

GENERATION IS REQUIRED FOR DATA-EFFICIENT PERCEPTION

Anonymous authors

Paper under double-blind review

ABSTRACT

Visual perception in the human brain is often thought to result from inverting a generative *decoder* that maps latents to images. In contrast, today’s most successful vision models are non-generative, relying on an *encoder* that maps images to latents without inverting an image decoder. This raises the question of whether generation is required for machines to achieve human-level visual perception. In this work, we approach this question from the perspective of data efficiency, a core feature of human perception. Specifically, we investigate whether *compositional generalization* is achievable, both in theory and practice, using generative and non-generative methods. We first formalize the inductive biases required to guarantee compositional generalization in generative (decoder-based) and non-generative (encoder-based) methods. We then provide theoretical results suggesting that such inductive biases cannot be enforced on an encoder through practical means such as regularization or architectural constraints. In contrast, we show that enforcing the inductive biases on a decoder is straightforward, enabling compositional generalization through inverting the decoder. We highlight how this inversion can be performed efficiently, either online through gradient-based search or offline through generative replay. Empirically, we train a range of non-generative methods on photorealistic image datasets, finding they often fail to generalize compositionally and require large-scale pretraining to improve generalization. By comparison, generative methods yield significant improvements in compositional generalization, without requiring additional data, by leveraging suitable inductive biases on a decoder along with search and replay.

1 INTRODUCTION

Perceiving the visual world requires forming internal representations of sensory input. Two opposing views exist for how these representations should be acquired. The **generative view** posits that representations result from inverting an internal generative model, or *decoder*, to infer the latent variables that give rise to the visual input (Friston and Stephan, 2007; Hinton, 2007; Olshausen, 2014; von Helmholtz, 1867). Conversely, the **non-generative view** holds that representations result from a feedforward *encoder* that directly maps visual input to latent variables without inverting a decoder (Gibson, 1979; LeCun, 2022; Yamins et al., 2014). A core problem in AI is to understand which of these paradigms should be adopted to build machines with *human-level visual perception*.

In recent years, consensus around this problem has shifted, following breakthroughs in non-generative methods for representation learning (Caron et al., 2021; Oquab et al., 2024; Radford et al., 2021; Tschannen et al., 2025). These methods, trained with self- or weak supervision, now enable unprecedented performance on perceptual tasks such as object recognition (Siméoni et al., 2025) and image captioning (Beyer et al., 2024; Fan et al., 2025). This progress has given rise to a common assumption that non-generative methods provide the most promising path toward human-level visual perception, while generative approaches are not necessary (Balestriero and LeCun, 2024).

Yet, despite their remarkable performance, current non-generative methods fall short in another key pillar of human visual perception: *data efficiency*. Specifically, these methods rely on web-scale datasets in which different visual concepts are encountered across diverse contexts, with high frequency (Udandarao et al., 2024), and often with language supervision (Zhang et al., 2024). In

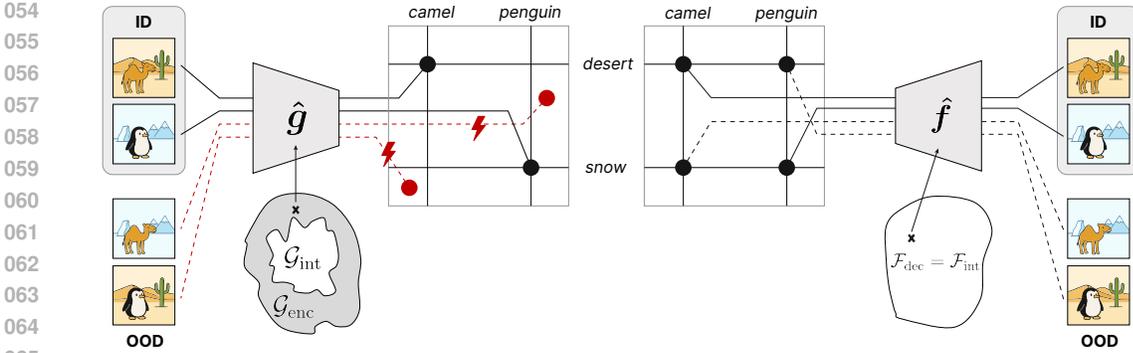


Figure 1: **Generative vs. non-generative compositional generalization.** We assume in-domain (ID) and out-of-domain (OOD) images arise from a latent variable model through an unknown generator $\mathbf{f} \in \mathcal{F}_{int}$, with inverse $\mathbf{g} \in \mathcal{G}_{int}$. Guaranteeing compositional generalization for a generative approach requires constraining a decoder such that $\hat{\mathbf{f}} \in \mathcal{F}_{int}$, and for a non-generative approach, an encoder such that $\hat{\mathbf{g}} \in \mathcal{G}_{int}$ (Sec. 2). We show theoretically in (Sec. 3) that placing such constraints on an encoder is generally infeasible with practical approaches while for a decoder it is straightforward. Empirically, this tends to manifest in an encoder yielding incorrect representations for OOD images (Sec. 5.2). In contrast, a decoder is able to correctly generate such images enabling compositional generalization through inversion (Sec. 4, 5.2).

contrast, human children achieve robust visual perception through much more constrained data, encountering concepts only a handful of times, mainly in the same settings (e.g., the home), and with little supervision (Lake et al., 2017; Tenenbaum et al., 2011). To reach this level of data efficiency, it has been conjectured across several disciplines (Kilbertus et al., 2018; Lake et al., 2015; Peters et al., 2024) that a generative approach may be necessary. This raises a key question: Can non-generative approaches to perception also achieve human-level data efficiency, or is generation required?

In this work, we approach this question by analyzing, both theoretically and empirically, whether generative and non-generative methods can achieve *compositional generalization* (Fodor and Pylyshyn, 1988; Greff et al., 2020). Compositional generalization is the ability to perceive out-of-domain (OOD) scenes containing unseen combinations of concepts, e.g., a dog in the park after only seeing dogs in the house. It is thus essential for realizing the data efficiency of human perception.

Structure and Contributions. We build upon Brady et al. (2025) to formalize the constraints required to guarantee compositional generalization for both generative (decoder-based) and non-generative (encoder-based) approaches. In Sec. 3, we show theoretically that enforcing such constraints on encoders is generally infeasible, as they depend on the geometry of out-of-domain regions of the data manifold, which is unknown. In contrast, we show that for generative models the constraints are not data-dependent and can be imposed directly through regularization or architectural design. These results suggest that inversion of a decoder is necessary to guarantee compositional generalization. In Sec. 4, we describe how such inversion can be implemented efficiently: in-distribution via an autoencoder, and out-of-distribution via gradient-based search (Sec. 4.1) and generative replay (Sec. 4.2). Finally, in Sec. 5, we empirically evaluate compositional generalization using photorealistic image data containing concepts such as animals and backgrounds (Bordes et al., 2023). We find that non-generative models frequently fail to generalize compositionally on this data, requiring large-scale pretraining to succeed (Sec. 5.2). In contrast, generative methods leveraging search and replay achieve substantial gains in OOD performance without requiring additional data.

2 PROBLEM SETUP

Perception. We begin by formalizing *visual perception*. To this end, we assume that images $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^{d_x}$ arise from a latent variable model. Specifically, we assume \mathbf{x} is generated from a latent vector $\mathbf{z} \in \mathcal{Z} := \mathbb{R}^{d_z}$ by a diffeomorphic generator $\mathbf{f} : \mathcal{Z} \rightarrow \mathcal{X}$, i.e., $\mathbf{x} = \mathbf{f}(\mathbf{z})$. Visual concepts in \mathbf{x} (e.g. “camel” and “desert” in Fig. 1) are modelled as K distinct *slots* of latents $z_k \in \mathbb{R}^m$ such that $\mathbf{z} = (z_1, \dots, z_K)$ (Brady et al., 2025). Now, assume we have a representation of an image $\hat{\mathbf{z}} = \phi(\mathbf{x})$, where $\phi : \mathbb{R}^{d_x} \rightarrow \mathcal{Z}$. We define perception as the ability to invert the generator \mathbf{f} via ϕ to recover the slots z_k that generated \mathbf{x} . In general, recovering z_k exactly is impossible. Thus,

we only require that ϕ inverts f up to re-parameterizations and permutations of the slots. Formally, let h_π be a function composed of slot-wise bijections $h_k : \mathbb{R}^m \rightarrow \mathbb{R}^m$ and permutations π , i.e., $h_\pi(z) := \{h_k(z_{\pi(k)})\}_{k=1}^K$. Perception on a set $\mathcal{Z}^S \subseteq \mathcal{Z}$ requires that there exist an h_π such that

$$\forall z \in \mathcal{Z}^S, \phi(f(z)) = h_\pi(z). \quad (2.1)$$

Eq. (2.1) takes the perspective of perception as an inverse problem (Tenenbaum et al., 2011), but with respect to the ground-truth generator f . This contrasts with a task-based view (Yamins and DiCarlo, 2016) where perception is defined with respect to solving a downstream task. We note that the task-based view can be framed as a special case of Eq. (2.1), by treating task-specific predictions as the latent variables to be recovered by ϕ . Moreover, if a representation satisfying Eq. (2.1) is learned, downstream tasks such as object classification can be solved via a simple readout applied independently to each inferred slot \hat{z}_k (see Sec. 5).

Generative and non-generative approaches. Using Eq. (2.1), we now characterize the *generative* and *non-generative* approaches to perception. For the generative approach, representations are obtained by inverting a learned *decoder* $\hat{f} : \mathcal{Z} \rightarrow \mathbb{R}^{d_x}$, i.e., $\phi(x) = \hat{f}^{-1}(x)$. For this to satisfy Eq. (2.1), the decoder \hat{f} must *identify* the ground-truth generator f such that for $z \in \mathcal{Z}$

$$\hat{f}(h_\pi(z)) = f(z). \quad (2.2)$$

Alternatively, for the non-generative approach, a representation is defined as $\phi(x) = \hat{g}(x)$, where $\hat{g} : \mathbb{R}^{d_x} \rightarrow \mathcal{Z}$ is a learned *encoder* which is *not* constructed to invert a decoder \hat{f} . For this to satisfy Eq. (2.1), \hat{g} must identify the inverse generator $g := f^{-1}$ such that for $x \in \mathcal{X}$

$$\hat{g}(x) = h_\pi(g(x)). \quad (2.3)$$

We emphasize that the difference between the generative and non-generative approaches is not whether an encoder or decoder is used. **Instead, it is whether a representation satisfying Eq. (2.1) is obtained by learning an approximation \hat{f} of the ground-truth generator (Eq. (2.2)) and then inverting this model, or by learning an approximation \hat{g} of the inverse generator directly (Eq. (2.3)).**

Compositional generalization. We now formalize *compositional generalization*. Informally, compositional generalization is the ability to perceive out-of-domain images containing unseen concept combinations (e.g. ‘‘penguin’’ and ‘‘desert’’ in Fig. 1). To formalize this, we assume observed images $\mathcal{X}_{\text{ID}} \subset \mathcal{X}$ arise from only a subset of possible concept combinations $\mathcal{Z}_{\text{ID}} \subset \mathcal{Z}$, i.e., $\mathcal{X}_{\text{ID}} := f(\mathcal{Z}_{\text{ID}})$ (see Fig. 2). OOD concept combinations \mathcal{Z}_{OOD} are defined as the set of all unseen combinations of slots

$$\mathcal{Z}_{\text{OOD}} := \{\mathcal{Z}_1 \times \mathcal{Z}_2 \times \dots \times \mathcal{Z}_K\} \setminus \mathcal{Z}_{\text{ID}} \quad \text{with} \quad \mathcal{Z}_k := \{z_k \in \mathbb{R}^m \mid z \in \mathcal{Z}_{\text{ID}}\}, \quad (2.4)$$

which give rise to OOD images $\mathcal{X}_{\text{OOD}} := f(\mathcal{Z}_{\text{OOD}})$ (Fig. 2). Compositional generalization is then achieved if Eq. (2.1) is satisfied both in-domain, for $z \in \mathcal{Z}_{\text{ID}}$, and out-of-domain, for all $z \in \mathcal{Z}_{\text{OOD}}$.

The problem of identifiability. Compositional generalization, using a generative or non-generative approach, requires identifying the ground-truth generator f or its inverse g , both in-domain and out-of-domain. In-domain identifiability is a well-studied problem (Hyvärinen et al., 2023). It can be solved using both approaches by leveraging observed images $x \in \mathcal{X}_{\text{ID}}$ together with self- (Gresele et al., 2021; von Kügelgen et al., 2021; Zimmermann et al., 2021) or weakly-supervised information (Hyvärinen and Morioka, 2016; Khemakhem et al., 2020; Locatello et al., 2020a) about the data-generating process. Out-of-domain identifiability, however, presents a different challenge: because $x \in \mathcal{X}_{\text{OOD}}$ is unobserved, the strategies above cannot be applied. Consequently, out-of-domain identifiability must be implied by in-domain identifiability (Wiedemer et al., 2024b). This is only possible if f belongs to a function class \mathcal{F} such that for all $f^1, f^2 \in \mathcal{F}$

$$\forall z \in \mathcal{Z}_{\text{ID}}, f^1(h_\pi(z)) = f^2(z) \implies \forall z \in \mathcal{Z}_{\text{OOD}}, f^1(h_\pi(z)) = f^2(z), \quad (2.5)$$

which equivalently implies that for all inverses $g^1, g^2 \in \mathcal{G} := \{f^{-1} \mid f \in \mathcal{F}\}$

$$\forall x \in \mathcal{X}_{\text{ID}}, h_\pi(g^1(x)) = g^2(x) \implies \forall x \in \mathcal{X}_{\text{OOD}}, h_\pi(g^1(x)) = g^2(x). \quad (2.6)$$

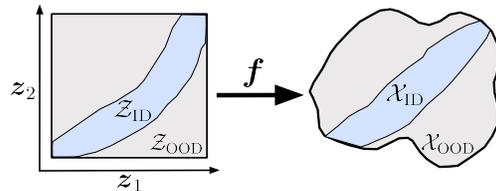


Figure 2: Visualization of a data generating process with in- and out-of-domain regions.

If these implications do not hold then the problem is *non-identifiable*, since there is no way to distinguish between f^1 and f^2 or g^1 and g^2 from observed data.

Further assumptions on f . Under our current assumptions, the ground-truth generator can be any diffeomorphism from \mathcal{Z} to \mathcal{X} . This function class is far too large to satisfy Eq. (2.5). Thus, further assumptions on f are required. Recently, Brady et al. (2025, Thm. 4.4) proved that diffeomorphisms (on their image) with the following form will satisfy Eq. (2.5) (when f is sufficiently nonlinear)

$$f(z) = \sum_{k=1}^K f^k(z_k) + \sum_{\alpha: |\alpha| \leq n} c_\alpha z^\alpha, \quad (2.7)$$

where $n \in \mathbb{N}$, $c_\alpha \in \mathbb{R}^{d_x}$, and $\alpha \in \mathbb{N}_0^{d_x}$ is a *multi-index*.¹ This function class, denoted \mathcal{F}_{int} , was introduced to model concepts with varying degrees of interaction n . For example, when $n = 1$, the second-sum on the RHS vanishes and concepts can only interact additively (Lachapelle et al., 2023). For $n > 1$ concepts can interact explicitly via polynomial functions of components from different slots. This aims to capture more complex concept interactions such as between objects and backgrounds. Such functions thus offer a flexible model of visual concepts, and are the largest function class shown to enable OOD identifiability (Eq. (2.5)). For these reasons, we assume that ground-truth generators f belong to \mathcal{F}_{int} , and inverse generators g to $\mathcal{G}_{\text{int}} := \{f^{-1} \mid f \in \mathcal{F}_{\text{int}}\}$.

Guaranteeing compositional generalization. We can now formalize what is required to guarantee compositional generalization using both a generative and non-generative approach. To this end, we assume ID identifiability holds for a decoder \hat{f} and encoder \hat{g} , i.e., Equations 2.2 and 2.3 are satisfied in-domain. Since ground-truth generators \mathcal{F}_{int} and inverses \mathcal{G}_{int} satisfy Equations 2.5 and 2.6, OOD identifiability is guaranteed if the decoder class $\hat{f} \in \mathcal{F}_{\text{dec}}$ is constrained to $\mathcal{F}_{\text{dec}} = \mathcal{F}_{\text{int}}$ and similarly if the encoder class $\hat{g} \in \mathcal{G}_{\text{enc}}$ is constrained to $\mathcal{G}_{\text{enc}} = \mathcal{G}_{\text{int}}$. Compositional generalization is thus possible *in theory* for both generative and non-generative approaches. This does not necessarily imply, however, that it can be guaranteed *in practice* for both approaches. Specifically, guaranteeing compositional generalization in practice depends on whether practical means exist to enforce $\hat{f} \in \mathcal{F}_{\text{int}}$ in the generative case and $\hat{g} \in \mathcal{G}_{\text{int}}$ in the non-generative case.

3 THEORETICAL ANALYSIS

In this section, we theoretically analyze the structure of \mathcal{F}_{int} and \mathcal{G}_{int} to understand whether a model can be constrained to these classes with practical means such as regularization or architecture design.

Structure of \mathcal{F}_{int} . Generators in \mathcal{F}_{int} are defined as diffeomorphisms which take the form of Eq. (2.7). Consequently, to enforce $\hat{f} \in \mathcal{F}_{\text{int}}$, we must constrain a decoder to match this form. This can be done in a straightforward manner via architecture design. For example, the first term on the RHS of Eq. (2.7) can be parameterized as the sum of slot-wise neural networks and the second term using learned coefficients for c_α . Furthermore, we highlight that functions of the form in Eq. (2.7) can equivalently be expressed as having block-diagonal derivative matrices $D^{n+1}f(z)$ (Brady et al., 2025; Lachapelle et al., 2023). Specifically, if $n = 1$, then the Hessian D^2f has the structure that for any two slots z_k and z_l ,

$$\forall 1 \leq k \neq l \leq K, \quad D_{z_k} D_{z_l} f(z) = 0. \quad (3.1)$$

For $n > 1$, analogous conditions hold for higher-order derivatives (Brady et al., 2025). Thus, we can also enforce that $\hat{f} \in \mathcal{F}_{\text{int}}$ for a decoder \hat{f} via regularization. For example, when $n = 1$, we can use the following regularizer (with similar expressions for higher-order derivatives when $n > 1$)

$$\mathcal{R}(\hat{f}, z) = \sum_{k \neq l \in [K]} \left\| D_{z_k, z_l}^2 \hat{f}(z) \right\|. \quad (3.2)$$

3.1 STRUCTURE OF \mathcal{G}_{int} .

We now investigate the structure of inverse generators in \mathcal{G}_{int} . For simplicity, we present results for $n = 1$; similar statements can in principle be derived for higher order derivatives for the case $n > 1$.

¹A *multi-index* is an ordered tuple $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_d)$ of non-negative integers $\alpha_i \in \mathbb{N}_0$, with operations $|\alpha| := \alpha_1 + \alpha_2 + \dots + \alpha_d$, and $z^\alpha := z_1^{\alpha_1} z_2^{\alpha_2} \dots z_d^{\alpha_d}$.

We first note that inverse generators $\mathbf{g} \in \mathcal{G}_{\text{int}}$ do not admit an analytical form similar to Eq. (2.7). Thus, understanding their structure requires analyzing finer-grained properties of these functions. To this end, we investigate their derivatives. We also study whether we can find architectures with an inductive bias towards \mathcal{G}_{int} , but delegate this to Appendix A.2 due to space constraints.

We will first assume that the observed dimension d_x equals the ground-truth latent dimension d_z such that $\mathcal{X} = \mathcal{Z}$. In this case, we show that, similar to generators in \mathcal{F}_{int} , inverse generators in \mathcal{G}_{int} have a structured Jacobian and Hessian. Specifically, we prove the following result.

Lemma 3.1. *Let $\mathbf{g} \in \mathcal{G}_{\text{int}}$ for $n = m = 1$ and $d_x = d_z$. Then \mathbf{g} has the property that for $\mathbf{x} \in \mathcal{X}$*

$$(D\mathbf{g})^{-\top}(\mathbf{x})D^2\mathbf{g}_s(\mathbf{x})(D\mathbf{g})^{-1}(\mathbf{x}) \in \text{Diag}(d_x) \quad (3.3)$$

is a diagonal matrix for $s \in [d_z]$. Further, if \mathbf{g} is a diffeomorphism satisfying Eq. (3.3) then $\mathbf{g} \in \mathcal{G}_{\text{int}}$.

Thus, when $\mathcal{X} = \mathcal{Z}$, enforcing that $\hat{\mathbf{g}} \in \mathcal{G}_{\text{int}}$ requires constraining an encoder according to Eq. (3.3). This is achievable, for instance, through regularization on the derivatives of $\hat{\mathbf{g}}$ analogous to Eq. (3.2).

This setting, however, is not applicable to image data since images typically lie in a manifold embedded in a higher-dimensional ambient space \mathbb{R}^{d_x} . We therefore consider the more practical case where d_x is larger than the ground-truth latent dimension d_z . Specifically, we assume $d_x \geq d_z^3$. In this case, we first prove that the aforementioned structure on $D\mathbf{g}$ and $D^2\mathbf{g}$ is no longer present.

Theorem 3.2. *Assume that $d_x \geq d_z^3$. Let $\mathbf{B}_l \in \mathbb{R}^{d_x \times d_x}$ be symmetric matrices for $1 \leq l \leq d_z$. Then there is for any $\mathbf{x}_0 \in \mathbb{R}^{d_x}$ and for almost every $\mathbf{A} \in \mathbb{R}^{d_z \times d_x}$ a generator $\mathbf{f} \in \mathcal{F}_{\text{int}}$ with a (left)-inverse $\mathbf{g} \in \mathcal{G}_{\text{int}}$, such that $\mathbf{f}(0) = \mathbf{x}_0$ and $D\mathbf{g}(\mathbf{x}_0) = \mathbf{A}$ and $D^2\mathbf{g}_l(\mathbf{x}_0) = \mathbf{B}_l$ for $1 \leq l \leq d_z$.*

Thus, when $d_x \gg d_z$, $D^2\mathbf{g}$ and $D\mathbf{g}$ can be arbitrary matrices (up to a set of measure 0). We emphasize that this result applies to \mathcal{F}_{int} with arbitrary interaction degree $n \geq 1$ and any slot dimensions. However, the structure expressed in Eq. (3.3) does not vanish entirely from \mathbf{g} . Instead, it persists, but only for the restriction of \mathbf{g} to the data manifold \mathcal{X} . Specifically, the constraint Eq. (3.3) holds more generally for $n = m = 1$ when $D\mathbf{g}$ is projected on the tangent space $T_{\mathbf{x}}\mathcal{X}$ of the data manifold, i.e.,

$$\left((D\mathbf{g}(\mathbf{x})\Pi_{T_{\mathbf{x}}\mathcal{X}})^+ \right)^\top (\mathbf{z})D^2\mathbf{g}_s(\mathbf{x})(D\mathbf{g}(\mathbf{x})\Pi_{T_{\mathbf{x}}\mathcal{X}})^+ \in \text{Diag}(d_z) \quad (3.4)$$

where $\Pi_{T_{\mathbf{x}}\mathcal{X}}$ denotes the orthogonal projection on the tangent space (see Lemma A.4 for details).

Constraining an encoder such that $\hat{\mathbf{g}} \in \mathcal{G}_{\text{int}}$ thus requires enforcing this structure on $\hat{\mathbf{g}}$. This is challenging because the constraints depend on the geometry of the data manifold \mathcal{X} . Enforcing such constraints is thus not only impractical but also ill-posed since the geometry of out-of-domain regions $\mathcal{X}_{\text{OOD}} \subset \mathcal{X}$ is unobserved. This suggests that constraining an encoder through approaches such as architectural design or regularization is infeasible, as any such method would necessarily be data-dependent as well as implicitly assume knowledge of \mathcal{X}_{OOD} .

We contrast this with the reverse direction for $\mathbf{f} \in \mathcal{F}_{\text{int}}$. In this case, the structure to be enforced (see Eq. (3.1)) is not manifold-dependent but is always aligned with the global coordinate axes (Fig. 3, right). This allows for a universal procedure to constrain a decoder to \mathcal{F}_{int} , rather than a manifold-dependent one (Fig. 3, left). Moreover, such constraints can also be applied in OOD regions, since the manifold \mathcal{Z}_{ID} extends in a Cartesian fashion and its structure is therefore known.

Special case of $n = 0$. We briefly discuss the case of functions in \mathcal{F}_{int} when $n = 0$. These functions, introduced by Brady et al. (2023), are a special case of $n = 1$ with the additional, more restrictive condition $|D_{z_k}\mathbf{f}_i(\mathbf{z})| \cdot |D_{z_l}\mathbf{f}_i(\mathbf{z})| = 0$ for each $i \in [d_x]$. In other words, each pixel i depends only on a single slot and no interactions (such as occlusions) between objects are possible. In this case, we can find a left inverse \mathbf{g} of $\mathbf{f} \in \mathcal{F}_{\text{int}}^{n=0}$ (for any $d_x \geq d_z$) whose Jacobian satisfies the sparsity constraint $|D_{x_i}\mathbf{g}_k| \cdot |D_{x_j}\mathbf{g}_l| = 0$ for $l \neq k$. This additional structure can thus be leveraged to restrict \mathcal{G}_{enc} . However, this remains challenging in practice because the sparsity pattern (i.e., which slots z_l depends on which pixel x_i) is not known a-priori. In Section 5, we study whether concepts satisfying $n = 0$ can empirically enable compositional generalization for non-generative approaches.

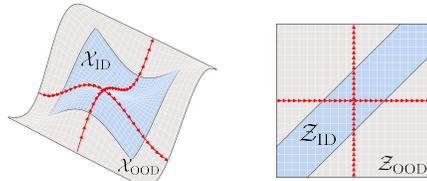


Figure 3: Structure of a data manifold \mathcal{X} and latent manifold \mathcal{Z} .

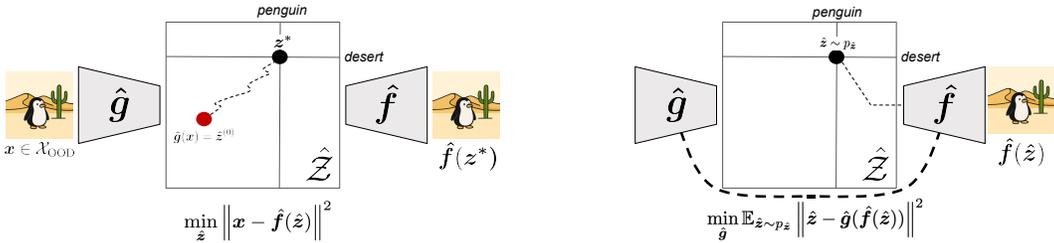


Figure 4: **Approaches for inverting a generator out-of-domain.** *Left.* Visualization of gradient-based search to invert a decoder \hat{f} out-of-domain, with initialization given by an encoder \hat{g} . *Right.* Visualization of generative replay in which an encoder is trained on OOD images generated by a decoder.

Takeaways. Our results suggest that it is generally not feasible to constrain an encoder class such that $\mathcal{G}_{\text{enc}} = \mathcal{G}_{\text{int}}$. This means that there can exist encoders $\hat{g} \in \mathcal{G}_{\text{enc}}$ that satisfy Eq. (2.3) in-domain, but fail to generalize this behavior OOD (see Fig. 1). Consequently, for non-generative methods, whether compositional generalization occurs depends on whether the optimization process happens to avoid converging to such a solution. In contrast, generative methods can avoid such solutions by construction through appropriate inductive biases on a decoder. In Sec. 5, we investigate the practical consequences of this asymmetry on compositional generalization for both classes of methods.

4 SEARCH AND REPLAY

Our results in Sec. 3 suggest that guaranteeing compositional generalization requires a generative approach, i.e., inverting a learned decoder \hat{f} . If a decoder admits an explicit inverse, this inversion is trivial. For image data, however, constructing such a decoder is challenging as this generally requires that $\mathcal{X} = \mathbb{R}^{d_x}$ (Papamakarios et al., 2021). Consequently, inverting \hat{f} requires solving an inference problem: given an image x , we must find a latent z^* such that

$$x = \hat{f}(z^*). \tag{4.1}$$

In this section, we explore strategies for solving this inference problem efficiently.

Inversion on \mathcal{X}_{ID} . For in-domain images, i.e. $x \in \mathcal{X}_{\text{ID}}$, inverting a decoder \hat{f} to obtain z^* can be done directly by training an *autoencoder*. Specifically, we can leverage an encoder \hat{g} to invert \hat{f} in-domain by minimizing the reconstruction objective

$$\min_{\hat{f}, \hat{g}} \mathbb{E}_{x \sim \mathcal{X}_{\text{ID}}} \left\| x - \hat{f}(\hat{g}(x)) \right\|^2. \tag{4.2}$$

Thus, for images $x \in \mathcal{X}_{\text{ID}}$, z^* (Eq. (4.1)) can be obtained directly as the output of the encoder. For out-of-domain images, however, minimizing Eq. (4.2) is not an option since $x \in \mathcal{X}_{\text{OOD}}$ is unobserved. Thus, to efficiently solve Eq. (4.1) on \mathcal{X}_{OOD} , other strategies are required. We explore two such strategies: gradient-based search (Sec. 4.1) and generative replay (Sec. 4.2).

4.1 GRADIENT-BASED SEARCH.

We note that the inference problem in Eq. (4.1) can be expressed as an optimization problem, i.e.,

$$z^* = \arg \min_{\hat{z}} \left\| x - \hat{f}(\hat{z}) \right\|^2. \tag{4.3}$$

Thus, for OOD images $x \in \mathcal{X}_{\text{OOD}}$, we can recover z^* *online* by solving Eq. (4.3) using *gradient-based optimization*. The efficiency of this, however, depends on the initialization $\hat{z}^{(0)}$. If $\hat{z}^{(0)}$ is far from the optimum, many gradient steps are required, leading to slow or suboptimal convergence. To mitigate this, we can leverage the encoder trained on \mathcal{X}_{ID} to provide an initial prediction for z^* such that $\hat{z}^{(0)} = \hat{g}(x)$ and then optimize Eq. (4.3) (see Fig. 4, left). Intuitively, the encoder gives a fast “System 1” guess that constrains the space for slower, “System 2” reasoning (Kahneman, 2011; Prabhudesai et al., 2023a), where “reasoning” corresponds to gradient-based search (LeCun, 2022) (for pseudocode, see Fig. 8, top).

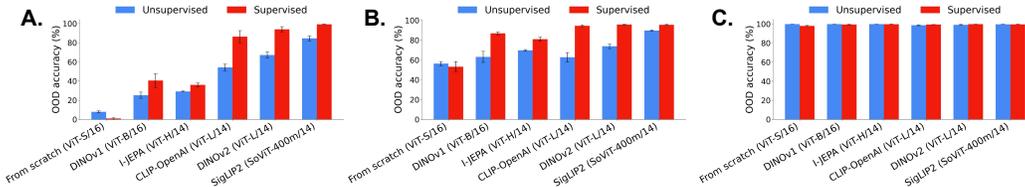


Figure 5: **OOD performance for non-generative methods.** We report OOD performance across three dataset splits for non-generative methods trained with and without supervision and with differing base encoders. On PUG-Background (A.), we see that strong OOD performance generally emerges only for base encoders with large scale pretraining such as SigLIP2 and is otherwise poor. We see a similar trend on PUG-Texture (B.) though OOD performance is generally higher across models. On PUG-Object (C.), concepts do not interact, such that \mathcal{G}_{int} is more constrained (Sec. 3.1). This structure is sufficient for all models to generalize OOD.

4.2 GENERATIVE REPLAY

For out-of-domain images, Eq. (4.1) can also be solved in an *offline* manner by leveraging *generative replay* (Kurth-Nelson et al., 2023; Schwartenbeck et al., 2023). Recall that images $x \in \mathcal{X}_{\text{OOD}}$ are generated by f as combinations of ground-truth slots z_k . Since the decoder \hat{f} identifies f up to slot-wise transformations, images $x \in \mathcal{X}_{\text{OOD}}$ can likewise be generated by re-combining inferred slots \hat{z}_k . Concretely, this can be achieved by sampling a latent \hat{z} from a distribution $p_{\hat{z}}$ with independent slot-wise marginals and decoding them with \hat{f} such that $\hat{f}(\hat{z}) \in \mathcal{X}_{\text{OOD}}$. We can then solve Eq. (4.1) out-of-domain by training an encoder \hat{g} on these samples such that $\hat{g}(\hat{f}(\hat{z})) = \hat{z}$ (see Fig. 4, right). This is captured by the following objective function (Wiedemer et al., 2024a) (see Fig. 8, bottom)

$$\min_{\hat{g}} \mathbb{E}_{\hat{z} \sim p_{\hat{z}}} \left\| \hat{z} - \hat{g}(\hat{f}(\hat{z})) \right\|^2. \quad (4.4)$$

5 EXPERIMENTS

In this section, we conduct an experimental study with two main components. First, we aim to assess the extent to which non-generative methods can achieve compositional generalization in practice without enforcing explicit constraints to this end. Second, we evaluate whether generative methods, which leverage search (Sec. 4.1) and replay (Sec. 4.2), can achieve superior compositional generalization. We describe our experimental setup below, further details can be found in App. B.

5.1 SETUP

Data. We are interested in evaluating compositional generalization for images in realistic settings. This is challenging, however, since web-scale image datasets do not provide explicit controllability over in- and out-of-domain regions. To address this, we leverage the PUG datasets (Bordes et al., 2023), which offer photorealistic images while remaining explicitly controllable. The images we consider are defined by a background and one or two animals, which can take on 10 and 32 different values, respectively. In addition, animals can vary in position and texture.

Using this dataset, we construct three different in- and out-of-domain splits (see Fig. 7, left). In *PUG-Background*, \mathcal{X}_{OOD} contains unseen combinations of animals and backgrounds. In *PUG-Texture*, \mathcal{X}_{OOD} contains unseen combinations of animals and textures. Finally, in *PUG-Object*, \mathcal{X}_{OOD} contains unseen combinations of animals. In this case, animals never occlude each other and therefore do not interact, meaning that concepts satisfy $n = 0$.

Evaluating compositional generalization. To evaluate compositional generalization on this data, we assume a model gives inferred latent slots \hat{z}_k . Each slot should encode either one of the two animals or the background, both on \mathcal{X}_{ID} and \mathcal{X}_{OOD} . To test this, we train a slot-wise readout in-domain to predict the category of the corresponding animal or background. We then report out-of-domain accuracy for these predictions.

Encoders. We consider encoder architectures with the following structure. Images are first divided into patches and processed by a *base encoder*, which produces a set of embeddings. These embeddings are then mapped to slots by a *slot encoder* (see Fig. 7, right). We implement the base

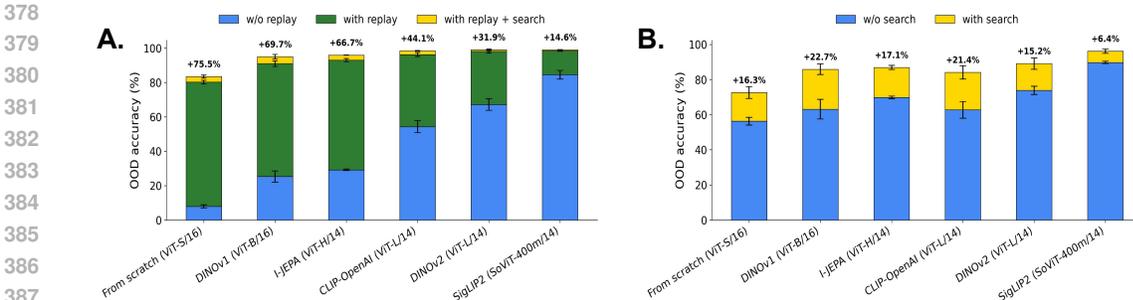


Figure 6: **OOD performance for generative methods.** We report OOD performance on PUG-Background and PUG-Texture for unsupervised autoencoders which leverage replay (Sec. 4.2) and search (Sec. 4.1) trained with differing base encoders. On PUG-Background (A.), we observe a significant increase in OOD performance using replay and additional gains through search. On PUG-Texture (B.) we also see a noticeable increase in OOD performance when using search.

encoder using a Vision Transformer (ViT) (Dosovitskiy et al., 2020), while the slot encoder is either a Transformer (Vaswani et al., 2017) or a Slot Attention module (Locatello et al., 2020b).

Ideally, we would train encoders from scratch using state-of-the-art non-generative methods. However, such methods rely on large-scale datasets, while our datasets are comparatively small (~ 20000 images). We thus leverage pretrained models. Concretely, for the base encoder, we use DINOv1 (Caron et al., 2021), I-JEPA (Assran et al., 2023), DINOv2 (Oquab et al., 2024), CLIP (Radford et al., 2021), and SigLIP2 (Tschannen et al., 2025). These models are optionally fine-tuned using a LoRA adapter (Hu et al., 2022), while the slot encoder is always trained from scratch. We note that the PUG datasets were not contained in the pretraining set for these models, thus data contamination is not an issue. Finally, we also include a ViT-Small base encoder trained from scratch.

Decoders. Brady et al. (2025) argued that constraining a decoder to \mathcal{F}_{int} can be done approximately using a regularized cross-attention Transformer. In this model, pixels query slots, and a regularization term is applied to the resulting attention weights to encourage pixels to specialize to a single slot. This model is also sufficiently flexible to capture complex images and concepts with varying degrees of interaction. For these reasons, we leverage such decoders in our experiments. In § C, we also report results when using unstructured decoders which are not designed to match \mathcal{F}_{int} .

Training objectives. To learn a representation \hat{z} , we train non-generative methods using both a supervised and unsupervised objective. In the supervised setting, the encoder is trained on \mathcal{X}_{ID} to predict the animal and background categories using a cross-entropy loss. In the unsupervised case, we train a variational autoencoder (VAE) (Kingma and Welling, 2014) with our regularized decoder architecture. This case is nevertheless non-generative since the encoder is only constructed to invert the decoder on \mathcal{X}_{ID} , and not on \mathcal{X}_{OOD} . For our generative methods, we take this learned decoder and invert it on \mathcal{X}_{OOD} using search and replay.

5.2 RESULTS

Non-generative methods. In Fig 5, we evaluate compositional generalization for non-generative methods trained on each PUG split. All methods achieve nearly perfect ID accuracy ($\sim 99\%$), thus we only visualize OOD accuracy. For each base encoder, we report the OOD accuracy obtained with the best-performing combination of slot encoder and fine-tuning choice. In Fig. 5 A. (PUG-Background), base encoders trained from scratch (ViT-Small) or pretrained on relatively small corpora (e.g., DINOv1 on ImageNet) fail to generalize OOD. OOD accuracy improves with encoders leveraging larger-scale pretraining, such as SigLIP2. In Fig. 5 B. (PUG-Texture), we observe a boost in OOD performance across models, though performance remains suboptimal overall. Again, models with larger-scale pretraining exhibit stronger OOD performance.

Finally, in Fig. 5 C., we report results for PUG-Object in which concepts do not interact. This corresponds to the special case of $n = 0$ in Sec. 3 in which \mathcal{G}_{int} is more structured. Although we do not explicitly enforce this structure on the models, they nevertheless achieve near-perfect OOD accuracy, indicating that such structure makes compositional generalization fundamentally easier.

Generative Methods. In Fig. 6, we take the autoencoders trained in Fig. 5 and report OOD accuracy after leveraging replay and search for inverting the decoder. On PUG-Background (Fig. 6 A.), we observe a significant increase in OOD accuracy when training encoders with replay across all models, with further improvement when additionally using search. On PUG-Texture (Fig. 6 B.), replay cannot be applied, since in our setup, slots are designed to capture objects and backgrounds, and therefore cannot be trivially recomposed to generate novel animal–texture combinations. However, leveraging search yields a clear improvement in OOD performance across all models. We do not report results on PUG-Object as all non-generative methods achieve near-perfect OOD performance on this dataset. Thus, further OOD gains through search and replay are not possible.

6 RELATED WORK

Limitations of non-generative methods for compositional generalization. Several empirical studies have shown limitations in compositional generalization for non-generative methods trained using natural language supervision (Assouel et al., 2025; Lewis et al., 2022; Ma et al., 2023; Tong et al., 2024; Yuksekgonul et al., 2022). These works generally posit that poor generalization arises from issues with standard contrastive language-image training objectives. In contrast, our theoretical and empirical contributions suggest that such issues are more fundamental, arising from the structure of the inverses of the unknown generator, i.e., \mathcal{G}_{int} .

Generative approaches for improving generalization. The idea that a generative approach can enable compositional generalization has long been advocated in the cognitive science community (Lake et al., 2015; 2017; Tenenbaum et al., 2011). Empirical realizations of this idea have recently been shown for diffusion models repurposed as classifiers (Jeong et al., 2025; Wang et al., 2025). Further (Prabhudesai et al., 2023a;b), showed that inverting a generative model with mechanisms similar to gradient-based search (Sec. 4.1) improves object-decomposition for OOD images and enhances the robustness of classifiers. Recent work explored training encoder-only models using synthetically generated data similar to Sec. 4.2, showing improvements in representations (Fan et al., 2025; Tian et al., 2023) and compositional generalization (Assouel et al., 2022; Jung et al., 2024; Wiedemer et al., 2024a). Our work provides a theoretical motivation for these approaches by highlighting challenges in achieving compositional generalization using non-generative methods.

Causal and anti-causal learning. Our theoretical contribution relates to ideas in the field of causality. A key heuristic in this area posits that the factorization $P(\text{cause})P(\text{effect}|\text{cause})$ is, in general, less complex than the reverse factorization $P(\text{effect})P(\text{cause}|\text{effect})$ (Janzing and Schölkopf, 2010; Sun et al., 2006; 2008). It was conjectured by Kilbertus et al. (2018) that this principle indicates generalization is typically easier to achieve in the causal direction than in the anti-causal direction. Moreover, they propose an abstract version of the search procedure (Sec. 4.1). The present paper can be seen as providing a formal justification for these ideas through theoretical insights on the structure of generators f (the causal direction) and their inverses g (the anti-causal direction).

7 DISCUSSION

Limitations. Our theory is limited to generators which belong to \mathcal{F}_{int} . We studied this function class as it provides a suitable model of visual data and is the largest class which enables OOD identifiability. However, these results may, in principle, fail to generalize to function classes associated with other settings, where non-generative strategies may be effective. Additionally, while our experiments leverage photorealistic data, they focus on concepts in simple settings which do not fully capture the complexity of real world data. To this end, an important future question is to understand how to create benchmarks to evaluate compositional generalization in a rigorous manner on data at a more realistic scale. For further discussions on limitations, see § D.

Conclusion. In this work, we sought a principled understanding of whether compositional generalization should be pursued through generative or non-generative approaches. Theoretically, we showed that for non-generative methods, enforcing the structure needed to guarantee compositionality tends to be infeasible. As a result, generalization is determined largely by the optimization process rather than by principled guarantees. Empirically, we observed that methods optimized from scratch or with little pretraining data tend to fail at compositional generalization, while larger-scale pretrained models improve OOD performance at the cost of data efficiency. By contrast, generative

486 approaches can directly enforce constraints for compositional generalization which manifest in sig-
487 nificant gains in OOD performance in practice. While scaling such generative approaches to more
488 challenging settings remains an open problem, we hope our findings will inspire renewed interest in
489 this direction.

490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539

REFERENCES

- 540
541
542 R. Assouel, P. Rodriguez, P. Taslakian, D. Vazquez, and Y. Bengio. Object-centric compositional
543 imagination for visual abstract reasoning. In *ICLR2022 Workshop on the Elements of Reasoning:
544 Objects, Structure and Causality*, 2022. [Cited on p. 9.]
- 545 R. Assouel, P. Astolfi, F. Bordes, M. Drozdal, and A. Romero-Soriano. Object-centric binding in
546 contrastive language-image pretraining. *arXiv preprint arXiv:2502.14113*, 2025. [Cited on p. 9.]
- 547 M. Assran, Q. Duval, I. Misra, P. Bojanowski, P. Vincent, M. Rabbat, Y. LeCun, and N. Ballas. Self-
548 supervised learning from images with a joint-embedding predictive architecture. In *Proceedings
549 of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15619–15629,
550 2023. [Cited on p. 8.]
- 551 R. Balestriero and Y. LeCun. How learning by reconstruction produces uninformative features for
552 perception. In *Forty-first International Conference on Machine Learning*, 2024. [Cited on p. 1.]
- 553 L. Beyer, A. Steiner, A. S. Pinto, A. Kolesnikov, X. Wang, D. Salz, M. Neumann, I. Alabdulmohsin,
554 M. Tschannen, E. Bugliarello, et al. Paligemma: A versatile 3b vlm for transfer. *arXiv preprint
555 arXiv:2407.07726*, 2024. [Cited on p. 1.]
- 556 F. Bordes, S. Shekhar, M. Ibrahim, D. Bouchacourt, P. Vincent, and A. Morcos. Pug: Photorealistic
557 and semantically controllable synthetic data for representation learning. *Advances in Neural
558 Information Processing Systems*, 36:45020–45054, 2023. [Cited on p. 2, 7, and 23.]
- 559 J. Brady, R. S. Zimmermann, Y. Sharma, B. Schölkopf, J. von Kügelgen, and W. Brendel. Provably
560 learning object-centric representations. In *International Conference on Machine Learning*, pages
561 3038–3062. PMLR, 2023. [Cited on p. 5.]
- 562 J. Brady, J. von Kügelgen, S. Lachapelle, S. Buchholz, T. Kipf, and W. Brendel. Interaction asym-
563 metry: A general principle for learning composable abstractions. In *The Thirteenth International
564 Conference on Learning Representations*, 2025. [Cited on p. 2, 4, 8, 23, 24, and 27.]
- 565 M. Caron, H. Touvron, I. Misra, H. Jégou, J. Mairal, P. Bojanowski, and A. Joulin. Emerging
566 properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international
567 conference on computer vision*, pages 9650–9660, 2021. [Cited on p. 1 and 8.]
- 568 A. Dittadi, S. S. Papa, M. De Vita, B. Schölkopf, O. Winther, and F. Locatello. Generalization
569 and robustness implications in object-centric learning. In *International Conference on Machine
570 Learning*, pages 5221–5285. PMLR, 2022. [Cited on p. 24 and 25.]
- 571 A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani,
572 M. Minderer, G. Heigold, S. Gelly, et al. An image is worth 16x16 words: Transformers for image
573 recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. [Cited on p. 8.]
- 574 D. Fan, S. Tong, J. Zhu, K. Sinha, Z. Liu, X. Chen, M. Rabbat, N. Ballas, Y. LeCun, A. Bar, et al.
575 Scaling language-free visual representation learning. *arXiv preprint arXiv:2504.01017*, 2025.
576 [Cited on p. 1 and 9.]
- 577 J. A. Fodor and Z. W. Pylyshyn. Connectionism and cognitive architecture: A critical analysis.
578 *Cognition*, 28(1):3–71, 1988. [Cited on p. 2.]
- 579 K. J. Friston and K. E. Stephan. Free-energy and the brain. *Synthese*, 159(3):417–458, 2007. [Cited
580 on p. 1.]
- 581 J. J. Gibson. The ecological approach to visual perception. 1979. [Cited on p. 1.]
- 582 Y. Grandvalet and Y. Bengio. Semi-supervised learning by entropy minimization. In L. Saul,
583 Y. Weiss, and L. Bottou, editors, *Advances in Neural Information Processing Systems*, volume 17.
584 MIT Press, 2004. [Cited on p. 24.]
- 585 K. Greff, S. Van Steenkiste, and J. Schmidhuber. On the binding problem in artificial neural net-
586 works. *arXiv preprint arXiv:2012.05208*, 2020. [Cited on p. 2.]
- 587 L. Gresele, J. von Kügelgen, V. Stimper, B. Schölkopf, and M. Besserve. Independent mechanism
588 analysis, a new concept? In *Advances in Neural Information Processing Systems*, volume 34,
589 pages 28233–28248, 2021. [Cited on p. 3.]
- 590
591
592
593

- 594 G. E. Hinton. To recognize shapes, first learn to generate images. *Progress in brain research*, 165:
595 535–547, 2007. [Cited on p. 1.]
- 596
597 E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, W. Chen, et al. Lora: Low-rank
598 adaptation of large language models. *ICLR*, 1(2):3, 2022. [Cited on p. 8.]
- 599 A. Hyvärinen and H. Morioka. Unsupervised feature extraction by time-contrastive learning and
600 nonlinear ICA. In *Advances in Neural Information Processing Systems*, pages 3765–3773, 2016.
601 [Cited on p. 3.]
- 602 A. Hyvärinen, I. Khemakhem, and H. Morioka. Nonlinear independent component analysis for
603 principled disentanglement in unsupervised deep learning. *Patterns*, 4(10), 2023. [Cited on p. 3.]
- 604 A. Jaegle, S. Borgeaud, J.-B. Alayrac, C. Doersch, C. Ionescu, D. Ding, S. Koppula, D. Zoran,
605 A. Brock, E. Shelhamer, et al. Perceiver IO: A general architecture for structured inputs & outputs.
606 In *International Conference on Learning Representations*, 2022. [Cited on p. 24.]
- 607
608 P. Jaini, K. Clark, and R. Geirhos. Intriguing properties of generative classifiers. In *The Twelfth*
609 *International Conference on Learning Representations*, 2024. [Cited on p. 27.]
- 610 D. Janzing and B. Schölkopf. Causal inference using the algorithmic markov condition. *IEEE*
611 *Transactions on Information Theory*, 56(10):5168–5194, Oct. 2010. [Cited on p. 9.]
- 612
613 Y. Jeong, A. Uselis, S. J. Oh, and A. Rohrbach. Diffusion classifiers understand compositionality,
614 but conditions apply. *arXiv preprint arXiv:2505.17955*, 2, 2025. [Cited on p. 9.]
- 615 W. Jung, J. Yoo, S. Ahn, and S. Hong. Learning to compose: Improving object centric learning by
616 injecting compositionality. In *The Twelfth International Conference on Learning Representations*,
617 2024. [Cited on p. 9.]
- 618 D. Kahneman. Thinking, fast and slow. *Farrar, Straus and Giroux*, 2011. [Cited on p. 6.]
- 619
620 I. Khemakhem, D. Kingma, R. Monti, and A. Hyvärinen. Variational autoencoders and nonlinear
621 ICA: A unifying framework. In *International Conference on Artificial Intelligence and Statistics*,
622 pages 2207–2217. PMLR, 2020. [Cited on p. 3.]
- 623 N. Kilbertus, G. Parascandolo, and B. Schölkopf. Generalization in anti-causal learning. *arXiv*
624 *preprint arXiv:1812.00524*, 2018. [Cited on p. 2 and 9.]
- 625 D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *International Conference*
626 *on Learning Representations*, 2015. [Cited on p. 24.]
- 627
628 D. P. Kingma and M. Welling. Auto-encoding variational bayes. In *International Conference on*
629 *Learning Representations*, 2014. [Cited on p. 8.]
- 630 Z. Kurth-Nelson, T. Behrens, G. Wayne, K. Miller, L. Luettgau, R. Dolan, Y. Liu, and P. Schwarten-
631 beck. Replay and compositional computation. *Neuron*, 111(4):454–469, 2023. [Cited on p. 7.]
- 632 S. Lachapelle, D. Mahajan, I. Mitliagkas, and S. Lacoste-Julien. Additive decoders for latent vari-
633 ables identification and cartesian-product extrapolation. In *Advances in Neural Information Pro-*
634 *cessing Systems*, volume 36, 2023. [Cited on p. 4.]
- 635
636 B. M. Lake, R. Salakhutdinov, and J. B. Tenenbaum. Human-level concept learning through proba-
637 bilistic program induction. *Science*, 350(6266):1332–1338, 2015. [Cited on p. 2 and 9.]
- 638 B. M. Lake, T. D. Ullman, J. B. Tenenbaum, and S. J. Gershman. Building machines that learn and
639 think like people. *Behavioral and brain sciences*, 40:e253, 2017. [Cited on p. 2 and 9.]
- 640 Y. LeCun. A path towards autonomous machine intelligence version 0.9.2, 2022-06-27. *OpenRe-*
641 *view*, pages 1–62, 2022. [Cited on p. 1 and 6.]
- 642
643 M. Lewis, N. V. Nayak, P. Yu, Q. Yu, J. Merullo, S. H. Bach, and E. Pavlick. Does clip bind
644 concepts? probing compositionality in large image models. *arXiv preprint arXiv:2212.10537*,
645 2022. [Cited on p. 9.]
- 646 F. Locatello, B. Poole, G. Rätsch, B. Schölkopf, O. Bachem, and M. Tschannen. Weakly-supervised
647 disentanglement without compromises. In *International Conference on Machine Learning*, pages
6348–6359. PMLR, 2020a. [Cited on p. 3.]

- 648 F. Locatello, D. Weissenborn, T. Unterthiner, A. Mahendran, G. Heigold, J. Uszkoreit, A. Dosovit-
649 ski, and T. Kipf. Object-centric learning with slot attention. In *Advances in Neural Information*
650 *Processing Systems*, volume 33, pages 11525–11538, 2020b. [Cited on p. 8 and 24.]
- 651 Z. Ma, J. Hong, M. O. Gul, M. Gandhi, I. Gao, and R. Krishna. Crepe: Can vision-language
652 foundation models reason compositionally? In *Proceedings of the IEEE/CVF Conference on*
653 *Computer Vision and Pattern Recognition*, pages 10910–10921, 2023. [Cited on p. 9.]
- 654 B. A. Olshausen. Perception as an inference problem. *The cognitive neurosciences*, pages 295–304,
655 2014. [Cited on p. 1.]
- 656 M. Oquab, T. Darcet, T. Moutakanni, H. V. Vo, M. Szafraniec, V. Khalidov, P. Fernandez, D. HAZ-
657 IZA, F. Massa, A. El-Nouby, M. Assran, N. Ballas, W. Galuba, R. Howes, P.-Y. Huang, S.-W. Li,
658 I. Misra, M. Rabbat, V. Sharma, G. Synnaeve, H. Xu, H. Jegou, J. Mairal, P. Labatut, A. Joulin,
659 and P. Bojanowski. DINOv2: Learning robust visual features without supervision. *Transactions*
660 *on Machine Learning Research*, 2024. ISSN 2835-8856. [Cited on p. 1 and 8.]
- 661 G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan. Normal-
662 izing flows for probabilistic modeling and inference. *Journal of Machine Learning Research*, 22
663 (57):1–64, 2021. [Cited on p. 6.]
- 664 W. Peebles and S. Xie. Scalable diffusion models with transformers. In *Proceedings of the*
665 *IEEE/CVF international conference on computer vision*, pages 4195–4205, 2023. [Cited on
666 p. 27.]
- 667 B. Peters, J. J. DiCarlo, T. Gureckis, R. Haefner, L. Isik, J. Tenenbaum, T. Konkle, T. Naselaris,
668 K. Stachenfeld, Z. Tavares, et al. How does the primate brain combine generative and discrimi-
669 native computations in vision? *ArXiv*, pages arXiv–2401, 2024. [Cited on p. 2.]
- 670 M. Prabhudesai, A. Goyal, S. Paul, S. Van Steenkiste, M. S. Sajjadi, G. Aggarwal, T. Kipf, D. Pathak,
671 and K. Fragkiadaki. Test-time adaptation with slot-centric models. In *International Conference*
672 *on Machine Learning*, pages 28151–28166. PMLR, 2023a. [Cited on p. 6 and 9.]
- 673 M. Prabhudesai, T.-W. Ke, A. Li, D. Pathak, and K. Fragkiadaki. Diffusion-tta: Test-time adaptation
674 of discriminative models via generative feedback. *Advances in Neural Information Processing*
675 *Systems*, 36:17567–17583, 2023b. [Cited on p. 9.]
- 676 A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin,
677 J. Clark, et al. Learning transferable visual models from natural language supervision. In *Inter-*
678 *national conference on machine learning*, pages 8748–8763. PmLR, 2021. [Cited on p. 1 and 8.]
- 679 R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis
680 with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and*
681 *pattern recognition*, pages 10684–10695, 2022. [Cited on p. 27.]
- 682 M. S. Sajjadi, H. Meyer, E. Pot, U. Bergmann, K. Greff, N. Radwan, S. Vora, M. Lučić, D. Duck-
683 worth, A. Dosovitskiy, et al. Scene representation transformer: Geometry-free novel view syn-
684 thesis through set-latent scene representations. In *Proceedings of the IEEE/CVF Conference on*
685 *Computer Vision and Pattern Recognition*, pages 6229–6238, 2022. [Cited on p. 24.]
- 686 P. Schwartenbeck, A. Baram, Y. Liu, S. Mark, T. Muller, R. Dolan, M. Botvinick, Z. Kurth-
687 Nelson, and T. Behrens. Generative replay underlies compositional inference in the hippocampal-
688 prefrontal circuit. *Cell*, 186(22):4885–4897, 2023. [Cited on p. 7.]
- 689 O. Siméoni, H. V. Vo, M. Seitzer, F. Baldassarre, M. Oquab, C. Jose, V. Khalidov, M. Szafraniec,
690 S. Yi, M. Ramamonjisoa, et al. Dinov3. *arXiv preprint arXiv:2508.10104*, 2025. [Cited on p. 1.]
- 691 X. Sun, D. Janzing, and B. Schölkopf. Causal inference by choosing graphs with most plausible
692 markov kernels. In *Proceedings of the 9th International Symposium on Artificial Intelligence and*
693 *Mathematics*, pages 1–11. Max-Planck-Gesellschaft, Jan. 2006. [Cited on p. 9.]
- 694 X. Sun, D. Janzing, and B. Schölkopf. Causal reasoning by evaluating the complexity of conditional
695 densities with kernel methods. *Neurocomputing*, 71(7-9):1248–1256, Mar. 2008. [Cited on p. 9.]
- 696 J. B. Tenenbaum, C. Kemp, T. L. Griffiths, and N. D. Goodman. How to grow a mind: Statistics,
697 structure, and abstraction. *Science*, 331(6022):1279–1285, 2011. [Cited on p. 2, 3, and 9.]

- 702 Y. Tian, L. Fan, P. Isola, H. Chang, and D. Krishnan. Stablerep: Synthetic images from text-to-image
703 models make strong visual representation learners. *Advances in Neural Information Processing*
704 *Systems*, 36:48382–48402, 2023. [Cited on p. 9.]
- 705 S. Tong, Z. Liu, Y. Zhai, Y. Ma, Y. LeCun, and S. Xie. Eyes wide shut? exploring the visual
706 shortcomings of multimodal llms. In *Proceedings of the IEEE/CVF Conference on Computer*
707 *Vision and Pattern Recognition*, pages 9568–9578, 2024. [Cited on p. 9.]
- 708 M. Tschannen, A. Gritsenko, X. Wang, M. F. Naeem, I. Alabdulmohsin, N. Parthasarathy, T. Evans,
709 L. Beyer, Y. Xia, B. Mustafa, et al. Siglip 2: Multilingual vision-language encoders with improved
710 semantic understanding, localization, and dense features. *arXiv preprint arXiv:2502.14786*, 2025.
711 [Cited on p. 1 and 8.]
- 712 V. Udandarao, A. Prabhu, A. Ghosh, Y. Sharma, P. Torr, A. Bibi, S. Albanie, and M. Bethge. No”
713 zero-shot” without exponential data: Pretraining concept frequency determines multimodal model
714 performance. *Advances in Neural Information Processing Systems*, 37:61735–61792, 2024.
715 [Cited on p. 1.]
- 716 S. van Steenkiste, D. Zoran, Y. Yang, Y. Rubanova, R. Kabra, C. Doersch, D. Gokay, E. Pot, K. Greff,
717 D. Hudson, et al. Moving off-the-grid: Scene-grounded video representations. *Advances in*
718 *Neural Information Processing Systems*, 37:124319–124346, 2024. [Cited on p. 26.]
- 719 A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polo-
720 sukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, vol-
721 ume 30, 2017. [Cited on p. 8.]
- 722 H. von Helmholtz. *Handbuch der physiologischen Optik*, volume 3. Voss, 1867. [Cited on p. 1.]
- 723 J. von Kügelgen, Y. Sharma, L. Gresele, W. Brendel, B. Schölkopf, M. Besserve, and F. Locatello.
724 Self-supervised learning with data augmentations provably isolates content from style. In *Ad-*
725 *vances in Neural Information Processing Systems*, volume 34, pages 16451–16467, 2021. [Cited
726 on p. 3.]
- 727 Y. Wang, J. Dauwels, and Y. Du. Compositional scene understanding through inverse generative
728 modeling. In *Forty-second International Conference on Machine Learning*, 2025. [Cited on p. 9.]
- 729 J. C. R. Whittington, W. Dorrell, S. Ganguli, and T. Behrens. Disentanglement with biological con-
730 straints: A theory of functional cell types. In *The Eleventh International Conference on Learning*
731 *Representations*, 2023. [Cited on p. 27.]
- 732 T. Wiedemer, J. Brady, A. Panfilov, A. Juhos, M. Bethge, and W. Brendel. Provable compositional
733 generalization for object-centric learning. In *International Conference on Learning Representa-*
734 *tions*, 2024a. [Cited on p. 7, 9, and 24.]
- 735 T. Wiedemer, P. Mayilvahanan, M. Bethge, and W. Brendel. Compositional generalization from first
736 principles. In *Advances in Neural Information Processing Systems*, volume 36, 2024b. [Cited on
737 p. 3.]
- 738 D. Yamins, H. Hong, C. F. Cadieu, E. A. Solomon, D. Seibert, and J. J. DiCarlo. Performance-
739 optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of*
740 *the National Academy of Sciences*, 111:8619 – 8624, 2014. [Cited on p. 1.]
- 741 D. L. Yamins and J. J. DiCarlo. Using goal-driven deep learning models to understand sensory
742 cortex. *Nature neuroscience*, 19(3):356–365, 2016. [Cited on p. 3.]
- 743 M. Yuksekgonul, F. Bianchi, P. Kalluri, D. Jurafsky, and J. Zou. When and why vision-language
744 models behave like bags-of-words, and what to do about it? *arXiv preprint arXiv:2210.01936*,
745 2022. [Cited on p. 9.]
- 746 J. Zhang, J. Huang, S. Jin, and S. Lu. Vision-language models for vision tasks: A survey. *IEEE*
747 *transactions on pattern analysis and machine intelligence*, 46(8):5625–5644, 2024. [Cited on
748 p. 1.]
- 749 R. S. Zimmermann, Y. Sharma, S. Schneider, M. Bethge, and W. Brendel. Contrastive learning
750 inverts the data generating process. In *International Conference on Machine Learning*, pages
751 12979–12990. PMLR, 2021. [Cited on p. 3.]

Appendices

Table of Contents

A Proofs	15
A.1 Local regularization of encoders	15
A.2 Constraining \mathcal{G}_{enc} by architecture	19
B Experimental details	23
B.1 Data	23
B.2 Models	24
B.3 Training Objectives	24
C Additional Experiments	25
D Extended Discussion	26

Use of Language Models Large language models (LLMs) were employed exclusively during the final stages of manuscript preparation for the purpose of refining language, grammar, and readability. They were not used for generating ideas, conducting analysis, or contributing to the substantive content of this work.

A PROOFS

In this section we collect the proofs of the results in the paper and some additional background material. First, in Section A.1 we investigate the local restrictions that $\mathbf{g} \in \mathcal{G}_{\text{int}}$ need to satisfy. Similarly we investigate in Section A.2 whether we can enforce $\mathbf{g} \in \mathcal{G}_{\text{int}}$ by architectural constraints. Let us, however, first introduce a notation for a subset of \mathcal{F}_{int} .

Definition A.1 (Additive functions). We denote the function class of coordinate-wise additive functions $\mathbf{f} : \mathbb{R}^{d_z} \rightarrow \mathbb{R}^{d_x}$ by \mathcal{F}_{add} . They can be expressed as

$$\mathbf{f}(\mathbf{x}) = \sum_{i=1}^{d_z} \mathbf{f}_i(\mathbf{x}_i) \quad (\text{A.1})$$

where $\mathbf{f}_i : \mathbb{R} \rightarrow \mathbb{R}^{d_x}$.

Clearly \mathcal{F}_{add} agrees with \mathcal{F}_{int} for $n = m = 1$, i.e., interactions of first order and blocks of dimension 1 and generally $\mathcal{F}_{\text{add}} \subset \mathcal{F}_{\text{int}}$ for $n \geq 1$ (higher order interactions and larger blocks are more flexible).

A.1 LOCAL REGULARIZATION OF ENCODERS

As discussed in the main text, we can enforce $\mathbf{f} \in \mathcal{F}_{\text{int}}$ by enforcing that certain derivatives of \mathbf{f} vanish (see equation 3.2). We now study to what extent this generalizes to functions $\mathbf{g} \in \mathcal{G}_{\text{int}}$ that are left inverses of such functions.

The key relation. The key relation that we need for the proofs below is that if $\mathbf{g} \circ \mathbf{f}(\mathbf{z}) = \mathbf{z}$ for two functions $\mathbf{f} : \mathbb{R}^{d_z} \rightarrow \mathbb{R}^{d_x}$ and $\mathbf{g} : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_z}$, then for every $s \in [d_z]$

$$D\mathbf{f}^\top(\mathbf{z})D^2\mathbf{g}_s(\mathbf{f}(\mathbf{z}))D\mathbf{f}(\mathbf{z}) + \sum_{k=1}^{d_x} (\partial_k \mathbf{g}_s)(\mathbf{f}(\mathbf{z}))D^2\mathbf{f}_k(\mathbf{z}) = 0. \quad (\text{A.2})$$

This relation follows by straightforward calculation, indeed we find using the chain rule

$$\begin{aligned} \partial_i \partial_j \mathbf{g}_s(\mathbf{f}(\mathbf{z})) &= \partial_i \left(\sum_{k=1}^{d_x} \partial_j \mathbf{f}_k(\mathbf{z}) (\partial_k \mathbf{g}_s)(\mathbf{f}(\mathbf{z})) \right) \\ &= \sum_{k,l=1}^{d_x} \partial_j \mathbf{f}_k(\mathbf{z}) \partial_i \mathbf{f}_l(\mathbf{z}) (\partial_k \partial_l \mathbf{g}_s)(\mathbf{f}(\mathbf{z})) + \sum_{k=1}^{d_x} \partial_i \partial_j \mathbf{f}_k(\mathbf{z}) (\partial_k \mathbf{g}_s)(\mathbf{f}(\mathbf{z})) \end{aligned} \quad (\text{A.3})$$

which is equation A.2 after rewriting the relation in matrix form.

Restrictions for $d_x = d_z$. We now prove Lemma 3.1 showing that for $d_x = d_z$, i.e., for equal dimension of latent space and data it is possible to find a local constraint for the inverses of additive functions $\mathbf{f} \in \mathcal{F}_{\text{add}}$.

Lemma 3.1. *Let $\mathbf{g} \in \mathcal{G}_{\text{int}}$ for $n = m = 1$ and $d_x = d_z$. Then \mathbf{g} has the property that for $\mathbf{x} \in \mathcal{X}$*

$$(D\mathbf{g})^{-\top}(\mathbf{x}) D^2 \mathbf{g}_s(\mathbf{x}) (D\mathbf{g})^{-1}(\mathbf{x}) \in \text{Diag}(d_x) \quad (\text{3.3})$$

is a diagonal matrix for $s \in [d_z]$. Further, if \mathbf{g} is a diffeomorphism satisfying Eq. (3.3) then $\mathbf{g} \in \mathcal{G}_{\text{int}}$.

Remark A.2. For higher dimensional slots there is a natural generalization, namely, the expression $D\mathbf{g}^{-\top} D^2 \mathbf{g}_s D\mathbf{g}^{-1}$ has a block diagonal structure.

Proof. Note that $\mathbf{g} \circ \mathbf{f}(\mathbf{z}) = \mathbf{z}$ implies $\text{Id}_{d_z} = (D\mathbf{g} \circ \mathbf{f}) D\mathbf{f}$ and thus $D\mathbf{f}(\mathbf{z}) = (D\mathbf{g})^{-1}(\mathbf{f}(\mathbf{z}))$. Therefore, we find using equation A.2

$$(D\mathbf{g})^{-\top}(\mathbf{f}(\mathbf{z})) D^2 \mathbf{g}_s(\mathbf{f}(\mathbf{z})) (D\mathbf{g})^{-1}(\mathbf{f}(\mathbf{z})) = - \sum_{k=1}^{d_x} (\partial_k \mathbf{g}_s)(\mathbf{f}(\mathbf{z})) D^2 \mathbf{f}_k(\mathbf{z}) \in \text{Diag}(d_z). \quad (\text{A.4})$$

where we used that $\mathbf{f} \in \mathcal{F}_{\text{add}}$ implies that the off-diagonal entries of $D^2 \mathbf{f}$ vanish. This implies the first part of the statement. For the reverse statement, we apply equation A.2 to $\mathbf{f} \circ \mathbf{g}(\mathbf{x}) = \mathbf{x}$ (here we use $d_z = d_x$) and we find that

$$0 = (D\mathbf{g})^{\top}(\mathbf{x}) D^2 \mathbf{f}_s(\mathbf{g}(\mathbf{x})) D\mathbf{g}(\mathbf{x}) + \sum_{k=1}^{d_z} (\partial_k \mathbf{f}_s)(\mathbf{g}(\mathbf{x})) D^2 \mathbf{g}_k(\mathbf{x}). \quad (\text{A.5})$$

We multiply this relation from the left and right by $(D\mathbf{g})^{-\top}(\mathbf{x})$ and $(D\mathbf{g})^{-1}(\mathbf{x})$ respectively (the inverses exist by assumption) and we find

$$D^2 \mathbf{f}_s(\mathbf{g}(\mathbf{x})) = - \sum_{k=1}^{d_z} (\partial_k \mathbf{f}_s)(\mathbf{g}(\mathbf{x})) (D\mathbf{g})^{-\top}(\mathbf{x}) D^2 \mathbf{g}_k(\mathbf{x}) (D\mathbf{g})^{-1}(\mathbf{x}) \in \text{Diag}(d_z). \quad (\text{A.6})$$

Here we used the assumption equation 3.3 to conclude that the right hand side is diagonal. Therefore \mathbf{f} has a diagonal Hessian which implies that it is additive. \square

The previous statement can be generalized to the general case $d_x > d_z$. The crucial ingredient is the following simple and standard lemma.

Lemma A.3. *Let $\mathbf{A} \in \mathbb{R}^{d_1 \times d_2}$ and $\mathbf{B} \in \mathbb{R}^{d_2 \times d_1}$ two matrices with $d_2 \geq d_1$ and assume that $\mathbf{A}\mathbf{B} = \mathbf{1}_{d_1 \times d_1}$. Then $\mathbf{B} = (\mathbf{A}\Pi)^+$ where Π denotes the orthogonal projection onto $\text{Range}(\mathbf{B})$ and $(\cdot)^+$ the Moore-Penrose inverse of a matrix.*

Proof. We check that \mathbf{B} satisfies the Moore-Penrose axioms ($\mathbf{M}\mathbf{M}^+\mathbf{M} = \mathbf{M}$, $\mathbf{M}^+\mathbf{M}\mathbf{M}^+ = \mathbf{M}^+$, $\mathbf{M}^+\mathbf{M}$ and $\mathbf{M}\mathbf{M}^+$ are Hermitian). We find

$$\mathbf{A}\Pi\mathbf{B}\mathbf{A}\Pi = \mathbf{A}\mathbf{B}\mathbf{A}\Pi = \mathbf{A}\Pi \quad (\text{A.7})$$

where we used $\Pi\mathbf{B} = \mathbf{B}$ by definition of Π . Similarly, we obtain

$$\mathbf{B}\mathbf{A}\Pi\mathbf{B} = \Pi\mathbf{B} = \mathbf{B}. \quad (\text{A.8})$$

Next we claim that

$$\mathbf{B}\mathbf{A}\mathbf{\Pi} = \mathbf{\Pi} \quad (\text{A.9})$$

which is Hermitian. Consider $v \in \mathbb{R}^{d_2}$ then by definition of $\mathbf{\Pi}$ there is $w \in \mathbb{R}^{d_1}$ such that $\mathbf{B}w = \mathbf{\Pi}v$ and thus

$$\mathbf{B}\mathbf{A}\mathbf{\Pi}v = \mathbf{B}\mathbf{A}\mathbf{B}w = \mathbf{B}w = \mathbf{\Pi}v. \quad (\text{A.10})$$

Finally, we find

$$\mathbf{A}\mathbf{\Pi}\mathbf{B} = \mathbf{A}\mathbf{B} = \mathbf{1}_{d_1 \times d_1}. \quad (\text{A.11})$$

□

We have the following generalization of Lemma 3.1.

Lemma A.4. *Let $\mathbf{f} \in \mathcal{F}_{\text{add}}$ and \mathbf{g} a left-inverse of \mathbf{f} . Then \mathbf{g} has the property that for $\mathbf{x} \in \mathcal{X}$*

$$\left((D\mathbf{g}(\mathbf{x})\mathbf{\Pi}_{T_{\mathbf{x}}\mathcal{X}})^+ \right)^\top (\mathbf{z}) D^2\mathbf{g}_s(\mathbf{x}) (D\mathbf{g}(\mathbf{x})\mathbf{\Pi}_{T_{\mathbf{x}}\mathcal{X}})^+ \in \text{Diag}(d_z) \quad (\text{A.12})$$

is a diagonal matrix for $s \in [d_z]$. Here, we denote by $\mathbf{\Pi}_{T_{\mathbf{x}}\mathcal{X}}$ the orthogonal projection on the tangent space at \mathbf{x} .

Proof. Starting from equation A.2 we find that for $\mathbf{f} \in \mathcal{F}_{\text{add}}$ we get

$$D\mathbf{f}^\top(\mathbf{z}) D^2\mathbf{g}_s(\mathbf{f}(\mathbf{z})) D\mathbf{f}(\mathbf{z}) \in \text{Diag}(d_z). \quad (\text{A.13})$$

Applying Lemma A.3 we find

$$D\mathbf{f}(\mathbf{z}) = (D\mathbf{g}(\mathbf{f}(\mathbf{z}))\mathbf{\Pi}_{T_{\mathbf{f}(\mathbf{z})}\mathcal{X}})^+ \quad (\text{A.14})$$

because the range of $D\mathbf{f}$ is the tangent space of the data manifold. Therefore we conclude that for $\mathbf{x} \in \mathcal{X}$ the relation equation A.12 indeed holds. □

Regularization for $d_x > d_z$. In this paragraph we investigate the local restrictions that $\mathbf{g} \in \mathcal{G}_{\text{int}}$ need to satisfy, and in particular we prove Theorem 3.2. The proof of Theorem 3.2 requires two lemmas as a key ingredient, which state that the crucial constraint on the second derivative stated in equation A.2 can be satisfied for a suitable choice of $\mathbf{M} = D\mathbf{f}(0)$ and $D^2\mathbf{f}(0)$ for given matrices \mathbf{B}_s corresponding to the Hessian of \mathbf{g} and almost every matrix \mathbf{A} (corresponding to the Jacobian of \mathbf{g}). The first lemma establishes the existence of \mathbf{M} such that first term in equation A.2 (given by $\mathbf{M}^\top \mathbf{B}_s \mathbf{M}$) is diagonal for all s . The second lemma constructs suitable second derivatives $D^2\mathbf{f}$ so that the relation equation A.2 also holds for the diagonal entries.

Lemma A.5. *Assume $d_x \geq d_z^3$. For all symmetric matrices $\mathbf{B}_s \in \mathbb{R}^{d_x \times d_x}$ for $s \in [d_z]$, and almost every $\mathbf{A} \in \mathbb{R}^{d_z \times d_x}$ there is a matrix $\mathbf{M} \in \mathbb{R}^{d_x \times d_z}$ such that $\mathbf{M}^\top \mathbf{B}_s \mathbf{M} \in \text{Diag}(d_z)$ for $s \in [d_z]$ and $\mathbf{A}\mathbf{M} = \text{Id}_{d_z}$.*

Remark A.6. 1. Counting parameters and equations, we find that \mathbf{M} has $d_z d_x$ parameters and (by symmetry of \mathbf{B}_s) there are

$$d_z \cdot \frac{d_z(d_z - 1)}{2} + d_z^2 = \frac{d_z^2(d_z + 1)}{2} \quad (\text{A.15})$$

equations. So, generally, we expect the result to hold for $d_x \geq d_z(d_z + 1)/2$.

2. On the other hand, the result does not hold for every \mathbf{A} with maximal rank. Indeed, there can be a non-trivial null set of full rank matrices \mathbf{A} such that the result does not hold. E.g., consider $d_z = 2$, $\mathbf{A} \in \mathbb{R}^{d_z \times d_x}$ such that all entries of \mathbf{A} are zero except $\mathbf{A}_{1,1} = \mathbf{A}_{2,2} = 1$. Moreover, \mathbf{B}_1 has all entries zero except $(\mathbf{B}_1)_{1,2} = (\mathbf{B}_1)_{2,1} = 1$. Then $\mathbf{A}\mathbf{M} = \text{Id}_{d_z}$ implies that $\mathbf{M}_{1,1} = \mathbf{M}_{2,2} = 1$, and $\mathbf{M}_{1,2} = \mathbf{M}_{2,1} = 0$. But then we find $\mathbf{M}_{:,1}^\top \mathbf{B}_1 \mathbf{M}_{:,2} = (\mathbf{B}_1)_{1,2} = 1 \neq 0$.

918 *Proof.* We inductively construct d_z linear subspaces $V_i \subset \mathbb{R}^{d_x}$ such that $\dim(V_i) = d_z$ and

$$919 \quad (\mathbf{v}^i)^\top \mathbf{B}_s \mathbf{v}^j = 0 \quad (\text{A.16})$$

921 for $\mathbf{v}^i \in V_i$, $\mathbf{v}^j \in V_j$ and $i \neq j$. We pick V_1 arbitrarily. Then, given a basis $\mathbf{v}^{i,1}, \dots, \mathbf{v}^{i,d_z}$ of V_i
 922 for $i \leq j$ we select $V_{j+1} \subset \ker \mathbf{T}_j$ where $\mathbf{T}_j : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_z \cdot j}$ given by $(\mathbf{T}_j \mathbf{v})_{s,(k,i)} = (\mathbf{v}^{i,k})^\top \mathbf{B}_s \mathbf{v}$
 923 (here it is convenient to identify $[d_z^2 \cdot j]$ with $[d_z] \times ([d_z] \times [j])$). By assumption $d_x - d_z^2 \cdot j \geq$
 924 $d_x - d_z^2 \cdot (d_z - 1) \geq d_z$ and therefore $\dim \ker \mathbf{T}_j \geq d_z$ and we can find a suitable subspace
 925 $V_{j+1} \subset \ker \mathbf{T}_j$. Given a matrix $\mathbf{A} = (\mathbf{a}^1, \dots, \mathbf{a}^{d_z})^\top \in \mathbb{R}^{d_z \times d_x}$, we want to find $\mathbf{w}^i \in V_i$ so that
 926 $\mathbf{M} = (\mathbf{w}^1, \dots, \mathbf{w}^{d_z})$ satisfies $\mathbf{A}\mathbf{M} = \text{Id}_{d_z}$. Equivalently $\mathbf{A}\mathbf{w}^i = \mathbf{e}^i$, where \mathbf{e}^i denotes the i -th
 927 standard basis vector. We expand into the basis of V_i , i.e., $\mathbf{w}^i = \sum_j \lambda_j^i \mathbf{v}^{i,j}$ and find the equivalent
 928 relation

$$929 \quad \mathbf{A}\mathbf{w}^i = (\mathbf{a}^1, \dots, \mathbf{a}^{d_z})^\top (\mathbf{v}^{i,1}, \dots, \mathbf{v}^{i,d_z}) \boldsymbol{\lambda}^i = \mathbf{e}^i. \quad (\text{A.17})$$

931 Since the second matrix has maximal rank ($(\mathbf{v}^{i,k})_{1 \leq k \leq d_z}$ is a basis of V_i), we find that for almost
 932 all \mathbf{A} the matrix product is invertible, and a solution $\boldsymbol{\lambda}^i$ exists and thus a suitable \mathbf{w}^i exists. To see
 933 this, we can assume that the basis $\mathbf{v}^{i,\cdot}$ is an orthonormal basis and expand \mathbf{a}_i in this basis (and an
 934 irrelevant orthogonal complement). We conclude that for almost all \mathbf{A} such a \mathbf{w}^i exists. Since the
 935 union of null-sets is a null-set the same statement holds for almost all \mathbf{A} for all i at the same time
 936 and therefore we find a matrix \mathbf{M} such that $\mathbf{A}\mathbf{M} = \text{Id}_{d_z}$ and, moreover, $(\mathbf{w}^i)^\top \mathbf{B}_s \mathbf{w}^j = 0$ because
 937 this holds for all $\mathbf{w}^i \in V_i$ and $\mathbf{w}^j \in V_j$. \square

938 We now construct the diagonal matrices that will later correspond to $D^2 \mathbf{f}_s$.

939 **Lemma A.7.** Assume $d_x \geq d_z$. Given $\mathbf{A} \in \mathbb{R}^{d_z \times d_x}$ of maximal rank and diagonal matrices
 940 $\mathbf{D}^1, \dots, \mathbf{D}^{d_z} \in \mathbb{R}^{d_z \times d_z}$ we can find diagonal matrices $\boldsymbol{\Lambda}^1, \dots, \boldsymbol{\Lambda}^{d_z} \in \mathbb{R}^{d_z \times d_z}$ such that for all
 941 $s \in [d_z]$

$$942 \quad \mathbf{D}^s = - \sum_{i=1}^{d_x} \mathbf{A}_{s,i} \boldsymbol{\Lambda}^i. \quad (\text{A.18})$$

943 *Proof.* The proof is straightforward as soon as one observes that this is a linear equation for
 944 the diagonal entries of $\boldsymbol{\Lambda}^i$. Indeed, denoting by $\boldsymbol{\lambda} = (\boldsymbol{\Lambda}_{11}^1, \dots, \boldsymbol{\Lambda}_{d_z, d_z}^1, \dots, (\boldsymbol{\Lambda}_{d_z, d_z}^{d_z})^\top)^\top \in$
 945 $\mathbb{R}^{d_z \cdot d_x}$ the vector containing all diagonal entries of the matrices $\boldsymbol{\Lambda}^i$ and similarly $\mathbf{d} =$
 946 $(\mathbf{D}_{11}^1, \dots, \mathbf{D}_{d_z, d_z}^1, \dots, \mathbf{D}_{d_z, d_z}^{d_z})^\top \in \mathbb{R}^{d_z^2}$ for the diagonal entries of \mathbf{D}^s . Then we can rewrite equa-
 947 tion A.18 as follows using the Kronecker product \otimes

$$948 \quad (\mathbf{A} \otimes \text{Id}_{d_z}) \boldsymbol{\lambda} = -\mathbf{d}. \quad (\text{A.19})$$

949 Now the rank of the matrix $\mathbf{A} \otimes \text{Id}_{d_z}$ is the product of the ranks, i.e., $d_z \min(d_x, d_z) = d_z^2 \leq d_x d_z$
 950 and thus a solution $\boldsymbol{\lambda}$ exists. \square

951 With these technical lemmas at hand, we can prove the theorem which we now restate for conve-
 952 nience of the reader.

953 **Theorem 3.2.** Assume that $d_x \geq d_z^3$. Let $\mathbf{B}_l \in \mathbb{R}^{d_x \times d_x}$ be symmetric matrices for $1 \leq l \leq d_z$.
 954 Then there is for any $\mathbf{x}_0 \in \mathbb{R}^{d_x}$ and for almost every $\mathbf{A} \in \mathbb{R}^{d_z \times d_x}$ a generator $\mathbf{f} \in \mathcal{F}_{\text{int}}$ with a
 955 (left)-inverse $\mathbf{g} \in \mathcal{G}_{\text{int}}$, such that $\mathbf{f}(0) = \mathbf{x}_0$ and $D\mathbf{g}(\mathbf{x}_0) = \mathbf{A}$ and $D^2 \mathbf{g}_l(\mathbf{x}_0) = \mathbf{B}_l$ for $1 \leq l \leq d_z$.
 956

957 *Proof of Theorem 3.2.* Clearly we can assume that $\mathbf{x}_0 = 0$. The key idea is that if we can ensure
 958 that equation A.2 holds for $\mathbf{z} = 0$ we can extend \mathbf{f} and \mathbf{g} such that $\mathbf{g} \circ \mathbf{f}(\mathbf{z}) = \mathbf{z}$ and $\mathbf{f} \in \mathcal{F}_{\text{add}}$.
 959 To achieve this, we first apply Lemma A.5 and then find a matrix \mathbf{M} such that $\mathbf{A}\mathbf{M} = \text{Id}_{d_z}$ and
 960 $\mathbf{M}^\top \mathbf{B}_s \mathbf{M} \in \text{Diag}(d_z)$. Then we apply Lemma A.7 and find matrices $\boldsymbol{\Lambda}^i$ such that

$$961 \quad \mathbf{M}^\top \mathbf{B}_s \mathbf{M} + \sum_{i=1}^{d_x} \mathbf{A}_{s,i} \boldsymbol{\Lambda}^i = 0. \quad (\text{A.20})$$

962 Now we pick a function $\mathbf{f} \in \mathcal{F}_{\text{add}}$ such that $\mathbf{f}(0) = 0$, $D\mathbf{f}(0) = \mathbf{M}$ and $D^2 \mathbf{f}_i = \boldsymbol{\Lambda}^i$. Clearly,
 963 this is possible because $\boldsymbol{\Lambda}^i$ are diagonal, e.g., we can locally use a quadratic polynomial to achieve

972 this. The next step is to construct a function \mathbf{g} such that $\mathbf{g} \circ \mathbf{f}(z) = z$. Using standard techniques
 973 (partition of unity) it is sufficient to construct this locally and then extend it globally. First, we
 974 consider $\bar{\phi} : \Omega \subset \mathbb{R}^{d_x} \times \mathbb{R}^{d_x}$ such that $\bar{\phi}(\mathbf{f}(z)) = (z, 0)$ for $z \in \Omega$ (e.g., by the existence of tubular
 975 neighbourhoods). We call the composition of $\bar{\phi}$ with the projection on the first d_z components ϕ .
 976 Then we define

$$977 \mathbf{g}(\mathbf{x}) = \mathbf{A}\mathbf{x} + \frac{1}{2} \begin{pmatrix} \mathbf{x}^\top \mathbf{B}_1 \mathbf{x} \\ \vdots \\ \mathbf{x}^\top \mathbf{B}_{d_z} \mathbf{x} \end{pmatrix} + \mathbf{h}(\phi(\mathbf{x})) \quad (\text{A.21})$$

982 where \mathbf{h} is given by

$$983 \mathbf{h}(z) = \mathbf{h}(\phi(\mathbf{f}(z))) = z - \mathbf{A}\mathbf{f}(z) - \frac{1}{2} \begin{pmatrix} \mathbf{f}(z)^\top \mathbf{B}_1 \mathbf{f}(z) \\ \vdots \\ \mathbf{f}(z)^\top \mathbf{B}_{d_z} \mathbf{f}(z) \end{pmatrix} \quad (\text{A.22})$$

988 We can calculate

$$989 \mathbf{g}(\mathbf{f}(z)) = z \quad (\text{A.23})$$

991 so \mathbf{g} is indeed a left-inverse of \mathbf{f} . Taking the derivative of this equation at 0 we obtain

$$992 D\mathbf{h}(0) = \text{Id}_{d_z} - \mathbf{A}(D\mathbf{f}(0)) + 0 = \text{Id}_{d_z} - \mathbf{A}\mathbf{M} = 0. \quad (\text{A.24})$$

994 For the second derivative we get

$$995 D^2\mathbf{h}_s(0) = - \sum_{i=1}^{d_x} \mathbf{A}_{s,i} D^2 \mathbf{f}_i(z) - D\mathbf{f}(0)^\top \mathbf{B}_s D\mathbf{f}(0) = - \sum_{i=1}^{d_x} \mathbf{A}_{s,i} \Lambda^i + \mathbf{M}^\top \mathbf{B}_s \mathbf{M} = 0. \quad (\text{A.25})$$

1000 Here we used for the derivative of the quadratic term that the contribution where the derivative hits
 1001 one $\mathbf{f}(z)$ twice vanishes since $\mathbf{f}(0) = 0$. Finally, we can now evaluate

$$1002 D\mathbf{g}(0) = \mathbf{A} + D\mathbf{h}(\phi(0)) = \mathbf{A} + D\mathbf{h}(0) = \mathbf{A} \quad (\text{A.26})$$

1004 and

$$1005 D^2\mathbf{g}_s(0) = \mathbf{B}_s + D^2\mathbf{h}_s \circ \phi = \mathbf{B}_s \quad (\text{A.27})$$

1006 where $D^2\mathbf{h}_s \circ \phi = 0$ follows from the chain rule and $D\mathbf{h}(0) = 0$ and $D^2\mathbf{h}(0) = 0$. \square

1010 A.2 CONSTRAINING \mathcal{G}_{enc} BY ARCHITECTURE

1011 In this section we discuss results showing that it is challenging to construct practical function classes
 1012 \mathcal{G}_{enc} which are sufficiently expressive so that they contain a left-inverse for each $\mathbf{f} \in \mathcal{F}_{\text{int}}$. As
 1013 explained before, the main challenge is that setting $\mathcal{G}_{\text{enc}} = \mathcal{G}_{\text{int}}$ is in principle sufficient to ensure
 1014 identifiability and out of distribution generalization. So we need to make additional assumptions
 1015 on \mathcal{G}_{enc} that function classes used in widely applied algorithms satisfy which then ensure that \mathcal{G}_{enc}
 1016 is very expressive preventing that equation 2.6 holds. We will make the assumption that \mathcal{G}_{enc} has a
 1017 linear structure, i.e., $\mathbf{g}_1 + \mathbf{g}_2 \in \mathcal{G}_{\text{enc}}$ if $\mathbf{g}_1, \mathbf{g}_2 \in \mathcal{G}_{\text{enc}}$. This is clearly satisfied when \mathcal{G}_{enc} is a vector
 1018 space (e.g., this assumption is satisfied for linear or kernel methods, or when learning a linear head
 1019 on a fixed representation). For functions implemented by neural networks with fixed architecture this
 1020 is in general not true. However, it does apply to infinite width limits of fixed architectures (this does
 1021 not generally imply universal approximation properties when the architecture is sparse, e.g., we use
 1022 slot-wise neural networks for the forward direction which cannot approximate interactions $\mathbf{x}_1 \mathbf{x}_2$
 1023 even at infinite width). Note that large width is also generally required to make neural networks
 1024 sufficiently expressive because for fixed width neural networks implement a parametric function
 1025 class while \mathcal{F}_{int} is non-parametric. We then show that such a function class \mathcal{G}_{enc} does not have a
 useful inductive bias.

Architecture constraints for $d_x = d_z$. We first consider the simpler case $d_x = d_z$ where \mathbf{f} is bijective on the codomain (and not only on its image).

Our main result here is that \mathcal{G}_{enc} has the universal approximation property when $\mathcal{G}_{\text{enc}} \supset \mathcal{F}_{\text{add}}^{-1}$ and \mathcal{G}_{enc} is closed under addition.

Theorem A.8. *Assume $d_x = d_z = d$. Consider an encoder function class \mathcal{G}_{enc} with the following two properties:*

1. *The class \mathcal{G}_{enc} is closed under addition, i.e., for $\mathbf{g}_1, \mathbf{g}_2 \in \mathcal{G}_{\text{enc}}$ also $\mathbf{g}_1 + \mathbf{g}_2 \in \mathcal{G}_{\text{enc}}$.*
2. *The function class \mathcal{G}_{enc} is expressive enough such that it contains all inverses of additive functions, i.e., $\mathcal{F}_{\text{add}}^{-1} \subset \mathcal{G}_{\text{enc}}$.*

Then \mathcal{G}_{enc} is dense in the space of all continuous functions on all compact subset of \mathbb{R}^{d_x} .

Since $\mathcal{F}_{\text{add}} \subset \mathcal{F}_{\text{int}}$ for $n \geq 1$ and any m we directly get the following corollary.

Corollary A.9. *Assume $d_x = d_z = d$ and the encoder function class \mathcal{G}_{enc} is closed under addition and satisfies $\mathcal{G}_{\text{int}} \subset \mathcal{G}_{\text{enc}}$. Then \mathcal{G}_{enc} is dense in the space of all continuous functions on all compact subset of \mathbb{R}^{d_x} .*

The takeaway from these results is that it is challenging to find natural function classes \mathcal{G}_{enc} so that $\mathcal{G}_{\text{enc}} \supset \mathcal{G}_{\text{int}}$ (sufficient expressivity) but \mathcal{G}_{enc} is not much larger than \mathcal{G}_{int} . Therefore, learning only encoders from \mathcal{G}_{enc} does not provide a strong inductive bias towards the inverse of the ground truth decoder and out of distribution generalization.

Proof of Theorem A.8. The general strategy is to prove that the conditions imply that all maps \mathbf{g} where \mathbf{g}_j (the j -th coordinate of \mathbf{g}) is any polynomial and $\mathbf{g}_i = 0$ for $i \neq j$ are contained in \mathcal{G}_{enc} . This will end the proof because polynomials are dense in the scalar valued continuous functions, and we can then apply this result coordinate-wise using the additive structure.

Step 1: Vector space structure of \mathcal{G}_{enc} . We now show that we can scale certain functions in \mathcal{G}_{enc} . Denote for $\lambda \neq 0$ by M_λ the multiplication map $z \rightarrow \lambda z$. Then $\mathbf{f} \circ M_\lambda \in \mathcal{F}_{\text{add}}$ if $\mathbf{f} \in \mathcal{F}_{\text{add}}$. Since $(\mathbf{f} \circ M_\lambda)^{-1} = \lambda^{-1} \mathbf{f}^{-1}$ we conclude that scalar multiples of \mathbf{f}^{-1} are in $\mathcal{F}_{\text{add}}^{-1}$ and the first assumption then implies that the vector space V generated by \mathbf{f}^{-1} for $\mathbf{f} \in \mathcal{F}_{\text{add}}$ is contained in \mathcal{G}_{enc} .

Step 2: We show that the monomials x_i^k are contained in \mathcal{G}_{enc} . Consider the map $\mathbf{f} \in \mathcal{F}_{\text{add}}$ where $\mathbf{x} = \mathbf{f}(z)$ has coordinates

$$\begin{aligned} x_1 &= z_2^k + z_1 \\ x_i &= z_i \quad \text{for } d \geq i \geq 2. \end{aligned} \tag{A.28}$$

This is clearly an additive function with inverse

$$\begin{aligned} z_1 &= x_1 - x_2^k \\ z_i &= x_i \quad \text{for } d \geq i \geq 2. \end{aligned} \tag{A.29}$$

Similarly, we consider

$$\begin{aligned} x_1 &= -(-z_2)^k - z_1 \\ x_i &= -z_i \quad \text{for } d \geq i \geq 2. \end{aligned} \tag{A.30}$$

with inverse

$$\begin{aligned} z_1 &= -x_1 - x_2^k \\ z_i &= -x_i \quad \text{for } d \geq i \geq 2. \end{aligned} \tag{A.31}$$

Summing these two functions, we find that the function \mathbf{g} with

$$\mathbf{g}_i(x) = -2\delta_{1i}x_1^k \tag{A.32}$$

satisfies $\mathbf{g} \in \mathcal{G}_{\text{enc}}$. By permutation of the outputs and inputs (and scaling) we find that all functions \mathbf{g} with $\mathbf{g}_i(x) = \delta_{ij}x_l^k$ are in \mathcal{G}_{enc} for all $j, l \in [d]$ and $k \in \mathbb{N}$.

Step 3: Now we show with a similar argument that more generally functions of the form $g_j(\mathbf{x}) = \delta_{jl} (\sum_{i=1}^d \alpha_i x_i)^k$ for all coefficients α_i and all $1 \leq l \leq d$ are in \mathcal{G}_{enc} . If only one α_i is non-zero we have shown this before, so we can assume that at least two α_i are non-zero and without loss of generality we assume that α_i for $1 \leq i \leq k$ are non-zero where $2 \leq k \leq d$. Then we consider the additive map g which satisfies for $\mathbf{x} = g(\mathbf{z})$

$$\begin{aligned} \mathbf{x}_1 &= \frac{1}{\alpha_1} \left(\mathbf{z}_1^k + \mathbf{z}_1 - \sum_{i=2}^k \mathbf{z}_i \right), \\ \mathbf{x}_2 &= \frac{1}{\alpha_2} (\mathbf{z}_2 - \mathbf{z}_1^k), \\ \mathbf{x}_i &= \frac{1}{\alpha_i} \mathbf{z}_i \quad \text{for } 3 \leq i \leq k, \\ \mathbf{x}_i &= \mathbf{z}_i \quad \text{for } k < i \leq d. \end{aligned} \tag{A.33}$$

Then we observe that

$$\sum_{i=1}^k \alpha_i \mathbf{x}_i = \mathbf{z}_1 \tag{A.34}$$

and thus the inverse is given by

$$\begin{aligned} \mathbf{z}_1 &= \sum_{i=1}^k \alpha_i \mathbf{x}_i, \\ \mathbf{z}_2 &= \alpha_2 \mathbf{x}_2 + \left(\sum_{i=1}^k \alpha_i \mathbf{x}_i \right)^k, \\ \mathbf{z}_i &= \alpha_i \mathbf{x}_i \quad \text{for } 3 \leq i \leq k \\ \mathbf{z}_i &= \mathbf{x}_i \quad \text{for } d \geq i > k. \end{aligned} \tag{A.35}$$

Similarly, we find that the inverse of the additive map given in coordinates by

$$\begin{aligned} \mathbf{x}_1 &= \frac{1}{\alpha_1} \left(-(-\mathbf{z}_1)^k - \mathbf{z}_1 + \sum_{i=2}^k \mathbf{z}_i \right), \\ \mathbf{x}_2 &= \frac{1}{\alpha_2} (-\mathbf{z}_2 + (-\mathbf{z}_1)^k), \\ \mathbf{x}_i &= -\frac{1}{\alpha_i} \mathbf{z}_i \quad \text{for } 3 \leq i \leq k, \\ \mathbf{x}_i &= \mathbf{z}_i \quad \text{for } k < i \leq d. \end{aligned} \tag{A.36}$$

can be written as (note $\sum_{i=1}^k \alpha_i \mathbf{x}_i = -\mathbf{z}_1$)

$$\begin{aligned} \mathbf{z}_1 &= -\sum_{i=1}^k \alpha_i \mathbf{x}_i, \\ \mathbf{z}_2 &= -\alpha_2 \mathbf{x}_2 + \left(\sum_{i=1}^k \alpha_i \mathbf{x}_i \right)^k, \\ \mathbf{z}_i &= -\alpha_i \mathbf{x}_i \quad \text{for } 3 \leq i \leq k \\ \mathbf{z}_i &= -\mathbf{z}_i \quad \text{for } d \geq i > k. \end{aligned} \tag{A.37}$$

Summing the two inverses in equation A.33 and equation A.37 we find that the map g given by $g_j(\mathbf{x}) = 2\delta_{j2} \left(\sum_{i=1}^k \alpha_i x_i \right)^k$ is in \mathcal{G}_{enc} and by permuting the indices and scaling we find that all maps of the form

$$g_j(x) = \delta_{jl} \left(\sum_{i=1}^k \alpha_i x_i \right)^k \tag{A.38}$$

are in \mathcal{G}_{enc} . Using Lemma A.10 stated below we infer that indeed all multinomial polynomials are in \mathcal{G}_{enc} and this ends the proof in light of the Stone-Weierstrass Theorem. \square

The following technical but standard lemma was used in the proof of Theorem A.8.

Lemma A.10. Consider the space of functions $g_{\alpha} : \mathbb{R}^d \rightarrow \mathbb{R}$ for $\alpha \in \mathbb{R}^d$ given by

$$g_{\alpha}(\mathbf{x}) = \left(\sum_{i=1}^d \alpha_i x_i \right)^k. \quad (\text{A.39})$$

Then the vector space generated by the functions g_{α} is the space of all k -homogeneous polynomials.

Proof. This is a general version of the well known polarization identity, namely

$$(\mathbf{x}_1 + \mathbf{x}_2)^2 - (\mathbf{x}_1 - \mathbf{x}_2)^2 = 4\mathbf{x}_1\mathbf{x}_2. \quad (\text{A.40})$$

For completeness we sketch the full proof. Denote the generated space by V . Let $\phi_j(\mathbf{x})$ be linear functions for $1 \leq j \leq k$, i.e., $\phi_j(\mathbf{x}) = \sum_{i=1}^d \alpha_i^j x_i$ for some α_i^j . Then using the multinomial expansion we find

$$\begin{aligned} & \sum_{(\epsilon_1, \dots, \epsilon_k) \in \{-1, 1\}^k} \left(\prod_{j=1}^k \epsilon_j \right) \left(\sum_{j=1}^k \epsilon_j \phi_j \right)^k \\ &= \sum_{(\epsilon_1, \dots, \epsilon_k) \in \{-1, 1\}^k} \left(\prod_{j=1}^k \epsilon_j \right) \sum_{\gamma_1 + \dots + \gamma_k = k} \frac{k!}{\gamma_1! \dots \gamma_k!} \prod_{i=j}^k \phi_j^{\gamma_j} \\ &= \sum_{(\epsilon_1, \dots, \epsilon_k) \in \{-1, 1\}^k} \left(\prod_{i=j}^k \epsilon_j \right) \sum_{\gamma_1 + \dots + \gamma_k = k} \frac{k!}{\gamma_1! \dots \gamma_k!} \prod_{j=1}^k (\epsilon_j \phi_j)^{\gamma_j} \\ &= \sum_{\gamma_1 + \dots + \gamma_k = k} \frac{k!}{\gamma_1! \dots \gamma_k!} \prod_{j=1}^k \phi_j^{\gamma_j} \prod_{j=1}^k \sum_{\epsilon_j \in \{-1, 1\}} \epsilon_j^{\gamma_j + 1}. \end{aligned} \quad (\text{A.41})$$

Now the last sum is 0 for γ_j even and 2 for γ_j odd. So the only non-zero term corresponds to all γ_j odd and thus $\gamma_j = 1$ for all j and therefore

$$\sum_{(\epsilon_1, \dots, \epsilon_k) \in \{-1, 1\}^k} \left(\prod_{j=1}^k \epsilon_j \right) \left(\sum_{j=1}^k \epsilon_j \alpha_j \right)^k = 2^k k! \prod_{j=1}^k \phi_j \in V. \quad (\text{A.42})$$

Clearly this implies that the monomials $\prod_{i=1}^d x_i^{\beta_i}$ with $\beta_i \geq 0$ and $\sum_{i=1}^d \beta_i = k$ are generated by the functions g_{α} (pick $\phi_j(\mathbf{x}) = x_i$ for β_i of the ϕ_j). \square

Architectural constraints for $d_x > d_z$. The case $d_x > d_z$ is more challenging because the data manifold is then a submanifold and even if we know that there is a $g \in \mathcal{G}_{\text{enc}}$ inverting f on the data manifold (i.e., \mathcal{G}_{enc} is sufficiently expressive) this provides (essentially) no information about g away from $\mathcal{X} = f(\mathcal{Z})$ and the data manifolds for different generators are unrelated. However, we can leverage the result for $d_x = d_z$ to obtain a weaker version in the general case. Here we make the additional assumption that \mathcal{G}_{enc} is closed under coordinate projections in the sense that $\tilde{g} \in \mathcal{G}_{\text{enc}}$ if $\tilde{g}(\mathbf{x}) = g(\mathbf{x}_I, \mathbf{0}_{[d_x] \setminus I})$ for some index set $I \subset [d_x]$ and $g \in \mathcal{G}_{\text{enc}}$. Note that this is naturally satisfied for neural networks where we can remove the influence of a coordinate by zeroing its outgoing weights.

Corollary A.11. Assume that \mathcal{G}_{enc} is a class of encoder functions such that \mathcal{G}_{enc} is closed under addition and coordinate projections and sufficiently expressive, i.e., for every $f \in \mathcal{F}_{\text{int}}$ there is $g \in \mathcal{G}_{\text{enc}}$ such that $g \circ f = \text{id}$. Let $f \in \mathcal{F}_{\text{int}}$ be such that f_I is a diffeomorphism (on its image) for some I with $|I| = d_z$. Then $\mathcal{G}_{\text{enc}} \circ f$ is dense in all continuous functions $C(K, \mathbb{R}^{d_z})$ for every compact $K \subset \mathbb{R}^{d_x}$, i.e., essentially arbitrary representations can be learned using function in \mathcal{G}_{enc} .

Proof. Consider a set $I \subset [d_z]$. Then the restrictions f_I of functions $f \in \mathcal{F}_{\text{int}}$ such that $f_{I^c}(z) = \mathbf{0}$ (i.e., functions that vanish in all but d_z coordinates) are in bijection to functions in \mathcal{F}_{int} mapping $\mathbb{R}^{d_z} \rightarrow \mathbb{R}^{d_z}$. Applying Theorem A.8 we therefore find that the set of functions $z_I \rightarrow g(z_I, \mathbf{0}_{I^c})$ for $g \in \mathcal{G}_{\text{enc}}$ is dense in the continuous functions defined on any compact set K' . It is convenient to introduce the shorthand \tilde{g} for the function $z_I \rightarrow g(z_I, \mathbf{0}_{I^c})$ by \tilde{g} . Then we can restate the density statement before as follows: Given any continuous function $h : K \rightarrow \mathbb{R}^{d_z}$ we can find for any $\epsilon > 0$ a $g \in \mathcal{G}_{\text{enc}}$ so that $\|\tilde{g} - h \circ (f_I)^{-1}\| < \epsilon$ on the compact set $K' = f_I(K)$ (here we use that f_I is bijective on its image to invert it). Using that \mathcal{G}_{enc} is closed under coordinate projections we can find \tilde{g} is in \mathcal{G}_{enc} and satisfies

$$\begin{aligned} \max_{z \in K} \|\tilde{g}f(z) - h(z)\|_\infty &= \|g(f_I(z), \mathbf{0}_{I^c}) - h \circ (f_I)^{-1} \circ f_I(z)\|_\infty \\ &= \|\tilde{g}(f_I(z)) - h \circ (f_I)^{-1}(f_I(z))\|_\infty \\ &\leq \max_{x_I \in K'} \|\tilde{g}(x_I) - h \circ (f_I)^{-1}(x_I)\|_\infty. \end{aligned} \tag{A.43}$$

This ends the proof. □

While the previous corollary makes the strong assumption that f_I is globally bijective we note that this is generally true at least locally. Moreover, we can patch such representations due to the additive structure. Therefore, it seems unlikely that a function class \mathcal{G}_{enc} satisfying the following three constraints exists: First, the function class \mathcal{G}_{enc} is expressive enough, i.e., it contains left inverses for all $f \in \mathcal{F}_{\text{int}}$. Secondly, \mathcal{G}_{enc} is not too expressive so it provides a useful inductive bias towards \mathcal{G}_{int} . And finally, \mathcal{G}_{enc} can be efficiently parametrized and used for optimization.

B EXPERIMENTAL DETAILS

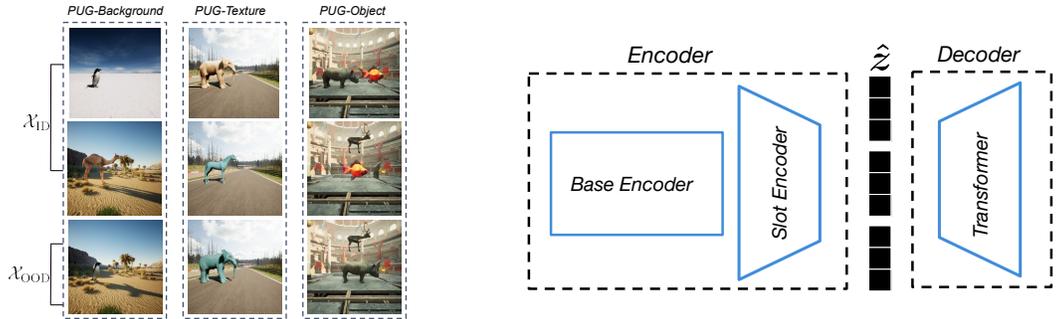


Figure 7: Left. Overview of the data splits used in the experiments. PUG-Background contains unseen combinations of background and object in its OOD split \mathcal{X}_{OOD} , PUG-Texture contains unseen object-texture combinations in \mathcal{X}_{OOD} , and PUG-Object contains unseen object combinations in \mathcal{X}_{OOD} . Right. General structure of the employed models. A *base encoder* (pretrained for most experiments) is used to extract features from the images which are then mapped through a *slot encoder* which leverages a cross-attention mechanism and potentially self-attention. For our decoders we use the regularized cross-attention Transformer architecture from Brady et al. (2025)

B.1 DATA

We create datasets for our experiments in Sec. 5 based on the PUG: Animals dataset (Bordes et al., 2023). This data consists of 43,520 high-resolution images which we resize to $224 \times 224 \times 3$. To create PUG-Background, we create an OOD set containing 32,000 images which consist of unseen combinations of animal category and backgrounds, e.g., penguin in a desert in Fig. 7, and a corresponding ID set containing 11,520 images. For PUG-Texture, the OOD set contains 16,000 images consisting of unseen combinations of animal and texture/color, e.g. blue elephant in Fig. 7, and the ID set contains 27,520 images. Lastly, for PUG-Object, the ID and OOD set both contain 21,760

1242 images. The OOD set here consist of unseen combinations of animal categories, e.g., rhinoceros
1243 and caribou in in Fig. 7.

1244 B.2 MODELS

1245 **Base encoders.** We use six different pretrained base encoders along with a ViT small for our
1246 experiments in Sec. 5. The specific sizes of each model can be seen in Fig. 5. When fine tuning these
1247 models with a LORA adapter, we use a rank of 16, a scaling factor of 32, and a dropout value of 0.1.

1248 **Slot encoders.** We use either a Transformer or Slot Attention model for the slot encoder in Sec. 5.
1249 The transformer model consist of both self and cross-attention layers. Both models consist of 5
1250 layers. We use 3 slots for each model, with dimensions of 64.

1251 **Decoders.** All decoders in our experiments use the cross-attention Transformer from Brady et al.
1252 (2025); Jaegle et al. (2022); Sajjadi et al. (2022). In this model, slots are first projected by a 2 layer
1253 slot-wise MLP and then passed through 2 layers of a cross-attention Transformer with pixel queries.
1254 Pixels are tokenized using a 2 layer MLP. The pixel outputs of the cross-attention Transformer are
1255 mapped to have a channel dimension of 3 using a 3 layer MLP. The attention weights in this model
1256 are regularized using the regularizer introduced by Brady et al. (2025). For all experiments, we use
1257 a value of 0.01 for this loss.

1258 **Readout.** We use a single layer linear readout shared across slots to predict animal or background
1259 categories from slots.

1260 B.3 TRAINING OBJECTIVES

1261 **Supervised models.** We train all supervised models for 100000 iterations across 3 random seeds
1262 using a batch size of 64, with the Adam optimizer (Kingma and Ba, 2015) and a learning rate of
1263 $1e-4$.

1264 **VAEs.** We train all unsupervised VAE models for 300000 iterations across 3 random seeds using a
1265 batch size of 32, with the Adam optimizer (Kingma and Ba, 2015) and a learning rate of $5e-4$, which
1266 is decayed by a factor of .1 throughout training and warmed up for the first 10000 iterations. We use
1267 a value of either 0.005 or 0.001 for the hyperparameter β on the KL loss.

1268 **Readout.** In our unsupervised experiments, we train a linear readout on learned slots for 7500
1269 iterations. To resolve the permutation between inferred and ground-truth slots we rely on on the
1270 Hungarian matching procedure used in Dittadi et al. (2022); Locatello et al. (2020b).

1271 **Gradient-based search.** When performing gradient based search in our experiments, we optimize
1272 Eq. 4.3 using Adam with a learning rate of .001. We optimize for either 300 or 500 iterations on
1273 PUG-Background and 700 iterations on PUG-Texture. To further aid in optimization we add an
1274 additional regularizer to the optimization procedure which minimizes the entropy of the logits under
1275 the classifier. This aims to ensure that the search procedure yields latent slots which are within the
1276 set of slots which the classifier has already observed. We use a value of either 10 or 50 for this loss.
1277 We note that a similar loss was used for semi-supervised learning in Grandvalet and Bengio (2004).

1278 **Generative replay.** For our experiments using generative replay, we generate OOD data by follow-
1279 ing the procedure in Wiedemer et al. (2024a) in which ID slots are randomly shuffled to create novel
1280 compositions. We train an encoder on batches of 64 of OOD samples for 15000 iterations with a
1281 learning rate of $5e-4$.

1282 **Compute.** We train all models using 2 NVIDIA A100 GPUs. Total training time was approximately
1283 1500h.

1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349

```

import torch
from torch.optim import Adam

def latent_search(x_ood, encoder, decoder, learning_rate, num_iters):
    # get encoding for ood image
    zh = encoder(x_ood)
    optimizer = Adam(params=[zh],
                      lr=learning_rate)

    # optimize encoding to minimize mse under decoder
    for _ in range(num_iters):
        xh = decoder(zh)
        mse_loss = (x_ood - xh).square().mean()
        mse_loss.backward()
        optimizer.step()

    return zh

def generative_replay(zh_ood, encoder, decoder, learning_rate):
    optimizer = Adam(params=encoder.parameters(),
                     lr=learning_rate)

    # ensure gradients not computed for decoder
    for param in decoder.parameters():
        param.requires_grad = False

    # generate OOD image
    xh_ood = decoder(zh_ood)

    # re-encode image with encoder
    zh_ood_recon = encoder(xh_ood)

    # ensure re-encoding matches original encoding
    mse_loss = (zh_ood - zh_ood_recon).square().mean()
    mse_loss.backward()
    optimizer.step()

    return encoder

```

Figure 8: Top. PyTorch pseudocode for gradient-based search using a decoder for a given OOD image x . Bottom. PyTorch pseudocode for one gradient step of generative replay.

C ADDITIONAL EXPERIMENTS

Setup.

Our experiments in § 5 highlight that generative methods can yield significant gains in compositional generalization. However, these results are only shown for one decoder architecture. Consequently, it remains unclear whether gains in OOD performance come from our specific setup, which leverages a structured slot-based decoder, or whether a less-structured decoder can yield similar gains. To test this, we train an autoencoder on PUG-Background with the following structure: images are mapped by a CNN encoder to a latent representation with the same total dimensionality as our slot-based model, but which does not have a slot structure. This latent is then decoded by a CNN decoder which is not slot-based. We train this model on PUG-Background across 3 seeds using the same training settings as § 5 and with a VAE loss with $\beta = 0.005$. For the CNN encoder and decoder, we leverage the same architecture used in Dittadi et al. (2022). To evaluate classification performance, we train 3 separate linear readouts on the latents of this model to predict each animals and background. To evaluate whether this decoder offers gains in OOD performance, we perform 750 iterations of search § 4.1 under the trained decoder and re-evaluate OOD accuracy.

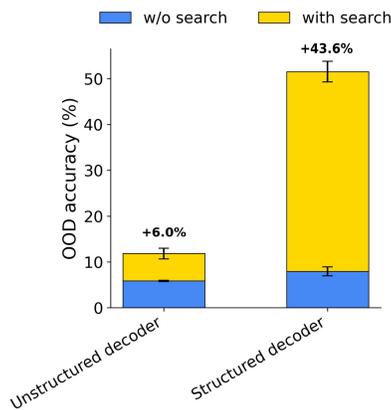


Figure 9: OOD accuracy on PUG-Background after applying search § 4.1 using a structured slot-based Transformer decoder designed according to § 3 vs. an unstructured CNN decoder.

We benchmark the performance of this unstructured model against our structured Transformer decoder from § 5 which uses a ViT-Small base encoder trained from scratch. Similarly, we perform 750 iterations of search under this decoder, and re-evaluate OOD performance.

Results.

For both models, we observe near-perfect ID classification performance. Thus, similar to § 5, we only report OOD performance. Results can be seen in Fig. 9. We find that both models achieve similar OOD accuracy before leveraging search. However, after applying search, we observe significantly higher gains in OOD performance when using the structured Transformer decoder compared to the unstructured CNN. These results suggest that non-trivial gains in compositional generalization using generative methods require constraining a decoder’s structure as highlighted by our theory in § 3.

D EXTENDED DISCUSSION

More complex datasets. One limitation of our experimental study is that we do not evaluate compositional generalization on real-world image datasets. The main issue in leveraging such datasets is that evaluating compositional generalization requires data in which all ground-truth latent factors in each image (e.g. objects, backgrounds, textures, etc.) are explicitly known, thereby enabling us to construct ID and OOD splits in a principled way. For real-world, unstructured data the underlying latents are unknown making such datasets unsuitable for rigorously testing compositional generalization. Consequently, we elected to use PUG, which to the best of our knowledge is the most visually complex data in which the underlying latents are known. Yet, this dataset is still limited in its visual complexity relative to real world data. We thus believe that an important direction for future work is to create a large scale, visually realistic image dataset which offers access to the ground-truth latents, perhaps by leveraging a synthetic or neural network based image renderer.

Compositional generalization via data scale. A key result of our experiments in § 5 is that non-generative methods leveraging encoders with large-scale pretraining achieve substantial gains in OOD performance. This suggests that, as the pretraining size of the base encoder increases, the full encoder (base encoder plus slot encoder) becomes initialized in a region of the loss landscape where all reachable optima correspond to models that also generalize OOD. A limitation of our current experiments is that we cannot investigate this phenomenon in a more systematic manner, since we rely on pretrained encoders that differ in architecture, training objective, and pretraining dataset. A more principled approach would require an empirical study in which an encoder’s architecture and training objective are held fixed, while the model is trained on gradually increasing amounts of pretraining data. We leave such an investigation as an important direction for future work.

Computational cost of generative methods. One drawback of leveraging generative methods is that training a decoder as well as using search or replay add computational overhead relative to encoder-only methods. For example, in our experiments in § 5.2, a single iteration for a supervised ViT-Small base encoder took ~ 55 ms for a batch size of 32, while adding a Transformer decoder and training with a VAE loss increased this time to ~ 366 ms. This overhead can be reduced by only decoding a subset of pixels similar to van Steenkiste et al. (2024), but we leave this for future work. Additionally, using generative replay incurs approximately ~ 170 ms per iteration, while search requires ~ 215 ms with a batch size of 32. We note that we train for only 15,000 replay iterations, making this procedure relatively computationally scalable. Search, on the other hand, requires 750 iterations per batch. Thus, further work is needed to scale this procedure more effectively.

Limitations of assumptions in theory. Our theoretical results in § 3, rely on two main assumptions on the ground-truth generator f : (i) that f is a diffeomorphism and (ii) that the interactions between slots under f are restricted according to Eq. (2.7). While these assumptions are the most general which have been shown to enable compositional generalization, it is possible that they may fail to accurately model real-world data in certain circumstances. For example, the assumption that f is a diffeomorphism may not hold for images in which objects exhibit strong occlusions such that f is

1404 no longer invertible. Additionally, it is possible that certain interactions between concepts may not
1405 be captured by Eq. (2.7) such as complex shadows or reflections.
1406

1407 **Limitations of assumptions in practice.** Enforcing our theoretical assumptions on a model in prac-
1408 tice can potentially pose scalability challenges. For example, enforcing that a model is a diffeomor-
1409 phism via an autoencoder is computationally costly and thus can be challenging to scale. We note,
1410 however, that modern generative diffusion models and classifiers (Jaini et al., 2024; Peebles and Xie,
1411 2023; Rombach et al., 2022) rely on learning a diffeomorphic mapping between image and latent
1412 space, highlighting the feasibility of enforcing this assumption at scale. Furthermore, enforcing that
1413 a decoder takes the form of Eq. (2.7) exactly is not scalable in practice (Brady et al., 2025). To this
1414 end, we aim to approximate this form in our experiments using a structured Transformer decoder
1415 and restricting interactions via a KL penalty and a sparsity regularizer on the decoder’s attention
1416 weights (Brady et al., 2025). Importantly, we find this model yields significant gains for composi-
1417 tional generalization highlighting that enforcing that a decoder exactly matches Eq. (2.7) may not
1418 be necessary in practice. Furthermore, we suspect that biologically plausible regularizers such as
1419 in Whittington et al. (2023) which aim to minimize activation energy through sparsity regularizers
1420 on the latents and the decoder’s weight matrix could serve a similar role as our regularizers.
1421

1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457