et al., 2023; Almudévar et al., 2024), (ii) fairer decision-making by mitigating or removing biases in sensitive attributes such as gender or race (Calders & Žliobaitė, 2013; Barocas & Selbst, 2016; Kumar et al., 2017; Locatello et al., 2019a; Sarhan et al., 2020; Chen et al., 2023), and (iii) empowering generative models to synthesize data with controllable attributes (Yan et al., 2016; Paige et al., 2017; Huang et al., 2018; Dupont, 2018; Kazemi et al., 2019; Antoran & Miguel, 2019; Zhou et al., 2020; Shen et al., 2022; Almudévar et al., 2024).

The definition in the previous paragraph provides an intuitive, though imprecise, understanding of disentangled representations. It raises two fundamental questions: (i) what does it precisely mean to separate the factors, and (ii) how can we evaluate whether the factors are indeed separated? In Section 2, we review several works that attempt to address these two questions. Most of these approaches assume that the factors of variation are independent—both from one another and from the nuisances. Consequently, the corresponding definitions and evaluation metrics are grounded in this assumption of independence. Many of the benchmark datasets used to validate these methods also exhibit (approximately) independent factors (LeCun et al., 2004; Liu et al., 2015; Reed et al., 2015; Matthey et al., 2017; Burgess & Kim, 2018; Gondal et al., 2019), thereby making the independence-based definitions and metrics largely applicable in those contexts.

However, in real-world scenarios, factors of variation are rarely statistically independent—either from each other or from nuisances. Returning to the earlier example, the shape of a fruit is not independent of attributes such as color, size, or shadow: encountering a yellow banana-shaped fruit is common, whereas a yellow strawberry-shaped fruit is highly unlikely. When factors are independent, separating them in a representation is relatively straightforward. In contrast, when factors are dependent, it becomes less clear what it means for a representation to "separate" them. As illustrated in Figure 1, even with a perfectly disentangled encoder, individual components of the representation can still carry information about multiple factors. In this work, we focus on defining and quantifying disentanglement in the setting where the factors of variation are not independent. Our key contributions can be summarized as follows:

- We introduce a set of four information-theoretic properties that characterize desirable behavior for disentangled representations in the presence of dependencies and nuisances.
- We prove that these properties hold if and only if the representations are both *minimal* and *sufficient*. Moreover, we argue that evaluating minimality and sufficiency provides a more meaningful assessment of disentanglement than analyzing the properties in isolation.
- We propose a method to quantify the degree of *minimality* and *sufficiency*, enabling the evaluation of disentanglement in arbitrary representations. [1].

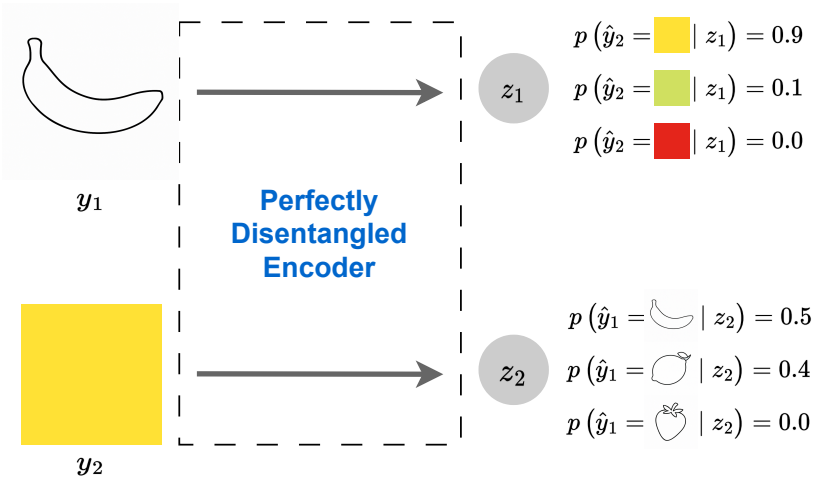

Figure 1: If the factors are dependent, then—even with a perfectly disentangled encoder—each sub-representation will still contain information about other factors. For example, if z_1 encodes the shape of a banana, the fruit's color is most likely yellow, possibly green if unripe, and almost never red. Likewise, if z_2 encodes the color yellow, the shape is most likely that of a banana or a lemon, but very unlikely to be that of a strawberry.

2 Related Work

Definition of Disentanglement Despite being a topic of great interest, there is no widely agreed-upon definition of disentangled representations. Intuitively, a disentangled representation separates the different factors of variation underlying the data (Desjardins et al., 2012; Bengio et al., 2013; Cohen & Welling, 2014; Kulkarni et al., 2015; Ridgeway, 2016). Early evaluations of disentanglement were largely based on visual inspection (Kingma & Welling, 2013). However, more precise definitions are necessary to understand the limitations and practical relevance of disentangled representations in downstream applications. One of the earliest and most widely cited definitions is given in Bengio et al. (2013), and has been adopted by several subsequent works (Higgins et al., 2017; Kim & Mnih, 2018; Kim et al., 2019; Suter et al., 2019). This

[1] `https://github.com/link/to/project`.

definition states that a disentangled representation is one in which a change in a single dimension of the representation corresponds to a change in a single factor of variation, while remaining invariant to changes in other factors. An alternative view argues that disentanglement should be defined such that a change in a single factor of variation leads to a change in only one dimension of the representation (Kumar et al., 2017; Locatello et al., 2019b). More recent efforts define disentanglement based on a set of desirable properties. Notably, Ridgeway & Mozer (2018) and Eastwood & Williams (2018) independently proposed three key properties that a representation should satisfy to be considered disentangled. While the terminology differs slightly between the two works, the underlying ideas align: (i) *Modularity* (referred to as *disentanglement* in Eastwood & Williams (2018)): each representation variable captures at most one factor of variation; (ii) *Compactness* (*completeness* in Eastwood & Williams (2018)): each factor of variation is captured by a single representation variable; and (iii) *Explicitness* (*informativeness* in Eastwood & Williams (2018)): the representation retains all task-relevant information about the factors present in the input. All of the above definitions present two key limitations: (i) they do not account for invariance to nuisance factors that may be present in the input; and (ii) they assume that the factors of variation are independent of one another and of the nuisances—which generally does not hold in real-world data. We propose a set of properties that a representation must satisfy to be considered disentangled, explicitly addressing both of these aspects.

Metrics for Disentanglement Due to the lack of consensus on the definition of disentanglement, there is also no agreement on how it should be measured. Below, we provide an overview of common disentanglement metrics, organized by their underlying principles, following Carbonneau et al. (2022).

- **Intervention-based Metrics.** These methods construct subsets of data samples that share a common factor of variation while differing in others. Representations for these subsets are then computed and compared to produce a disentanglement score. Representative methods in this category include β-VAE (Higgins et al., 2017), FactorVAE (Kim & Mnih, 2018), and R-FactorVAE (Kim et al., 2019). These metrics primarily aim to assess modularity but do not account for compactness or explicitness (Sepliarskaia et al., 2019).

- **Information-based Metrics.** These approaches estimate mutual information between the factors and individual variables in the representation. Widely used examples include the Mutual Information Gap (MIG) (Chen et al., 2018) and Robust MIG (Do & Tran, 2019). These metrics are mainly designed to capture compactness.

- **Predictor-based Metrics.** These methods train regressors or classifiers to predict the factors from the learned representations. The resulting predictors are then analyzed to determine the contribution of each representation variable to each factor. Separated Attribute Predictability (SAP) (Kumar et al., 2017), focuses exclusively on compactness. Subsequently, more comprehensive frameworks have been proposed to evaluate multiple properties of disentanglement. Notably, Eastwood & Williams (2018) introduced DCI (Disentanglement, Completeness, and Informativeness), while Ridgeway & Mozer (2018) proposed an equivalent set: Modularity, Compactness, and Explicitness.

Disentanglement under Dependent Factors of Variation As previously mentioned, disentanglement under dependent factors is a relatively underexplored area. However, several works have addressed related ideas, among which the following stand out: Suter et al. (2019) offers a causal perspective on representation learning that connects disentanglement with robustness to domain shift; Choi et al. (2020) proposes a method to disentangle variations that are shared across all classes from those that are specific to each class; Montero et al. (2020; 2022) examine the effects of excluding certain factor combinations during training and evaluate models on their ability to reconstruct unseen combinations at test time; Träuble et al. (2021) introduces synthetic dependencies between pairs of factors to study disentanglement in more realistic settings; Roth et al. (2022) presents the Hausdorff Factorized Support (HFS) criterion, which accommodates arbitrary (including correlated) distributions over factor supports; Ahuja et al. (2023) explores how interventional data can support causal representation learning by revealing geometric patterns in latent variables; and von Kügelgen et al. (2024) proposes a nonparametric method for discovering latent causal variables and their relations from high-dimensional observations, without relying on strong parametric assumptions.

Information Bottleneck in Representation Learning The Information Bottleneck (IB) framework (Tishby et al., 2000) formalizes a trade-off between fidelity and compression in learned representations. A representation is deemed *sufficient* if it retains all information relevant to the target task, and *minimal* if it excludes all information irrelevant to it. Formally, given an input x, target variable y, and representation z, sufficiency corresponds to $I(z;y) = I(x;y)$, while minimality requires $I(z;x) = I(z;y)$. The IB principle has been leveraged in a range of works to design regularizers for representation learning (Alemi et al., 2016; 2018; Achille & Soatto, 2018b;a; Almudévar et al., 2025), and has also been applied to promote disentanglement (Vera et al., 2018; Yamada et al., 2020; Jeon et al., 2021; Gao et al., 2021). To the best of our knowledge, however, this is the first work to formally define disentanglement in terms of minimal sufficient representations. In our formulation, each factor of variation is treated as a distinct target variable, and the overall representation is decomposed into sub-representations. We define a representation as disentangled if each sub-representation is both minimal and sufficient with respect to its corresponding factor.

3 Desirable Properties of a Disentangled Representation

In our representation learning setup, we consider raw data $\boldsymbol{x}$, which is fully explained by a set of task-relevant factors $\boldsymbol{y} = \{y_i\}_{i=1}^n$ and a set of nuisances $\boldsymbol{n}$. The factors $\boldsymbol{y}$ correspond to the underlying sources or causes that influence the observed data and are relevant to the task at hand. We assume that these factors are conditionally independent[2] given $\boldsymbol{x}$, i.e., $p(y_i, y_j \mid \boldsymbol{x}) = p(y_i \mid \boldsymbol{x})p(y_j \mid \boldsymbol{x})$ for $i \neq j$. The nuisances $\boldsymbol{n}$ refer to sources of variation in $\boldsymbol{x}$ that are irrelevant to the task.

We define a representation $\boldsymbol{z}$ as a random variable governed by the conditional distribution $p(\boldsymbol{z} \mid \boldsymbol{x})$. This setup induces the Markov chains $\boldsymbol{y} \leftrightarrow \boldsymbol{x} \leftrightarrow \boldsymbol{z}$ and $y_i \leftrightarrow \boldsymbol{x} \leftrightarrow \boldsymbol{z}$ for all $i = 1, \dots, n$. We further assume that $\boldsymbol{z}$ can be decomposed into a set of variables, i.e., $\boldsymbol{z} = \{z_j\}_{j=1}^m$, with each z_j also governed by a conditional distribution $p(z_j \mid \boldsymbol{x})$. From this point onward, we refer to the z_j as the variables of $\boldsymbol{z}$.

Next, we define four desirable properties for a disentangled representation. These properties are conceptually related to those discussed in Section 2, but differ in a crucial aspect: while previous definitions assume independence among factors and from nuisances, our formulation explicitly accounts for dependencies. For the following definitions, we assume that z_j is the representation variable intended to describe the factor y_i.

(i) **Factors-Invariance**: We say that a variable $z_j \in \boldsymbol{z}$ is *factors-invariant* with respect to a factor y_i when it satisfies the following Markov chain:

$$\tilde{\boldsymbol{y}}_i \leftrightarrow y_i \leftrightarrow z_j \tag{1}$$

where $\tilde{\boldsymbol{y}}_i = \{y_k\}_{k \neq i}$ denotes the set of all other factors. This is equivalent to the conditional independence $I(z_j; \tilde{\boldsymbol{y}}_i \mid y_i) = 0$, meaning that once y_i is known, z_j provides no additional information about the remaining factors $\tilde{\boldsymbol{y}}_i$. In other words, the distribution of z_j is fully determined by the value of y_i alone. This property is directly connected to modularity (Ridgeway & Mozer, 2018) (or disentanglement in Eastwood & Williams (2018)): a representation z_j is said to be modular (or disentangled) if it captures at most one factor y_i. However, such definitions implicitly assume that the factors are independent. When dependencies exist among the factors—as is often the case in real-world scenarios—this definition becomes problematic. For instance, if two factors y_i and y_k are dependent, then y_k inherently contains information about y_i. As a result, under the original definition, it would be impossible for z_j to simultaneously be modular and fully informative about y_i, since it would necessarily encode some information about y_k as well.

(ii) **Nuisances-Invariance**: We say that a variable $z_j \in \boldsymbol{z}$ is *nuisances-invariant* with respect to a factor y_i and a set of nuisances $\boldsymbol{n}$ when it satisfies the following Markov chain:

$$\boldsymbol{n} \leftrightarrow y_i \leftrightarrow z_j \tag{2}$$

This is equivalent to the conditional independence $I(z_j; \boldsymbol{n} \mid y_i) = 0$, meaning that once y_i is known, z_j remains unaffected by variations in the nuisances. This property is typically neither discussed nor

[2]Most prior work assumes marginal independence between factors, i.e., $p(y_i, y_j) = p(y_i)p(y_j)$ for $i \neq j$, which rarely holds in real-world scenarios.

measured in prior work—possibly due to the difficulty of estimating the distribution of $\boldsymbol{n}$. Despite this, it is an important criterion for disentangled representations. While factors-invariance ensures that z_j is unaffected by other factors y_k ($k \neq i$), we would also expect z_j to be invariant to aspects of the input that are not factors of variation—i.e., the nuisances—if the representation is truly disentangled (Carbonneau et al., 2022).

(iii) **Representations-Invariance**: We say that a variable z_j of the representation $\boldsymbol{z}$ is *representations-invariant* for a factor y_i when it satisfies the following Markov chain:

$$\tilde{\boldsymbol{z}}_j \leftrightarrow z_j \leftrightarrow y_i \tag{3}$$

where $\tilde{\boldsymbol{z}}_j = \{z_k\}_{k \neq j}$ denotes all other variables in the representation. This is equivalent to the conditional independence $I(y_i; \tilde{\boldsymbol{z}}_j \mid z_j) = 0$, which implies that y_i can be entirely predicted from z_j alone, without any additional information from the rest of the representation. This property is useful in downstream applications: for instance, if the goal is to predict y_i, one could use z_j in isolation and ignore the remaining components of $\boldsymbol{z}$. Similarly, in controllable generative models, it would suffice to manipulate z_j to control the value of y_i in the generated output $\boldsymbol{x}$. This notion is closely related to compactness as defined in (Ridgeway & Mozer, 2018) (or completeness in Eastwood & Williams (2018)), where a representation z_j is said to be compact (or complete) if it is the sole component encoding information about a factor y_i. As with factors-invariance, this property assumes independence among factors. In realistic scenarios where factors such as y_i and y_k are dependent, and represented by z_j and z_l respectively, it is natural for both z_j and z_l to carry information about both factors.

(iv) **Explicitness** We say that a representation $\boldsymbol{z}$ is *explicit* for a factor y_i when it satisfies the following Markov chain:

$$\boldsymbol{x} \leftrightarrow \boldsymbol{z} \leftrightarrow y_i \tag{4}$$

This means that $I(y_i; \boldsymbol{x} | \boldsymbol{z}) = 0$, i.e., $\boldsymbol{x}$ provides no information about y_i when $\boldsymbol{z}$ is known or, equivalently, $\boldsymbol{z}$ contains all the information about y_i. This property is referred to as explicitness in Ridgeway & Mozer (2018) and as informativeness in Eastwood & Williams (2018). Unlike the other properties discussed, explicitness is not specific to disentangled representations—it is a desirable quality for representations more generally (Bengio et al., 2013).

4 Measuring Disentanglement via Minimality and Sufficiency

In this section, we define minimality and sufficiency, and argue that directly measuring these properties is more convenient and informative than evaluating the four previously discussed properties independently. We also provide practical methods for estimating these quantities from data. We now present the formal definitions:

Definition 1. Given a representation z of an input x and a target variable y, we say that z is *minimal* with respect to y if it satisfies $I(z; x|y) = 0$ or, equivalently, if $I(z; y) = I(z; x)$, i.e., z contains information *only* about y.

Definition 2. Given a representation z of an input x and a target variable y, we say that z is *sufficient* with respect to y if it satisfies $I(y; x|z) = 0$ or, equivalently, if $I(z; x) = I(y; x)$, i.e., z contains *all* information about y.

4.1 Connection between Disentanglement and Minimality and Sufficiency

We connect below the properties introduced in Section 3 to the concepts of sufficiency and minimality.

Theorem 1. *Let $\boldsymbol{y} = \{y_l\}_{l=1}^n$ denote a set of factors and $\boldsymbol{z} = \{z_j\}_{j=1}^m$ a representation. Then, z_j is a minimal representation of y_i if and only if z_j is factors-invariant with respect to y_i and nuisances-invariant. Equivalently, we have that:* [3]

$$(\boldsymbol{x} \leftrightarrow y_i \leftrightarrow z_j) \iff (\tilde{\boldsymbol{y}}_i \leftrightarrow y_i \leftrightarrow z_j) \wedge (\boldsymbol{n} \leftrightarrow y_i \leftrightarrow z_j)$$

[3]Note that $\iff$ denotes "if and only if", whereas $\leftrightarrow$ represents a dependency structure in a Markov chain.

where $\tilde{\boldsymbol{y}}_i = \{y_l\}_{l \neq i}$. Therefore, satisfying minimality is equivalent to jointly satisfying factors-invariance and nuisances-invariance. Intuitively, if z_j is fully determined by y_i (i.e., minimal), then neither the remaining factors $\tilde{\boldsymbol{y}}_i$ nor the nuisances $\boldsymbol{n}$ influence z_j once y_i is known—captured by factors-invariance and nuisances-invariance, respectively. Conversely, since the input $\boldsymbol{x}$ is completely defined by $\boldsymbol{y}$ and $\boldsymbol{n}$, the implication also holds in the reverse direction. A formal proof of this theorem is provided in Appendix A.

Theorem 2. *Let y_i denote a factor and $\boldsymbol{z} = \{z_l\}_{l=1}^m$ a representation. Then, z_j is a sufficient representation of y_i if and only if z_j is representations-invariant for y_i and $\boldsymbol{z}$ is explicit for y_i. Equivalently, we have that:*

$$(\boldsymbol{x} \leftrightarrow z_j \leftrightarrow y_i) \Longleftrightarrow (\tilde{\boldsymbol{z}}_j \leftrightarrow z_j \leftrightarrow y_i) \wedge (\boldsymbol{x} \leftrightarrow \boldsymbol{z} \leftrightarrow y_i)$$

where $\tilde{\boldsymbol{z}}_j = \{z_l\}_{l \neq j}$. Thus, satisfying sufficiency is equivalent to jointly satisfying representations-invariance and explicitness. Intuitively, if y_i is fully determined by z_j independently of $\boldsymbol{x}$ (i.e., sufficiency), then there is no information about y_i in $\boldsymbol{z}$ that is not already contained in z_j (representations-invariance), and since $z_j \in \boldsymbol{z}$, this also implies that y_i is fully captured by $\boldsymbol{z}$ (explicitness). Conversely, if all the information in $\boldsymbol{z}$ relevant to y_i is already present in z_j (representations-invariance), and $\boldsymbol{z}$ captures all information about y_i (explicitness), then z_j must fully describe y_i—that is, z_j is sufficient. A formal proof of this theorem is provided in Appendix A.

4.2 Why Measuring Disentanglement via Minimality and Sufficiency?

As previously demonstrated, a representation is disentangled—according to the properties defined in Section 3—if and only if it is both minimal and sufficient. Next, we argue that it is more practical and principled to directly assess sufficiency and minimality, rather than evaluating the four individual properties separately.

- **Minimality vs. Factors-Invariance + Nuisances-Invariance**: When evaluating whether a representation is disentangled, the true objective is to determine whether each of its variables is influenced by a single factor—regardless of the presence of other factors or nuisances. It is not sufficient for a variable to be aligned with a specific factor if it is also significantly affected by nuisance variables. Similarly, a variable that is invariant to nuisances but still influenced by multiple factors cannot be considered disentangled. Therefore, we argue that these two properties should be assessed jointly, and—as established in Theorem 1—this can be effectively achieved by measuring minimality.
- **Sufficiency vs. Representations-Invariance + Explicitness**: When evaluating whether a representation is disentangled, we aim to determine whether each factor can be both exclusively and fully described by a single variable in the representation. A representation in which a single variable influences a factor but captures only a small portion of its information fails to satisfy the fundamental purpose of a representation—namely, to describe the underlying factors. Conversely, if a factor is well described by the representation as a whole but is equally influenced by all variables, we cannot claim that the representation is disentangled. For this reason, we argue that it is more meaningful to assess these two aspects jointly rather than separately. As established in Theorem 2, this can be effectively achieved by measuring sufficiency.

4.3 Defining metrics for Minimality and Sufficiency

At the beginning of this section, we defined when a representation is minimal or sufficient. In practice, however, representations are rarely perfectly minimal or sufficient. It is therefore desirable to design continuous metrics that quantify how close a representation is to these ideals. Moreover, while minimality and sufficiency can be defined for each pair (z_j, y_i), it is often useful to summarize these values into a single score that characterizes the entire representation $\boldsymbol{z}$. We next introduce metrics that meet these requirements.

Minimality Recall that a variable z_j is minimal with respect to a factor y_i if $I(z_j; \boldsymbol{x} \mid y_i) = I(z_j; x) - I(z_j; y_i) = 0$. To improve interpretability and facilitate comparison across metrics, we normalize this value by $I(z_j; \boldsymbol{x})$ to ensure it lies in $[0, 1]$. We thus define the minimality of z_j with respect to y_i as

$$m_{ij} = 1 - \frac{I(z_j; \boldsymbol{x} \mid y_i)}{I(z_j; \boldsymbol{x})} = \frac{I(z_j; y_i)}{I(z_j; x)}. \tag{5}$$

Here, m_{ij} reflects the fraction of information in z_j that is attributable to factor y_i. While these values provide factor-wise insights, it is common to report a single scalar as a disentanglement score. Assuming, as in Section 3, a one-to-one correspondence between z_j and a specific y_i, we define $\hat{m}_j = \max_i m_{ij}$, measuring the minimality of z_j relative to its most associated factor. The overall minimality score is then obtained by averaging across all representation variables:

$$\bar{m} = \frac{1}{m} \sum_j \hat{m}_j \tag{6}$$

Sufficiency Analogously, a variable z_j is sufficient for a factor y_i if $I(y_i; \boldsymbol{x} \mid z_j) = I(y_i; \boldsymbol{x}) - I(y_i; z_j) = 0$. Applying the same normalization yields the sufficiency of z_j with respect to y_i:

$$s_{ij} = 1 - \frac{I(y_i; \boldsymbol{x} \mid z_j)}{I(y_i; \boldsymbol{x})} = \frac{I(y_i; z_j)}{I(y_i; x)}. \tag{7}$$

Thus, s_{ij} measures the proportion of information about y_i captured by z_j. Following the same reasoning as for minimality, we compute $\hat{s}_i = \max_j s_{ij}$, which evaluates how well the most informative representation variable predicts y_i. Averaging across factors yields the overall sufficiency score:

$$\bar{s} = \frac{1}{n} \sum_i \hat{s}_i \tag{8}$$

4.4 Computing Minimality and Sufficiency in Practice

Calculating the mutual information terms from the previous section can be intractable, particularly when the factor y_i or the representation z_j is continuous. To address this, we proceed as follows: (i) we discretize z_j by binning, a common practice in representation learning (Shwartz-Ziv & Tishby, 2017; Locatello et al., 2019b; Carbonneau et al., 2022); and (ii) in the less common case where $y_i$is not categorical, we also discretize it by binning, while leaving it unchanged when it is categorical. Once both y_i and z_j are categorical, we obtain: (i) $I(z_j; y_i)$ can be computed directly; (ii) $I(z_j; x) = H(z_j)$, since z_j is fully defined by x and therefore $H(z_j|x) = 0$; and (iii) $I(y_i; x) = H(y_i)$, by the same reasoning.

5 Toy Experiments

Since our goal is to evaluate whether the proposed metrics reliably capture disentanglement, we manually construct the representations in this section's experiments. This design provides full access to the ground truth and, consequently, clear expectations about the metrics' behavior.

We generate two toy experiments where a set of factors is sampled and the corresponding representations are explicitly constructed from them. This controlled setup enables a systematic comparison between our metrics and existing disentanglement metrics under conditions that are particularly challenging for prior approaches, namely: (i) dependence between factors of variation, and (ii) the presence of nuisances in the representation. Accordingly, we design two simple experiments, each addressing one of these scenarios.

5.1 Disentanglement under Dependence between Factors of Variation

Definition of the factors We define a set of n factors $\boldsymbol{y} = \{y_i\}_{i=1}^n$, where each y_i is constructed from a corresponding set of latent variables $\boldsymbol{\epsilon} = \{\epsilon_i\}_{i=1}^n$, with $\epsilon_i \sim \mathcal{U}[0, 1]$ for $i = 1, \ldots, n$. Each factor is constructed as follows:

$$y_i' = \delta \epsilon_i + \frac{1-\delta}{n-1} \sum_{l \neq i} \epsilon_l, \qquad y_i = \mathcal{B}(y_i')$$

for $i = 1, \ldots, n$, where $\delta \in \left[\frac{1}{n}, 1\right]$ is a hyperparameter that controls the level of independence among the factors, and the function $\mathcal{B} : \mathbb{R} \to \{0, 1, \ldots, K\}$ bins the continuous values into K discrete classes. In particular: (i) when $\delta = 1$, each y_i depends solely on ϵ_i, resulting in independent factors; and (ii) when $\delta = \frac{1}{n}$, each y_i is equally influenced by all elements of $\boldsymbol{\epsilon}$, yielding maximal dependence.

Definition of the representations For simplicity, we define a representation $\boldsymbol{z} = \{z_j\}_{j=1}^n$ with the same number of variables as factors, where each z_j is intended to correspond to y_j. We define each representation variable as:

$$z'_j = \alpha y_j + \frac{1-\alpha}{n-1} \sum_{i \neq j} y_i, \qquad z_j = \cos\left(\frac{\pi z'_j}{K}\right)$$

for $j = 1, \ldots, n$, where $\alpha \in \left[\frac{1}{n}, 1\right]$ is a hyperparameter that controls the degree of disentanglement. Specifically: (i) when $\alpha = 1$, each z_j depends only on y_j, resulting in a perfectly disentangled representation; (ii) when $\alpha = \frac{1}{n}$, each z_j depends equally on all factors in $\boldsymbol{y}$, yielding a fully entangled representation. The cosine function serves to introduce a non-linearity in the mapping from factors to representations. We note that, since its argument lies in $[0, \pi]$, the transformation remains bijective, ensuring that no information about the original factor is lost.

What results should we expect? Under the described setup, a disentanglement metric should satisfy the following desiderata:

(i) The metric should attain a value of 1 when $\alpha = 1$, regardless of the value of δ. In this case, there is a bijective mapping between each factor y_j and its corresponding representation variable z_j, implying perfect disentanglement.

(ii) For a fixed value of α, the metric should increase as δ decreases. Consider a scenario in which the factors are highly dependent (i.e., low δ), and each representation variable z_j is influenced by all factors (i.e., low α). In such a case, most of the information in z_j can still be attributed to its corresponding factor y_j, resulting in higher factors-invariance. This is because the other factors contributing to z_j are themselves strongly correlated with y_j. For example, in Figure 1, suppose the encoder is not perfectly disentangled and z_1, intended to represent fruit shape, also encodes color. Since banana-shaped fruits are typically yellow, z_1 still contains information almost exclusively about shape, despite also carrying color information. Similarly, if z_1 does not fully capture shape because of imperfect disentanglement, the fact that it also encodes color indirectly reinforces shape information: if the fruit is yellow, its shape is much more likely to be that of a banana than that of a strawberry.

What results do we observe? Figure 2 shows the results of various disentanglement metrics for $n = 4$ and $K = 5$ across different values of α and δ. We observe that *Minimality* and *Sufficiency*[4] are the only metrics that consistently satisfy the desirable properties outlined above. More specifically, we can see that: (i) Metrics such as *MIG*, *SAP*, *Disentanglement*, and *Completeness* fail to reach a value of one when $\alpha = 1$ and $\delta = 1$, i.e., even in the fully disentangled and independent case; (ii) Metrics such as *FactorVAE*, *MIG*, *SAP*, and *Modularity* never attain their maximum values as $\delta = \frac{1}{n}$, meaning they fail to correctly reflect disentanglement when correlations are present; and (iii) *Disentanglement* and *Completeness* remain largely invariant with respect to δ, contrary to the expected behavior discussed earlier.

5.2 Disentanglement under the Presence of Nuisances

Definition of the factors In this case, we define the factors using the same formulation as in Section 5.1, but restrict our analysis to the case $\delta = 1$, i.e., complete independence between factors.

Definition of the representations We define the representation $\{z_j\}_{j=1}^n$ such that each z_j is intended to correspond to the factor y_j. Specifically:

$$z'_j = y_j + \beta\epsilon, \qquad z_j = \cos\left(\frac{\pi z'_j}{K}\right)$$

[4]As described in Section 4.4, we compute these metrics by discretizing the representations into 15 bins.

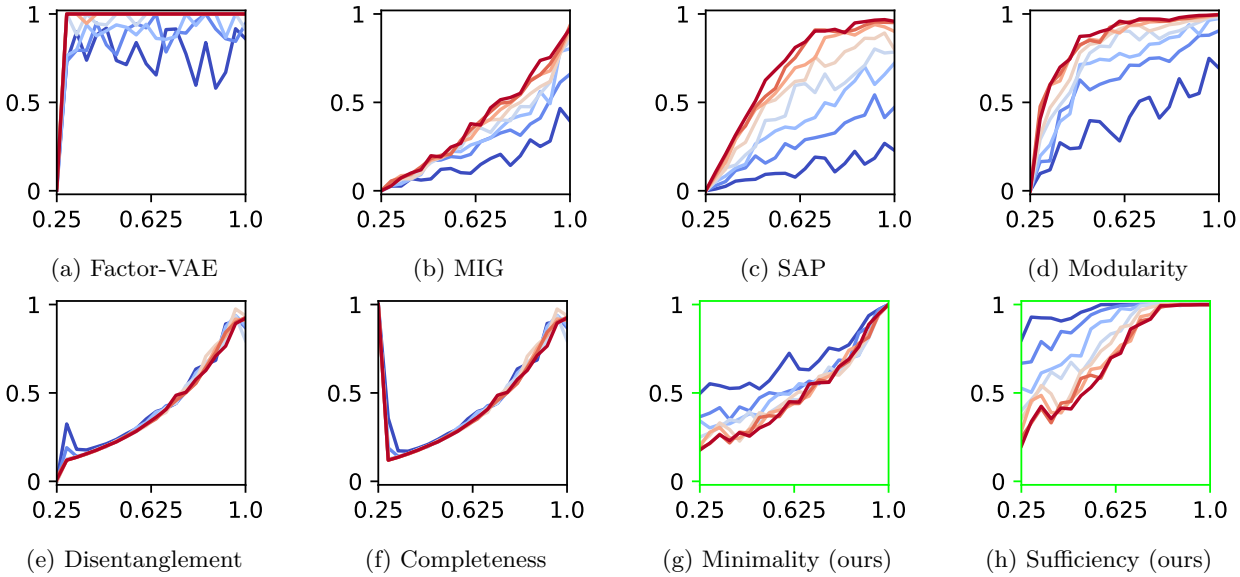

(a) Factor-VAE (b) MIG (c) SAP (d) Modularity

(e) Disentanglement (f) Completeness (g) Minimality (ours) (h) Sufficiency (ours)

Figure 2: Comparison of different metrics (y-axis) across varying values of α (x-axis) and δ (color). Colors range from dark blue (highest dependence, $\delta = \frac{1}{n}$) to dark red (lowest dependence, $\delta = 1$).

where $\epsilon \sim \mathcal{U}[0, 1]$ and $\beta \in \left[0, 1 - \frac{1}{K}\right]$ is a hyperparameter that controls the amount of nuisance information in the representation. As before, the argument of the cosine lies in $[0, \pi]$, ensuring that no information about y_j is lost in the transformation.

What results should we expect? Under this setup, we expect a disentanglement metric that is sensitive to the presence of nuisances to attain its maximum value when $\beta = 0$, and to decrease as β increases.

What results do we observe? Figure 3 shows how different metrics vary with respect to β. We observe that *Minimality* is the only one that is sensitive to the presence of nuisances. While our *Sufficiency* metric remains invariant to changes in β, this is the expected behavior, as shown in Theorem 2, which demonstrates that sufficiency is connected to *representations-invariance* and *explicitness*, but not to *nuisances-invariance*. Finally, we note that *Minimality* does not reach zero even at the highest value of β, because it jointly reflects both *factors-invariance* and *nuisances-invariance* and, in this setup, the former remains maximal.

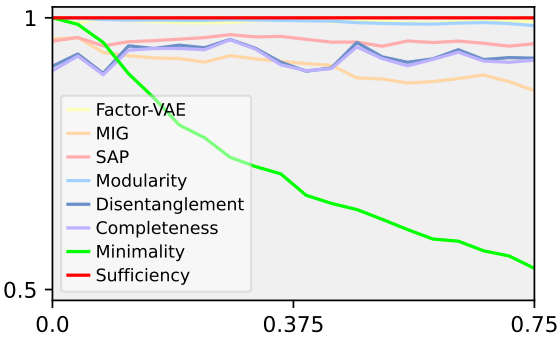

Figure 3: Comparison of different metrics (y-axis) for different values of β (x-axis)

6 Neural Representations Experiments

In this section, we compare our proposed metrics against existing ones on representations obtained from neural network–based encoders. Unlike in earlier experiments, ground-truth disentanglement levels are unavailable, precluding a direct analysis. Instead, as noted in Section 1, the practical value of disentangled representations often lies in their ability to enable strong performance on downstream tasks. We therefore evaluate the metrics by examining their relationship with representation quality across applications typically associated with disentanglement.

To this end, we train models known to induce disentangled representations on three widely used datasets: DSprites (Matthey et al., 2017), MPI3D (Gondal et al., 2019), and Shapes3D (Burgess & Kim, 2018). For

each dataset, we explore six encoder architectures (detailed in Appendix C) combined with six disentanglement losses: β-VAE (Higgins et al., 2017), AnnealedVAE (Burgess et al., 2018), β-TCVAE (Chen et al., 2018), FactorVAE (Kim & Mnih, 2018), Ada-GVAE (Locatello et al., 2020), and HFSVAE (Roth et al., 2022), yielding 36 encoders per dataset. All models are trained using the hyperparameters listed in Appendix B.

After training, we evaluate how each metric correlates with downstream performance under varying levels of factor dependence. Following Roth et al. (2022), we introduce controlled dependencies among factors, even though the original datasets are independent by design. Specifically, we analyze cases with correlations between one, two, or three pairs of factors, as well as a confounded setting in which a single factor depends on all others. A metric suitable for dependent-factor scenarios should satisfy two criteria: (i) it should maintain a strong correlation with downstream performance, and (ii) this correlation should remain stable across different levels of factor dependence. The second condition is particularly critical: since the encoder itself does not change with the degree of factor dependence, a reliable metric should consistently reflect encoder performance irrespective of the underlying factor correlations.

6.1 Accuracy under Data and Computational Constraints

A central motivation for disentangled representations is their expected utility in downstream tasks, particularly under constraints of limited data and computational resources (Schölkopf et al., 2012; Bengio et al., 2013). Building on Locatello et al. (2019b), we examine whether the proposed metrics correlate with predictive performance—measured by factor prediction accuracy—under two restrictive conditions: (i) predictions rely on a single dimension of the representation, and (ii) training is performed with only 1000 samples.

What results should we expect? In this setting, metrics designed to assess *representations-invariance* and *explicitness* are expected to correlate strongly with predictive accuracy. The intuition is that when all the information about a factor y_i is encoded within a single representation dimension z_j, the mapping from z_j to y_i is straightforward. This implies that even a classifier with limited capacity can recover y_i reliably from z_j, underscoring the practical value of such representations for downstream prediction.

What results do we observe? Figure 4 reports the rank correlation between disentanglement metrics and predictive accuracy under the low-resource setting described above. When the factors are independent, *Sufficiency* achieves the strongest correlation with accuracy, while most other metrics also exhibit moderate to high correlations. This indicates that, in the absence of dependencies, several metrics can serve as reasonable proxies for downstream performance. As factor dependence increases, however, a marked difference emerges: *Sufficiency* remains stable, whereas the correlations of the other metrics decline sharply. These findings demonstrate that *Sufficiency* provides the most reliable and robust predictor of downstream accuracy among the metrics considered, especially in scenarios where factors are dependent.

6.2 Interventional Robustness

Another motivation for disentangled representations is their robustness to interventions on the underlying factors. Specifically, if a factor y_i is modified while all others remain fixed, the resulting change in the representation should be confined to its corresponding dimension z_j. This property is particularly critical for controlled generative models: if it holds, modifying z_j ensures that the generated data changes exclusively in the associated factor y_i. To quantify this, Suter et al. (2019) introduced the Interventional Robustness Score (IRS), which quantifies the robustness of a representation under changes in the factors. We next compute the rank correlation between IRS and all considered metrics.

What results should we expect? In this setting, metrics designed to capture *factors-invariance* are expected to correlate positively with IRS. The underlying intuition is that if a representation dimension z_j is entirely determined by a factor y_i, i.e., $p(z_j \mid y_i) = p(z_j \mid x)$, then an intervention on y_i—while keeping all other factors fixed—will affect z_j exclusively through y_i. In such cases, z_j can be regarded as robust to interventions, which is precisely the property that IRS is intended to measure.

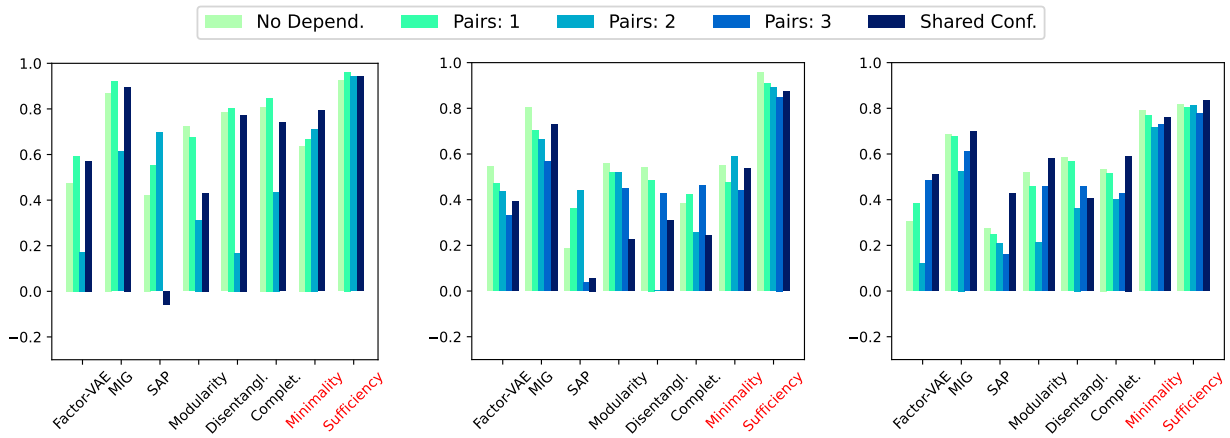

(a) DSprites (b) MPI3D (c) Shapes3D

Figure 4: Rank correlation between disentanglement metrics and accuracy of a random forest classifier trained on 1000 samples across different levels of factor dependence.

What results do we observe? In contrast to the accuracy results, Figure 5 shows that most metrics have weak—or even negative—correlations with IRS, including when factors are independent, and their correlations vary substantially across dependence levels, sometimes changing sign. By contrast, *Minimality* exhibits a consistently strong, positive correlation with IRS across all datasets and for every dependence regime considered. The combination of stability and magnitude indicates that *Minimality* tracks interventional robustness far more faithfully than competing metrics.

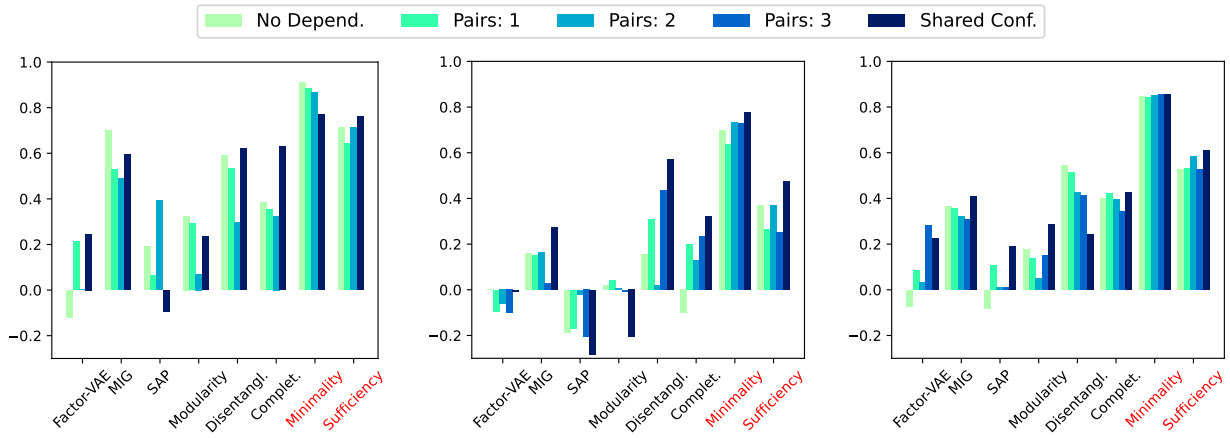

(a) DSprites (b) MPI3D (c) Shapes3D

Figure 5: Rank correlation between disentanglement metrics and Interventional Robustness Score across different levels of factor dependence.

7 Conclusions

In this work, we identified two fundamental limitations of existing definitions and metrics of disentanglement: (i) the assumption that factors of variation are independent, and (ii) the lack of consideration for invariance to nuisances in the representation. These limitations restrict the applicability of standard metrics in realistic scenarios.

To address them, we proposed four information-theoretic properties that characterize disentanglement in the presence of both dependent factors and nuisances. We showed that these properties naturally connect to

the concepts of *minimality* and *sufficiency*, and argued that directly measuring these two quantities provides a more principled and interpretable characterization of disentanglement. To this end, we introduced novel metrics for quantifying *minimality* and *sufficiency*, along with tractable estimators.

Our experiments demonstrate that the proposed metrics capture disentanglement more reliably across a wider range of scenarios than existing alternatives. Moreover, unlike prior metrics, *minimality* and *sufficiency* consistently exhibit positive correlation with downstream performance in tasks typically associated with disentangled representations, even in the presence of dependent factors of variation.

Acknowledgments

This work has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 101007666, MCIN/AEI/10.13039/501100011033 under Grant PID2021-126061OB-C44, and the Government of Aragón (Grant Group T36 23R).

References

Alessandro Achille and Stefano Soatto. Emergence of invariance and disentanglement in deep representations. *Journal of Machine Learning Research*, 19(50):1–34, 2018a.

Alessandro Achille and Stefano Soatto. Information dropout: Learning optimal representations through noisy computation. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2897–2905, 2018b.

Kartik Ahuja, Divyat Mahajan, Yixin Wang, and Yoshua Bengio. Interventional causal representation learning. In *International conference on machine learning*, pp. 372–407. PMLR, 2023.

Alexander Alemi, Ben Poole, Ian Fischer, Joshua Dillon, Rif A Saurous, and Kevin Murphy. Fixing a broken elbo. In *International conference on machine learning*, pp. 159–168. PMLR, 2018.

Alexander A Alemi, Ian Fischer, Joshua V Dillon, and Kevin Murphy. Deep variational information bottleneck. *arXiv preprint arXiv:1612.00410*, 2016.

Antonio Almudévar, Théo Mariotte, Alfonso Ortega, Marie Tahon, Luis Vicente, Antonio Miguel, and Eduardo Lleida. Predefined prototypes for intra-class separation and disentanglement. *arXiv preprint arXiv:2406.16145*, 2024.

Antonio Almudévar, José Miguel Hernández-Lobato, Sameer Khurana, Ricard Marxer, and Alfonso Ortega. Aligning multimodal representations through an information bottleneck. *arXiv preprint arXiv:2506.04870*, 2025.

Antonio Almudévar, Théo Mariotte, Alfonso Ortega, and Marie Tahon. Unsupervised multiple domain translation through controlled disentanglement in variational autoencoder. In *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 7010–7014, 2024. doi: 10.1109/ICASSP48485.2024.10446649.

Javier Antoran and Antonio Miguel. Disentangling and learning robust representations with natural clustering. In *2019 18th IEEE International Conference On Machine Learning And Applications (ICMLA)*, pp. 694–699. IEEE, 2019.

Solon Barocas and Andrew D Selbst. Big data's disparate impact. *Calif. L. Rev.*, 104:671, 2016.

Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.

Chris Burgess and Hyunjik Kim. 3d shapes dataset. https://github.com/deepmind/3dshapes-dataset/, 2018.

Christopher P Burgess, Irina Higgins, Arka Pal, Loic Matthey, Nick Watters, Guillaume Desjardins, and Alexander Lerchner. Understanding disentangling in β-vae. *arXiv preprint arXiv:1804.03599*, 2018.

Toon Calders and Indrė Žliobaitė. Why unbiased computational processes can lead to discriminative decision procedures. In *Discrimination and Privacy in the Information Society: Data mining and profiling in large databases*, pp. 43–57. Springer, 2013.

Marc-André Carbonneau, Julian Zaidi, Jonathan Boilard, and Ghyslain Gagnon. Measuring disentanglement: A review of metrics. *IEEE transactions on neural networks and learning systems*, 2022.

Richard J Chen, Judy J Wang, Drew FK Williamson, Tiffany Y Chen, Jana Lipkova, Ming Y Lu, Sharifa Sahai, and Faisal Mahmood. Algorithmic fairness in artificial intelligence for medicine and healthcare. *Nature biomedical engineering*, 7(6):719–742, 2023.

Ricky TQ Chen, Xuechen Li, Roger B Grosse, and David K Duvenaud. Isolating sources of disentanglement in variational autoencoders. *Advances in neural information processing systems*, 31, 2018.

Jaewoong Choi, Geonho Hwang, and Myungjoo Kang. Discond-vae: disentangling continuous factors from the discrete. *arXiv preprint arXiv:2009.08039*, 2020.

Taco Cohen and Max Welling. Learning the irreducible representations of commutative lie groups. In *International Conference on Machine Learning*, pp. 1755–1763. PMLR, 2014.

Guillaume Desjardins, Aaron Courville, and Yoshua Bengio. Disentangling factors of variation via generative entangling. *arXiv preprint arXiv:1210.5474*, 2012.

Kien Do and Truyen Tran. Theory and evaluation metrics for learning disentangled representations. *arXiv preprint arXiv:1908.09961*, 2019.

Emilien Dupont. Learning disentangled joint continuous and discrete representations. *Advances in neural information processing systems*, 31, 2018.

Cian Eastwood and Christopher KI Williams. A framework for the quantitative evaluation of disentangled representations. In *International conference on learning representations*, 2018.

Gege Gao, Huaibo Huang, Chaoyou Fu, Zhaoyang Li, and Ran He. Information bottleneck disentanglement for identity swapping. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 3404–3413, 2021.

Leilani H Gilpin, David Bau, Ben Z Yuan, Ayesha Bajwa, Michael Specter, and Lalana Kagal. Explaining explanations: An overview of interpretability of machine learning. In *2018 IEEE 5th International Conference on data science and advanced analytics (DSAA)*, pp. 80–89. IEEE, 2018.

Muhammad Waleed Gondal, Manuel Wuthrich, Djordje Miladinovic, Francesco Locatello, Martin Breidt, Valentin Volchkov, Joel Akpo, Olivier Bachem, Bernhard Schölkopf, and Stefan Bauer. On the transfer of inductive bias from simulation to the real world: a new disentanglement dataset. *Advances in Neural Information Processing Systems*, 32, 2019.

Irina Higgins, Loic Matthey, Arka Pal, Christopher P Burgess, Xavier Glorot, Matthew M Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. *ICLR (Poster)*, 3, 2017.

Wei-Ning Hsu, Yu Zhang, and James Glass. Unsupervised learning of disentangled and interpretable representations from sequential data. *Advances in neural information processing systems*, 30, 2017.

Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 172–189, 2018.

Insu Jeon, Wonkwang Lee, Myeongjang Pyeon, and Gunhee Kim. Ib-gan: Disentangled representation learning with information bottleneck generative adversarial networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 7926–7934, 2021.

Hadi Kazemi, Seyed Mehdi Iranmanesh, and Nasser Nasrabadi. Style and content disentanglement in generative adversarial networks. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 848–856. IEEE, 2019.

Hyunjik Kim and Andriy Mnih. Disentangling by factorising. In *International conference on machine learning*, pp. 2649–2658. PMLR, 2018.

Minyoung Kim, Yuting Wang, Pritish Sahu, and Vladimir Pavlovic. Relevance factor vae: Learning and identifying disentangled factors. *arXiv preprint arXiv:1902.01568*, 2019.

Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

Tejas D Kulkarni, William F Whitney, Pushmeet Kohli, and Josh Tenenbaum. Deep convolutional inverse graphics network. *Advances in neural information processing systems*, 28, 2015.

Abhishek Kumar, Prasanna Sattigeri, and Avinash Balakrishnan. Variational inference of disentangled latent concepts from unlabeled observations. *arXiv preprint arXiv:1711.00848*, 2017.

Brenden M Lake, Tomer D Ullman, Joshua B Tenenbaum, and Samuel J Gershman. Building machines that learn and think like people. *Behavioral and brain sciences*, 40:e253, 2017.

Yann LeCun, Fu Jie Huang, and Leon Bottou. Learning methods for generic object recognition with invariance to pose and lighting. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, volume 2, pp. II–104. IEEE, 2004.

Wenqian Liu, Runze Li, Meng Zheng, Srikrishna Karanam, Ziyan Wu, Bir Bhanu, Richard J Radke, and Octavia Camps. Towards visually explaining variational autoencoders. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 8642–8651, 2020.

Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.

Francesco Locatello, Gabriele Abbati, Thomas Rainforth, Stefan Bauer, Bernhard Schölkopf, and Olivier Bachem. On the fairness of disentangled representations. *Advances in neural information processing systems*, 32, 2019a.

Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Raetsch, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem. Challenging common assumptions in the unsupervised learning of disentangled representations. In *international conference on machine learning*, pp. 4114–4124. PMLR, 2019b.

Francesco Locatello, Ben Poole, Gunnar Rätsch, Bernhard Schölkopf, Olivier Bachem, and Michael Tschannen. Weakly-supervised disentanglement without compromises. In *International conference on machine learning*, pp. 6348–6359. PMLR, 2020.

Loic Matthey, Irina Higgins, Demis Hassabis, and Alexander Lerchner. dsprites: Disentanglement testing sprites dataset. https://github.com/deepmind/dsprites-dataset/, 2017.

Milton Montero, Jeffrey Bowers, Rui Ponte Costa, Casimir Ludwig, and Gaurav Malhotra. Lost in latent space: Examining failures of disentangled models at combinatorial generalisation. *Advances in Neural Information Processing Systems*, 35:10136–10149, 2022.

Milton Llera Montero, Casimir JH Ludwig, Rui Ponte Costa, Gaurav Malhotra, and Jeffrey Bowers. The role of disentanglement in generalisation. In *International Conference on Learning Representations*, 2020.

Meike Nauta, Jan Trienes, Shreyasi Pathak, Elisa Nguyen, Michelle Peters, Yasmin Schmitt, Jörg Schlötterer, Maurice van Keulen, and Christin Seifert. From anecdotal evidence to quantitative evaluation methods: A systematic review on evaluating explainable ai. *ACM Computing Surveys*, 55(13s):1–42, 2023.

Brooks Paige, Jan-Willem Van De Meent, Alban Desmaison, Noah Goodman, Pushmeet Kohli, Frank Wood, Philip Torr, et al. Learning disentangled representations with semi-supervised deep generative models. *Advances in neural information processing systems*, 30, 2017.

Scott E Reed, Yi Zhang, Yuting Zhang, and Honglak Lee. Deep visual analogy-making. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015.

Karl Ridgeway. A survey of inductive biases for factorial representation-learning. *arXiv preprint arXiv:1612.05299*, 2016.

Karl Ridgeway and Michael C Mozer. Learning deep disentangled embeddings with the f-statistic loss. *Advances in neural information processing systems*, 31, 2018.

Karsten Roth, Mark Ibrahim, Zeynep Akata, Pascal Vincent, and Diane Bouchacourt. Disentanglement of correlated factors via hausdorff factorized support. *arXiv preprint arXiv:2210.07347*, 2022.

Mhd Hasan Sarhan, Nassir Navab, Abouzar Eslami, and Shadi Albarqouni. Fairness by learning orthogonal disentangled representations. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIX 16*, pp. 746–761. Springer, 2020.

Jürgen Schmidhuber. Learning factorial codes by predictability minimization. *Neural computation*, 4(6): 863–879, 1992.

Bernhard Schölkopf, Dominik Janzing, Jonas Peters, Eleni Sgouritsa, Kun Zhang, and Joris Mooij. On causal and anticausal learning. *arXiv preprint arXiv:1206.6471*, 2012.

Anna Sepliarskaia, Julia Kiseleva, Maarten de Rijke, et al. Evaluating disentangled representations. *arXiv preprint arXiv:1910.05587*, 2019.

Xinwei Shen, Furui Liu, Hanze Dong, Qing Lian, Zhitang Chen, and Tong Zhang. Weakly supervised disentangled generative causal representation learning. *Journal of Machine Learning Research*, 23(241): 1–55, 2022.

Ravid Shwartz-Ziv and Naftali Tishby. Opening the black box of deep neural networks via information. *arXiv preprint arXiv:1703.00810*, 2017.

Raphael Suter, Djordje Miladinovic, Bernhard Schölkopf, and Stefan Bauer. Robustly disentangled causal mechanisms: Validating deep representations for interventional robustness. In *International Conference on Machine Learning*, pp. 6056–6065. PMLR, 2019.

Naftali Tishby, Fernando C Pereira, and William Bialek. The information bottleneck method. *arXiv preprint physics/0004057*, 2000.

Frederik Träuble, Elliot Creager, Niki Kilbertus, Francesco Locatello, Andrea Dittadi, Anirudh Goyal, Bernhard Schölkopf, and Stefan Bauer. On disentangled representations learned from correlated data. In *International conference on machine learning*, pp. 10401–10412. PMLR, 2021.

Matias Vera, Pablo Piantanida, and Leonardo Rey Vega. The role of the information bottleneck in representation learning. In *2018 IEEE international symposium on information theory (ISIT)*, pp. 1580–1584. IEEE, 2018.

Julius von Kügelgen, Michel Besserve, Liang Wendong, Luigi Gresele, Armin Kekić, Elias Bareinboim, David Blei, and Bernhard Schölkopf. Nonparametric identifiability of causal representations from unknown interventions. *Advances in Neural Information Processing Systems*, 36, 2024.

Nicholas Watters, Loic Matthey, Christopher P Burgess, and Alexander Lerchner. Spatial broadcast decoder: A simple architecture for learning disentangled representations in vaes. *arXiv preprint arXiv:1901.07017*, 2019.

Daniel E Worrall, Stephan J Garbin, Daniyar Turmukhambetov, and Gabriel J Brostow. Interpretable transformations with encoder-decoder networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 5726–5735, 2017.

Masanori Yamada, Heecheol Kim, Kosuke Miyoshi, Tomoharu Iwata, and Hiroshi Yamakawa. Disentangled representations for sequence data using information bottleneck principle. In *Asian Conference on Machine Learning*, pp. 305–320. PMLR, 2020.

Xinchen Yan, Jimei Yang, Kihyuk Sohn, and Honglak Lee. Attribute2image: Conditional image generation from visual attributes. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*, pp. 776–791. Springer, 2016.

Quan-shi Zhang and Song-Chun Zhu. Visual interpretability for deep learning: a survey. *Frontiers of Information Technology & Electronic Engineering*, 19(1):27–39, 2018.

Sharon Zhou, Eric Zelikman, Fred Lu, Andrew Y Ng, Gunnar Carlsson, and Stefano Ermon. Evaluating the disentanglement of deep generative models through manifold topology. *arXiv preprint arXiv:2006.03680*, 2020.

Xinqi Zhu, Chang Xu, and Dacheng Tao. Where and what? examining interpretable disentangled representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5861–5870, 2021.

A Proofs of Theorems

Theorem 1. *Let $\boldsymbol{y} = \{y_l\}_{l=1}^n$ be some factors and $\boldsymbol{z} = \{z_j\}_{j=1}^m$ a representation. If z_j is a minimal representation of y_i, then z_j is factors-invariant for y_i and nuisances-invariant. Equivalently, we have that:*

$$(\boldsymbol{x} \leftrightarrow y_i \leftrightarrow z_j) \Longleftrightarrow (\tilde{\boldsymbol{y}}_i \leftrightarrow y_i \leftrightarrow z_j) \wedge (\boldsymbol{n} \leftrightarrow y_i \leftrightarrow z_j)$$

Proof. First, we demonstrate that $(\boldsymbol{x} \leftrightarrow y_i \leftrightarrow z_j) \Longrightarrow (\tilde{\boldsymbol{y}}_i \leftrightarrow y_i \leftrightarrow z_j)$: Given this Markov chain, by the DPI, we have that $I(z_j;\boldsymbol{x}) \leq I(z_j;y_i)$. Since z_j is a representation of $\boldsymbol{x}$, we also have by DPI that $I(z_j;\boldsymbol{y}) \leq I(z_j;\boldsymbol{x})$. Therefore, it follows that $I(z_j;\boldsymbol{y}) \leq I(z_j;y_i)$ On the other hand, by the chain rule of mutual information, we have that $I(z_j;\boldsymbol{y}) = I(z_j;y_i,\tilde{\boldsymbol{y}}_{\boldsymbol{i}}) = I(z_j;y_i) + I(z_j;\tilde{\boldsymbol{y}}_{\boldsymbol{i}}|y_i)$. By the non-negativity of mutual information, we are left with $I(z_j;\boldsymbol{y}) \geq I(z_j;y_i)$. Therefore, we have that $I(z_j;y_i) = I(z_j;\boldsymbol{y}) = I(z_j;y_i,\tilde{\boldsymbol{y}}_{\boldsymbol{i}})$ or, equivalently, $\tilde{\boldsymbol{y}}_i \leftrightarrow y_i \leftrightarrow z_j$.

Second, we demonstrate that $(\boldsymbol{x} \leftrightarrow y_i \leftrightarrow z_j) \Longrightarrow (\boldsymbol{n} \leftrightarrow y_i \leftrightarrow z_j)$: Given this Markov chain, we know from the DPI that $I(z_j;\boldsymbol{x}) \leq I(z_j;y_i)$. On the other hand, since z_j is a representation of $\boldsymbol{x}$, we have the Markov chain $(\boldsymbol{n},y_i) \leftrightarrow \boldsymbol{x} \leftrightarrow z_j$. Equivalently, we have that $I(z_j;\boldsymbol{n},y_i) \leq I(z_j;\boldsymbol{x})$. The chain rule of mutual information tells us that $I(z_j;\boldsymbol{n}|y_i) = I(z_j;\boldsymbol{n},y_i) - I(z_j|y_i)$. According to the above two points, we are left with $I(z_j;\boldsymbol{n}|y_i) \leq I(z_j;\boldsymbol{x}) - I(z_j|y_i)$. Therefore, because of the non-negativity of mutual information, it is only possible that $I(z_j;\boldsymbol{n}|y_i) = 0$ or, equivalently, $\boldsymbol{n} \leftrightarrow y_i \leftrightarrow z_j$.

Finally, we demonstrate that $(\tilde{\boldsymbol{y}}_i \leftrightarrow y_i \leftrightarrow z_j) \wedge (\boldsymbol{n} \leftrightarrow y_i \leftrightarrow z_j) \Longrightarrow (\boldsymbol{x} \leftrightarrow y_i \leftrightarrow z_j)$: We know by the Bayes' theorem that $p(z_j,\boldsymbol{y},\boldsymbol{n}) = p(z_j|\boldsymbol{y},\boldsymbol{n})p(\boldsymbol{n}|\boldsymbol{y})p(\boldsymbol{y}) = p(\boldsymbol{n}|z_j,\boldsymbol{y})p(z_j|\boldsymbol{y})p(\boldsymbol{y}) = p(\boldsymbol{n}|z_j,\boldsymbol{y})p(z_j|y_i)p(\boldsymbol{y})$. Thus, $p(z_j|x) = p(z_j|\boldsymbol{y},\boldsymbol{n}) = p(z_j|y_i)$ if and only if $p(\boldsymbol{n}|\boldsymbol{y}) = p(\boldsymbol{n}|z_j,\boldsymbol{y})$, which is equivalent to having $p(z_j|\boldsymbol{y}) = p(z_j|\boldsymbol{y},\boldsymbol{n}) = p(z_j|\boldsymbol{x})$ and since $p(z_j|\boldsymbol{y}) = p(z_j|y_i)$, we are left with $p(z_j|\boldsymbol{x}) = p(z_j|y_i)$, and thus, $(\boldsymbol{x} \leftrightarrow y_i \leftrightarrow z_j)$. □

Theorem 2. *Let y_i be a factor and $\boldsymbol{z} = \{z_l\}_{l=1}^m$ a representation. Then, z_j is a sufficient representation of y_i if and only if z_j is representations-invariant for y_i and $\boldsymbol{z}$ is explicit for y_i. Equivalently, we have the that:*

$$(\boldsymbol{x} \leftrightarrow z_j \leftrightarrow y_i) \Longleftrightarrow (\tilde{\boldsymbol{z}}_j \leftrightarrow z_j \leftrightarrow y_i) \wedge (\boldsymbol{x} \leftrightarrow \boldsymbol{z} \leftrightarrow y_i)$$

Proof. First, we demonstrate that $(\boldsymbol{x} \leftrightarrow z_j \leftrightarrow y_i) \Longrightarrow (\tilde{\boldsymbol{z}}_j \leftrightarrow z_j \leftrightarrow y_i)$: Given this Markov chain, by the DPI, we have that $I(y_i;\boldsymbol{x}) \leq I(y_i;z_j)$. Since $\boldsymbol{z}$ is a representation of $\boldsymbol{x}$, we also have that $I(y_i;\boldsymbol{z}) \leq I(y_i;\boldsymbol{x})$. Therefore, it follows that $I(y_i;\boldsymbol{z}) \leq I(y_i;z_j)$ On the other hand, by the mutual information chain rule, we have that $I(y_i;\boldsymbol{z}) = I(y_i;z_j,\tilde{\boldsymbol{z}}_{\boldsymbol{j}}) = I(y_i;z_j) + I(y_i;\tilde{\boldsymbol{z}}_{\boldsymbol{j}}|z_j)$. By the non-negativity of mutual information, we are left with $I(y_i;\boldsymbol{z}) \geq I(y_i;z_j)$. Therefore, we have that $I(y_i;z_j) = I(y_i;\boldsymbol{z}) = I(y_i;z_j,\tilde{\boldsymbol{z}}_{\boldsymbol{j}})$ or, equivalently, $\tilde{\boldsymbol{z}}_j \leftrightarrow z_j \leftrightarrow y_i$.

Second, we demonstrate that $(\boldsymbol{x} \leftrightarrow z_j \leftrightarrow y_i) \Longrightarrow (\boldsymbol{x} \leftrightarrow \boldsymbol{z} \leftrightarrow y_i)$: Since z_j is a representation of $\boldsymbol{x}$, we have that $p(y_i|\boldsymbol{x},z_j) = p(y_i|\boldsymbol{x})$. Moreover, since z_j is sufficient, we have that $p(y_i|\boldsymbol{x},z_j) = p(y_i|z_j)$. Putting the above two terms together, we are left with $p(y_i|\boldsymbol{x}) = p(y_i|z_j)$. As we have already shown, if z_j is sufficient, then we have the Markov chain $(\tilde{\boldsymbol{z}}_j \leftrightarrow z_j \leftrightarrow y_i)$, so $p(y_i|\boldsymbol{z}) = p(y_i|z_j)$. Therefore, we are left with $p(y_i|\boldsymbol{z}) = p(y_i|\boldsymbol{x})$. Finally, since $\boldsymbol{z}$ is a representation of $\boldsymbol{x}$, we know that $p(y_i|\boldsymbol{x}) = p(y_i|\boldsymbol{x},\boldsymbol{z})$. Putting all of the above together, we are left with $p(y_i|\boldsymbol{x},\boldsymbol{z}) = p(y_i|\boldsymbol{z})$ or, equivalently, $\boldsymbol{x} \leftrightarrow \boldsymbol{z} \leftrightarrow y_i$.

Finally, we demonstrate that $(\tilde{\boldsymbol{z}}_j \leftrightarrow z_j \leftrightarrow y_i) \wedge (\boldsymbol{x} \leftrightarrow \boldsymbol{z} \leftrightarrow y_i) \Longrightarrow (\boldsymbol{x} \leftrightarrow z_j \leftrightarrow y_i)$: Since z_j and $\boldsymbol{z}$ are representations of $\boldsymbol{x}$, we have that $p(y_i|\boldsymbol{x},z_j) = p(y_i|\boldsymbol{x},\boldsymbol{z}) = p(y_i|\boldsymbol{x})$. Furthermore, given the Markov chain $\boldsymbol{x} \leftrightarrow \boldsymbol{z} \leftrightarrow y_i$, we know that $p(y_i|\boldsymbol{x},\boldsymbol{z}) = p(y_i|\boldsymbol{z})$. Equivalently, given $\tilde{\boldsymbol{z}}_j \leftrightarrow z_j \leftrightarrow y_i$, we have that $p(y_i|z_j) = p(y_i|z_j,\tilde{\boldsymbol{z}}_{\boldsymbol{j}}) = p(y_i|\boldsymbol{z})$. Recapitulating the above, we are left with $p(y_i|z_j) = p(y_i|\boldsymbol{z}) = p(y_i|\boldsymbol{x},\boldsymbol{z}) = p(y_i|\boldsymbol{x},z_j)$. Therefore, we have that $\boldsymbol{x} \leftrightarrow z_j \leftrightarrow y_i$. □

B Training Hyperparameters

All training hyperparameters are listed in Table 1.

Table 1: Training hyperparameters.

Hyperparameter	**Value**
Batch size	256
Number of epochs	50
Optimizer	Adam
Base batch size	64
Learning rate	1×10^{-4}
Weight decay	4×10^{-5}
Scheduler	StepLR
Step size	20
Gamma	0.1

C VAE Architectures

All our architectures are based on VAE approaches. Below, we summarize the six architectures considered, with example output shapes provided for a 64×64 RGB input image (C=3).

C.1 Architecture Following Burgess et al. (2018)

The encoder described in Table 2 and the decoder in Table 3 are used.

Table 2: Encoder architecture (Burgess et al., 2018).

Layer	Type	Parameters	Output shape	Activation
Input	–	Image size = (C, H, W)	(B, 3, 64, 64)	–
Conv1	Conv2d	$C \rightarrow 32$, kernel=4, stride=2, pad=1	(B, 32, 32, 32)	ReLU
Conv2	Conv2d	$32 \rightarrow 32$, kernel=4, stride=2, pad=1	(B, 32, 16, 16)	ReLU
Conv3	Conv2d	$32 \rightarrow 32$, kernel=4, stride=2, pad=1	(B, 32, 8, 8)	ReLU
Conv4*	Conv2d	$32 \rightarrow 32$, kernel=4, stride=2, pad=1	(B, 32, 4, 4)	ReLU
Flatten	–	–	(B, 512)	–
FC1	Linear	$512 \rightarrow 256$	(B, 256)	ReLU
FC2	Linear	$256 \rightarrow 256$	(B, 256)	ReLU
FC3	Linear	$256 \rightarrow$ latent_dim $\times$ 2	(B, latent_dim $\times$ 2)	–
Output	Reshape	–	(B, latent_dim, 2)	–

Table 3: Decoder architecture (Burgess et al., 2018).

Layer	Type	Parameters	Output shape	Activation
Input	–	Latent vector $\in \mathbb{R}^{\text{latent_dim}}$	(B, latent_dim)	–
FC1	Linear	latent_dim $\rightarrow 256$	(B, 256)	ReLU
FC2	Linear	$256 \rightarrow 256$	(B, 256)	ReLU
FC3	Linear	$256 \rightarrow 32 \times 4 \times 4$	(B, 512)	ReLU
Reshape	–	–	(B, 32, 4, 4)	–
ConvT_64*	ConvTranspose2d	$32 \rightarrow 32$, kernel=4, stride=2, pad=1	(B, 32, 8, 8)	ReLU
ConvT1	ConvTranspose2d	$32 \rightarrow 32$, kernel=4, stride=2, pad=1	(B, 32, 16, 16)	ReLU
ConvT2	ConvTranspose2d	$32 \rightarrow 32$, kernel=4, stride=2, pad=1	(B, 32, 32, 32)	ReLU
ConvT3	ConvTranspose2d	$32 \rightarrow C$, kernel=4, stride=2, pad=1	(B, C, 64, 64)	Sigmoid

C.2 Architecture Following Chen et al. (2018)

The encoder described in Table 4 and the decoder in Table 5 are used.

Table 4: MLP encoder architecture (Chen et al., 2018).

Layer	Type	Parameters	Output shape	Activation
Input	–	Image flattened	(B, $64 \times 64 \times 3 = 12288$)	–
FC1	Linear	$12288 \to 1200$	(B, 1200)	ReLU
FC2	Linear	$1200 \to 1200$	(B, 1200)	ReLU
FC3	Linear	$1200 \to$ latent_dim $\times 2$	(B, latent_dim $\times$ 2)	–
Output	Reshape	–	(B, latent_dim, 2)	–

Table 5: MLP decoder architecture (Chen et al., 2018).

Layer	Type	Parameters	Output shape	Activation
Input	–	Latent vector $\in \mathbb{R}^{\text{latent_dim}}$	(B, latent_dim)	–
FC1	Linear	latent_dim $\to 1200$	(B, 1200)	Tanh
FC2	Linear	$1200 \to 1200$	(B, 1200)	Tanh
FC3	Linear	$1200 \to 1200$	(B, 1200)	Tanh
FC4	Linear	$1200 \to 64 \times 64 \times 3 = 12288$	(B, 12288)	Sigmoid
Output	Reshape	–	(B, 3, 64, 64)	–

C.3 Architecture Following Locatello et al. (2020)

The encoder described in Table 6 and the decoder in Table 7 are used.

Table 6: Encoder architecture (Locatello et al., 2020).

Layer	Type	Parameters	Output shape	Activation
Input	–	Image size = (C, 64, 64)	(B, C, 64, 64)	–
Conv1	Conv2d	$C \to 32$, kernel=4, stride=2, pad=1	(B, 32, 32, 32)	ReLU
Conv2	Conv2d	$32 \to 32$, kernel=4, stride=2, pad=1	(B, 32, 16, 16)	ReLU
Conv3	Conv2d	$32 \to 64$, kernel=4, stride=2, pad=1	(B, 64, 8, 8)	ReLU
Conv4	Conv2d	$64 \to 64$, kernel=4, stride=2, pad=1	(B, 64, 4, 4)	ReLU
Flatten	–	–	(B, 1024)	–
FC1	Linear	$1024 \to 256$	(B, 256)	ReLU
FC2	Linear	$256 \to$ latent_dim $\times 2$	(B, latent_dim $\times$ 2)	–
Output	Reshape	–	(B, latent_dim, 2)	–

Table 7: Decoder architecture (Locatello et al., 2020).

Layer	Type	Parameters	Output shape	Activation
Input	–	Latent vector $\in \mathbb{R}^{\text{latent_dim}}$	(B, latent_dim)	–
FC1	Linear	latent_dim $\to 256$	(B, 256)	ReLU
FC2	Linear	$256 \to 64 \times 4 \times 4 = 1024$	(B, 1024)	ReLU
Reshape	–	–	(B, 64, 4, 4)	–
ConvT1	ConvTranspose2d	$64 \to 64$, kernel=4, stride=2, pad=1	(B, 64, 8, 8)	ReLU
ConvT2	ConvTranspose2d	$64 \to 32$, kernel=4, stride=2, pad=1	(B, 32, 16, 16)	ReLU
ConvT3	ConvTranspose2d	$32 \to 32$, kernel=4, stride=2, pad=1	(B, 32, 32, 32)	ReLU
ConvT4	ConvTranspose2d	$32 \to C$, kernel=4, stride=2, pad=1	(B, C, 64, 64)	Sigmoid

C.4 Architecture Combining Locatello et al. (2020) and Watters et al. (2019)

The encoder described in Table 6 and the decoder in Table 8 are used.

Table 8: Spatial Broadcast Decoder (Watters et al., 2019).

Layer	Parameters	Output shape	Activation
Input	Latent vector	(B, latent_dim)	–
Spatial Broadcast	Tile latent + concat XY mesh	(B, latent_dim+2, 64, 64)	–
Conv1	(latent_dim + 2) $\rightarrow$ 64, kernel=5, stride=1, pad=2	(B, 64, 64, 64)	ReLU
Conv2	64 $\rightarrow$ 64, kernel=5, stride=1, pad=2	(B, 64, 64, 64)	ReLU
Conv3	64 $\rightarrow$ 64, kernel=5, stride=1, pad=2	(B, 64, 64, 64)	ReLU
Conv4	64 $\rightarrow$ 64, kernel=5, stride=1, pad=2	(B, 64, 64, 64)	ReLU
Conv5	64 $\rightarrow$ C, kernel=5, stride=1, pad=2	(B, C, 64, 64)	Sigmoid

C.5 Small Architecure Following Montero et al. (2022)

The encoder described in Table 9 and the decoder in Table 10 are used.

Table 9: Small encoder architecture (Montero et al., 2022).

Layer	Type	Parameters	Output shape	Activation
Input	–	Image size = (C, 64, 64)	(B, C, 64, 64)	–
Conv1	Conv2d	$C \rightarrow 32$, kernel=4, stride=2, pad=1	(B, 32, 32, 32)	ReLU
Conv2	Conv2d	32 $\rightarrow$ 32, kernel=4, stride=2, pad=1	(B, 32, 16, 16)	ReLU
Conv3	Conv2d	32 $\rightarrow$ 64, kernel=4, stride=2, pad=1	(B, 64, 8, 8)	ReLU
Conv4	Conv2d	64 $\rightarrow$ 64, kernel=4, stride=2, pad=1	(B, 64, 4, 4)	ReLU
Conv5	Conv2d	64 $\rightarrow$ 128, kernel=4, stride=2, pad=1	(B, 128, 2, 2)	ReLU
Flatten	–	–	(B, 512)	–
FC1	Linear	512 $\rightarrow$ 256	(B, 256)	ReLU
FC2	Linear	256 $\rightarrow$ latent_dim $\times$ 2	(B, latent_dim $\times$ 2)	–
Output	Reshape	–	(B, latent_dim, 2)	–

Table 10: Small decoder architecture (Montero et al., 2022).

Layer	Type	Parameters	Output shape	Activation
Input	–	Latent vector $\in \mathbb{R}^{\text{latent_dim}}$	(B, latent_dim)	–
FC1	Linear	latent_dim $\rightarrow$ 256	(B, 256)	ReLU
FC2	Linear	$256 \rightarrow 128 \times 2 \times 2 = 512$	(B, 512)	ReLU
Reshape	–	–	(B, 128, 2, 2)	–
ConvT1	ConvTranspose2d	128 $\rightarrow$ 64, kernel=4, stride=2, pad=1	(B, 64, 4, 4)	ReLU
ConvT2	ConvTranspose2d	64 $\rightarrow$ 64, kernel=4, stride=2, pad=1	(B, 64, 8, 8)	ReLU
ConvT3	ConvTranspose2d	64 $\rightarrow$ 32, kernel=4, stride=2, pad=1	(B, 32, 16, 16)	ReLU
ConvT4	ConvTranspose2d	32 $\rightarrow$ 32, kernel=4, stride=2, pad=1	(B, 32, 32, 32)	ReLU
ConvT5	ConvTranspose2d	32 $\rightarrow$ C, kernel=4, stride=2, pad=1	(B, C, 64, 64)	Sigmoid

C.6 Large Architecure Following Montero et al. (2022)

The encoder described in Table 11 and the decoder in Table 12 are used.

Table 11: Large encoder architecture (Montero et al., 2022)).

Layer	Type	Parameters	Output shape	Activation
Input	–	Image size = (C, 64, 64)	(B, C, 64, 64)	–
Conv1	Conv2d	$C \to 64$, kernel=4, stride=2, pad=1	(B, 64, 32, 32)	ReLU
Conv2	Conv2d	$64 \to 64$, kernel=4, stride=2, pad=1	(B, 64, 16, 16)	ReLU
Conv3	Conv2d	$64 \to 128$, kernel=4, stride=2, pad=1	(B, 128, 8, 8)	ReLU
Conv4	Conv2d	$128 \to 128$, kernel=4, stride=2, pad=1	(B, 128, 4, 4)	ReLU
Conv5	Conv2d	$128 \to 256$, kernel=4, stride=2, pad=1	(B, 256, 2, 2)	ReLU
Flatten	–	–	(B, 1024)	–
FC1	Linear	$1024 \to 256$	(B, 256)	ReLU
FC2	Linear	$256 \to$ latent_dim $\times$ 2	(B, latent_dim $\times$ 2)	–
Output	Reshape	–	(B, latent_dim, 2)	–

Table 12: Large decoder architecture (Montero et al., 2022).

Layer	Type	Parameters	Output shape	Activation
Input	–	Latent vector $\in \mathbb{R}^{\text{latent_dim}}$	(B, latent_dim)	–
FC1	Linear	latent_dim $\to 256$	(B, 256)	ReLU
FC2	Linear	$256 \to 256 \times 2 \times 2 = 1024$	(B, 1024)	ReLU
Reshape	–	–	(B, 256, 2, 2)	–
ConvT1	ConvTranspose2d	$256 \to 128$, kernel=4, stride=2, pad=1	(B, 128, 4, 4)	ReLU
ConvT2	ConvTranspose2d	$128 \to 128$, kernel=4, stride=2, pad=1	(B, 128, 8, 8)	ReLU
ConvT3	ConvTranspose2d	$128 \to 64$, kernel=4, stride=2, pad=1	(B, 64, 16, 16)	ReLU
ConvT4	ConvTranspose2d	$64 \to 64$, kernel=4, stride=2, pad=1	(B, 64, 32, 32)	ReLU
ConvT5	ConvTranspose2d	$64 \to C$, kernel=4, stride=2, pad=1	(B, C, 64, 64)	Sigmoid