

# INTERACTION ASYMMETRY: A GENERAL PRINCIPLE FOR LEARNING COMPOSABLE ABSTRACTIONS

Jack Brady<sup>\*1,2</sup> Julius von Kügelgen<sup>3</sup> Sébastien Lachapelle<sup>4</sup>

Simon Buchholz<sup>1,2</sup> Thomas Kipf<sup>†5</sup> Wieland Brendel<sup>†1,2,6</sup>

<sup>1</sup> Max Planck Institute for Intelligent Systems, Tübingen <sup>2</sup> Tübingen AI Center <sup>3</sup> ETH Zürich

<sup>4</sup> Samsung - SAIT AI Lab, Montreal <sup>5</sup> Google DeepMind <sup>6</sup> ELLIS Institute, Tübingen

## ABSTRACT

Learning disentangled representations of concepts and re-composing them in unseen ways is crucial for generalizing to out-of-domain situations. However, the underlying properties of concepts that enable such disentanglement and compositional generalization remain poorly understood. In this work, we propose the principle of *interaction asymmetry* which states: “Parts of the same concept have more complex interactions than parts of different concepts”. We formalize this via block diagonality conditions on the  $(n+1)$ <sup>th</sup> order derivatives of the generator mapping concepts to observed data, where different orders of “complexity” correspond to different  $n$ . Using this formalism, we prove that interaction asymmetry enables *both* disentanglement and compositional generalization. Our results unify recent theoretical results for learning concepts of objects, which we show are recovered as special cases with  $n = 0$  or  $1$ . We provide results for up to  $n = 2$ , thus extending these prior works to more flexible generator functions, and conjecture that the same proof strategies generalize to larger  $n$ . Practically, our theory suggests that, to disentangle concepts, an autoencoder should penalize its latent capacity and the interactions between concepts during decoding. We propose an implementation of these criteria using a flexible Transformer-based VAE, with a novel regularizer on the attention weights of the decoder. On synthetic image datasets consisting of objects, we provide evidence that this model can achieve comparable object disentanglement to existing models that use more explicit object-centric priors.

## 1 INTRODUCTION

A core feature of human cognition is the ability to use abstract conceptual knowledge to generalize far beyond direct experience (Behrens et al., 2018; Mitchell, 2021; Murphy, 2004; Tenenbaum et al., 2011). For example, by applying abstract knowledge of the concept “chair”, we can easily infer how to use a “chair on a beach”, even if we have not yet observed this combination of concepts. This feat is non-trivial and requires solving two key problems. Firstly, one must acquire an abstract, internal model of different concepts in the world. This implies learning a *separate* internal representation of each concept from sensory observations. Secondly, these representations must remain valid when observations consist of novel compositions of concepts, e.g., “chair” and “beach”. In machine learning, these two problems are commonly referred to as learning *disentangled representations* (Bengio et al., 2013; Higgins et al., 2018; Schölkopf et al., 2021) and *compositional generalization* (Fodor and Pylyshyn, 1988; Goyal and Bengio, 2022; Greff et al., 2020; Lake et al., 2017).

Both problems are known to be challenging due to the issue of *non-identifiability* (Hyvärinen et al., 2023). Namely, many models can explain the same data equally well, but only some will learn representations of concepts which are disentangled and generalize compositionally. To guarantee *identifiability* with respect to (w.r.t.) these criteria, it is necessary to incorporate suitable inductive biases into a model (Hyvärinen and Pajunen, 1999; Lachapelle et al., 2023; Locatello et al., 2019). These inductive biases, in turn, must reflect some underlying properties of the concepts which give rise to observed data. This raises a fundamental question: What properties of concepts enable learning models which provably achieve disentanglement and compositional generalization?

\*Correspondence to: jack.brady@tue.mpg.de. †Joint senior author.  
Code available at: [github.com/JackBrady/interaction-asymmetry](https://github.com/JackBrady/interaction-asymmetry)

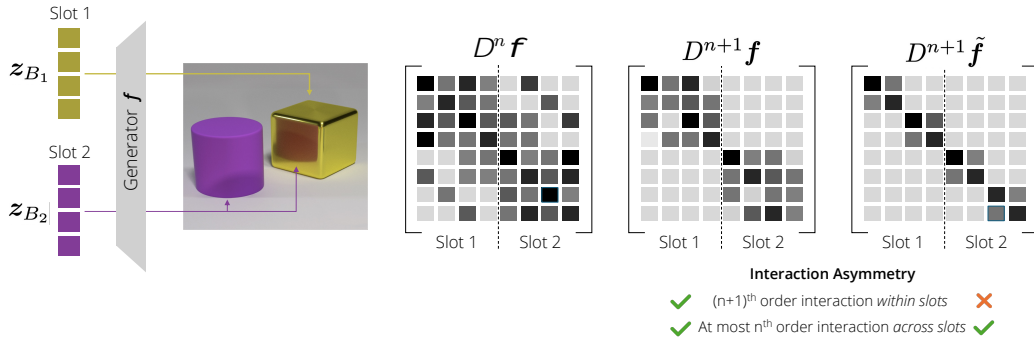


Figure 1: **Illustration of Interaction Asymmetry.** (Left) Observations  $\mathbf{x}$  result from a generator  $\mathbf{f}$  applied to latent slots  $\mathbf{z}_{B_k}$  that represent separate concepts. As indicated by the reflection of the cylinder upon the cube, slots can interact during generation. Our key assumption, interaction asymmetry, states that these interactions across slots must be less complex than interactions within the same slot. (Right) This is formalized by assuming block-diagonality *across* but not *within* slots for the  $(n+1)^{\text{th}}$  order derivatives of the generator, i.e.,  $D^{n+1} \tilde{\mathbf{f}}$ .

Many works aim to answer this question by studying properties enabling either disentanglement or compositional generalization *in isolation*. This is insufficient, however, as disentanglement alone does not imply compositional generalization (Montero et al., 2022a; 2021; Schott et al., 2022), while compositional generalization requires first disentangling the concepts to be composed. Only a few studies investigate properties enabling *both* disentanglement and compositional generalization (Brady et al., 2023; Lachapelle et al., 2023; Wiedemer et al., 2024a). Yet, the properties proposed in these works are rather restrictive and specific to objects in simple visual scenes. There is growing evidence, however, that the principles humans use to learn conceptual knowledge are not concept-specific, but shared across different concepts (objects, attributes, events, etc.) (Behrens et al., 2018; Constantinescu et al., 2016; Hawkins et al., 2018). This suggests there exist some *general* properties of concepts which enable both disentanglement and compositional generalization.

In this work, we seek to formulate such a general property for disentangling and composing concepts. We begin by aiming to deduce, from first principles, properties which are fundamental to concepts (§ 3). From this, we arrive at the guiding principle of *interaction asymmetry* (Principle 3.1) stating: “Parts of the same concept have more complex interactions than parts of different concepts”. As illustrated in Fig. 1 (left), we define concepts as distinct groups, or *slots*, of latent variables which generate the observed data (§ 2). Interaction asymmetry is then formalized as a block-diagonality condition *across* but *not within* slots of  $D^{n+1} \mathbf{f}$ , the tensor of  $(n+1)^{\text{th}}$  order partial derivatives of the generator function (Asm. 3.5), where  $n$  determines the complexity of interactions, see Fig. 1 (right).

**Theory.** Using this formulation, we prove that interaction asymmetry dually enables *both* disentanglement (Thm. 4.3) and compositional generalization (Thm. 4.4). We also show that our formalism provides a unifying framework for prior results of Brady et al. (2023) and Lachapelle et al. (2023), by proving that the properties studied in these works for visual objects are special cases of our assumptions for  $n=0$  and 1, respectively. We provide results for up to  $n=2$ , thereby extending these prior works to more general function classes, and conjecture that our results generalize to arbitrary  $n \geq 0$ .

**Method.** Our theory suggests that to disentangle concepts, a model should (i) enforce invertibility, without using more latent dimensions than necessary, and (ii) penalize interactions across slots during decoding. To translate these insights into a practical method, we leverage a VAE loss (Kingma and Welling, 2014) for (i), and observe that the Transformer architecture (Vaswani et al., 2017) offers an approximate means to achieve (ii) since interactions are determined by the attention weights of the model. To this end, we introduce an inexpensive interaction regularizer for a cross-attention mechanism, which we incorporate, with the VAE loss, into a flexible Transformer-based model (§ 5).

**Empirical Results.** We test this model’s ability to disentangle concepts of visual objects on a Sprites dataset (Watters et al., 2019a) and on CLEVR6 (Johnson et al., 2017). We find that the model reliably learns disentangled representations of objects, improving performance over an unregularized Transformer (§ 6). Furthermore, we provide preliminary evidence that our regularized Transformer can achieve comparable performance to models with more explicit object-centric priors such as Slot Attention (Locatello et al., 2020b) and Spatial Broadcast Decoders (Watters et al., 2019b).

**Notation.** We write scalars in lowercase ( $z$ ), vectors in lowercase bold ( $\mathbf{z}$ ), and matrices in capital bold ( $\mathbf{M}$ ).  $[K]$  stands for  $\{1, 2, \dots, K\}$ .  $D_i$  and  $D_{i,j}^2$  stand for the first- and second-order partial derivatives with respect to (w.r.t.)  $z_i$  and  $(z_i, z_j)$ , respectively. If  $B \subseteq [n]$  and  $\mathbf{z} \in \mathbb{R}^n$ ,  $\mathbf{z}_B$  denotes the subvector  $(z_i)_{i \in B}$  indexed by  $B$ . A function is  $C^n$  if it is  $n$ -times continuously differentiable.

## 2 BACKGROUND

We begin with formalizing the core ideas of concepts, disentanglement, and compositional generalization, mostly following the setup of Lachapelle et al. (2023). To begin, we assume that the observed data  $x \in X \subseteq \mathbb{R}^{d_x}$  results from applying a diffeomorphic generator:  $Z \rightarrow X$  to latent vectors  $z \in Z := \mathbb{R}^{d_z}$ , sampled from some distribution  $p_z$ . Concepts underlying (objects, attributes, events, etc.) are then modelled as disjoint groups or slots of latents  $z_{B_k}$  such that  $z = (z_{B_1}, \dots, z_{B_K})$ , where  $B_k \subseteq [d_z]$ . We assume that  $z$  is only supported on a subset  $Z_{\text{supp}} \subseteq Z$  which gives rise to observed data  $x_{\text{supp}} := f(Z_{\text{supp}})$ . This generative process can be summarized as:

$$x = f(z); \quad z \sim p_z; \quad \text{supp}(p_z) = Z_{\text{supp}}. \quad (2.1)$$

Next, consider a model  $\hat{f}: Z \rightarrow \mathbb{R}^{d_x}$  trained to be invertible from  $X_{\text{supp}}$  to  $Z_{\text{supp}} := \hat{f}^{-1}(X_{\text{supp}})$ , whose inverse  $\hat{f}^{-1}$  maps a representation  $x \in X_{\text{supp}}$  to  $z \in Z_{\text{supp}}$ . This model is said to learn a disentangled representation  $z \in Z_{\text{supp}}$  if each model slot  $z_{B_j}$  captures exactly one concept  $c_{B_j}$ .

**Definition 2.1 (Disentanglement)** Let  $f: Z \rightarrow X$  be a diffeomorphism and  $\hat{f}: Z \rightarrow \mathbb{R}^{d_x}$ . A model  $\hat{f}$  disentangles  $z$  on  $Z$  w.r.t.  $f$  if there exist a permutation of  $[K]$  and slot-wise diffeomorphism  $h = (h_1, \dots, h_K)$  with  $h_k: \mathbb{R}^{B_{(k)j}} \rightarrow \mathbb{R}^{B_{kj}}$  and  $j_{B_{(k)j}} = j_{B_{kj}}$  such that for all  $z \in Z$ :

$$\hat{f}(h_1(z_{B_{(1)j}}), \dots, h_K(z_{B_{(K)j}})) = f(z). \quad (2.2)$$

In other words, a representation is disentangled if the model inverts the generator up to permutation and reparametrization of the slots. For compositional generalization we would like this to hold not only on  $Z_{\text{supp}}$  but also for arbitrary combinations of the slots therein. Namely, also on the set

$$Z_{\text{CPE}} := Z_1 \times Z_2 \times \dots \times Z_K; \quad \text{with} \quad Z_k := f(z_{B_{kj}} | z \in Z_{\text{supp}}) \quad (2.3)$$

where  $Z_k$  denote the marginal supports of  $f$  and  $Z_{\text{CPE}}$  the Cartesian-product extension (Lachapelle et al., 2023) of  $Z_{\text{supp}}$ . In general  $Z_{\text{supp}}$  is a subset of  $Z_{\text{CPE}}$ . Thus, to generalize compositionally, a model must also achieve disentanglement “out-of-domain” on novel compositions of slots in

**Definition 2.2 (Compositional Generalization)** Let  $f: Z \rightarrow X$  be a diffeomorphism. A model  $\hat{f}$  that disentangles  $z$  on  $Z_{\text{supp}}$  w.r.t.  $f$  (Defn. 2.1) generalizes compositionally if it also disentangles  $z$  on  $Z_{\text{CPE}}$  w.r.t.  $f$ .

**On the Necessity of Inductive Biases.** It is well known that only a small subset of invertible models achieve disentanglement on  $Z_{\text{supp}}$  (Hyvärinen and Pajunen, 1999; Locatello et al., 2019) or generalize compositionally on  $Z_{\text{CPE}}$  (Lachapelle et al., 2023). To provably achieve these goals (without explicit supervision), we thus need to further restrict the space of permissible models, i.e., place additional assumptions on the generative process in Eq. (2.1). Such assumptions then translate into inductive biases on a model. To this end, the core challenge is formulating assumptions on  $f$  which faithfully reflect properties of concepts, while sufficiently restricting the problem.

**Assumptions on  $p_z$ .** To guarantee disentanglement, several assumptions have been proposed, such as conditional independence of latents given an auxiliary variable (Hyv et al., 2019; Khemakhem et al., 2020); particular temporal (Vä and Hyvarinen, 2020; Hyvärinen and Morioka, 2016; 2017; Klindt et al., 2021), spatial (Vä et al., 2021; 2024), or other latent structures (Kivva et al., 2022; Kori et al., 2024); multiple views (Ahuja et al., 2022; Brehmer et al., 2022; Gresele et al., 2020; Locatello et al., 2020a; von Kügelgen et al., 2021; Yao et al., 2024; Zimmermann et al., 2021); or interventional information (Buchholz et al., 2023; Lachapelle et al., 2022; 2024; Lippe et al., 2022; 2023; Varici et al., 2024; von Kügelgen et al., 2023). While sufficient for disentanglement, such assumptions do not guarantee compositional generalization. The latter requires that the behavior of the generator on  $Z_{\text{CPE}}$  can be determined solely from its behavior on  $Z_{\text{supp}}$  (see Defn. 2.2). In the most extreme case, where the values of each slot are seen only once on  $Z_{\text{supp}}$  will be a one-dimensional manifold embedded in  $\mathbb{R}^{d_z}$  while  $Z_{\text{CPE}}$  is always  $d_z$ -dimensional. This highlights that generalizing from  $Z_{\text{supp}}$  to  $Z_{\text{CPE}}$  is only possible if the form of the generator is restricted.

**Assumptions on  $f$ .** Restrictions of  $f$  which enable compositional generalization have been proposed by Dong and Ma (2022); Lippl and Stachenfeld (2024); Wiedemer et al. (2024b). Yet, these results rely on quite limited function classes and do not address disentanglement, assuming it to be solved a priori. Conversely, several works explore restriction through orthogonality (Buchholz et al., 2022; Gresele et al., 2021; Horan et al., 2021) or sparsity (Leemann et al., 2023; Moran

et al., 2022; Zheng and Zhang, 2023) which address disentanglement but not compositional generalization. More recently, Brady et al. (2023) and Lachapelle et al. (2023) proposed assumptions on which enable both disentanglement and compositional generalization (Wiedemer et al., 2024a). Yet, these assumptions are overly restrictive such that only model limited types of concepts, e.g., non-interacting objects, and not more general concepts. We discuss these two works further in § 4.3.

### 3 THE INTERACTION ASYMMETRY PRINCIPLE

In this section, we attempt to formulate assumptions that enable disentanglement and compositional generalization, while capturing more general properties of concepts. To approach this, we take a step back and try to understand what are the defining properties of concepts. Specifically, we consider the question: Why are some structures in the world recognized as different concepts (e.g., apple vs. dog) and others as part of the same concept? We propose an answer to this for concepts grounded in sensory data, such as objects (e.g., “car”), events (e.g., “making coffee”), or attributes (e.g., “color”).

Sensory-grounded concepts correspond to reoccurring visual or temporal patterns that follow an abstract template. They tend to be modular, such that independently changing one concept generally leaves the structure of other concepts intact (Greff et al., 2015, § 4.1.1; Peters et al., 2017). For example, a car can change position without affecting the structure of the street, buildings, or people around it. Thus, different concepts appear, in some sense, to interact.

On the other hand, parts of the same concept do not seem to possess this modularity. Namely, arbitrarily changing one part of a concept without adjusting other parts is generally not possible without destroying its inherent structure. For example, it is not possible to change the position of the front half of a car, while maintaining something we would still consider a car, without also changing the back half’s position. Thus, parts of the same concept seem to interact.

This may then lead us to answer our initial question with: Parts of the same concept interact, while different concepts do not. However, this is an oversimplified view, as parts of different concepts can, in fact, interact. For example, in Fig. 1 we see the purple cylinder rests upon and thus interacts with the golden cube. However, such interactions across concepts appear somehow simpler than interactions within a concept: whereas the latter can alter the concept’s structure, the former generally will not. In other words, the complexity of interaction within and across concepts appears to be asymmetric. We formulate this as the following principle (see Appx. G.1 for related principles).

**Principle 3.1 (Interaction Asymmetry)** Parts of the same concept have more complex interactions than parts of different concepts.

To investigate the implications of Principle 3.1 for disentanglement and compositional generalization, we must first give it a precise formalization. To this end, we need a mathematical definition of the “complexity of interaction” between parts of concepts, i.e., groups of latents from the same or different slots. This can be formalized either through assumptions on the latent distribution on the generator. Since the latter are essential for compositional generalization, this is our focus.

Let us start by imagining what it would mean if two groups of latent components  $z_A$  and  $z_B$  interact with no complexity, i.e., have no interaction within  $f$ . A natural way to formalize this is that  $z_A$  and  $z_B$  affect distinct output components. Mathematically, this is captured as follows.

**Definition 3.2 (At most 0<sup>th</sup> order/No interaction)** Let  $f : Z \rightarrow X$  be  $C^1$ , and let  $A, B \subseteq [d_z]$  be non-empty.  $z_A$  and  $z_B$  have no interaction within  $f$  if for all  $z \in Z$ , and all  $i \in A, j \in B$ :

$$D_i f(z) \cdot D_j f(z) = 0: \quad (3.1)$$

To define the next order of interaction complexity, we assume that  $z_A$  and  $z_B$  do interact, i.e., they affect the same output, such that  $D_i f(z)$  and  $D_j f(z)$  are non-zero for some  $i \in A, j \in B$ . This interaction, however, should have the lowest possible complexity. A natural way to capture this is to say that  $z_i$  can affect the same output as  $z_j$  but cannot affect the way in which  $f$  depends on  $z_j$ . Since the latter is captured by  $D_j f(z)$ , this amounts to a question about the order derivative  $D_{ij}^2 f$ . We thus arrive at the following definition for the next order of interaction complexity.

**Definition 3.3 (At most 1<sup>st</sup> order interaction)** Let  $f : Z \rightarrow X$  be  $C^2$ , and let  $A, B \subseteq [d_z]$  be non-empty.  $z_A$  and  $z_B$  have at most 1<sup>st</sup> order interaction within  $f$  if for all  $z \in Z$ , and all  $i \in A, j \in B$ :

$$D_{ij}^2 f(z) = 0: \quad (3.2)$$

Using the same line of reasoning, we can continue to define interactions at increasing orders of complexity. For example, for at most  $2^{nd}$  order interaction,  $z_i$  can affect the derivative  $D_j f_i$ , such that  $D_{ij}^2 f_i(z) \neq 0$ , but cannot affect the way in which  $D_j f_i$  depends on any other  $z_k$ , i.e.,  $D_{ij;k}^3 f_i(z) = 0$ . This leads to a general definition of interactions with most<sup>th</sup> order complexity.

**Definition 3.4 (At most<sup>n</sup>th order interaction)** Let  $n \geq 1$  be an integer. Let  $f : Z \rightarrow X$  be  $C^{n+1}$ . Let  $A; B \subseteq [d_z]$  be non-empty,  $z_A$  and  $z_B$  have at most<sup>n</sup>th order interaction within  $f$  if for all  $z \in Z$ , all  $i \in A; j \in B$ , and all multi-indices  $\alpha \in N_0^{d_z}$  with  $|\alpha| = n + 1$  and  $1 \leq i; j$ :

$$D^\alpha f(z) = 0: \tag{3.3}$$

In other words,  $z_A$  and  $z_B$  have at most<sup>n</sup>th order interaction with  $f$  if all higher-than<sup>n</sup>th order cross partial derivatives w.r.t. at least one component of  $z_A$  and of  $z_B$  are zero everywhere. Otherwise, if the statement in Defn. 3.4 does not hold for some  $z$ ,  $i \in A$ , and  $j \in B$ , we say that  $z_A$  and  $z_B$  have  $(n+1)^{th}$  order interaction at  $z$  (and similarly for 1<sup>st</sup> order interaction if Defn. 3.2 does not hold). With these definitions, we can now provide a precise formalization of Principle 3.1.

**Assumption 3.5 (Interaction asymmetry (formal))** There exists  $n \in N_0$  such that (i) any two distinct slots  $z_{B_i}$  and  $z_{B_j}$  have at most<sup>n</sup>th order interaction within  $f$ ; and (ii) for all  $z \in Z$ , all slots  $z_{B_k}$  and all non-empty  $A; B$  with  $B_k = A \cup B$ ,  $z_A$  and  $z_B$  have  $(n+1)^{th}$  order interaction within  $f$  at  $z$ .

We emphasize that Asm. 3.5 (ii) does not state that subsets of latents within a slot must have  $(n+1)^{th}$  order interaction, but only that a slot cannot split into two parts with at most<sup>n</sup>th order interaction, see Fig. 1 (right). We also note that condition (ii) must hold uniformly over  $Z$ , which resembles the notion of “uniform statistical dependence” among latents introduced by [Byrd and Morioka \(2017, Defn. 1\)](#). For further discussions of Asm. 3.5, see Appx. H.1.

## 4 THEORETICAL RESULTS

We now explore the theoretical implications of Asm. 3.5 for disentanglement and compositional generalization to  $\mathcal{Z}_{CPE}$ . We provide our results for up to at most<sup>2nd</sup> order interaction across slots. All results—i.e., at most<sup>1st</sup> (no interaction), at most<sup>1st</sup>, and at most<sup>2nd</sup> order interaction—use a uni ed proof strategy. Thus, we conjecture this strategy can also be used to obtain results for  $(n + 1) \geq 4$  derivatives of compositions of multivariate functions, which becomes very tedious as  $n$  grows. General<sup>n</sup>th order results are thus left for future work.

### 4.1 DISENTANGLEMENT

We start by proving disentanglement  $\mathcal{Z}_{supp}$  for which we will need two additional assumptions.

**Basis-Invariant Interactions.** First, one issue we must address is that our formalization of interaction asymmetry (Asm. 3.5) is not basis invariant. Specifically, it is possible that all splits of a slot  $z_{B_k}$  have  $(n+1)^{th}$  order interactions while for  $M_k \in R^{d \times d}$ , with  $M_k$  a slot-wise change of basis matrix, they have at most<sup>n</sup>th order interactions. Since  $M_k$  need not affect interactions across slots, interaction asymmetry may no longer hold in the new basis. This makes it ambiguous whether interaction asymmetry is truly satisfied, as  $z_{B_k}$  and  $M_k z_{B_k}$  contain the same information. To address this, we assume interaction asymmetry holds for all slot-wise basis changes via equivalent generators

**Definition 4.1 (Equivalent Generators)** A function  $f : R^{d_z} \rightarrow R^{d_x}$  is said to be equivalent to a generator  $g$  if for all  $k \in [K]$  there exists an invertible matrix  $M_k \in R^{d \times d}$  such that

$$g(z \in R^{d_z} : (M_1 z_{B_1}; \dots; M_K z_{B_K})) = f(z_{B_1}; \dots; z_{B_K}): \tag{4.1}$$

**Sufficient Independence.** We require one additional assumption for which we call sufficient independence. This assumption amounts to a linear independence condition on blocks of higher-order derivatives of  $f$ . Its main purpose is to remove redundancy in the derivatives across slots, which can be interpreted as further constraining the interaction across slots during generation. In the case

A multi-index is an ordered tuple  $\alpha = (\alpha_1; \alpha_2; \dots; \alpha_d)$  of non-negative integers,  $\alpha_i \in N_0$ , with operations  $|\alpha| := \alpha_1 + \alpha_2 + \dots + \alpha_d$ ,  $z^\alpha := z_1^{\alpha_1} z_2^{\alpha_2} \dots z_d^{\alpha_d}$ , and  $D^\alpha := \frac{\partial^{\alpha_1}}{\partial z_1^{\alpha_1}} \frac{\partial^{\alpha_2}}{\partial z_2^{\alpha_2}} \dots \frac{\partial^{\alpha_d}}{\partial z_d^{\alpha_d}}$ , see Appx. B for details.

of  $n = 0$  (i.e., no interaction across slots), sufficient independence reduces to linear independence between slot-wise Jacobians of (Defn. A.8). This is satisfied automatically since  $f$  is a diffeomorphism. When  $n > 0$ , we require an analogous linear independence condition on higher order derivatives of  $f$ . Below, we present this for the case  $n = 2$ , while for  $n = 1$ , it is presented in Defn. A.9.

**Definition 4.2 (Sufficient Independence (2<sup>nd</sup> Order)).** A  $C^3$  function  $f : \mathbb{R}^{d_z} \rightarrow \mathbb{R}^{d_x}$  with at most 2<sup>nd</sup> order interactions across slots is said to have sufficiently independent derivatives if  $\forall z \in \mathbb{R}^{d_z}$ :

$$\begin{aligned} & \text{rank} \begin{bmatrix} D_i f(z) \\ [D_{i_1 i_2}^2 \text{ of } f(z)]_{i_1, i_2 \in B_k, i_1 \neq i_2} \\ [D_{i_1 i_2 i_3}^3 \text{ of } f(z)]_{(i_1, i_2, i_3) \in \mathcal{B}_k^3} \end{bmatrix}_{k \in [K]} \\ &= \sum_{k \in [K]} \text{rank} \begin{bmatrix} D_i f(z) \\ [D_{i_1 i_2}^2 \text{ of } f(z)]_{i_1, i_2 \in B_k} \end{bmatrix} + \text{rank} \begin{bmatrix} [D_{i_1 i_2 i_3}^3 \text{ of } f(z)]_{(i_1, i_2, i_3) \in \mathcal{B}_k^3} \end{bmatrix} : \end{aligned}$$

With Defns. 4.1 and 4.2, we can now state our theoretical results; see Appx. A for complete proofs.

**Theorem 4.3 (Disentanglement on  $Z_{\text{supp}}$ ).** Let  $n \in \{0, 1, 2\}$ . Let  $f : Z \rightarrow X$  be a  $C^{n+1}$  diffeomorphism satisfying interaction asymmetry (Asm. 3.5) for all equivalent generators (Defn. 4.1) and sufficient independence (Appx. A.2). Let  $Z_{\text{supp}}$  be regular closed (Defn. A.3), path-connected (Defn. A.14) and aligned-connected (Defn. A.16). A model  $\hat{f} : \mathbb{R}^{d_x} \rightarrow Z_{\text{supp}}$  disentangles  $z$  on  $Z_{\text{supp}}$  w.r.t.  $f$  (Defn. 2.1) if it is (i) a  $C^{n+1}$  diffeomorphism between  $Z_{\text{supp}}$  and  $X_{\text{supp}}$  with (ii) at most  $n^{\text{th}}$  order interactions across slots (Defn. 3.4) on  $Z_{\text{supp}}$ .

**Intuition.** Assume for a contradiction that  $\hat{f}^{-1} \circ f$  entangles a ground-truth slot  $B_k$ , i.e.,  $D_{B_k} h(z)$  has multiple non-zero blocks. Because  $\hat{f}$  and  $f$  are invertible,  $h$  must encode all of  $B_k$  in  $\hat{z} := h(z)$ . Further, because  $f$  satisfies interaction asymmetry,  $B_k$  cannot be split into two parts with less than  $(n+1)^{\text{th}}$  order interaction. Taken together, this implies that  $\hat{f}$  entangles  $z_{B_k}$ , then there exist parts  $z_A$  and  $z_B$  of  $z_{B_k}$ , with  $(n+1)^{\text{th}}$  order interaction, encoded in different model slots. Since the model  $\hat{f}$  is constrained to have at most  $n^{\text{th}}$  order interactions across slots, it cannot capture this interaction. Thus, the only way that  $\hat{f}$  can satisfy (i) and (ii) without achieving disentanglement is if reparametrizing  $\hat{f}$  via  $h$  removed the interaction between  $z_A$  and  $z_B$ . This situation is prevented by assuming sufficient independence and that Asm. 3.5 holds for all equivalent generators.

**Conditions on  $Z_{\text{supp}}$ .** The regular closed condition on  $Z_{\text{supp}}$  in Thm. 4.3 ensures that equality between two functions on  $Z_{\text{supp}}$  implies equality of their derivatives, while the path-connectedness condition prevents the one-to-one correspondence between the slots of  $Z$  from changing across different  $z$  (Lachapelle et al., 2023). The aligned-connectedness condition is novel and allows one to take integrals to go from local to global disentanglement (see Appx. A.3 for more details).

## 4.2 COMPOSITIONAL GENERALIZATION

We now show how Asm. 3.5 also enables learning a model that generalizes compositionally (Defn. 2.2), i.e., that equality of  $\hat{f}$  and  $f \circ h$  on  $Z_{\text{supp}}$  also implies their equality on  $Z_{\text{CPE}}$ . As discussed in § 2, such generalization is non-trivial and requires specific restrictions on a function class. A key restriction imposed by interaction asymmetry is that interactions across slots are limited to at most  $n^{\text{th}}$  order. In Thm. 4.3, this prevents  $\hat{f} \circ h$  from modelling interactions between parts of the same ground-truth slot in different model slots. We now aim to show that limiting the interactions across slots serves the dual role of making  $\hat{f} \circ h$  “predictable”, such that their behavior on  $Z_{\text{CPE}}$  can be determined from  $Z_{\text{supp}}$ . To do this, we will require a characterization of the form of functions with at most  $n^{\text{th}}$  order interactions across slots, which we prove in Thm. C.2 to be:

$$f(z) = \sum_{k=1}^K f^k(z_{B_k}) + \sum_{j_1, \dots, j_n \in [K]} c_{j_1, \dots, j_n} z : \quad (4.2)$$

where  $c \in \mathbb{R}^{d_x}$ . In the first sum, slots are processed separately by functions  $f^k$ , while in the second, they can interact more explicitly via polynomial functions of components from different slots, with degree determined by the order of interaction. With this, we can now state our result.

**Theorem 4.4 (Compositional Generalization)** Let  $n \in \{0, 1, 2\}$ . Let  $Z_{\text{supp}}$  be regular closed (Defn. A.3). Let  $f : Z \rightarrow X$  and  $\hat{f} : Z \rightarrow \mathbb{R}^{d_x}$  be  $C^3$  diffeomorphisms with at most  $n^{\text{th}}$  order interactions across slots on  $Z_{\text{supp}}$ . If  $\hat{f}$  disentangles  $z$  on  $Z_{\text{supp}}$  w.r.t.  $f$  (Defn. 2.1), then it generalizes compositionally (Defn. 2.2).

Intuition. Consider the red dotted line in Fig. 2 (left) corresponding to  $z \in \mathbb{R}^2 \mid z_1 = z_1^g$ . To generalize compositionally, the behavior of the partial derivative  $\frac{\partial f}{\partial z} (z_1; z_2)$  on this line must be predictable from the behavior of  $f$  on  $Z_{\text{supp}}$ , and similarly for  $f^{\wedge} h$ . Because  $f^{\wedge} h$  have at most  $n^{\text{th}}$  order interactions across slots of  $\mathbb{R}^{d_z}$ , the form of this derivative is constrained to be a fixed-degree polynomial, see Eq. (4.2) and Fig. 2 (right). Thus, its global behavior on the dotted line can be determined from its derivative locally in a region  $\tilde{z}_{\text{supp}}$ . Applying this reasoning to all such line segments intersecting  $Z_{\text{supp}}$ , we can show that the behavior of  $f^{\wedge} h$  on  $Z_{\text{CPE}}$  can be determined from  $\tilde{z}_{\text{supp}}$ .

Figure 2: See intuition for Theorem 4.4.

### 4.3 UNIFYING AND EXTENDING PRIOR RESULTS

We now show that our theory also recovers the results of Brady et al. (2023) and Lachapelle et al. (2023) as special cases for  $\alpha = 0$  and  $n = 1$ , and extends them to more flexible generative processes.

**At most 0<sup>th</sup> Order Interaction.** Brady et al. (2023) proposed two properties for which enable disentanglement and compositional generalization (Wiedemer et al., 2024): compositionality (Defn. E.1) and irreducibility (Defn. E.2). Compositionality states that different slots affect distinct output components such that  $f(z)$  has a block-like structure. This is equivalent to having at most 0<sup>th</sup> order interaction across slots (Defn. 3.2). Irreducibility is a rank condition on  $f(z)$  which Brady et al. (2023) interpreted as parts of the same object sharing information. In Thm. E.3, we show that irreducibility is equivalent to having 1<sup>st</sup> order interaction within slots for all equivalent generators. Thus, the assumptions in Brady et al. (2023) are equivalent to interaction asymmetry for all equivalent generators when  $\alpha = 0$ . Further, we recover their disentanglement result using a novel proof strategy, unified with proofs for at most 1<sup>st</sup> / 2<sup>nd</sup> order interaction across slots (Thm. A.20).

**At most 1<sup>st</sup> Order Interaction.** Lachapelle et al. (2023) also proposed two assumptions on  $f$  for disentanglement and compositional generalization: additivity (Defn. E.4) and sufficient non-linearity (Defn. E.5). Additivity is equivalent to  $f$  having a block-diagonal Hessian for all  $l \in [1, d_x]$  (Lachapelle et al., 2023). This is the same as having at most 1<sup>st</sup> order interaction across slots (Defn. 3.4). Sufficient nonlinearity is a linear independence condition on column-wise 2<sup>nd</sup> derivatives of  $f$ . In Thm. E.6, we show that sufficient nonlinearity implies that  $f$  satisfies sufficient independence for  $\alpha = 1$  and has 2<sup>nd</sup> order interaction within slots for all equivalent generators. Further, we conjecture that the reverse implication does not hold. Thus, the assumptions of Lachapelle et al. (2023) imply, and are conjectured to be stronger than, our assumptions when  $\alpha = 1$ . We also recover their same disentanglement result using a unified proof strategy (Thm. A.22).

**Allowing More Complex Interactions.** Our theory not only unifies but also extends these prior results to more general function classes. This is clear from considering the form of functions with at most  $n^{\text{th}}$  order interactions across slots in Eq. (4.2). For at most 0<sup>th</sup> order interactions (Brady et al., 2023) or 1<sup>st</sup> order interactions (Lachapelle et al., 2023), the sum of polynomials on the RHS of (4.2) vanishes. Consequently,  $f$  reduces to an additive function. Such generators can only model concepts with trivial interactions such as non-occluding objects. In contrast, we are able to go beyond additive interactions via the polynomial terms in (4.2). This formally corroborates the “generality” of interaction asymmetry, in that it enables more flexible generative processes where concepts can explicitly interact.

## 5 METHOD: ATTENTION-REGULARIZED TRANSFORMER VAE

We now explore how our theoretical results in § 4 can inform the design of a practical estimation method. Our theory puts forth two key properties that a model should satisfy: (i) invertibility and (ii) limited interactions across slots of at most  $n^{\text{th}}$  order. To achieve disentanglement on  $Z_{\text{supp}}$ , (i) and (ii) must hold only “in-domain” on  $Z_{\text{supp}}$  and  $X_{\text{supp}}$  (Thm. 4.3), while for compositional generalization, they must also hold out-of-domain, on all  $Z$  and  $X$  (Thm. 4.4). We will focus on approaches for achieving (i) and (ii) in-domain. Achieving them out-of-domain requires addressing separate practical challenges, which are out of the scope of this work. We discuss this in detail in Appx. H.3.

On Scalability. Approaches that enforce (i) and (ii) will generally only be computationally tractable in low-dimensional settings. Such computational issues are typical when translating a disentanglement result into an empirical method, often resulting in methods which directly adhere to theory but cannot scale beyond toy data (e.g., [Brady et al., 2023](#); [Gresele et al., 2021](#)). Our core motivation, however, is learning representations of concepts underlying high-dimensional sensory data, such as images. Thus, to formulate a method which scales to such settings, we do not restrict ourselves to approaches which exactly enforce (i) and (ii) and also approximate approaches.

(i) Invertibility. Our theory requires invertibility between  $X_{\text{supp}} \in \mathbb{R}^{d_x}$  and  $Z_{\text{supp}} \in \mathbb{R}^{d_z}$ . For most settings of interest, the observed dimension  $d_x$  exceeds the ground-truth latent dimension  $d_z$ . Thus, we generally cannot use models which are invertible by construction such as normalizing flows ([Papamakarios et al., 2021](#)). An alternative is to use an autoencoder in which  $f^{-1}$  and  $f$  are parameterized separately by an encoder  $g : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_z}$  and a decoder  $h : \mathbb{R}^{d_z} \rightarrow \mathbb{R}^{d_x}$ , which are trained to invert each other ( $Z_{\text{supp}}$  and  $X_{\text{supp}}$ ) by minimizing a reconstruction loss  $L_{\text{rec}} := \mathbb{E}_x \|g(x) - h(g(x))\|^2$ . Minimizing  $L_{\text{rec}}$  alone, however, does not suffice unless the inferred latent dimension equals the ground-truth  $d_z$ . Yet, in practice  $d_z$  is unknown. Moreover, choosing  $d_z > d_x$  is important for scalability ([Sajjadi et al., 2022a](#)). A viable alternative is thus to employ a soft constraint where  $d_z > d_x$ , but the model is encouraged to encode using minimal latent dimensions. To achieve this, we leverage the well known VAE loss ([Kingma and Welling, 2014](#)), which couples with a KL divergence loss  $L_{\text{KL}}$  between a factorized posterior  $q(z|x)$  and prior distribution  $p(z)$ , i.e.,  $L_{\text{KL}} := \mathbb{E}_{z \sim q(z|x)} \sum_{i=1}^{d_z} D_{\text{KL}}(q(z_i|x) \parallel p(z_i))$ . This loss encourages each dimension to be insensitive to changes in  $x$  such that unnecessary dimensions should contain no information ([Boutin et al., 2019](#)).

(ii) At Most  $n^{\text{th}}$  Order Interactions. One approach to enforce at most  $n^{\text{th}}$  order interactions across slots would be to parameterize the decoder to match the form of such functions (see [Thm. C.2](#)) for some  $x \in \mathbb{R}^n$ . However, this can result in an overly restrictive inductive bias and limit scalability. Moreover,  $n$  is generally unknown. Thus, a more promising approach is to regularize interactions to be minimal. Doing this naively though using gradient descent would require computing gradients of high-order derivatives, which is intractable beyond toy data. This leads to the question: Is there a scalable architecture which permits efficient regularization of the interactions across slots?

Transformers for Interaction Regularization. We make the observation that the Transformer architecture ([Vaswani et al., 2017](#)) provides an efficient means to approximately regularize interactions. In a Transformer, slots are only permitted to interact via an attention mechanism. We will focus on a cross-attention mechanism, which maps a latent vector  $z$  to an output  $x_1$  (e.g., a pixel) via:

$$K = W^K [z_{B_1} \dots z_{B_K}]; \quad V = W^V [z_{B_1} \dots z_{B_K}]; \quad Q = W^Q [o_1 \dots o_{d_x}]; \quad (5.1)$$

$$A_{l;k} = \frac{\exp(Q_{l;k}^> K_{:;k})}{\sum_{i \in [K]} \exp(Q_{l;i}^> K_{:;i})}; \quad x_1 = A_{l;:} V^>; \quad \hat{x}_1 = \sigma(x_1); \quad (5.2)$$

In [Eq. \(5.1\)](#), all slots are assumed to have equal size, and key and value  $v_{:;k}$  vectors are computed for each slot  $k \in [K]$ . Query vectors are computed for output dimensions  $o \in [d_x]$  (e.g., pixel coordinates) and each is assigned a fixed vector  $o_l$ . In [Eq. \(5.2\)](#), queries and keys are used to compute attention weights  $A_{l;k}$ . These weights determine the slots pixel attends to when generating pixel token  $x_1$ , which is mapped to a pixel  $\hat{x}_1$  by nonlinear function  $\sigma$ ; see [Appx. F](#) for further details.

Within cross-attention, interactions across slots occur if the query vector for a pixel attends to multiple slots, i.e., if  $A_{l;k}$  is non-zero for more than one  $k$ . Conversely, if  $A_{l;k}$  is non-zero for only one  $k$ , then, intuitively, no interactions should occur. This intuition can be corroborated formally by computing the Jacobian of cross-attention w.r.t. each slot (see [Appx. F.1](#)). Thus, an approximate means to minimize interactions across slots is to regularize towards having only one non-zero entry for each row  $A_{l;:}$ . To this end, we propose to minimize the sum of all pairwise products  $A_{l;j} A_{l;k}$ , where  $j \neq k$  (see [Fig. 4](#)). This quantity is non-negative and will only be zero when each row of  $A$  has exactly one non-zero entry. This resembles the positional contrast of [Brady et al. \(2023\)](#), but computed on  $A$ , which can be efficiently optimized, as opposed to the Jacobian of which is intractable to optimize. We refer to this regularizer as  $L_{\text{interact}}$ ; see [Eq. \(F.9\)](#).

Model. Combining these different objectives leads us to the following weighted three-part-loss:

$$L_{\text{disent}}(f^{\wedge}; g; x) = L_{\text{rec}} + \lambda L_{\text{interact}} + \mu L_{\text{KL}}; \quad (5.3)$$



Figure 3: (A) Sprites Normalized slot-wise Jacobians for an unregularized ( $\lambda = 0$ ) and a regularized ( $\lambda > 0$ ;  $\lambda > 0$ ) Transformer and a Spatial Broadcast Decoder (SBD). The unregularized model encodes objects across multiple slots, while the regularized model matches the disentanglement of the SBD. (B) Slot-wise Jacobians for a regularized Transformer and a SBD on objects in CLEVR6 which interact via reflections. As can be seen in reconstructions and Jacobians, the regularized Transformer models reflections, while mostly removing unnecessary interactions, while the SBD fails to model reflections due to its restricted architecture.

We apply this loss to a flexible Transformer-based autoencoder, similar to the models of [Jabri et al. \(2023\)](#); [Jaegle et al. \(2022\)](#); [Sajjadi et al. \(2022b\)](#). For the encoder, we first map data to features using the CNN of [Locatello et al. \(2020b\)](#). These features are processed by a Transformer, which has both self- and cross-attention at every layer, yielding a representation  $\hat{z}$ . Our decoder then maps  $\hat{z}$  to an output  $\hat{y}$  using a cross-attention Transformer regularized with  $\mathcal{L}_{\text{interact}}$  see Appx. J for details.

**Relationship to Models In Object-Centric Learning.** Existing models for learning disentangled representations of concepts, particularly for disentangling objects without supervision, typically rely on architectural priors rather than regularization ([Greff et al., 2019](#); [Locatello et al., 2020b](#); [Seitzer et al., 2023](#); [Singh et al., 2022a](#)). While such priors promote disentanglement, they are often too restrictive. For example, Spatial Broadcast Decoders ([Watters et al., 2019b](#)) decode slots separately and only allow for weak interaction through a softmax function, which prevents modelling real-world data where objects exhibit more complex interactions ([Singh et al., 2022a](#)). While some works have shown success in disentangling objects using more powerful Transformer decoders ([Sajjadi et al., 2022a](#); [Singh et al., 2022a;b](#)), they rely on encoders that use Slot Attention ([Locatello et al., 2020b](#)) as an architectural component, which differs from current large-scale models, typically based on Transformers ([Anil et al., 2023](#)). In contrast, we explore the more flexible approach of starting with a very general Transformer-based model and regularizing it towards a more constrained model.

## 6 EXPERIMENTS

We now apply our attention-regularized Transformer-VAE (§ 5) for learning representations of concepts. Since this model is designed to enforce the criteria outlined in Thm. 4.3 for disentanglement on  $Z_{\text{supp}}$ , we focus on evaluating disentanglement, as opposed to compositional generalization. To this end, we focus on disentangling objects in visual scenes, and leave an empirical study of a wider range of concepts (e.g., attributes, object-parts, events) for future work (see Appx. J for details).

**Data.** We consider two multi-object datasets in our experiments. The first, which we refer to as Sprites ([Brady et al., 2023](#); [Watters et al., 2019a](#); [Wiedemer et al., 2024b](#)), consist of images with 4 objects set against a black background. The second is the dataset ([Johnson et al., 2017](#)), consisting of images with 2–6 objects. In Sprites, objects do not have reflections and rarely occlude such that slots have essentially have no interaction. In CLEVR6, however, objects can cast shadows and reflect upon each other (see Fig. 1 for an example), introducing more complex interactions.

**Metrics.** A common metric for object disentanglement is the Adjusted-Rand Index (ARI; [Hubert and Arabie, 1985](#)). The ARI measures the similarity between the set of pixels encoded by a model slot, and the set of ground-truth pixels for a given object in a scene, yielding an optimal score if each slot corresponds to exactly one object. To assign a pixel to a unique model slot, prior works typically choose the slot with the largest attention score (from, e.g., Slot Attention) for that pixel ([Seitzer et al., 2023](#)). However, using attention scores can make model comparisons challenging and is also somewhat unprincipled (see Appx. J.2). We thus consider an alternative and compute the ARI using the Jacobian of a decoder (J-ARI). Specifically, we assign a pixel to the slot with the largest  $L_1$  norm for the slot-wise Jacobian  $\mathbb{D}_{B_k} \hat{f}_1(\mathbf{z})$  (see Fig. 3 for a visualization of these Jacobians).

While J-ARI indicates which slots are most responsible for encoding each object, it does not indicate if additional slots affect the same object, i.e.  $\mathbb{D}_{B_k} \hat{f}_1(\mathbf{z})_{k_1} \neq 0$  for more than one  $k$ . To measure

Table 1: Empirical Results. We show the mean std. dev. for J-ARI and JIS (in %) over 3 seeds for different choices of encoders and decoders and weights of the loss terms in Eq. (5.3) on Sprites and CLEVR6.

Model			Sprites				CLEVR6			
Encoder	Decoder	Loss	J-ARI(%)		JIS (%)		J-ARI (%)		JIS (%)	
Slot Attention	Spatial-Broadcast	$\lambda = 0; \mu = 0$	89:3	1:5	91:4	0:8	97:0	0:2	95:3	0:7
Slot Attention	Transformer	$\lambda = 0; \mu = 0$	90:1	1:4	73:6	1:5	95:5	1:0	63:1	1:0
Transformer	Transformer	$\lambda = 0; \mu = 0$	80:5	4:1	57:0	8:0	92:7	3:3	54:8	3:5
Transformer	Transformer	$\lambda = 0:05; \mu = 0$	82:8	3:6	73:8	4:0	79:2	12:8	51:6	5:9
Transformer	Transformer	$\lambda = 0; \mu = 0:05$	92:6	2:0	92:8	0:9	96:6	0:3	80:3	0:4
Transformer	Transformer	$\lambda = 0:05; \mu = 0:05$ (Ours)	93:7	0:6	95:0	1:7	96:5	0:4	83:8	1:1

this, we also introduce the Jacobian Interaction Score (JIS). JIS is computed by taking the maximum of  $\| \mathbf{D}_{B_k} \hat{f}_1(\mathbf{z}) \|_1$  across slots after normalization, averaged over all pixels. If each pixel is affected by only one slot, JIS is 1. For datasets where objects essentially do not interact like Sprites, JIS should be close to 1, whereas for CLEVR6, it should be as high as possible while maintaining invertibility.

## 6.1 RESULTS

**Enables Object Disentanglement.** In Tab. 1, we compare the J-ARI and JIS of our regularized Transformer-based model ( $\lambda > 0; \mu > 0$ ) trained with  $L_{\text{disent}}$  (Eq. (5.3)) to the same model trained without regularization ( $\lambda = 0; \mu = 0$ ), i.e., with only  $L_{\text{rec}}$ . On Sprites, the regularized model achieves notably higher scores for both J-ARI and JIS. This is corroborated by visualizing the slot-wise Jacobians in Fig. 3A, where we see the regularized model cleanly disentangles objects, whereas the unregularized model often encodes objects across multiple slots. Similarly, on CLEVR6, the regularized model achieves superior disentanglement, as indicated by the higher values for both metrics.

**Comparison to Existing Object-Centric Autoencoders.** In Tab. 1, we also compare our model to existing models using encoders with Slot Attention and Spatial Broadcast Decoders (SBDs). On Sprites, our model achieves higher J-ARI and JIS than these models, despite using a weaker architectural prior. On CLEVR6, our model outperforms Slot Attention with a Transformer decoder in terms of J-ARI and JIS. Models using a SBD, however, achieve a higher and nearly perfect JIS, i.e., the learned slots essentially never affect the same pixel. In Fig 3B, we see this comes at the cost of SBDs failing to model reflections between objects, while our model captures this interaction. This highlights that regularizing a flexible architecture with  $L_{\text{disent}}$  can enable a better balance between restricting interactions and model expressivity.

**Ablation Over Losses.** Lastly, in Tab. 1, we ablate the impact of the regularizer  $L_{\text{disent}}$ . Training without  $L_{\text{KL}}$  ( $\lambda > 0; \mu = 0$ ) can in some cases give improvements in J-ARI and JIS over an unregularized model ( $\lambda = 0; \mu = 0$ ). However, across datasets this loss yields worse disentanglement than  $L_{\text{disent}}$  ( $\lambda > 0; \mu > 0$ ). This highlights that penalizing latent capacity via  $L_{\text{KL}}$  is important for object disentanglement. Training with  $L_{\text{interac}}$  ( $\lambda = 0; \mu > 0$ ) generally yields a drop across both metrics compared to  $L_{\text{disent}}$  though on CLEVR6 this loss achieves a comparable J-ARI. We found that training with  $L_{\text{KL}}$  can, in some cases, implicitly minimize  $L_{\text{interac}}$  explaining this result (Fig. 5).

**More Complex Data.** Tab. 2 in Appx. I presents additional results for the visually complex CLEVR-Text dataset (Karazija et al., 2021). For these experiments, we follow Seitzer et al. (2023) and reconstruct image representations based on a pre-trained encoder rather than the original images. We find our model to achieve superior J-ARI compared to an unregularized Transformer and a Slot Attention baseline, but slot-wise MLP decoders yield higher JIS. For further details, see Appx. I.

## 7 CONCLUSION

In this work, we proposed interaction asymmetry as a general principle for learning disentangled and composable representations. Formalizing this idea led to a constraint on the partial derivatives of the generator function, which unifies assumptions from prior efforts and extends their results to a more flexible class of generators that allow for non-trivial interactions. These theoretical insights inspired the development of a flexible estimation method based on the Transformer architecture with a novel cross-attention regularizer, which can be efficiently implemented at scale, and which shows promising results on object-centric learning datasets. Future work should seek to further extend our theoretical results, address the empirical challenges for achieving compositional generalization, and test our method on more large-scale data involving not only objects but also other types of concepts.

## ACKNOWLEDGEMENTS

JB would like to thank Nicolò Zottino for providing the initial idea (expressed in Fig. 2) that compositional generalization could be extended to functions with at most order interactions. The authors thank Luigi Gresele and the Robust Machine Learning Group for helpful discussions.

This work was supported by the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039A, 01IS18039B. WB acknowledges financial support via an Emmy Noether Grant funded by the German Research Foundation (DFG) under grant no. BR 6382/1-1 and via the Open Philanthropy Foundation funded by the Good Ventures Foundation. WB is a member of the Machine Learning Cluster of Excellence, EXC number 2064/1 – Project number 390727645. This work utilized compute resources at the Tübingen Machine Learning Cloud, DFG FKZ INST 37/1057-1 FUGG. Part of this work was done while JvK was affiliated with the Max Planck Institute for Intelligent Systems in Tübingen.

## AUTHOR CONTRIBUTIONS

JB initiated and led the project. JB wrote the manuscript in close collaboration with JvK, and with feedback from all authors. JB conceived of the framework of interaction asymmetry and then formalized it together with SL and JvK. JB proved the disentanglement results in Appx. A in collaboration with SL, who improved the clarity of several results and simplified some arguments. SL extended these results from local to global disentanglement in Appx. A.3. SB proved the characterization of functions with at most<sup>th</sup> order interactions in Appx. C, based on an earlier weaker result for analytic functions by JvK. SB proved the compositional generalization results in Appx. D after discussions with JB. JB proved the results for unifying prior works in Appx. E. SB carried out the Jacobian computation in Appx. F.1. JB conceived of the regularizers and method in § 5 and implemented and executed all experiments in § 6 with advising from TK. WB created Fig. 1 and Fig. 3 with insight from JB and feedback from JvK. SL created Fig. 2 based on an outline provided by JB.

## REFERENCES

- K. Ahuja, J. S. Hartford, and Y. Bengio. Weakly supervised representation learning with sparse perturbations. In *Advances in Neural Information Processing Systems*, volume 35, pages 15516–15528, 2022. [Cited on p. 3.]
- K. Ahuja, D. Mahajan, Y. Wang, and Y. Bengio. Interventional causal representation learning. In *International Conference on Machine Learning*, pages 372–407. PMLR, 2023. [Cited on p. 45.]
- A. A. Alemi, I. Fischer, J. V. Dillon, and K. Murphy. Deep variational information bottleneck. arXiv preprint arXiv:1612.00410, 2016. [Cited on p. 44.]
- R. Anil, S. Borgeaud, Y. Wu, J.-B. Alayrac, J. Yu, R. Soricut, J. Schalkwyk, A. M. Dai, A. Hauth, et al. Gemini: a family of highly capable multimodal models. arXiv preprint arXiv:2312.11805, 2023. [Cited on p. 9.]
- J. L. Ba. Layer normalization. arXiv preprint arXiv:1607.06450, 2016. [Cited on p. 42.]
- D. Baldwin, A. Andersson, J. Saffran, and M. Meyer. Segmenting dynamic human action via statistical structure. *Cognition*, 106(3):1382–1407, 2008. [Cited on p. 44.]
- T. E. J. Behrens, T. H. Muller, J. C. R. Whittington, S. Mark, A. Baram, K. L. Stachenfeld, and Z. Kurth-Nelson. What is a cognitive map? organizing knowledge for flexible behavior. *Neuron*, 100:490–509, 2018. [Cited on p. 1, 2, and 45.]
- Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013. [Cited on p. 1.]
- O. Biza, S. Van Steenkiste, M. S. Sajjadi, G. F. Elsayed, A. Mahendran, and T. Kipf. Invariant slot attention: Object discovery with slot-centric reference frames. arXiv preprint arXiv:2302.04973, 2023. [Cited on p. 49 and 50.]
- J. Brady, R. S. Zimmermann, Y. Sharma, B. Schöopf, J. von Kugelgen, and W. Brendel. Provably learning object-centric representations. In *International Conference on Machine Learning*, pages 3038–3062. PMLR, 2023. [Cited on p. 2, 4, 7, 8, 9, 40, 43, 44, and 47.]

- J. Brehmer, P. De Haan, P. Lippe, and T. Cohen. Weakly supervised causal representation learning. In *Advances in Neural Information Processing Systems* volume 35, pages 38319–38331, 2022. [Cited on p. 3.]
- T. Brooks, A. Holynski, and A. A. Efros. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* pages 18392–18402, 2023. [Cited on p. 45.]
- S. Buchholz, M. Besserve, and B. Schölkopf. Function classes for identifiable nonlinear independent component analysis. *Advances in Neural Information Processing Systems* volume 35, pages 16946–16961, 2022. [Cited on p. 3.]
- S. Buchholz, G. Rajendran, E. Rosenfeld, B. Aragam, B. Schölkopf, and P. Ravikumar. Learning linear causal representations from interventions under general nonlinear mixtures. *Advances in Neural Information Processing Systems* 2023. [Cited on p. 3.]
- M. Caron, H. Touvron, I. Misra, H. Jou, J. Mairal, P. Bojanowski, and A. Joulin. Emerging properties in self-supervised vision transformers. *Proceedings of the IEEE/CVF international conference on computer vision* pages 9650–9660, 2021. [Cited on p. 49.]
- M. Chang, T. Griffiths, and S. Levine. Object representations as fixed points: Training iterative refinement algorithms with implicit differentiation. *Advances in Neural Information Processing Systems* 35:32694–32708, 2022. [Cited on p. 50.]
- A. O. Constantinescu, J. X. O'Reilly, and T. E. Behrens. Organizing conceptual knowledge in humans with a gridlike code. *Science* 352(6292):1464–1468, 2016. [Cited on p. 2.]
- A. Dittadi, S. S. Papa, M. De Vita, B. Schölkopf, O. Winther, and F. Locatello. Generalization and robustness implications in object-centric learning. *International Conference on Machine Learning* pages 5221–5285. PMLR, 2022. [Cited on p. 47.]
- K. Dong and T. Ma. First steps toward understanding the extrapolation of nonlinear models to unseen domains. In *International Conference on Learning Representations*, 2022. [Cited on p. 3.]
- A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. [Cited on p. 49.]
- B. Esmaeili, H. Wu, S. Jain, A. Bozkurt, N. Siddharth, B. Paige, D. H. Brooks, J. G. Dy, and J.-W. van de Meent. Structured disentangled representation. *International Conference on Artificial Intelligence and Statistics*, 2018. [Cited on p. 45.]
- J. A. Fodor and Z. W. Pylyshyn. Connectionism and cognitive architecture: A critical analysis. *Cognition* 28(1):3–71, 1988. [Cited on p. 1.]
- S. Fortunato and D. Hric. Community detection in networks: A user guide. *Physics reports* 659: 1–44, 2016. [Cited on p. 44.]
- A. Gopalakrishnan, K. Irie, J. Schmidhuber, and S. van Steenkiste. Unsupervised learning of temporal abstractions with slot-based transformers. *Neural Computation* 35(4):593–626, 2023. [Cited on p. 48.]
- A. Goyal and Y. Bengio. Inductive biases for deep learning of higher-level cognition. *Proceedings of the Royal Society*, 478(2266):20210068, 2022. [Cited on p. 1.]
- A. Goyal, A. Lamb, J. Hoffmann, S. Sodhani, S. Levine, Y. Bengio, and B. Schölkopf. Recurrent independent mechanisms. *International Conference on Learning Representations*, 2021. [Cited on p. 45.]
- K. Greff, R. K. Srivastava, and J. Schmidhuber. Binding via reconstruction clustering. preprint arXiv:1511.06418, 2015. [Cited on p. 4 and 44.]
- K. Greff, R. L. Kaufman, R. Kabra, N. Watters, C. Burgess, D. Zoran, L. Matthey, M. Botvinick, and A. Lerchner. Multi-object representation learning with iterative variational inference. In *International Conference on Machine Learning* pages 2424–2433. PMLR, 2019. [Cited on p. 9.]

- K. Greff, S. Van Steenkiste, and J. Schmidhuber. On the binding problem in artificial neural networks. arXiv preprint arXiv:2012.05208, 2020. [Cited on p. 1 and 44.]
- L. Gresele, P. K. Rubenstein, A. Mehrjou, F. Locatello, and B. Schölkopf. The incomplete rosetta stone problem: Identifiability results for multi-view nonlinear ICA. *Uncertainty in Artificial Intelligence*, pages 217–227. PMLR, 2020. [Cited on p. 3.]
- L. Gresele, J. von Kgelgen, V. Stimper, B. Schölkopf, and M. Besserve. Independent mechanism analysis, a new concept? *Advances in Neural Information Processing Systems*, volume 34, pages 28233–28248, 2021. [Cited on p. 3 and 8.]
- H. Hälvä and A. Hyvärinen. Hidden markov nonlinear ICA: Unsupervised learning from nonstationary time series. In *Conference on Uncertainty in Artificial Intelligence*, pages 939–948. PMLR, 2020. [Cited on p. 3.]
- H. Hälvä, S. Le Corff, L. Lelercy, J. So, Y. Zhu, E. Gassiat, and A. Hyvärinen. Disentangling identifiable features from noisy data with structured nonlinear ICA. *Advances in Neural Information Processing Systems*, 34:1624–1633, 2021. [Cited on p. 3.]
- H. Hälvä, J. So, R. E. Turner, and A. Hyvärinen. Identifiable feature learning for spatial data with nonlinear ICA. In *International Conference on Artificial Intelligence and Statistics*, pages 3331–3339. PMLR, 2024. [Cited on p. 3.]
- J. Hawkins, M. Lewis, M. Klukas, S. Purdy, and S. Ahmad. A framework for intelligence and cortical function based on grid cells in the neocortex. *Frontiers in Neural Circuits* 12, 2018. [Cited on p. 2.]
- I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner. beta-VAE: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations*, 2017. [Cited on p. 45.]
- I. Higgins, D. Amos, D. Pfau, S. Racanière, L. Matthey, D. J. Rezende, and A. Lerchner. Towards a definition of disentangled representations. arXiv preprint arXiv:1812.02230, 2018. [Cited on p. 1.]
- S. Hochreiter and J. Schmidhuber. Feature extraction through local neural computation. *Neural Computation* 11(3):679–714, 1999. [Cited on p. 44.]
- D. Horan, E. Richardson, and Y. Weiss. When is unsupervised disentanglement possible? In *Advances in Neural Information Processing Systems*, volume 34, 2021. [Cited on p. 3.]
- L. J. Hubert and P. Arabie. Comparing partitions. *Journal of Classification* 2:193–218, 1985. [Cited on p. 9.]
- A. Hyvärinen and H. Morioka. Unsupervised feature extraction by time-contrastive learning and nonlinear ICA. In *Advances in Neural Information Processing Systems*, pages 3765–3773, 2016. [Cited on p. 3.]
- A. Hyvärinen and H. Morioka. Nonlinear ICA of temporally dependent stationary sources. In *Artificial Intelligence and Statistics*, pages 460–469. PMLR, 2017. [Cited on p. 3 and 5.]
- A. Hyvärinen and P. Pajunen. Nonlinear independent component analysis: Existence and uniqueness results. *Neural networks* 12(3):429–439, 1999. [Cited on p. 1 and 3.]
- A. Hyvärinen and J. Perkio. Learning to segment any random vector. In *IEEE International Joint Conference on Neural Network Proceedings*, pages 4167–4172. IEEE, 2006. [Cited on p. 44.]
- A. Hyvärinen, H. Sasaki, and R. Turner. Nonlinear ICA using auxiliary variables and generalized contrastive learning. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 859–868. PMLR, 2019. [Cited on p. 3 and 21.]
- A. Hyvärinen, I. Khemakhem, and H. Morioka. Nonlinear independent component analysis for principled disentanglement in unsupervised deep learning. *Patterns* 4(10), 2023. [Cited on p. 1.]
- A. Jabri, D. J. Fleet, and T. Chen. Scalable adaptive computation for iterative generative. In *International Conference on Machine Learning*, pages 14569–14589. PMLR, 2023. [Cited on p. 9.]

- A. Jaegle, S. Borgeaud, J.-B. Alayrac, C. Doersch, C. Ionescu, D. Ding, S. Koppula, D. Zoran, A. Brock, E. Shelhamer, et al. Perceiver IO: A general architecture for structured inputs & outputs. In *International Conference on Learning Representations*, 2022. [Cited on p. 9 and 46.]
- J. Jiang, F. Deng, G. Singh, and S. Ahn. Object-centric slot diffusion. *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. [Cited on p. 47.]
- Y. Jiang, E. Liu, B. Eysenbach, J. Z. Kolter, and C. Finn. Learning options via compression. *Advances in Neural Information Processing Systems*, 35:21184–21199, 2022. [Cited on p. 44.]
- J. Johnson, B. Hariharan, L. Van Der Maaten, L. Fei-Fei, C. Lawrence Zitnick, and R. Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2901–2910, 2017. [Cited on p. 2, 9, and 50.]
- L. Karazija, I. Laina, and C. Rupprecht. ClevrTex: A Texture-Rich Benchmark for Unsupervised Multi-Object Segmentation. *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks*, 2021. [Cited on p. 10 and 49.]
- I. Khemakhem, D. Kingma, R. Monti, and A. Hyvriinen. Variational autoencoders and nonlinear ICA: A unifying framework. In *International Conference on Artificial Intelligence and Statistics*, pages 2207–2217. PMLR, 2020. [Cited on p. 3.]
- D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 2015. [Cited on p. 49 and 50.]
- D. P. Kingma and M. Welling. Auto-encoding variational bayes. *International Conference on Learning Representations*, 2014. [Cited on p. 2 and 8.]
- T. Kipf, Y. Li, H. Dai, V. Zambaldi, A. Sanchez-Gonzalez, E. Grefenstette, P. Kohli, and P. Battaglia. Compile: Compositional imitation learning and execution. *International Conference on Machine Learning*, pages 3418–3428. PMLR, 2019. [Cited on p. 48.]
- A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, et al. Segment anything. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023. [Cited on p. 45.]
- B. Kivva, G. Rajendran, P. Ravikumar, and B. Aragam. Identifiability of deep generative models without auxiliary information. *Advances in Neural Information Processing Systems*, 35:15687–15701, 2022. [Cited on p. 3.]
- D. A. Klindt, L. Schott, Y. Sharma, I. Ustyuzhaninov, W. Brendel, M. Bethge, and D. Paiton. Towards nonlinear disentanglement in natural data with temporal sparse coding. *International Conference on Learning Representations*, 2021. [Cited on p. 3.]
- A. Kori, F. Locatello, A. Santhirasekaram, F. Toni, B. Glocker, and F. D. S. Ribeiro. Identifiable object-centric representation learning via probabilistic slot attention. *arXiv preprint arXiv:2406.07141*, 2024. [Cited on p. 3 and 45.]
- A. Kraskov, H. Stogbauer, R. G. Andrzejak, and P. Grassberger. Hierarchical clustering using mutual information. *Europhysics Letters*, 70(2):278, 2005. [Cited on p. 44.]
- S. Lachapelle, P. Rodriguez, Y. Sharma, K. E. Everett, R. Le Priol, A. Lacoste, and S. Lacoste-Julien. Disentanglement via mechanism sparsity regularization: A new principle for nonlinear ICA. In *Conference on Causal Learning and Reasoning*, pages 428–484. PMLR, 2022. [Cited on p. 3.]
- S. Lachapelle, D. Mahajan, I. Mitliagkas, and S. Lacoste-Julien. Additive decoders for latent variables identification and cartesian-product extrapolation. *Advances in Neural Information Processing Systems*, volume 36, 2023. [Cited on p. 1, 2, 3, 4, 6, 7, 19, 21, 22, 23, 27, 35, 36, 39, 41, 42, 45, 46, and 48.]
- S. Lachapelle, P. R. López, Y. Sharma, K. Everett, R. L. Priol, A. Lacoste, and S. Lacoste-Julien. Nonparametric partial disentanglement via mechanism sparsity: Sparse actions, interventions and sparse temporal dependencies. *arXiv preprint arXiv:2401.04890*, 2024. [Cited on p. 3.]
- B. M. Lake, T. D. Ullman, J. B. Tenenbaum, and S. J. Gershman. Building machines that learn and think like people. *Behavioral and brain sciences*, 40:e253, 2017. [Cited on p. 1.]

- A. Lamb, D. He, A. Goyal, G. Ke, C.-F. Liao, M. Ravanelli, and Y. Bengio. Transformers with competitive ensembles of independent mechanisms. *arXiv preprint arXiv:2103.00336*, 2021. [Cited on p. 45.]
- Y. LeCun. A path towards autonomous machine intelligence version 0.9.2, 2022-06-27. Review, pages 1–62, 2022. [Cited on p. 45.]
- T. Leemann, M. Kirchhof, Y. Rong, E. Kasneci, and G. Kasneci. When are post-hoc conceptual explanations identifiable? *Uncertainty in Artificial Intelligence* pages 1207–1218. PMLR, 2023. [Cited on p. 3.]
- P. Lippe, S. Magliacane, S. Löwe, Y. M. Asano, T. Cohen, and S. Gavves. Citris: Causal identifiability from temporal intervened sequences. *International Conference on Machine Learning* pages 13557–13603. PMLR, 2022. [Cited on p. 3.]
- P. Lippe, S. Magliacane, S. Löwe, Y. M. Asano, T. Cohen, and E. Gavves. Causal representation learning for instantaneous and temporal effects in interactive systems. *International Conference on Learning Representations*, 2023. [Cited on p. 3.]
- S. Lippl and K. Stachenfeld. When does compositional structure yield compositional generalization? a kernel theory. *arXiv preprint arXiv:2405.16391*, 2024. [Cited on p. 3.]
- F. Locatello, S. Bauer, M. Lucic, G. Raetsch, S. Gelly, B. Schölkopf, and O. Bachem. Challenging common assumptions in the unsupervised learning of disentangled representations. *International Conference on Machine Learning*, pages 4114–4124. PMLR, 2019. [Cited on p. 1 and 3.]
- F. Locatello, B. Poole, G. Raetsch, B. Schölkopf, O. Bachem, and M. Tschannen. Weakly-supervised disentanglement without compromises. *International Conference on Machine Learning*, pages 6348–6359. PMLR, 2020a. [Cited on p. 3.]
- F. Locatello, D. Weissenborn, T. Unterthiner, A. Mahendran, G. Heigold, J. Uszkoreit, A. Dosovitskiy, and T. Kipf. Object-centric learning with slot attention. *Advances in Neural Information Processing Systems*, volume 33, pages 11525–11538, 2020b. [Cited on p. 2, 9, 47, 50, and 51.]
- L. Mahon, L. Shah, and T. Lukasiewicz. Correcting flaws in common disentanglement metrics. *arXiv preprint arXiv:2304.02335*, 2023. [Cited on p. 45.]
- C. Meo, L. Mahon, A. Goyal, and J. Dauwels.  $\beta$ -STC-VAE: On the relationship between disentanglement and diversity. *The Twelfth International Conference on Learning Representations* 2024. [Cited on p. 44.]
- M. Mitchell. Abstraction and analogy-making in artificial intelligence. *Annals of the New York Academy of Sciences*, 505(1):79–101, 2021. [Cited on p. 1.]
- M. Montero, J. Bowers, R. Ponte Costa, C. Ludwig, and G. Malhotra. Lost in latent space: Examining failures of disentangled models at combinatorial generalisation. *Advances in Neural Information Processing Systems*, volume 35, pages 10136–10149, 2022a. [Cited on p. 2.]
- M. L. Montero, C. J. Ludwig, R. P. Costa, G. Malhotra, and J. Bowers. The role of disentanglement in generalisation. *International Conference on Learning Representations*, 2021. [Cited on p. 2 and 45.]
- M. L. Montero, J. Bowers, R. P. Costa, C. J. Ludwig, and G. Malhotra. Lost in latent space: Examining failures of disentangled models at combinatorial generalisation. In A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, editors, *Advances in Neural Information Processing Systems*, 2022b. [Cited on p. 45.]
- G. E. Moran, D. Sridhar, Y. Wang, and D. Blei. Identifiable deep generative models via sparse decoding. *Transactions on Machine Learning Research*, 2022. [Cited on p. 3.]
- J. Munkres. *Analysis On Manifolds*. Basic Books, 1991. [Cited on p. 19.]
- G. Murphy. *The big book of concepts*. The MIT Press, 2004. [Cited on p. 1.]
- G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research*, 22(57):1–64, 2021. [Cited on p. 8.]

- J. Peters, D. Janzing, and B. Schölkopf. Elements of causal inference: foundations and learning algorithms The MIT Press, 2017. [Cited on p. 4.]
- A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen, and I. Sutskever. Zero-shot text-to-image generation. *International Conference on Machine Learning*, pages 8821–8831. Pmlr, 2021. [Cited on p. 45.]
- A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*(2):3, 2022. [Cited on p. 45.]
- N. Ravi, V. Gabeur, Y.-T. Hu, R. Hu, C. Ryali, T. Ma, H. Khedr, P. Ré, C. Rolland, L. Gustafson, et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. [Cited on p. 45.]
- J. R. Reynolds, J. M. Zacks, and T. S. Braver. A computational model of event segmentation from perceptual prediction. *Cognitive science*, 31(4):613–643, 2007. [Cited on p. 44.]
- M. Rolinek, D. Zietlow, and G. Martius. Variational autoencoders pursue pca directions (by accident). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* pages 12406–12415, 2019. [Cited on p. 8.]
- N. Ruiz, Y. Li, V. Jampani, Y. Pritch, M. Rubinstein, and K. Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22500–22510, 2023. [Cited on p. 45.]
- C. Saharia, W. Chan, S. Saxena, L. Li, J. Whang, E. L. Denton, K. Ghasemipour, R. Gontijo Lopes, B. Karagol Ayan, T. Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, volume 35, pages 36479–36494, 2022. [Cited on p. 45.]
- M. S. Sajjadi, D. Duckworth, A. Mahendran, S. Van Steenkiste, F. Pavetic, M. Lucic, L. J. Guibas, K. Greff, and T. Kipf. Object scene representation transformer. *Advances in Neural Information Processing Systems*, volume 35, pages 9512–9524, 2022a. [Cited on p. 8, 9, and 46.]
- M. S. Sajjadi, H. Meyer, E. Pot, U. Bergmann, K. Greff, N. Radwan, S. Vora, M. L. Duckworth, A. Dosovitskiy, et al. Scene representation transformer: Geometry-free novel view synthesis through set-latent scene representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6229–6238, 2022b. [Cited on p. 9 and 46.]
- A. C. Schapiro, T. T. Rogers, N. I. Cordova, N. B. Turk-Browne, and M. M. Botvinick. Neural representations of events arise from temporal community structure. *Nature neuroscience*, 16(4):486–492, 2013. [Cited on p. 44.]
- J. Schmidhuber. Towards compositional learning in dynamic networks. Technical University of Munich (Technical Report FKI-129-90) 990. [Cited on p. 44.]
- J. Schmidhuber. Learning factorial codes by predictability minimization. *Neural computation*, 4(6):863–879, 1992. [Cited on p. 44.]
- B. Schölkopf, F. Locatello, S. Bauer, N. R. Ke, N. Kalchbrenner, A. Goyal, and Y. Bengio. Toward causal representation learning. *Proceedings of the IEEE*, 109(5):612–634, 2021. [Cited on p. 1.]
- L. Schott, J. von Kugelgen, F. Taubert, P. V. Gehler, C. Russell, M. Bethge, B. Schölkopf, F. Locatello, and W. Brendel. Visual representation learning does not generalize strongly within the same domain. In *International Conference on Learning Representations*, 2022. [Cited on p. 2 and 45.]
- M. Seitzer, M. Horn, A. Zadaianchuk, D. Zietlow, T. Xiao, C.-J. Simon-Gabriel, T. He, Z. Zhang, B. Schölkopf, T. Brox, and F. Locatello. Bridging the gap to real-world object-centric learning. In *International Conference on Learning Representations*, 2023. [Cited on p. 9, 10, 46, 47, 48, 49, and 51.]
- G. Singh, F. Deng, and S. Ahn. Illiterate DALL-E learns to compose. *International Conference on Learning Representations*, 2022a. [Cited on p. 9, 46, 47, and 50.]



- G. Singh, Y.-F. Wu, and S. Ahn. Simple unsupervised object-centric learning for complex and naturalistic videos. In *Advances in Neural Information Processing Systems*, volume 35, pages 18181–18196, 2022b. [Cited on p. 9.]
- C. Summer eld. *Natural General Intelligence: How understanding the brain can help us build AI*. Oxford university press, 2022. [Cited on p. 45.]
- J. B. Tenenbaum, C. Kemp, T. L. Griffiths, and N. D. Goodman. How to grow a mind: Statistics, structure, and abstraction. *Science* 331(6022):1279–1285, 2011. [Cited on p. 1.]
- N. Tishby, F. C. Pereira, and W. Bialek. The information bottleneck method. *arXiv preprint physics/0004057*, 2000. [Cited on p. 44.]
- E. C. Tolman. Cognitive maps in rats and men. *Psychological review* 55 4:189–208, 1948. [Cited on p. 45.]
- A. Vani, B. Nguyen, S. Lavoie, R. Krishna, and A. Courville. Sparo: Selective attention for robust and compositional transformer encodings for vision. *arXiv preprint arXiv:2404.15721*, 2024. [Cited on p. 45.]
- B. Varici, E. Acarürk, K. Shanmugam, and A. Tajer. General identifiability and achievability for causal representation learning. *International Conference on Artificial Intelligence and Statistics*, pages 2314–2322. PMLR, 2024. [Cited on p. 3.]
- A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems*, volume 30, 2017. [Cited on p. 2, 8, and 42.]
- J. von Kügelgen, Y. Sharma, L. Gresele, W. Brendel, B. Schölkopf, M. Besserve, and F. Locatello. Self-supervised learning with data augmentations provably isolates content from style. *Advances in Neural Information Processing Systems*, volume 34, pages 16451–16467, 2021. [Cited on p. 3.]
- J. von Kügelgen, M. Besserve, L. Wendong, L. Gresele, A. Keesen, E. Bareinboim, D. M. Blei, and B. Schölkopf. Nonparametric identifiability of causal representations from unknown interventions. In *Advances in Neural Information Processing Systems*, 2023. [Cited on p. 3.]
- Y. Wang, L. Liu, and J. Dauwels. Slot-vae: Object-centric scene generation with slot attention. In *International Conference on Machine Learning*, pages 36020–36035. PMLR, 2023. [Cited on p. 45.]
- N. Watters, L. Matthey, S. Borgeaud, R. Kabra, and A. Lerchner. Spriteworld: A flexible, configurable reinforcement learning environment, 2019a. [Cited on p. 2, 9, and 50.]
- N. Watters, L. Matthey, C. P. Burgess, and A. Lerchner. Spatial broadcast decoder: A simple architecture for learning disentangled representations in VAEs. *arXiv preprint arXiv:1901.07017*, 2019b. [Cited on p. 2, 9, and 50.]
- T. Wiedemer, J. Brady, A. Panfili, A. Juhos, M. Bethge, and W. Brendel. Provable compositional generalization for object-centric learning. *International Conference on Learning Representations*, 2024a. [Cited on p. 2, 4, 45, and 49.]
- T. Wiedemer, P. Mayilvahanan, M. Bethge, and W. Brendel. Compositional generalization from first principles. In *Advances in Neural Information Processing Systems*, volume 36, 2024b. [Cited on p. 3, 7, and 9.]
- D. Yao, D. Xu, S. Lachapelle, S. Magliacane, P. Taslakian, G. Martius, J. von Kügelgen, and F. Locatello. Multi-view causal representation learning with partial observability. *International Conference on Learning Representations*, 2024. [Cited on p. 3.]
- J. M. Zacks, C. A. Kurby, M. L. Eisenberg, and N. Haroutunian. Prediction error associated with the perceptual segmentation of naturalistic events. *Journal of Cognitive Neuroscience* 23:4057–4066, 2011. [Cited on p. 44.]
- Y. Zheng and K. Zhang. Generalizing nonlinear ICA beyond structural sparsity. *Advances in Neural Information Processing Systems*, volume 36, pages 13326–13355, 2023. [Cited on p. 4.]
- R. S. Zimmermann, Y. Sharma, S. Schneider, M. Bethge, and W. Brendel. Contrastive learning inverts the data generating process. *International Conference on Machine Learning*, pages 12979–12990. PMLR, 2021. [Cited on p. 3.]

# Appendices

## Table of Contents

---

A	Disentanglement Proofs	19
A.1	Additional Definitions and Lemmas . . . . .	19
A.2	Sufficient Independence Assumptions . . . . .	20
A.3	From Local to Global Disentanglement . . . . .	21
A.4	Disentanglement (At Most <sup>1st</sup> Order/No Interaction) . . . . .	23
A.5	Disentanglement (At Most <sup>1st</sup> Order Interaction) . . . . .	26
A.6	Disentanglement (At Most <sup>2nd</sup> Order interaction) . . . . .	30
B	Multi-Index Notation	34
C	Characterization of Functions With At Most $n^{\text{th}}$ Order Interactions	34
D	Compositional Generalization Proofs	37
E	Unifying Assumptions from Prior Work	40
E.1	At Most <sup>0th</sup> Order Interaction Across Slots . . . . .	40
E.2	At Most <sup>1st</sup> Order Interaction Across Slots . . . . .	41
F	Transformers for Interaction Regularization	42
F.1	Jacobian of Cross-Attention Mechanism . . . . .	43
F.2	Interaction Regularizer . . . . .	43
G	Extended Related Work	44
G.1	Theory . . . . .	44
G.2	Method and Experiments . . . . .	45
H	Extended Discussion	46
H.1	Theoretical Assumptions . . . . .	46
H.2	Method and Experiments . . . . .	46
H.3	Enforcing Theoretical Criteria Out-of-Domain. . . . .	48
I	Experiments on CLEVRText	49
I.1	Experimental Setup . . . . .	49
I.2	Results . . . . .	49
J	Experimental Details	50
J.1	Data, Model, and Training Details . . . . .	50
J.2	Metrics and Evaluation . . . . .	51
J.3	Additional Figures . . . . .	51

---

## A DISENTANGLEMENT PROOFS

### A.1 ADDITIONAL DEFINITIONS AND LEMMAS

**Definition A.1** ( $C^k$ -diffeomorphism) Let  $A \subseteq \mathbb{R}^n$  and  $B \subseteq \mathbb{R}^m$ . A map  $f : A \rightarrow B$  is said to be a  $C^k$ -diffeomorphism if it is bijective,  $C^k$ , and has a  $C^k$  inverse.

**Remark A.2.** The property of being differentiable is usually defined only for functions with an open domain of  $\mathbb{R}^n$ . Note that, in the definition above, both  $A$  and  $B$  might not be open sets in their respective topologies. For an arbitrary domain  $A \subseteq \mathbb{R}^n$ , we say that a function is  $C^k$  if it can be extended to a  $C^k$  function defined on an open set containing  $A$ . More precisely  $f : A \rightarrow B$  is  $C^k$  if there exists a function  $g : U \rightarrow \mathbb{R}^m$  such that 1)  $U$  is an open set containing  $A$ , 2)  $g$  is  $C^k$ , and 3)  $g(a) = f(a)$  for all  $a \in A$ . See p.199 of [Munkres \(1991\)](#) for details about such constructions.

**Definition A.3** (Regular closed sets) A set  $Z_{\text{supp}} \subseteq \mathbb{R}^d$  is regular closed if  $Z_{\text{supp}} = \overline{Z_{\text{supp}}^\circ}$  i.e. if it is equal to the closure of its interior (in the standard topology of  $\mathbb{R}^d$ ).

**Lemma A.4** ([Lachapelle et al. \(2023\)](#)) Let  $A, B \subseteq \mathbb{R}^n$  and suppose there exists a homeomorphism  $f : A \rightarrow B$ . If  $A$  is regular closed in  $\mathbb{R}^n$ , we have that  $B$  is regular closed in  $\mathbb{R}^n$ .

The way we define  $C^k$  functions with arbitrary domain is such that a function can be differentiable without having a uniquely defined derivative everywhere on its domain. This happens when the derivative of two distinct extensions differ. The following Lemma states that the derivative of a function is uniquely defined on the closure of the interior of its domain.

**Lemma A.5** ([Lachapelle et al. \(2023\)](#)) Let  $A \subseteq \mathbb{R}^n$  and  $f : A \rightarrow \mathbb{R}^m$  be a  $C^k$  function. Then, its  $k$ -th derivatives are uniquely defined on  $\overline{A^\circ}$  in the sense that they do not depend on the specific choice of  $C^k$  extension.

**Notation** For a subset  $S \subseteq [d_z]$  and a matrix  $A \in \mathbb{R}^{m \times n}$ ,  $A_S$  will denote the sub-matrix consisting of the columns in  $A$  indexed by  $S$  i.e.  $A_S = [A_{:,i}]_{i \in S}$ . Similarly, for a vector  $z$ ,  $z_S$  will denote the sub-vector of  $z$  consisting of components indexed by  $S$  i.e.  $z_S := (z_i)_{i \in S}$ .

**Lemma A.6.** Let  $A \in \mathbb{R}^{m \times n}$  and let  $B$  be a partition of  $[n]$ . If

$$\text{rank}(A) = \sum_{S \in B} \text{rank}(A_S) \quad (\text{A.1})$$

Then  $\exists z \in \mathbb{R}^n$  s.t.  $Az = 0$ ,  $A_S z_S = 0$ , for any  $S \in B$ .

**Proof.** Assume for a contradiction that there exists  $z \in \mathbb{R}^n$ , s.t.  $Az = 0$ , and there exists  $S_1 \in B$  s.t.  $A_{S_1} z_{S_1} \neq 0$ .

Now construct the matrix, denoted  $A_{-S_1}$ , consisting of all columns in  $A$  except those indexed by  $S_1$ , i.e.

$$A_{-S_1} := [A_{:,i}]_{i \in [n] \setminus S_1} \quad (\text{A.2})$$

By using (A.1) and the property that  $\text{rank}(B; C) = \text{rank}(B) + \text{rank}(C)$ , we get

$$\text{rank}(A) = \sum_{S \in B \setminus S_1} \text{rank}(A_S) + \text{rank}(A_{S_1}) \quad (\text{A.3})$$

$$\text{rank}(A_{-S_1}) + \text{rank}(A_{S_1}) \quad (\text{A.4})$$

$$\text{rank}(A) \quad (\text{A.5})$$

Consequently, we have that:

$$\text{rank}(A) = \text{rank}(A_{-S_1}) + \text{rank}(A_{S_1}) \quad (\text{A.6})$$

This implies that the column spaces of both matrices denoted  $(\text{range}(A_{-S_1}); \text{range}(A_{S_1}))$  respectively, do not intersect, except at the zero vector.

A simple example of such a situation is the trivial function  $f : \mathbb{R} \rightarrow \mathbb{R}$  which is differentiable at 0 but does not have a well defined derivative because  $g(x) = x$  and  $h(x) = -x$  are both differentiable extensions of  $f$  but have different derivatives at  $x=0$ .

Now we know that  $Az = 0$

$$0 = Az \tag{A.7}$$

$$= A_{S_1} z_{S_1} + A_{S_1^c} z_{S_1^c} \tag{A.8}$$

Consequently,

$$A_{S_1} z_{S_1} = -A_{S_1^c} z_{S_1^c} \tag{A.9}$$

and by our assumed contradiction we know that

$$A_{S_1} z_{S_1} \neq 0 \tag{A.10}$$

This implies that the column spaces of  $A_{S_1}$  and  $A_{S_1^c}$  must intersect at a point other than the zero vector, which is a contradiction.  $\square$

**Lemma A.7.** Let  $A \in \mathbb{R}^{d \times d}$  be an invertible matrix and  $B_1, \dots, B_K$  be a partition of  $[d]$ . Assume there are  $k_1, k_2, k \in [K]$  such that:

$$A_{B_k; B_{k_1}} \neq 0 \neq A_{B_k; B_{k_2}} \tag{A.11}$$

Then there exists a subset  $S \subseteq [d]$  with cardinality  $|S| = k$  that has the following properties:

1. The sub-block  $A_{B_k; S}$  is invertible.
2.  $S \cap B_{k^0} = \emptyset$ , for any  $k^0 \in [K]$

*Proof.* We first prove that there must exist a satisfying point 1. Since  $A$  is invertible, each subset of rows is linearly independent and thus  $\text{rank}(A_{S, :}) = |S|$ . This implies that there exist a set  $S \subseteq [d]$  with cardinality  $|S| = k$  such that  $B_k \cap S = \emptyset$ ;  $A_{B_k; S}$  are linearly independent, and thus form a basis of  $\mathbb{R}^{B_k}$ .

If  $S \cap B_{k^0} = \emptyset$  for all  $k^0 \in [K]$ , we are done.

We consider the case where there exists  $S$  such that  $S \cap B_{k^0} \neq \emptyset$ . We will show that we can construct a different  $S$  from  $S$  which satisfies both conditions.

We know by (A.11) that there exist a second block  $k^0 \in [K]$  such that for some  $j \in B_{k^0}$ ,  $A_{B_k; j} \neq 0$ . Since  $A_{B_k; S}$  forms a basis of  $\mathbb{R}^{B_k}$ , the vector  $A_{B_k; j}$  can be represented uniquely as

$$A_{B_k; j} = \sum_{i \in S} a_i A_{B_k; i} \tag{A.12}$$

where  $a_i \in \mathbb{R}$  for all  $i$ . Because  $A_{B_k; j} \neq 0$ , there exists  $j \in S$  such that  $a_j \neq 0$ . Because this representation is unique, we know that  $A_{B_k; j}$  is outside the span of  $\{A_{B_k; i} \mid i \in S \setminus \{j\}\}$ . This means that, by taking  $S' := (S \setminus \{j\}) \cup \{j^0\}$ , we have that  $A_{B_k; S'}$  is a basis for  $\mathbb{R}^{B_k}$  or, in other words,  $A_{B_k; S'}$  is invertible. Also,  $S'$  is not included in a single block since  $S \cap B_{k^0} = \emptyset$  and  $j^0 \in B_{k^0}$  with  $k^0 \neq k$ .  $\square$

## A.2 SUFFICIENT INDEPENDENCE ASSUMPTIONS

**Definition A.8 (Sufficient Independence (0<sup>th</sup> Order)).** Let  $f : \mathbb{R}^{d_z} \rightarrow \mathbb{R}^{d_x}$  be a  $C^1$  function with 0<sup>th</sup> order interactions between slots (Def. 3.2). The function is said to have sufficiently independent derivatives if  $\forall z \in \mathbb{R}^{d_z}$ :

$$\text{rank} [D_i f(z)]_{i \in B_k, k \in [K]} = \sum_{k \in [K]} \text{rank} [D_i f(z)]_{i \in B_k} \tag{A.13}$$

**Definition A.9 (Sufficient Independence (1<sup>st</sup> Order)).** Let  $f : \mathbb{R}^{d_z} \rightarrow \mathbb{R}^{d_x}$  be a  $C^2$  function with at most 1<sup>st</sup> order interactions between slots (Def. 3.3). The function is said to have sufficiently independent derivatives if  $\forall z \in \mathbb{R}^{d_z}$ :

$$\begin{aligned} \text{rank} [D_i f(z)]_{i \in B_k} &= \text{rank} [D_{i,j}^2 f(z)]_{(i,j) \in B_k^2} \\ &= \sum_{k \in [K]} \text{rank} [D_i f(z)]_{i \in B_k} + \sum_{i \in B_k} \text{rank} [D_{i,j}^2 f(z)]_{(i,j) \in B_k^2} \end{aligned}$$

Definition 4.2 (Sufficient Independence<sup>2nd</sup> Order). A  $C^3$  function  $f : \mathbb{R}^{d_z} \rightarrow \mathbb{R}^{d_x}$  with at most  $2^{\text{nd}}$  order interactions across slots is said to have sufficiently independent derivatives if  $\forall z \in \mathbb{R}^{d_z}$ :

$$\text{rank} \left[ D_i f(z) \right]_{i \in B_k} \left[ D_{i_1}^2 f(z) \right]_{i_1 \in B_k, i_2 \in B_k} \left[ D_{i_1 i_2}^3 f(z) \right]_{(i_1, i_2) \in B_k^3} \geq \sum_{k \in [K]} \text{rank} \left[ D_i f(z) \right]_{i \in B_k} \left[ D_{i_1}^2 f(z) \right]_{i_1 \in B_k, i_2 \in B_k} + \text{rank} \left[ D_{i_1 i_2}^3 f(z) \right]_{(i_1, i_2) \in B_k^3} :$$

### A.3 FROM LOCAL TO GLOBAL DISENTANGLEMENT

This section takes care of technical subtleties when one has to go from local to global disentanglement. The disentanglement guarantee of this work is proven by first showing that the Jacobian of  $h := f^{-1} \circ f^\wedge$ , has a block-permutation structure everywhere, and from there showing that  $h$  can be written as  $h(z) = (h_1(z_{B_{(1)}}); h_2(z_{B_{(2)}}); \dots; h_K(z_{B_{(K)}}))$  (see Definition 2.1). Lachapelle et al. (2023) refers to the first condition as local disentanglement and the second condition as global disentanglement, the latter of which corresponds to the definition of disentanglement employed in the present work. The authors also show that going from local to global disentanglement requires special care when considering very general supports like we do in this work, as opposed to the more common assumption  $Z_{\text{supp}} := \mathbb{R}^{d_z}$  which makes this step more direct (e.g., see Harinen et al. (2019)). This section reuses definitions and lemmata taken from Lachapelle et al. (2023) and introduces a novel sufficient condition on the support of the latent factors, we named aligned-connectedness, to guarantee that the jump from local to global disentanglement can be made.

Definition A.10 (Partition-respecting permutations). Let  $B := \{B_1; B_2; \dots; B_K\}$  be a partition of  $\{1; \dots; d\}$ . A permutation  $\sigma$  over  $\{1; \dots; d\}$  respects  $B$  if, for all  $B \in B; (B) \in B$ .

Definition A.11 (B-block permutation matrices). A matrix  $A \in \mathbb{R}^{d \times d}$  is a B-block permutation matrix if it is invertible and can be written as  $A = CP$  where  $P$  is the matrix representing the B-respecting permutation (Definition A.10), i.e.  $P e_i = e_{\sigma(i)}$ , and  $C \in \mathbb{R}^{d \times d}$  is such that for all distinct blocks  $B; B^0 \in B, C_{B; B^0} = 0$ .

Proposition A.12. The inverse of a B-block permutation matrix is also a B-block permutation matrix.

Proof. First note that  $C$  must be invertible, otherwise  $A$  is not. Also,  $C^{-1}$  must also be such that  $(C^{-1})_{B; B^0} = 0$  for all distinct blocks  $B; B^0 \in B$ . This is because, without loss of generality, we can assume the blocks  $B$  are contiguous which implies that  $A$  is a block diagonal matrix so that  $C^{-1}$  is also block diagonal. Since  $\sigma$  preserves  $B$ , we have that  $P^T C^{-1} P$  is also block diagonal since, for all distinct  $B; B^0 \in B, (P^T C^{-1} P)_{B; B^0} = (C^{-1})_{(B); (B^0)} = 0$ , where we used the fact that the blocks  $(B)$  and  $(B^0)$  are in  $B$ , because  $\sigma$  is B-preserving, and are distinct, because  $\sigma$  is a bijection. We can thus see that

$$\begin{aligned} A^{-1} &= P^T C^{-1} P \\ &= P^T C^{-1} P P^T \\ &= C P^T \\ &= C P^{-1}; \end{aligned}$$

where  $C := P^T C^{-1} P$  is block diagonal and  $P^{-1}$  is block-preserving.  $\square$

Definition A.13 (Local disentanglement; Lachapelle et al. (2023)). A learned decoder  $f^\wedge : \mathbb{R}^{d_z} \rightarrow \mathbb{R}^{d_x}$  is said to be locally disentangled w.r.t. the ground-truth decoder  $f$  when  $f^\wedge \circ h(z) = f(z)$  for all  $z \in Z_{\text{supp}}$  where the mapping  $h$  is a diffeomorphism from  $Z_{\text{supp}}$  onto its image satisfying the following property: for all  $z \in Z_{\text{supp}}, Dh(z)$  is a block-permutation matrix respecting  $B := \{B_1; \dots; B_K\}$ .

Note that, in the above definition, the permutation of the blocks might change from one to another (see Example 5 in Lachapelle et al. (2023)). To prevent this possibility, we will assume  $Z_{\text{supp}}$  is path-connected.

**Definition A.14 (Path-connected sets)** A set  $Z_{\text{supp}} \subset \mathbb{R}^d$  is path-connected if for all pairs of points  $z^0, z^1 \in Z_{\text{supp}}$  there exists a continuous map:  $[0; 1] \rightarrow Z_{\text{supp}}$  such that  $(0) = z^0$  and  $(1) = z^1$ . Such a map is called a path between  $z^0$  and  $z^1$ .

The following Lemma from [Lachapelle et al. \(2023\)](#) can be used to show that when a diffeomorphism  $h$  and  $Z_{\text{supp}}$  is path-connected, the block structure cannot change. This is due to the fact that  $Dh(z)$  is invertible everywhere and a continuous function of  $z$ . We restate the Lemma without proof.

**Lemma A.15** ([Lachapelle et al. \(2023\)](#)) Let  $C$  be a path-connected topological space and  $M: C \rightarrow \mathbb{R}^{d \times d}$  be a continuous function. Suppose that, for all  $c \in C$ ,  $M(c)$  is an invertible  $B$ -block permutation matrix (Definition A.11). Then, there exists a  $B$ -respecting permutation such that for all  $c \in C$  and all distinct  $B, B' \in \mathcal{B}$ ,  $M(c)_{(B'), B} = 0$ .

It turns out that, in general, having that  $Dh$  has a constant block-permutation structure across its support  $Z_{\text{supp}}$  is not enough to make the jump to global disentanglement. See Example 7 from [Lachapelle et al. \(2023\)](#). We now propose a novel condition on the support and will show it is sufficient to guarantee global disentanglement in Lemma A.18.

**Definition A.16 (Aligned-connected sets)** A set  $A \subset \mathbb{R}^d$  is said to be aligned-connected w.r.t. a partition  $\mathcal{B} = \{B_1, B_2, \dots, B_K\}$  if, for all  $k \in [K]$  and all  $a^0 \in A$ , the set  $\{a \in A \mid a_{B_k} = a_{B_k}^0\}$  is path-connected.

**Remark A.17 (Relation to path-connectedness)** There exist sets that are path-connected but not aligned-connected and vice-versa. Example 7 from [Lachapelle et al. \(2023\)](#) presents a ‘‘U-shaped’’ support that is path-connected but not aligned-connected. Moreover, the set  $A^{(1)} \cup A^{(2)}$  where  $A^{(1)} := \{a \in \mathbb{R}^2 \mid a_1 \in [1; 2], a_2 \in [1; 1]\}$  and  $A^{(2)} := \{a \in \mathbb{R}^2 \mid a_1 \in [1; 2], a_2 \in [1; 1]\}$  is aligned-connected w.r.t. the partition  $\mathcal{B} = \{[1; 1], [2; 2]\}$  but not path-connected.

We now show how aligned-connectedness combined with path-connectedness is enough to guarantee global disentanglement from local disentanglement.

**Lemma A.18 (Local to global disentanglement)** Suppose  $h$  is a diffeomorphism from  $Z_{\text{supp}} \subset \mathbb{R}^d$  to its image and suppose  $Dh(z)$  is a  $B$ -block permutation matrix for all  $z \in Z_{\text{supp}}$  (local disentanglement). If  $Z_{\text{supp}}$  is path-connected (Defn. A.14) and aligned-connected set (Defn. A.16), then  $h(z) = (h_1(z_{B_{(1)}}); \dots; h_1(z_{B_{(K)}}))$  for all  $z \in Z_{\text{supp}}$  where the  $h_k$  are diffeomorphisms (global disentanglement).

**Proof.** Since  $h$  is a diffeomorphism  $Dh$  is continuous and  $Dh(z)$  is invertible for all  $z \in Z_{\text{supp}}$ . Since we also have that  $Z_{\text{supp}}$  is path-connected, we can apply Lemma A.15 to get that there exists a permutation  $\sigma: [K] \rightarrow [K]$  such that, for all  $z \in Z_{\text{supp}}$  and all distinct  $k, k' \in [K]$ , we have  $Dh(z)_{B_k; B_{(k')}} = 0$ . In other words,  $D_{B_{(k')}} h_{B_k}(z) = 0$ . We must now show that  $h_{B_k}(z)$  depends solely on  $z_{B_{(k)}}$ . Consider another point  $z^0 \in Z_{\text{supp}}$  such that  $z_{B_{(k)}} = z_{B_{(k)}}^0$ . We will now show that  $h_{B_k}(z) = h_{B_k}(z^0)$ , i.e. changing  $z_{B_{(k)'}}$  does not influence  $h_{B_k}(z)$ .

Because  $Z_{\text{supp}}$  is aligned-connected, there exists a continuous path  $[0; 1] \rightarrow Z_{\text{supp}}$  such that  $(0) = z^0$ ,  $(1) = z$  and  $z_{B_{(k)}}(t) = z_{B_{(k)}} = z_{B_{(k)}}^0$  for all  $t \in [0; 1]$ . By the fundamental

<sup>1</sup>This lemma also holds if  $C$  is connected.

theorem of calculus, we have that

$$\begin{aligned}
 h_{B_k}(z) - h_{B_k}(z^0) &= \int_{z^0}^z (h_{B_k}(\cdot))'(t) dt \\
 &= \int_{z^0}^z D h_{B_k}(\cdot)(t) dt \\
 &= \int_{z^0}^z \left( \sum_{k^0 \in K} D_{B_{(k^0)}} h_{B_k}(\cdot)(t) \right) dt + \sum_{k^0 \in K} \int_{z^0}^z D_{B_{(k^0)}} h_{B_k}(\cdot)(t) dt \\
 &= \int_{z^0}^z \left( \sum_{k^0 \in K} D_{B_{(k^0)}} h_{B_k}(\cdot)(t) \right) dt + \sum_{k^0 \in K} \int_{z^0}^z D_{B_{(k^0)}} h_{B_k}(\cdot)(t) dt \\
 &= 0;
 \end{aligned}$$

where we used the fact that  $\mathbb{1}_{B_{(k^0)}}(t)$  is a constant function and  $D_{B_{(k^0)}} h_{B_k}(z) = 0$  for distinct  $k; k^0$ .

We conclude that, for all  $k$ , we can write  $h_{B_k}(z) = h_{B_k}(z_{B_{(k)}})$ , which is the desired result.

Additionally, the functions  $h_{B_k}(z_{B_{(k)}})$  are diffeomorphisms because their Jacobians must be invertible otherwise the Jacobian of  $h$  (which is block diagonal) would not be invertible (which would violate the fact that it is a diffeomorphism).  $\square$

Contrasting with [Lachapelle et al. \(2023\)](#), instead of assuming aligned-connectedness, [Lachapelle et al. \(2023\)](#) assumed that the block-specific decoders, which would correspond to  $(z_k)$  in (4.2), are injective which, when combined with path-connectedness, is also enough to go from local to global disentanglement in the context of additive decoders (1). Whether a similar strategy could be adapted for more general decoders with at most order interactions is left as future work.

#### A.4 DISENTANGLEMENT (AT MOST 0<sup>TH</sup> ORDER/NO INTERACTION)

Lemma A.19. Let  $Z_{\text{supp}} \subset Z$  be a regular closed set (Defn. A.3). Let  $f: Z \rightarrow X$  be  $C^1$  and  $h: Z_{\text{supp}} \rightarrow Z_{\text{supp}}$  be a diffeomorphism. Let  $\hat{f} := f \circ h$ . If  $f$  has no interaction (Definition 3.4 with  $n = 0$ ), then, for all  $j; j^0 \in [d_z]$  and  $z \in Z_{\text{supp}}$  we have

$$D_j \hat{f}(z) - D_{j^0} \hat{f}(z) = W^f(h(z)) m^h(z; (j; j^0)); \tag{A.14}$$

where

$$\begin{aligned}
 W^f(z) &:= [W_k^f(z)]_{k \in [K]} \\
 W_k^f(z) &:= [D_{i_1} f(z) \quad D_{i_2} f(z)]_{(i_1, i_2) \in 2B_k^2} \\
 m^h(z; (j; j^0)) &:= [m_k^h(z; (j; j^0))]_{k \in [K]} \\
 m_k^h(z; (j; j^0)) &:= [D_j \circ h_{i_1}(z) D_{j^0} h_{i_2}(z)]_{(i_1, i_2) \in 2B_k^2};
 \end{aligned}$$

Proof. We have that

$$\hat{f}(z) = f \circ h(z); \quad z \in Z_{\text{supp}}$$

Following the same line of argument as [Lachapelle et al. \(2023\)](#), we can use Lemma A.5 to say that the function  $\hat{f}(z) = f \circ h(z)$  has well-defined derivatives on  $Z_{\text{supp}}$ . Since  $h^{-1}$  is a diffeomorphism from  $Z_{\text{supp}}$  (which is regular closed) to  $Z_{\text{supp}}$  Lemma A.4 implies that  $\hat{f}(z)$  has well-defined derivatives on  $Z_{\text{supp}}$ . This means that the function  $\hat{f}(z) = f \circ h(z)$  has well-defined derivatives for all  $z \in Z_{\text{supp}}$ .

By taking the derivative w.r.t.  $z_j$  on both sides of  $\hat{f}(z) = f \circ h(z)$ , we get

$$D_j \hat{f}(z) = \sum_{k \in [K]} \sum_{i \in 2B_k} D_i f(h(z)) D_j h_i(z) \tag{A.15}$$

We thus have that

$$\begin{aligned}
 D_j \hat{f}(z) D_j \circ \hat{f}(z) &= \prod_{k \in [K]} \left( \prod_{i \in \mathcal{B}_{k_1}} D_{i_1} f(h(z)) D_{j \circ h_{i_1}}(z) \right) \prod_{k \in [K]} \left( \prod_{i \in \mathcal{B}_{k_2}} D_{i_2} f(h(z)) D_{j \circ h_{i_2}}(z) \right) \\
 &= \prod_{k \in [K]} \left( \prod_{i \in \mathcal{B}_{k_1}} D_{i_1} f(h(z)) \right) \prod_{k \in [K]} \left( \prod_{i \in \mathcal{B}_{k_2}} D_{i_2} f(h(z)) D_{j \circ h_{i_2}}(z) \right) \\
 &= \prod_{k \in [K]} \left( \prod_{i \in \mathcal{B}_{k_1}} D_{i_1} f(h(z)) \right) \prod_{k \in [K]} \left( \prod_{i \in \mathcal{B}_{k_2}} D_{i_2} f(h(z)) D_{j \circ h_{i_2}}(z) \right);
 \end{aligned}$$

where the last equality used the fact that  $f$  has no interaction (Definition 3.2). We conclude by noticing

$$\begin{aligned}
 D_j \hat{f}(z) D_j \circ \hat{f}(z) &= \prod_{k \in [K]} \left( \prod_{i \in \mathcal{B}_{k_1}} D_{i_1} f(h(z)) \right) \prod_{k \in [K]} \left( \prod_{i \in \mathcal{B}_{k_2}} D_{i_2} f(h(z)) D_{j \circ h_{i_2}}(z) \right) \\
 &= W^f(h(z)) m^h(z; (j; j^0));
 \end{aligned}$$

□

**Theorem A.20.** Let  $f : Z \rightarrow X$  be a  $C^1$  diffeomorphism satisfying interaction asymmetry (Asm. 3.5) for all equivalent generators (Defn. 4.1) for  $\tau = 0$ . Let  $Z_{\text{supp}} \subseteq Z$  be regular closed (Defn. A.3), path-connected (Defn. A.14) and aligned-connected (Defn. A.16). A model  $\hat{f} : Z \rightarrow \mathbb{R}^{d_x}$  disentangles  $z$  on  $Z_{\text{supp}}$  w.r.t.  $f$  (Defn. 2.1) if it is (i) a  $C^1$  diffeomorphism between  $\hat{Z}_{\text{supp}}$  and  $X_{\text{supp}}$  with (ii) at most  $\Theta$  order interactions across slots (Defn. 3.4)  $\hat{Z}_{\text{supp}}$ .

**Proof.** As mentioned in Section A.3, the proofs will proceed in two steps: First, we show local disentanglement (Definition A.13) and then we show (global) disentanglement via Lemma A.18. We first show local disentanglement.

**Remark:** We will use the following notation below:

$$D_{i;j}^1 f(z) := D_j f(z) \prod_{i \in \mathcal{B}_k} D_i f(z) \in \mathbb{R}^m \quad (\text{A.16})$$

We first define the function  $h : \hat{Z}_{\text{supp}} \rightarrow Z_{\text{supp}}$  relating the latent spaces of these functions  $\hat{f}$  and  $f$ :

$$h := f^{-1} \circ \hat{f} \quad (\text{A.17})$$

The function  $\hat{f}$  can then be written in terms of  $h$  and  $f$  on  $\hat{Z}_{\text{supp}}$ :

$$\hat{f} = f \circ h \quad (\text{A.18})$$

Because  $f, \hat{f}$  are both  $C^1$  diffeomorphism between  $Z_{\text{supp}}, X_{\text{supp}}$  and  $\hat{Z}_{\text{supp}}, X_{\text{supp}}$ , respectively, we have that  $h$  is a  $C^1$  diffeomorphism.

By Lemma A.19, for all  $z \in \hat{Z}_{\text{supp}}, (j; j^0) \in [d_z]$ , we have:

$$D_j \hat{f}(z) D_j \circ \hat{f}(z) = W^f(h(z)) m^h(z; (j; j^0)) \quad (\text{A.19})$$

where  $W^f$  and  $m^h$  are defined in Lemma A.19.

Define the sets

$$D := \prod_{k \in [K]} B_k^2; \quad D^c := \{f_1, \dots, f_{d_z}\} \cap D \quad (\text{A.20})$$

Because  $\hat{f}$  has no interaction (Definition 3.2), we have that, for  $(j; j^0) \in D^c$

$$\begin{aligned}
 0 &= W^f(h(z)) m^h(z; (j; j^0)) \\
 &= \prod_{k \in [K]} W_k^f(h(z)) m_k^h(z; (j; j^0));
 \end{aligned}$$



Because  $f$  has no interaction, each row  $W_k^f(h(z))_{n_i}$  is non-zero for at most one  $k \in [K]$  (although this  $k$  can change for different values of  $z$ ). This implies that for all  $z \in Z_{\text{supp}}; (j; j^0) \in D^c, k \in [K]$ :

$$0 = W_k^f(h(z)) m_k^h(z; (j; j^0)) \quad (\text{A.21})$$

Case 1:  $|B_k| = 1$  (One-Dimensional Slots)

When  $|B_k| = 1$ , for all  $k \in [d_z]$ , the matrix  $W_k^f(h(z))$  can be written as:

$$W_k^f(h(z)) = [D_k f(z) \quad D_k f(z)] \quad (\text{A.22})$$

This matrix has a single column, which must be non-zero since  $f$  is a  $C^1$  diffeomorphism. Thus,  $W_k^f(h(z))$  has full column rank and thus has a null space equal to  $\{0\}$ . Using Eq. (A.21), we conclude that for all  $(j; j^0) \in D^c, k \in [d_z]$ :

$$0 = m_k^h(z; (j; j^0)) \quad (\text{A.23})$$

Applying the definition of  $m_k^h(z; (j; j^0))$ , this implies that for all  $(j; j^0) \in D^c, k \in [d_z]$ :

$$0 = D_j \circ h_k(z) D_j h_k(z) \quad (\text{A.24})$$

This means each row of the Jacobian matrix  $Dh(z)$  cannot have more than one nonzero value. Since the Jacobian is invertible, these nonzero values must all be different for different rows, otherwise a whole column would be zero. Hence  $Dh(z)$  is a permutation-scaling matrix, i.e. we have local disentanglement.

Case 2:  $|B_k| > 1$  (Multi-Dimensional Slots)

Assume for a contradiction that  $\hat{A}$  is not locally disentangled w.r.t.  $f$ . This implies that there exist  $z \in Z_{\text{supp}}; k, k^0 \in [K]$  for  $k \neq k^0$ , such that:

$$D_{B_{k^0}} h_{B_k}(z) \neq 0; \quad D_{B_{k^0}} h_{B_k}(z) \neq 0 \quad (\text{A.25})$$

Because  $f, \hat{A}$  are  $C^1$  diffeomorphisms, we know that  $\hat{A} \circ f$  is also a  $C^1$  diffeomorphism. Coupling this with Eq. (A.25), Lemma A.7 tells us that there exist  $S \subseteq [d_z]$  with cardinality  $|B_k|$  such that:

$$\forall B \subseteq B; S \subseteq B; \text{ and } \forall i \in S; D_i h_{B_k}(z) \text{ are linearly independent.} \quad (\text{A.26})$$

Now choose any  $B \subseteq B$  such that  $S_1 := S \setminus B \neq \emptyset$ . Furthermore, define the set  $S_2 := S \cap B$ . Because  $S \subseteq B$ , we know that  $S_2$  is non-empty. Further, by construction  $S_1 \cap S_2 = \emptyset$ . In other words,  $S_1$  and  $S_2$  are non-empty, form a partition of  $S$ , and do not contain any indices from the same slot.

Now construct the matrices, denoted  $A_{S_1}$  and  $A_{S_2}$  as follows:

$$A_{S_1} := D_{S_1} h_{B_k}(z); \quad A_{S_2} := D_{S_2} h_{B_k}(z) \quad (\text{A.27})$$

And the matrix denoted  $A_k$  as:

$$A_k := [A_{S_1}; A_{S_2}] \quad (\text{A.28})$$

Note that because  $\forall i \in S, D_i h_{B_k}(z)$  are linearly independent (Eq. (A.26)), we know that  $A_k$  is invertible.

Now, define the following block diagonal matrix  $A \in \mathbb{R}^{d_z \times d_z}$  as follows:

$$A := \begin{bmatrix} A_1 & 0 & \dots & 0 \\ 0 & A_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & A_K \end{bmatrix} \quad (\text{A.29})$$

where  $\forall i \in [K] n \neq k; A_i$  is the identity matrix, and thus invertible, while  $A_k$  is defined according to Eq. (A.28).

Define  $Z := A^{-1}Z$ , the function  $h : Z \rightarrow Z$  as  $h(z) := Az$  and the function  $f : Z \rightarrow X$  as  $f := f \circ h$ . By construction we have

$$\forall z \in Z; \quad f(A_1^{-1}z_{B_1}; \dots; A_K^{-1}z_{B_K}) = f(z_{B_1}; \dots; z_{B_K}) \quad (\text{A.30})$$

Because all  $A_i^{-1}$  are invertible, then  $m$  is equivalent to  $f$  in the sense of Def. (4.1).

We can now apply Lemma A.19 to  $f = f \circ h$  to obtain, for all  $j; j^0 \in [d_z]$ :

$$D_j f(z) - D_j \circ f(z) = W^f(h(z)) m^h(z; (j; j^0)) : \quad (\text{A.31})$$

Choose  $z \in Z$  such that  $h(z) = h(z')$ , which is possible because  $h(z') \in Z$  and  $h$  is a bijection from  $Z$  to  $Z$ . We can then write

$$D_j f(z) - D_j \circ f(z) = W^f(h(z')) m^h(z; (j; j^0)) : \quad (\text{A.32})$$

Let  $J; J^0 \subseteq B_k$  be a partition of  $B_k$  such that  $J$  is the set of columns of  $A$  corresponding to  $A_{S_1}$  and  $J^0$  be the set of columns of  $A$  corresponding to  $A_{S_2}$ . More formally, we have

$$A_{B_k; J} = A_{S_1} \quad \text{and} \quad A_{B_k; J^0} = A_{S_2}$$

Since  $A_{S_1} = D_{S_1} h_{B_k}(z)$  and  $A_{S_2} = D_{S_2} h_{B_k}(z)$ , we have that

$$A_{B_k; J} = D_{S_1} h_{B_k}(z) \quad \text{and} \quad A_{B_k; J^0} = D_{S_2} h_{B_k}(z)$$

Since  $Dh(z) = A$ , we have

$$D_J h_{B_k}(z) = D_{S_1} h_{B_k}(z) \quad \text{and} \quad D_{J^0} h_{B_k}(z) = D_{S_2} h_{B_k}(z) :$$

Choose some  $(j; j^0) \in J \times J^0$ . We know there must exist  $(s; s^0) \in S_1 \times S_2$  such that

$$D_j h_{B_k}(z) = D_s h_{B_k}(z) \quad \text{and} \quad D_{j^0} h_{B_k}(z) = D_{s^0} h_{B_k}(z) :$$

which implies

$$m_k^h(z; (j; j^0)) = m_k^h(z; (s; s^0)) : \quad (\text{A.33})$$

Moreover, since the Jacobian of  $f$  is block diagonal, we have that  $\partial_{k^0} h_k(z; (j; j^0)) = 0$  for all  $k^0 \notin k$  (recall that  $j^0 \in B_k$ ). This means we can rewrite (A.32) as

$$D_j f(z) - D_j \circ f(z) = W_k^f(h(z)) m_k^h(z; (j; j^0)) : \quad (\text{A.34})$$

Plugging (A.33) into the above equation yields

$$D_j f(z) - D_j \circ f(z) = W_k^f(h(z)) m_k^h(z; (s; s^0)) : \quad (\text{A.35})$$

Since  $(s; s^0) \in S_1 \times S_2 \subseteq D^c$ , we can apply (A.21) to get

$$D_j f(z) - D_j \circ f(z) = W_k^f(h(z)) m_k^h(z; (s; s^0)) = 0 : \quad (\text{A.36})$$

In other words, we found a partition  $J^0$  of the block  $B_k$  such that  $D_j f(z) - D_j \circ f(z) = 0$  for all  $(j; j^0) \in J \times J^0$ . This means that the blocks  $J$  and  $J^0$  have no interaction in  $f$  at  $z$ . This is a contradiction with Assm. 3.5. Hence, we have local disentanglement.

**Local to global disentanglement.** We showed that  $Dh(z)$  is a block-permutation matrix for all  $z \in \hat{Z}_{\text{supp}}$ , i.e. local disentanglement. Consider the inverse  $v := h^{-1}$ . The Jacobian of  $v$  is given by  $Dv(z) = Dh^{-1}(z) = Dh(v(z))^{-1}$ , by the inverse function theorem. By Proposition A.12, this means  $Dv(z)$  is also a block permutation matrix for all  $z \in \hat{Z}_{\text{supp}}$ . Since  $\hat{Z}_{\text{supp}}$  is aligned-connected (Definition A.16), Lemma A.18 guarantees that we can write  $v(z) = (v_1(z_{B_{(1)}}); \dots; v_K(z_{B_{(K)}}))$  for all  $z \in \hat{Z}_{\text{supp}}$  where the  $v_k$  are diffeomorphisms. This implies that  $h(z) = (v_1^{-1}(z_{B_{(1)}}); \dots; v_K^{-1}(z_{B_{(K)}}))$  for all  $z \in \hat{Z}_{\text{supp}}$  which concludes the proof.  $\square$

## A.5 DISENTANGLEMENT (AT MOST 1<sup>ST</sup> ORDER INTERACTION)

**Lemma A.21.** Let  $Z_{\text{supp}} \subseteq Z$  be a regular closed set (Defn. A.3). Let  $f: Z \rightarrow X$  be  $C^1$  and  $h: \hat{Z}_{\text{supp}} \rightarrow Z_{\text{supp}}$  be a diffeomorphism. Let  $\hat{f} := f \circ h$ . If  $f$  has at most 1<sup>st</sup> order interaction (Definition 3.4 with  $n = 1$ ), then, for all  $j; j^0 \in [d_z]$  and  $z \in \hat{Z}_{\text{supp}}$ , we have

$$D_{jj^0}^2 \hat{f}(z) = W^f(h(z)) m^h(z; (j; j^0)) ; \quad (\text{A.37})$$

where

$$\begin{aligned} W^f(z) &:= [W_k^f(z)]_{k \in [K]} \\ W_k^f(z) &:= [D_{i_1} f(z)]_{i_1 \in B_k}; [D_{i_1; i_2}^2 f(z)]_{(i_1; i_2) \in B_k^2} \\ m^h(z; (j; j^0)) &:= [m_k^h(z; (j; j^0))]_{k \in [K]} \\ m_k^h(z; (j; j^0)) &:= [D_{jj^0}^2 \circ h_{i_1}(z)]_{i_1 \in B_k}; [D_j \circ h_{i_2}(z) D_j h_{i_1}(z)]_{(i_1; i_2) \in B_k^2} \end{aligned}$$

Proof. The exact same argument as the one presented in Lemma A.19 (based on Lachapelle et al. (2023)) guarantees that  $f^{\wedge}$  and  $h$  have equal derivatives on  $Z_{\text{supp}}$ . We leverage this fact next.

By taking the derivative w.r.t.  $z_j$  on both sides of  $f^{\wedge}(z) = f(h(z))$ , we get

$$D_j f^{\wedge}(z) = \sum_{k \in [K]} \sum_{i \in B_k} D_i f(h(z)) D_j h_i(z) \quad (\text{A.38})$$

Now take another derivative w.r.t.  $z_{j^0}$  for some  $j^0 \in [d_z]$  to get

$$D_{jj^0}^2 f^{\wedge}(z) = \sum_{k_1 \in [K]} \sum_{i_1 \in B_{k_1}} D_{i_1} f(h(z)) D_{jj^0}^2 \circ h_{i_1}(z) + \sum_{k_2 \in [K]} \sum_{i_2 \in B_{k_2}} D_{i_1; i_2}^2 f(h(z)) D_j \circ h_{i_2}(z) D_j h_{i_1}(z) \quad (3)$$

Because we have at most first order interactions (Def. 3.4 with  $n = 1$ ), the second sum over  $[K]$  drops, and we are left with:

$$\begin{aligned} D_{jj^0}^2 f^{\wedge}(z) &= \sum_{k_1 \in [K]} \sum_{i_1 \in B_{k_1}} D_{i_1} f(h(z)) D_{jj^0}^2 \circ h_{i_1}(z) + \sum_{i_2 \in B_{k_1}} D_{i_1; i_2}^2 f(h(z)) D_j \circ h_{i_2}(z) D_j h_{i_1}(z) \quad (3) \\ &= \sum_{k_1 \in [K]} \sum_{i_1 \in B_{k_1}} D_{i_1} f(h(z)) D_{jj^0}^2 \circ h_{i_1}(z) + \sum_{(i_1; i_2) \in B_{k_1}^2} D_{i_1; i_2}^2 f(h(z)) D_j \circ h_{i_2}(z) D_j h_{i_1}(z) \quad (3) \\ &= W^f(h(z)) m^h(z; (j; j^0)); \end{aligned}$$

which concludes the proof.  $\square$

Theorem A.22. Let  $f : Z \rightarrow X$  be a  $C^2$  diffeomorphism satisfying interaction asymmetry (Asm. 3.5) for all equivalent generators (Defn. 4.1) for  $r = 1$  and sufficient independence (Defn. A.9). Let  $Z_{\text{supp}} \subset Z$  be regular closed (Defn. A.3), path-connected (Defn. A.14) and aligned-connected (Defn. A.16). A model  $\hat{z} : Z \rightarrow \mathbb{R}^{d_x}$  disentangles  $z$  on  $Z_{\text{supp}}$  w.r.t.  $f$  (Defn. 2.1) if it is (i) a  $C^2$  diffeomorphism between  $Z_{\text{supp}}$  and  $X_{\text{supp}}$  with (ii) at most first order interactions across slots (Defn. 3.4) on  $Z_{\text{supp}}$ .

Proof. As mentioned in Section A.3, the proofs will proceed in two steps: First, we show local disentanglement (Definition A.13) and then we show (global) disentanglement via Lemma A.18. We first show local disentanglement.

We first define the function  $h : Z_{\text{supp}} \rightarrow Z_{\text{supp}}$  relating the latent spaces of these functions  $z$  on  $Z_{\text{supp}}$ :

$$h := f^{-1} \circ f^{\wedge} \quad (\text{A.39})$$

The function  $f^{\wedge}$  can then be written in terms of  $f$  and  $h$  on  $Z_{\text{supp}}$ :

$$f^{\wedge} = f \circ h \quad (\text{A.40})$$

Because  $f; f^{\wedge}$  are both  $C^2$  diffeomorphism between  $Z_{\text{supp}}; X_{\text{supp}}$  and  $Z_{\text{supp}}; X_{\text{supp}}$ , respectively, we have that  $h$  is a  $C^2$  diffeomorphism.

Since  $f$  has at most 1st order interactions, we can apply Lemma A.21 to obtain, for all  $\mathcal{Z}_{\text{supp}}; j; j^0 \in \mathcal{D}_z$ ,

$$D_{jj^0}^2 f(z) = W^f(h(z))m^h(z; (j; j^0)) :$$

Since  $f$  has at most 1st order interaction, we have that, for all  $j \in \mathcal{D}_z$

$$0 = W^f(h(z))m^h(z; (j; j^0)) : \quad (\text{A.41})$$

By defining

$$\begin{aligned} W_k^{f;\text{rest}}(z) &:= [D_{i_1} f(z)]_{i_1 \in B_k} \\ W_k^{f;\text{high}}(z) &:= [D_{i_1; i_2}^2 f(z)]_{(i_1; i_2) \in B_k^2} \\ m_k^{h;\text{rest}}(z; (j; j^0)) &:= [D_{jj^0}^2 h_{i_1}(z)]_{i_1 \in B_k} \\ m_k^{h;\text{high}}(z; (j; j^0)) &:= [D_j h_{i_2}(z) D_j h_{i_1}(z)]_{(i_1; i_2) \in B_k^2} \end{aligned}$$

we can restate the sufficiently independent derivative assumption (Def. A.9) as, for all

$$\text{rank } W^f(z) = \sum_{k \in [K]} \text{rank } W_k^{f;\text{rest}}(z) + \text{rank } W_k^{f;\text{high}}(z)$$

This condition allows us to apply Lemma A.6 to go from (A.41) to, for all  $j \in \mathcal{D}_z, k \in [K]$ :

$$0 = W_k^{f;\text{high}}(h(z))m_k^{h;\text{high}}(z; (j; j^0)) \quad (\text{A.42})$$

**Case 1:  $|B_k| = 1$  (One-Dimensional Slots)** By Assumption 3.5.ii,  $(W_k^f = B = \text{fig}) D_{jj^0}^2 f(z) \in \mathbb{R}$ . Note that  $W_k^{f;\text{high}}(h(z)) = D_{k;k}^2 f(z)$ . Hence, (A.42) implies that  $m_k^{h;\text{high}}(z; (j; j^0)) = 0$  (which is a scalar). This means  $m_k^{h;\text{high}}(z; (j; j^0)) = D_j h_{i_2}(z) D_j h_{i_1}(z) = 0$ . Since this is true for all  $k$  and all distinct  $j; j^0$ , this means each row has at most one nonzero entry. Since  $Dh(z)$  is invertible, these nonzero entries must appear on different columns, otherwise a column will be filled with zeros. This means  $Dh(z)$  is a permutation-scaling matrix, i.e. we have local disentanglement (Definition A.13).

**Case 2:  $|B_k| > 1$  (Multi-Dimensional Slots)**

Assume for a contradiction that it is not locally disentangled on  $\mathcal{Z}_{\text{supp}}$  w.r.t.  $f$ . This implies that there exist  $\mathcal{Z} \in \mathcal{Z}_{\text{supp}}; k; k^0 \in [K]$  with  $k^0 \in k^0$  such that:

$$D_{B_{k^0}} h_{B_k}(z) \in \mathbb{R}; \quad D_{B_{k^0}} h_{B_k}(z) \in \mathbb{R} \quad (\text{A.43})$$

Because  $f; f^\wedge$  are  $C^1$  diffeomorphisms, we know that it is also a  $C^1$  diffeomorphism. Coupling this with Eq. (A.43), Lemma A.7 tells us that there exist  $\mathcal{S} \in \mathcal{D}_z$  with cardinality  $|B_k|$  such that:

$$B \in B; \quad \mathcal{S} \in B; \quad \text{and } \mathcal{S}_i \in \mathcal{S}; D_i h_{B_k}(z) \text{ are linearly independent.} \quad (\text{A.44})$$

Now choose any  $\mathcal{B} \in B$  such that  $\mathcal{S}_1 := f \setminus B \in \mathcal{B}$ ; . Furthermore, define the sets  $\mathcal{S}_2 := \mathcal{S} \cap \mathcal{S}_1$ . Because  $\mathcal{S} \in B$ , we know that  $\mathcal{S}_2$  is non-empty. Further, by construction  $\mathcal{S}_1 \cap \mathcal{S}_2 = \emptyset$ . In other words,  $\mathcal{S}_1$  and  $\mathcal{S}_2$  are non-empty, form a partition of  $\mathcal{S}$ , and do not contain any indices from the same slot.

Now construct the matrices, denoted  $A_{\mathcal{S}_1}$  and  $A_{\mathcal{S}_2}$  as follows:

$$A_{\mathcal{S}_1} := D_{\mathcal{S}_1} h_{B_k}(z); \quad A_{\mathcal{S}_2} := D_{\mathcal{S}_2} h_{B_k}(z) \quad (\text{A.45})$$

And the matrix denoted  $A_k$  as:

$$A_k := [A_{\mathcal{S}_1}; A_{\mathcal{S}_2}] \quad (\text{A.46})$$

Note that because  $\mathcal{S}_i \in \mathcal{S}, D_i h_{B_k}(z)$  are linearly independent (Eq. A.44), we know that  $A_k$  is invertible.

Now, define the following block diagonal matrix  $A \in \mathbb{R}^{d_z \times d_z}$  as follows:

$$A := \begin{pmatrix} A_1 & 0 & \dots & 0 \\ 0 & A_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & A_K \end{pmatrix} \quad (\text{A.47})$$

where  $\delta_i \in [K]$  and  $n \in \mathbb{N}$ ;  $A_i$  is the identity matrix, and thus invertible, while  $A_k$  is defined according to Eq. (A.46).

Define  $Z := A^{-1}Z$ , the function  $h : Z \rightarrow Z$  as  $h(z) := Az$  and the function  $f : Z \rightarrow X$  as  $f := f \circ h$ . By construction we have

$$\forall z \in Z; \quad f(A_1^{-1}z_{B_1}; \dots; A_K^{-1}z_{B_K}) = f(z_{B_1}; \dots; z_{B_K}) \quad (\text{A.48})$$

Because all  $A_i^{-1}$  are invertible, then  $h$  is equivalent to  $f$  in the sense of Def. (4.1).

We can now apply Lemma A.21 to  $f \circ h$  to obtain, for all  $(j; j^0) \in [d_z]$ :

$$D_{jj^0}^2 (f \circ h)(z) = W^f(h(z)) m^h(z; (j; j^0)) \quad (\text{A.49})$$

Choose  $z \in Z$  such that  $h(z) = h(z')$ , which is possible because  $h(z) \in Z$  and  $h$  is a bijection from  $Z$  to  $Z$ . We can then write

$$D_{jj^0}^2 (f \circ h)(z) = W^f(h(z')) m^h(z; (j; j^0)) \quad (\text{A.50})$$

Let  $J; J^0 \subseteq B_k$  be a partition of  $B_k$  such that  $J$  is the set of columns of  $A$  corresponding to  $A_{S_1}$  and  $J^0$  be the set of columns of  $A$  corresponding to  $A_{S_2}$ . More formally, we have

$$A_{B_k; J} = A_{S_1} \quad \text{and} \quad A_{B_k; J^0} = A_{S_2}$$

Since  $A_{S_1} = D_{S_1} h_{B_k}(z)$  and  $A_{S_2} = D_{S_2} h_{B_k}(z)$ , we have that

$$A_{B_k; J} = D_{S_1} h_{B_k}(z) \quad \text{and} \quad A_{B_k; J^0} = D_{S_2} h_{B_k}(z)$$

Since  $Dh(z) = A$ , we have

$$D_J h_{B_k}(z) = D_{S_1} h_{B_k}(z) \quad \text{and} \quad D_{J^0} h_{B_k}(z) = D_{S_2} h_{B_k}(z)$$

For all  $(j; j^0) \in J \times J^0$ , there must exist  $(s; s^0) \in S_1 \times S_2$  such that

$$D_j h_{B_k}(z) = D_s h_{B_k}(z) \quad \text{and} \quad D_{j^0} h_{B_k}(z) = D_{s^0} h_{B_k}(z)$$

This implies that, for all  $(j; j^0) \in J \times J^0$ , there exist  $(s; s^0) \in S_1 \times S_2$  such that

$$m_k^{h; \text{high}}(z; (j; j^0)) = m_k^{h; \text{high}}(z; (s; s^0)) \quad (\text{A.51})$$

Moreover, since  $h$  is a block-wise function we have that, for all  $(j; j^0) \in J \times J^0 \subseteq B_k$  and  $k \in [K]$  and  $n \in \mathbb{N}$ ,  $m_k^{h; \text{rest}}(z; (j; j^0)) = 0$ . We can thus write:

$$D_{jj^0}^2 (f \circ h)(z) = W_k^f(h(z)) m_k^h(z; (j; j^0)) \quad (\text{A.52})$$

Since  $h$  is linear, we have that  $m_k^{h; \text{rest}}(z; (j; j^0)) = 0$ , and thus

$$D_{jj^0}^2 (f \circ h)(z) = W_k^{f; \text{high}}(h(z)) m_k^{h; \text{high}}(z; (j; j^0)) \quad (\text{A.53})$$

Plug the (A.51) into the above to obtain that, for all  $(j; j^0) \in J \times J^0$ ,

$$D_{jj^0}^2 (f \circ h)(z) = W_k^{f; \text{high}}(h(z)) m_k^{h; \text{high}}(z; (s; s^0)) = 0 \quad (\text{A.54})$$

where the very last “= 0” is due to (A.42) (recall  $(s; s^0) \in S_1 \times S_2 \subseteq D^c$ ).

In other words, we found a partition  $J \times J^0$  of the block  $B_k$  and a value  $z$  such that  $D_{jj^0}^2 (f \circ h)(z) = 0$  for all  $(j; j^0) \in J \times J^0$ . This means that the blocks  $J$  and  $J^0$  have no second order interaction at  $z$ . This is a contradiction with Assm. 3.5. Hence, we have local disentanglement.

From local to global disentanglement. The same argument as in the proof of Theorem A.20 applies.

□

A.6 DISENTANGLEMENT (AT MOST  $2^{\text{ND}}$  ORDER INTERACTION)

Lemma A.23. Let  $Z_{\text{supp}} \subseteq Z$  be a regular closed set (Defn. A.3). Let  $f: Z \rightarrow \mathbb{R}^K$  be  $C^1$  and  $h: Z_{\text{supp}} \rightarrow Z_{\text{supp}}$  be a diffeomorphism. Let  $\hat{f} := f \circ h$ . If  $f$  has at most  $2^{\text{nd}}$  order interaction (Definition 3.4 with  $n = 2$ ), then, for all  $j \in [d_z]$  and  $z \in Z_{\text{supp}}$ , we have

$$D_{jj}^3 \hat{f}(z) = W^f(h(z)) m^h(z; (j; j^0; j^0)); \quad (\text{A.55})$$

where

$$\begin{aligned} W^f(z) &:= [W_k^f(z)]_{k \in [K]} \\ W_k^f(z) &:= [D_{i_1} f(z)]_{i_1 \in B_k}; \\ &\quad [D_{i_1, i_2}^2 f(z)]_{i_1, i_2 \in B_k}; \\ &\quad [D_{i_1, i_2, i_3}^3 f(z)]_{(i_1, i_2, i_3) \in B_k^3} \\ m^h(z; (j; j^0; j^0)) &:= [m_k^h(z; (j; j^0; j^0))]_{k \in [K]} \\ m_k^h(z; (j; j^0; j^0)) &:= [D_{jj}^3 \hat{f}(h(z))]_{i_1 \in B_k}; \\ &\quad [D_j h_{i_1}(z) D_{j^0}^2 \hat{f}(h(z)) + D_j h_{i_2}(z) D_{jj}^2 \hat{f}(h(z)) + D_j h_{i_3}(z) D_{jj}^2 \hat{f}(h(z))]_{i_1, i_2 \in B_k}; \\ &\quad [D_j h_{i_3}(z) D_j h_{i_2}(z) D_j h_{i_1}(z)]_{(i_1, i_2, i_3) \in B_k^3} \end{aligned}$$

Proof. As argued in Lemma A.21, differentiating  $\hat{f}(z) = f \circ h(z)$  w.r.t.  $z_j$  and  $z_{j^0}$  on both sides yields

$$D_{jj}^2 \hat{f}(z) = \sum_{k_1 \in [K]} \sum_{i_1 \in B_{k_1}} 4 D_{i_1} f(h(z)) D_{jj}^2 \hat{f}(h(z)) + \sum_{k_2 \in [K]} \sum_{i_2 \in B_{k_2}} D_{i_1, i_2}^2 f(h(z)) D_j h_{i_2}(z) D_j h_{i_1}(z)$$

Now take another derivative with respect to  $z_{j^0}$  to compute  $D_{jj}^3 \hat{f}(z)$ . For the first term in the sum, we have:

$$\sum_{k_1 \in [K]} \sum_{i_1 \in B_{k_1}} 4 \sum_{k_2 \in [K]} \sum_{i_2 \in B_{k_2}} D_{i_1, i_2}^2 f(h(z)) D_j h_{i_2}(z) D_{jj}^2 \hat{f}(h(z)) + D_{i_1} f(h(z)) D_{jj}^3 \hat{f}(h(z))$$

And for the second term in the sum (the nested sum), we have:

$$\sum_{k_1 \in [K]} \sum_{i_1 \in B_{k_1}} \sum_{k_2 \in [K]} \sum_{i_2 \in B_{k_2}} \sum_{k_3 \in [K]} \sum_{i_3 \in B_{k_3}} D_{i_1, i_2, i_3}^3 f(h(z)) D_j h_{i_3}(z) D_j h_{i_2}(z) D_j h_{i_1}(z) + D_{i_1, i_2}^2 f(h(z)) D_j h_{i_2}(z) D_{jj}^2 \hat{f}(h(z))$$

Because we have at most second order interactions (Def. 3.4 with  $n = 2$ ), this term can be rewritten as:

$$\sum_{k_1 \in [K]} \sum_{i_1 \in B_{k_1}} \sum_{k_2 \in [K]} \sum_{i_2 \in B_{k_2}} \sum_{k_3 \in [K]} \sum_{i_3 \in B_{k_3}} D_{i_1, i_2, i_3}^3 f(h(z)) D_j h_{i_3}(z) D_j h_{i_2}(z) D_j h_{i_1}(z) + \sum_{k_2 \in [K]} \sum_{i_2 \in B_{k_2}} D_{i_1, i_2}^2 f(h(z)) D_j h_{i_2}(z) D_{jj}^2 \hat{f}(h(z))$$

Combining the first and second terms, we get:

$$\begin{aligned}
D_{jj}^3 \circ_{j^0} \circ_{j^1} \circ_{j^2} \hat{f}(z) &= \prod_{k_1=2}^{K-1} \prod_{i_1=2}^{B_{k_1}} D_{i_1} f(h(z)) D_{jj}^3 \circ_{j^0} \circ_{j^1} \circ_{j^2} h_{i_1}(z) + \\
&\quad \prod_{k_2=2}^{K-1} \prod_{i_2=2}^{B_{k_2}} D_{i_1; i_2}^2 f(h(z)) D_j h_{i_1}(z) D_{j^0}^2 \circ_{j^1} \circ_{j^2} h_{i_2}(z) + D_j \circ h_{i_2}(z) D_{jj}^2 \circ_{j^0} \circ_{j^1} h_{i_1}(z) + D_j \circ h_{i_2}(z) D_{jj}^2 \circ_{j^0} h_{i_1}(z) + \\
&\quad \prod_{i_2=2}^{B_{k_1}} \prod_{i_3=2}^{B_{k_1}} D_{i_1; i_2; i_3}^3 f(h(z)) D_j \circ h_{i_3}(z) D_j \circ h_{i_2}(z) D_j h_{i_1}(z) \\
&= \prod_{k_1=2}^{K-1} \prod_{i_1=2}^{B_{k_1}} D_{i_1} f(h(z)) D_{jj}^3 \circ_{j^0} \circ_{j^1} \circ_{j^2} h_{i_1}(z) + \\
&\quad \prod_{i_1=2}^{B_{k_1}} \prod_{i_2=2}^{[d_2]} D_{i_1; i_2}^2 f(h(z)) D_j h_{i_1}(z) D_{j^0}^2 \circ_{j^1} \circ_{j^2} h_{i_2}(z) + D_j \circ h_{i_2}(z) D_{jj}^2 \circ_{j^0} \circ_{j^1} h_{i_1}(z) + D_j \circ h_{i_2}(z) D_{jj}^2 \circ_{j^0} h_{i_1}(z) + \\
&\quad \prod_{(i_1; i_2; i_3) \in \mathcal{B}_{k_1}^3} D_{i_1; i_2; i_3}^3 f(h(z)) D_j \circ h_{i_3}(z) D_j \circ h_{i_2}(z) D_j h_{i_1}(z) \\
&= W^f(h(z)) m^h(z; (j; j^0; j^1; j^2)) :
\end{aligned}$$

□

**Theorem A.24.** Let  $f : Z \rightarrow X$  be a  $C^3$  diffeomorphism satisfying interaction asymmetry (Asm. 3.5) for all equivalent generators (Defn. 4.1) for  $\ell=2$  and sufficient independence (Defn. 4.2). Let  $Z_{\text{supp}} \subset Z$  be regular closed (Defn. A.3), path-connected (Defn. A.14) and aligned-connected (Defn. A.16). A model  $\mathcal{Z} \rightarrow \mathbb{R}^{d_x}$  disentangles on  $Z_{\text{supp}}$  w.r.t.  $f$  (Defn. 2.1) if it is (i) a  $C^3$  diffeomorphism between  $\mathcal{Z}_{\text{supp}}$  and  $X_{\text{supp}}$  with (ii) at most  $2^d$  order interactions across slots (Defn. 3.4) on  $\mathcal{Z}_{\text{supp}}$ .

**Proof.** As mentioned in Section A.3, the proofs will proceed in two steps: First, we show local disentanglement (Definition A.13) and then we show (global) disentanglement via Lemma A.18. We first show local disentanglement.

We first define the function  $h : \mathcal{Z}_{\text{supp}} \rightarrow Z_{\text{supp}}$  relating the latent spaces of these functions  $\hat{f}$  on  $\mathcal{Z}_{\text{supp}}$ :

$$h := f^{-1} \circ \hat{f} \quad (\text{A.56})$$

The function  $\hat{f}$  can then be written in terms of  $h$  and  $h$  on  $\mathcal{Z}_{\text{supp}}$ :

$$\hat{f} = f \circ h \quad (\text{A.57})$$

Because  $f; \hat{f}$  are both  $C^2$  diffeomorphism between  $Z_{\text{supp}}; X_{\text{supp}}$  and  $\mathcal{Z}_{\text{supp}}; X_{\text{supp}}$  respectively, we have that  $h$  is a  $C^2$  diffeomorphism.

Since  $\hat{f}$  has at most  $2^d$  order interaction, we have that, for all  $\mathcal{Z}_{\text{supp}}; (j; j^0; j^1; j^2) \in \mathcal{D}^c[d_2]$ ,

$$0 = W^f(h(z)) m^h(z; (j; j^0; j^1; j^2)) : \quad (\text{A.58})$$

By definition

$$W_k^{f; \text{rest}}(z) := [D_{i_1} f(z)]_{i_1 \in B_k};$$

$$[D_{i_1, i_2}^2 f(z)]_{i_1 \in B_k, i_2 \in [d_z]}$$

$$W_k^{f; \text{high}}(z) := [D_{i_1, i_2, i_3}^3 f(z)]_{(i_1, i_2, i_3) \in B_k^3};$$

$$m_k^{h; \text{rest}}(z; (j; j^0; j^{00})) := [D_{jj^0j^{00}}^3 h_{i_1}(z)]_{i_1 \in B_k};$$

$$[D_j h_{i_1}(z) D_{j^0 j^{00}}^2 h_{i_2}(z) + D_j h_{i_2}(z) D_{jj^0j^{00}}^2 h_{i_1}(z) + D_j h_{i_2}(z) D_{jj^0j^{00}}^2 h_{i_1}(z)]_{i_1 \in B_k, i_2 \in [d_z]}$$

$$m_k^{h; \text{high}}(z; (j; j^0; j^{00})) := [D_j h_{i_3}(z) D_j h_{i_2}(z) D_j h_{i_1}(z)]_{(i_1, i_2, i_3) \in B_k^3};$$

we can restate the sufficiently independent derivative assumption (Def. 4.2) as, for all

$$\text{rank } W^f(z) = \sum_{k \in [K]} \text{rank } W_k^{f; \text{rest}}(z) + \text{rank } W_k^{f; \text{high}}(z)$$

This condition allows us to apply Lemma A.6 to go from (A.58) to, for all  $(j; j^0; j^{00}) \in D^c [d_z]$ ,  $k \in [K]$ :

$$0 = W_k^{f; \text{high}}(h(z)) m_k^{h; \text{high}}(z; (j; j^0; j^{00})) \quad (\text{A.59})$$

Case 1:  $|B_k| = 1$  (One-Dimensional Slots) By Assumption 3.5.ii (with  $\mathbf{w} = \mathbf{B} = \text{fig}$ ), we have that  $D_{i; i}^3 f(z) \notin 0$ . Note that  $W_k^{f; \text{high}}(h(z)) = D_{k; k; k}^3 f(z)$ . Hence, (A.42) implies that  $m_k^{h; \text{high}}(z; (j; j^0; j^{00})) = 0$  (which is a scalar). This means  $m_k^{h; \text{high}}(z; (j; j^0; j^{00})) = D_j h_k(z) D_j h_k(z) D_j h_k(z) = 0$  for all  $(j; j^0; j^{00}) \in D^c [d_z]$ . In particular, we have

$$D_j h_k(z)^2 D_j h_k(z) = 0;$$

for all  $(j; j^0) \in D^c$ . Since this is true for all  $k$  and all distinct  $j; j^0$ , this means each row of  $h(z)$  has at most one nonzero entry. Since  $Dh(z)$  is invertible, these nonzero entries must appear on different columns, otherwise a column would be filled with zeros. This means  $h(z)$  is a permutation-scaling matrix, i.e. we have local disentanglement (Definition A.13).

Case 2:  $|B_k| > 1$  (Multi-Dimensional Slots)

Assume for a contradiction that  $\hat{\mathbf{A}}$  does not disentangle on  $Z_{\text{supp}}$  w.r.t.  $f$ . This implies that there exist  $z \in Z_{\text{supp}}; k; k^0; k^{00} \in [K]$  with  $k^0 \notin k^{00}$  such that:

$$D_{B_{k^0}} h_{B_k}(z) \notin 0; \quad D_{B_{k^{00}}} h_{B_k}(z) \notin 0 \quad (\text{A.60})$$

Because  $f; \hat{\mathbf{A}}$  are  $\mathbb{C}^3$  diffeomorphisms, we know that  $\hat{\mathbf{A}}$  is also a  $\mathbb{C}^3$  diffeomorphism. Coupling this with Eq. (A.60), Lemma A.7 tells us that there exist  $S \subseteq [d_z]$  with cardinality  $|B_k|$  such that:

$$|S| \in B; \quad |S^c| \in B; \quad \text{and } \{D_i h_{B_k}(z)\}_{i \in S} \text{ are linearly independent.} \quad (\text{A.61})$$

Now choose any  $B \subseteq B$  such that  $S_1 := f(S) \cap B \notin \emptyset$ . Furthermore, define the set  $S_2 := S \cap S_1$ . Because  $|S| \in B$ , we know that  $S_2$  is non-empty. Further, by construction  $S = S_1 \cup S_2$ . In other words,  $S_1$  and  $S_2$  are non-empty, form a partition of  $S$ , and do not contain any indices from the same slot.

Now construct the matrices, denoted  $A_{S_1}$  and  $A_{S_2}$  as follows:

$$A_{S_1} := D_{S_1} h_{B_k}(z); \quad A_{S_2} := D_{S_2} h_{B_k}(z) \quad (\text{A.62})$$

And the matrix denoted  $A_k$  as:

$$A_k := [A_{S_1}; A_{S_2}] \quad (\text{A.63})$$

Note that because  $|S| \in B$ ,  $\{D_i h_{B_k}(z)\}_{i \in S}$  are linearly independent (Eq. (A.61)), we know that  $A_k$  is invertible.



Now, define the following block diagonal matrix  $A \in \mathbb{R}^{d_z \times d_z}$  as follows:

$$A := \begin{pmatrix} A_1 & 0 & \cdots & 0 \\ 0 & A_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & A_K \end{pmatrix} \quad (\text{A.64})$$

where  $A_i \in \mathbb{R}^{n_i \times n_i}$ ;  $A_i$  is the identity matrix, and thus invertible, while  $A_k$  is defined according to Eq. (A.63).

Define  $Z := A^{-1}Z$ , the function  $h : Z \rightarrow Z$  as  $h(z) := Az$  and the function  $f : Z \rightarrow X$  as  $f := f \circ h$ . By construction we have

$$D_{j;j^0;j^{00}}^3 f(z) = A_1^{-1} z_{B_1}; \dots; A_K^{-1} z_{B_K} = f(z_{B_1}; \dots; z_{B_K}) \quad (\text{A.65})$$

Because all  $A_i^{-1}$  are invertible,  $f$  is equivalent to  $f$  in the sense of Def. (4.1).

We can now apply Lemma A.23 to  $f \circ h$  to obtain, for all  $(j; j^0; j^{00}) \in [d_z]$ :

$$D_{j;j^0;j^{00}}^3 f \circ h(z) = W^f(h(z)) m^h(z; (j; j^0; j^{00})) \quad (\text{A.66})$$

Choose  $Z$  such that  $h(z) = h(z')$ , which is possible because  $h$  is a bijection from  $Z$  to  $Z$ . We can then write

$$D_{j;j^0;j^{00}}^3 f \circ h(z) = W^f(h(z')) m^h(z; (j; j^0; j^{00})) \quad (\text{A.67})$$

Let  $J; J^0 \subseteq B_k$  be a partition of  $B_k$  such that  $J$  is the set of columns of  $A$  corresponding to  $S_1$ , and  $J^0$  be the set of columns of  $A$  corresponding to  $S_2$ . More formally, we have

$$A_{B_k; J} = A_{S_1} \quad \text{and} \quad A_{B_k; J^0} = A_{S_2}$$

Since  $A_{S_1} = D_{S_1} h_{B_k}(z)$  and  $A_{S_2} = D_{S_2} h_{B_k}(z)$ , we have that

$$A_{B_k; J} = D_{S_1} h_{B_k}(z) \quad \text{and} \quad A_{B_k; J^0} = D_{S_2} h_{B_k}(z)$$

Since  $Dh(z) = A$ , we have

$$D_J h_{B_k}(z) = D_{S_1} h_{B_k}(z) \quad \text{and} \quad D_{J^0} h_{B_k}(z) = D_{S_2} h_{B_k}(z)$$

For all  $(j; j^0; j^{00}) \in J \cup J^0 \subseteq B_k$  there must exist  $(s; s^0; s^{00}) \in S_1 \cup S_2 \subseteq S$  such that

$$D_j h_{B_k}(z) = D_s h_{B_k}(z); \quad D_{j^0} h_{B_k}(z) = D_{s^0} h_{B_k}(z); \quad \text{and} \quad D_{j^{00}} h_{B_k}(z) = D_{s^{00}} h_{B_k}(z)$$

This implies that for all  $(j; j^0; j^{00}) \in J \cup J^0 \subseteq B_k$  there must exist  $(s; s^0; s^{00}) \in S_1 \cup S_2 \subseteq S$  such that

$$m_k^{h; \text{high}}(z; (j; j^0; j^{00})) = m_k^{h; \text{high}}(z; (s; s^0; s^{00})) \quad (\text{A.68})$$

Moreover, since  $h$  is a block-wise function, we have that, for all  $(j; j^0; j^{00}) \in J \cup J^0 \subseteq B_k \subseteq B_k^3$  and all  $k^0 \in [K]$ , we have  $m_k^{h; \text{high}}(z; (j; j^0; j^{00})) = 0$ , which allows us to rewrite (A.67) as

$$D_{j;j^0;j^{00}}^3 f \circ h(z) = W_k^f(h(z)) m_k^h(z; (j; j^0; j^{00})) \quad (\text{A.69})$$

Since  $h$  is linear, we have that  $m_k^{h; \text{rest}}(z; (j; j^0; j^{00})) = 0$ , and thus

$$D_{j;j^0;j^{00}}^3 f \circ h(z) = W_k^{f; \text{high}}(h(z)) m_k^{h; \text{high}}(z; (j; j^0; j^{00})) \quad (\text{A.70})$$

Plug (A.68) into the above to obtain that for all  $(j; j^0; j^{00}) \in J \cup J^0 \subseteq B_k$ , there exist  $(s; s^0; s^{00}) \in S_1 \cup S_2 \subseteq S$  such that

$$D_{j;j^0;j^{00}}^3 f \circ h(z) = W_k^{f; \text{high}}(h(z)) m_k^{h; \text{high}}(z; (s; s^0; s^{00})) = 0 \quad (\text{A.71})$$

where the very last “= 0” is due to (A.59) (recall  $(s; s^0; s^{00}) \in S_1 \cup S_2 \subseteq S \subseteq D^c[d_z]$ ).

In other words, we found a partition  $J^0$  of the block  $B_k$  and a value  $\alpha$  such that  $D_{j;j^0;j^0}^3 \text{ of } (z) = 0$  for all  $(j; j^0; j^0) \in J \setminus J^0 \times B_k$ . One can show that as no 3rd order interaction across blocks because it is equivalent to  $f$ , which also has no 3rd order interactions across blocks. We thus have that  $D_{j;j^0;j^0}^3 \text{ of } (z) = 0$  for all  $(j; j^0; j^0) \in J \setminus J^0 \times [d_z]$ . This means that the blocks and  $J^0$  have no third order interaction in  $f$  at  $z$ . This is a contradiction with Assm. 3.5.

From local to global disentanglement. The same argument as in the proof of Theorem A.20 applies. □

## B MULTI-INDEX NOTATION

Multi-index notation is a convenient shorthand to denote higher order derivatives. A multi-index of dimension  $d$  is an ordered tuple  $\alpha = (\alpha_1; \dots; \alpha_d) \in \mathbb{N}^d$ . We introduce the shorthands

$$j^\alpha = \prod_{i=1}^d j_i^{\alpha_i}; \quad \alpha! = \prod_{i=1}^d \alpha_i! \tag{B.1}$$

and we write  $\alpha \leq \beta$  if  $\alpha_i \leq \beta_i$  for all  $i$  and  $\alpha \pm \beta$  denotes the element wise sum (difference) of the entries. We write

$$D^\alpha = \frac{\partial^{\alpha_1}}{\partial z_1^{\alpha_1}} \dots \frac{\partial^{\alpha_d}}{\partial z_d^{\alpha_d}} \tag{B.2}$$

and

$$z^\alpha = \prod_{i=1}^d z_i^{\alpha_i} \tag{B.3}$$

We will need the important property that

$$D^\alpha z^\beta = \begin{cases} \frac{\alpha!}{(\alpha-\beta)!} z^{\beta-\alpha} & \text{if } \alpha \leq \beta \\ 0 & \text{otherwise.} \end{cases} \tag{B.4}$$

Consider now a partition  $[d_z]$  into slots  $B_1; \dots; B_k$ . We define the set of interaction multi-indices of order  $n$  for  $n \geq 2$  by

$$I_n = \{ \alpha \in \mathbb{N}^{d_z} : \sum_j \alpha_j = n, \exists i_1; i_2 \text{ s.t. } i_1 \in B_{k_1}, i_2 \in B_{k_2} \text{ with } k_1 \neq k_2 \text{ and } \alpha_{i_1}; \alpha_{i_2} > 0 \} \tag{B.5}$$

i.e., the set of all multi-indices such that the non-zero components are contained in at least two blocks. Clearly  $I_n$  depends on the block partition which we do not reflect in the notation. We also consider

$$I_n = \bigcup_{m=2}^n I_m \tag{B.6}$$

Clearly, if  $\alpha \in I_j$  and  $\beta$  is any multi-index, then  $\alpha + \beta \in I_{j+\beta_1}$ .

## C CHARACTERIZATION OF FUNCTIONS WITH AT MOST $n^{\text{TH}}$ ORDER INTERACTIONS

In this section we characterize functions with interaction of at most order  $n$  by proving Theorem C.2. Our characterization relies on the notion of aligned-connectedness introduced in Definition A.16 and the following topological notion.

Definition C.1. A topological space  $X$  is contractible if there is a continuous function  $F: X \times [0; 1] \rightarrow X$  such that  $F(x; 0) = x$  and  $F(x; 1) = x_0$  for a point  $x_0 \in X$ . We call a subset  $C \subset \mathbb{R}^d$  contractible if it is contractible as a topological space with respect to the induced subspace topology.

Roughly, contractibility means that we can transform a topological space continuously into a point, which is possible if the space has no holes. Note that, e.g., all one dimensional connected sets and all convex sets are contractible. Sets that are not contractible are, e.g., spheres and disconnected sets. Note that the characterization in the following theorem generalizes Proposition 7 in Lachapelle et al. (2023) by allowing higher order interactions and showing the result for more general domains. We denote, similar to (2.3), for  $\mathbb{R}^{d_x}$  by  $\pi_i = f_{z_{B_i}} : z \in Z$  the projections of  $\pi$  on the blocks.

**Theorem C.2 (Characterization of functions with at most  $n$ th order interactions across slots)** Let  $\pi$  be an open connected and aligned-connected set such that  $\pi$  is contractible. Let  $f(z) = f(z_{B_1}; z_{B_2}; \dots; z_{B_K})$  be a  $C^{n+1}$  function on  $\pi$  for an integer  $n \geq 1$ . Then any distinct slots  $z_{B_i}$  and  $z_{B_j}$  have at most  $n$ th order interaction with  $\pi$  (Defn. 3.4) if and only if, for some constants  $c \in \mathbb{R}^{d_x \times 2^{1-n}}$  and some  $C^{n+1}$  functions  $f^k : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_x}$  such that for all  $z \in Z$

$$f(z) = \sum_{k=1}^K f^k(z_{B_k}) + \sum_{i \in 2^{1-n}} c_i z_i \quad (\text{C.1})$$

**Remark C.3.** To avoid unnecessary complications we focus on the case where the ground is defined on  $Z = \mathbb{R}^{d_z}$ . Then  $\pi = Z$  clearly satisfies the assumptions and actually the proof is slightly simpler. The more general result here would allow us to handle  $Z \subset \mathbb{R}^{d_z}$  in Appendix D with minor changes.

The proof can be essentially decomposed in two steps: We show how to reduce from interaction of at most order  $n$  to interaction of at most order  $n-1$  and then we establish the induction base for  $n=2$ .

**Lemma C.4.** Suppose  $f : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_x}$  is a  $C^{n+1}$  function and  $\pi$  open and connected. Assume that  $f$  has interaction of at most order  $n$  between any two different slots for some  $n \geq 2$ . Let  $z_0 \in \pi$  be any point. Then the function

$$f(z) - \sum_{i \in 2^{1-n}} \frac{D_i f(z_0)}{|I|} z_i \quad (\text{C.2})$$

has interaction of order at most  $n-1$ .

**Proof.** First we observe that having interaction at most  $n$  implies that  $D_i f$  is constant in  $\pi$  for  $i \in 2^{1-n}$ . Indeed, since  $i \in 2^{1-n}$  we conclude  $i + e_j \in 2^{1-n+1}$  where  $e_j$  denotes the tuple with  $j$ th entry 1 and all other entries 0. Then, by definition of having interaction at most  $n$  in Definition 3.4, we conclude that

$$D_i D_j f(z) = D^{i+e_j} f(z) = 0 \quad (\text{C.3})$$

This implies that the total derivative  $D_i f$  vanishes on  $\pi$ , which implies that  $D_i f$  is constant because  $\pi$  is connected. Consider now any  $i \in 2^{1-n}$ . Then we find using (B.4)

$$D_i f(z) - \sum_{i \in 2^{1-n}} \frac{D_i f(z_0)}{|I|} z_i = D_i f(z) - \frac{D_i f(z_0)}{|I|} z_i = 0 \quad (\text{C.4})$$

where we used that  $D_i f$  is constant and  $D_i z = 0$  for  $i \in \mathbb{R}^n$  if  $j \neq i$ . This ends the proof.  $\square$

We now establish the functional form for interaction of at most order 1. This is essentially a similar statement as in Proposition 7 in Lachapelle et al. (2023) except that we consider more general domains so that their proof does not apply.

**Lemma C.5.** Assume  $\pi$  is an open connected and aligned-connected set such that  $\pi$  is contractible. If  $f$  is a function such that different slots have interaction at most of order 1 then there are functions  $f^k$  such that

$$f(z) = \sum_{k=1}^K f^k(z_{B_k}) \quad (\text{C.5})$$

Proof. Fix a  $1 \leq k \leq K$ . With slight abuse of notation we write  $z = (z_{B_k}; z_{B_k^c})$ . Fix now some value  $z_{B_k}$ . We claim that for all  $z_{B_k^c}; z_{B_k^c}^0$  such that  $(z_{B_k}; z_{B_k^c}); (z_{B_k}; z_{B_k^c}^0) \in \mathcal{Z}$

$$D_{z_{B_k}} f((z_{B_k}; z_{B_k^c})) = D_{z_{B_k}} f((z_{B_k}; z_{B_k^c}^0)) \tag{C.6}$$

By assumption we indeed know that

$$D_{z_{B_k^c}} D_{z_{B_k}} f((z_{B_k}; z_{B_k^c})) = 0 \tag{C.7}$$

Moreover, by aligned-connectedness we know that the set  $\mathcal{Z}_{z_{B_k}} = \{z_{B_k^c} : (z_{B_k}; z_{B_k^c}) \in \mathcal{Z}\}$  is connected so we conclude that the function

$$g^k(z_{B_k^c}) = D_{z_{B_k}} f((z_{B_k}; z_{B_k^c})) \tag{C.8}$$

is indeed constant. This implies that there is a function  $g^k$  depending on  $z_{B_k}$  such that

$$g^k(z_{B_k}) = D_{z_{B_k}} f((z_{B_k}; z_{B_k^c})) \tag{C.9}$$

for all  $z = (z_{B_k}; z_{B_k^c}) \in \mathcal{Z}$ . Locally  $g^k$  is the gradient of a function, but by assumption  $\mathcal{Z}$  is contractible and therefore, by the Poincaré lemma, there is a function  $f^k$  such that  $Df^k = g^k$ . Then we find

$$D_{z_{B_k}} f((z_{B_k}; z_{B_k^c})) = g^k(z_{B_k}) = D_{z_{B_k}} f^k(z_{B_k}) = D_{z_{B_k}} \sum_{k^0=1}^K f^{k^0}(z_{B_{k^0}}) \tag{C.10}$$

Thus the difference  $\sum_{k=1}^K f^k$  has vanishing derivative on  $\mathcal{Z}$  and since  $\mathcal{Z}$  is connected we conclude that it is constant. This implies (C.5) after shifting  $f$  by this constant.  $\square$

Based on these two lemmas the proof of Theorem C.2 is straightforward.

Proof of Theorem C.2 In the first step we show that if the at most  $s$ th order interaction condition holds then  $f$  can be written as in (C.1), i.e.  $f \in \mathcal{F}_s$ . Applying inductively Lemma C.4 we conclude that there are constants  $c_m \in \mathbb{R}^{d^m}$  such that

$$f(z) = \sum_{m=2}^s c_m \sum_{I_m} \prod_{i \in I_m} z_i \tag{C.11}$$

has interaction of order at most  $s$ . Thus, we can apply Lemma C.5 which implies that a representation as in (C.1) exists on  $\mathcal{Z}$ . For the reverse direction we observe that clearly the functional form implies for  $\mathcal{Z} \subseteq \mathbb{R}^{I_{n+1}}$  the relation

$$Df = 0 \tag{C.12}$$

Let us show through examples that the topological conditions on the set are necessary. The following examples shows that the condition that  $\mathcal{Z}$  is contractible cannot be dropped.

Example C.6. For every  $z \in \mathbb{R}^2 \setminus \{0\}$  we denote by  $(z) \in [0; 2\pi)$  the argument (i.e., the angle to the positive x-axis in radian) and by  $|z|$  the radius of  $z$ . We consider  $\mathcal{Z} \subseteq \mathbb{R}^4$  and  $B_1 = \{z : |z| \leq 1\}$ ,  $B_2 = \{z : |z| \leq 2\}$  given by

$$\mathcal{Z} = \{z : (z_{B_1}); (z_{B_2}) \in (0; 2\pi) \times (0; 2\pi)\} \tag{C.13}$$

and the function

$$f : \mathcal{Z} \rightarrow \mathbb{R}; f(z) = (z_{B_1}) - (z_{B_2}) \tag{C.14}$$

Then  $\mathcal{Z}$  is aligned-connected because the sets in questions are annular sectors and in particular path connected. Moreover,  $f$  is smooth because  $(z_{B_1}) - (z_{B_2}) \in (0; 2\pi)$  so it does not jump and  $D_{z_{B_1}} D_{z_{B_2}} f = 0$  because it is locally additive. However it is not globally additive as in (C.1).

The necessity of the aligned connectedness condition can be shown by an example that is similar to Example 7 in Lachapelle et al. (2023).

Example C.7. Consider  $\mathcal{Z} = ([-1; 0] \cup [2; 2]) \times ([0; 1] \cup [1; 2]) \times ([0; 1] \cup [2; 1])$  and  $f : \mathcal{Z} \rightarrow \mathbb{R}$  given by

$$f(z) = \begin{cases} z_1^3 & \text{if } z_1; z_2 > 0 \\ 0 & \text{otherwise} \end{cases} \tag{C.15}$$

Then  $\mathcal{Z}$  is  $C^2$ ,  $f$  has interaction of order at most 1 but cannot be written as in (C.1). Note that  $\mathcal{Z}$  is not aligned-connected because  $\mathcal{Z}_{z_2} = \{z_1 : (z_1; z_2) \in \mathcal{Z}\} = [-2; 1] \cup [1; 2]$  is not connected.

## D COMPOSITIONAL GENERALIZATION PROOFS

In this appendix we prove extrapolation result Theorem 4.4. Based on the functional form derived in Theorem C.2 we relate two different disentangled representations.

Lemma D.1. Let  $f : Z \rightarrow \mathbb{R}^{d_x}$  be a  $C^3$  diffeomorphism of the form:

$$f(z) = \sum_{k=1}^K f^k(z_{B_k}) + \sum_{2 \leq j \leq K} c_j z_j \quad (\text{D.1})$$

for some  $f^i$  in  $C^3$ . Let  $f^\wedge : Z \rightarrow \mathbb{R}^{d_x}$  be a diffeomorphism of the same functional form. Let  $h : Z_{\text{supp}} \rightarrow Z$  be such that  $f = f^\wedge \circ h$  on  $Z_{\text{supp}}$ . If  $h$  is a slot-wise function, i.e. for all  $2 \leq k \leq K$ ;  $h_k(z) = h_k(z_{B_k})$  and  $Z_{\text{supp}}$  is regularly closed then for all  $z \in Z_{\text{supp}}$

$$\sum_{k=1}^K f^k(z_{B_k}) = \sum_{k=1}^K f^{\wedge(k)}(h_k(z_{B_k})) + L(z) \quad (\text{D.2})$$

for some affine function  $L : \mathbb{R}^{d_z} \rightarrow \mathbb{R}^{d_x}$ .

Remark D.2. We note that it is not possible to remove the affine function from the statement. Indeed if all slots have dimension 1 and  $h_1(z_1) = z_1 + 1$ ,  $h_2(z_2) = z_2 + 1$  then  $h_1(z_1)h_2(z_2) = z_1z_2 = z_1 + z_2 + 1$  is an additive function. Moreover, we cannot in general prove that itself is slot-wise affine because the coefficients can be zero. In this case can be any slot-wise diffeomorphism.

Proof. First we remark that the polynomial part of the functional form in (D.1) contains all terms  $z_i z_j$  where  $i, j$  are in different slots, thus it can be equivalently written as

$$\sum_{2 \leq j \leq K} c_j z_j = \sum_{k=1}^K \sum_{k^0=k+1}^K z_{B_k} z_{B_{k^0}} A_{kk^0} \quad (\text{D.3})$$

for some constant matrices  $A_{kk^0} \in \mathbb{R}^{(|B_k| \times |B_{k^0}|)}$   $\mathbb{N}_{|K| \times |K|}$ , where  $\delta_{k^0 < k}$  denotes the Kronecker product (e.g.  $[z_1; z_2] [z_3; z_4] = [z_1z_3; z_1z_4; z_2z_3; z_2z_4]$ ).

We assume that the permutations is the identity. We know that  $f, f^\wedge$  are diffeomorphisms between the same spaces and can thus be related by the function

$$f = f^\wedge \circ h \quad (\text{D.4})$$

Inserting the functional forms for  $f, f^\wedge$  and leveraging that  $h$  is a slot-wise function and is the identity, we have for all  $z \in Z$

$$\begin{aligned} \sum_{k=1}^K f^k(z_{B_k}) + \sum_{k=1}^K \sum_{k^0=k+1}^K z_{B_k} z_{B_{k^0}} A_{kk^0} \\ = \sum_{k=1}^K f^{\wedge k}(h_{B_k}(z_{B_k})) + \sum_{k=1}^K \sum_{k^0=k+1}^K h^k(z_{B_k}) h^{k^0}(z_{B_{k^0}}) A_{kk^0} \end{aligned} \quad (\text{D.5})$$

To prove the claim we now consider the expression

$$\begin{aligned} L(z) &= \sum_{k=1}^K f^k(z_{B_k}) - \sum_{k=1}^K f^{\wedge k}(h_{B_k}(z_{B_k})) \\ &= \sum_{i=k}^K \sum_{k^0=k+1}^K h^k(z_{B_k}) h^{k^0}(z_{B_{k^0}}) A_{kk^0} - \sum_{k=1}^K \sum_{k^0=k+1}^K z_{B_k} z_{B_{k^0}} A_{kk^0} \end{aligned} \quad (\text{D.6})$$

and prove that  $L(z)$  is an affine function. To show this it is sufficient to prove that the second derivative  $D^2L$  vanishes because  $Z_{\text{supp}}$  is path-connected. Thus we consider all partial derivatives. Consider first the case where  $i \in B_k$  and  $i^0 \in B_{k^0}$  for  $k < k^0$ . Then we find that

$$D_i D_{i^0} L(z) = D_i D_{i^0} \prod_{k=1}^k f^k(z_{B_k}) \prod_{k=1}^{k^0} f^{k^0}(h_{B_k}(z_{B_k})) = 0 \quad (\text{D.7})$$

It remains to consider derivatives of the form  $D_i D_{i^0}$  where  $i \in B_k$  for some slot  $i$ . Then we clearly have

$$D_i D_{i^0} \prod_{k=1}^k z_{B_k} \prod_{k^0=k+1}^{k^0} A_{kk^0} = 0 \quad (\text{D.8})$$

because this is a linear expression in  $z_{B_k}$ . Next, we want to show that

$$D_i D_{i^0} h^k(z_{B_k}) h^{k^0}(z_{B_{k^0}}) A_{kk^0} = 0 \quad (\text{D.9})$$

for all  $k < k^0$ . To prove this we show the more general statement (that will be used in the proof of Theorem 4.4 below) that for any  $v \in \mathbb{R}^{B_{k^0}}$  and any vector  $z \in \mathbb{R}^{B_k}$  the functions

$$z_{B_k} \cdot (h^k(z_{B_k}) \cdot v) A_{kk^0} \quad (\text{D.10})$$

are affine on  $Z_k$  or equivalently that

$$D_i D_{i^0} h^k(z_{B_k}) \cdot v A_{kk^0} = 0 \quad (\text{D.11})$$

for every  $v \in \mathbb{R}^{B_{k^0}}$ . To prove this we consider any  $z \in B_{k^0}$  and apply the derivative  $D_i D_{i^0} D_j$  to (D.5) to get

$$0 = D_i D_{i^0} h^k(z_{B_k}) \cdot D_j h^{k^0}(z_{B_{k^0}}) A_{kk^0} \quad (\text{D.12})$$

for every  $z \in Z_{\text{supp}}$ . Now we use that by assumption  $h$  is a diffeomorphism. Using the block structure of  $h$  we find that also  $h^k$  are diffeomorphisms. In particular, this implies that for any  $z \in Z_{\text{supp}}$  the vectors  $(D_j h^{k^0}(z_{B_{k^0}}))_{j \in B_{k^0}}$  are linearly independent vectors in  $\mathbb{R}^{B_{k^0}}$  and they thus generate  $\mathbb{R}^{B_{k^0}}$ . Therefore we can find coefficients  $\alpha_j$  (depending on  $z_{B_{k^0}}$ ) such that

$$\sum_{j \in B_{k^0}} \alpha_j D_j h^{k^0}(z_{B_{k^0}}) = v \quad (\text{D.13})$$

Then we get using (D.12)

$$\begin{aligned} D_i D_{i^0} h^k(z_{B_k}) \cdot v A_{kk^0} &= D_i D_{i^0} \left( h^k(z_{B_k}) \cdot \sum_{j \in B_{k^0}} \alpha_j D_j h^{k^0}(z_{B_{k^0}}) \right) A_{kk^0} \\ &= \sum_{j \in B_{k^0}} \alpha_j D_i D_{i^0} h^k(z_{B_k}) \cdot D_j h^{k^0}(z_{B_{k^0}}) A_{kk^0} = 0 \end{aligned} \quad (\text{D.14})$$

So (D.10) holds and thus also (D.9) (we actually only get this for points  $z \in Z_k$  such that there is  $z \in Z_{\text{supp}}$  with  $z_k = z_{B_k}$  but by continuity and since  $Z_{\text{supp}}$  is regularly closed this actually holds on  $Z_k$ ). The same reasoning shows that this is also true if  $i \in B_{k^0}$  (instead of  $i \in B_k$ ). We then find that for  $i \in B_k$

$$\begin{aligned} D_i D_{i^0} \prod_{k=1}^k h^k(z_{B_k}) \prod_{k^0=k+1}^{k^0} A_{kk^0} &= \\ &= \prod_{k=1}^k D_i D_{i^0} h^k(z_{B_k}) \prod_{k^0=k+1}^{k^0} A_{kk^0} = 0 \end{aligned} \quad (\text{D.15})$$

The last display together with (D.8) and (D.7) imply  $D^2L = 0$  and thus  $L$  is affine. When  $L$  is not the identity the proof is similar.  $\square$

We also need the following simple lemma which states that we have unique Cartesian-product extension of functions with interaction of order at most  $\ell$  between different slots.

**Lemma D.3.** Let  $f : Z \rightarrow \mathbb{R}^{d_x}$  be a  $C^3$  diffeomorphism with interaction at most  $\ell$  between different slots such that  $Z_{\text{supp}}$  is regularly closed and for  $z \in Z_{\text{supp}}$

$$f(z) = \prod_{k=1}^K f^k(z_{B_k}) + \sum_{2 \leq m \leq n \leq 2l_m} X_{2 \dots m} X_{m+1 \dots n} c_{2 \dots n} z \quad (\text{D.16})$$

for some  $f^i$  in  $C^3$ . Then this relation holds on  $Z_{\text{CPE}}$ .

**Proof.** We know by Theorem C.2 that a representation as in (D.16) holds on  $\mathbb{R}^{d_z}$  and thus can be restricted to  $Z_{\text{CPE}}$ , however it might not be the same representation but involve functions and constants. Taking the difference and setting  $g^k = f^k - f^k$  and  $c = c - c$  we find that on  $Z_{\text{supp}}$

$$0 = \prod_{k=1}^K f^k(z_{B_k}) + \sum_{2 \leq m \leq n \leq 2l_m} X_{2 \dots m} X_{m+1 \dots n} c_{2 \dots n} z \quad (\text{D.17})$$

But by applying D for  $2 \leq l_m$  for  $m = n$  down to  $m = 2$  we find  $c = 0$  for all  $2 \leq l_n$  and thus the polynomial term vanishes. Next, we apply and find that  $f^k$  is constant on  $Z_k$  (because  $Z_{\text{supp}}$  is regularly closed). This implies that (D.17) holds on  $Z_{\text{CPE}}$  and thus (D.16) holds on  $Z_{\text{CPE}}$ .  $\square$

Using the previous lemmas we can prove Theorem 4.4.

**Theorem 4.4 (Compositional Generalization)** Let  $n \in \{0, 1, 2\}$ . Let  $Z_{\text{supp}}$  be regular closed (Defn. A.3). Let  $f : Z \rightarrow X$  and  $f^\wedge : Z \rightarrow \mathbb{R}^{d_x}$  be  $C^3$  diffeomorphisms with at most  $n$  order interactions across slots  $\alpha$ . If  $f^\wedge$  disentangles  $f$  on  $Z_{\text{supp}}$  w.r.t.  $f$  (Defn. 2.1), then it generalizes compositionally (Defn. 2.2).

**Proof of Theorem 4.4** Note that Corollary 3 in Lachapelle et al. (2023) already handles the case  $n = 0, 1$  but the proof below is more general, and also covers the case  $n = 2$ , since functions with at most  $\ell$  and  $\ell'$  order interactions are special cases of functions with at most  $\ell$  order interactions assuming  $f$  is a  $C^3$  diffeomorphism.

We can apply Theorem C.2 to which implies that  $f$  can be written on  $Z = \mathbb{R}^{d_z}$  as in (C.1) and as explained in Lemma D.1 an equivalent representation is

$$f(z) = \prod_{k=1}^K f^k(z_{B_k}) + \sum_{k=1}^K \sum_{k^0=k+1}^K X_{z_{B_k}} X_{z_{B_{k^0}}} A_{kk^0} \quad (\text{D.18})$$

and we have similarly

$$f^\wedge(z) = \prod_{k=1}^K f^{\wedge k}(z_{B_k}) + \sum_{k=1}^K \sum_{k^0=k+1}^K X_{z_{B_k}} X_{z_{B_{k^0}}} \hat{A}_{kk^0} \quad (\text{D.19})$$

By assumption we have  $f = f^\wedge \circ h$  on  $Z_{\text{supp}}$  where  $h(z) := (h_1(z_{B_{(1)}}), \dots, h_K(z_{B_{(K)}}))$  and the functions  $h_k : \mathbb{R}^{B_{(k)}} \rightarrow \mathbb{R}^{B_k}$  are diffeomorphisms. Our goal is to show that this relation actually holds on the Cartesian-product extension  $Z_{\text{CPE}}$ . Let  $U$  be the set of points such that  $f(z) = f^\wedge \circ h(z)$  for  $z \in U$ . We claim that if  $z = (z_{B_1}, \dots, z_{B_K}) \in U$  then  $z^0 = (z_{B_1}^0, \dots, z_{B_1}^0, \dots, z_{B_K}^0) \in U$  for any  $z_{B_1}^0 \in Z_1$ . Let us define the map  $\varphi : Z_1 \rightarrow Z$  given by  $\varphi(z_{B_1}^0) = z^0$ . We know by Lemma D.1 that the function

$$z \mapsto \prod_{k=1}^K f^k(z_{B_k}) - \prod_{k=1}^K f^{\wedge(k)}(h_k(z)_{B_k}) = L(z) \quad (\text{D.20})$$

is affine on  $Z_{\text{supp}}$ . Applying Lemma D.3 the same holds  $\bar{z}_{\text{PE}}$ . Thus we conclude that

$$z_{B_1}^0 \cdot \prod_{k=1}^K f^k(e^z(z_{B_1}^0)_{B_k}) = \prod_{k=1}^K f^{\wedge(k)}(h_k(e^z(z_{B_1}^0)_{B_k})) = L(e^z(z_{B_1}^0)) \quad (\text{D.21})$$

is affine on  $Z_1$ . Moreover,

$$z_{B_1}^0 \cdot \prod_{k=1}^K e^z(z_{B_1}^0)_{B_k} = \prod_{k^0=k+1}^K e^z(z_{B_1}^0)_{B_{k^0}} A_{kk^0} \quad (\text{D.22})$$

is clearly affine on  $Z_1$  and by (D.10) the same holds for

$$z_{B_1}^0 \cdot \prod_{k=1}^K h^k(e^z(z_{B_1}^0)_{B_k}) = \prod_{k^0=k+1}^K h^{k^0}(e^z(z_{B_1}^0)_{B_{k^0}}) A_{kk^0} \quad (\text{D.23})$$

The last three displays together imply that

$$z_{B_1}^0 \cdot f(e^z(z_{B_1}^0)) = f^{\wedge} h(e^z(z_{B_1}^0)) \quad (\text{D.24})$$

is affine on  $Z_1$  and since it is zero in a neighbourhood of  $z_{B_1}^0 = z_{B_1}$  (because  $\alpha \geq 2$ ) it is equal to zero on  $Z_1$ . Since this is true for any slot  $B_1$ , we can now conclude that  $\mathcal{A} = Z$ . Indeed, pick any open rectangle  $Z_1^0 \times Z_2^0 \times \dots \times Z_K^0 \times Z_{\text{supp}} \times U$ . We then infer that  $Z_1 \times Z_2^0 \times \dots \times Z_K^0 \times U$  and by inducting over the slots and applying continuity at the boundary we obtain the claim.

## E UNIFYING ASSUMPTIONS FROM PRIOR WORK

### E.1 AT MOST 0<sup>TH</sup> ORDER INTERACTION ACROSS SLOTS

To prove that the assumptions in Brady et al. (2023) are a special case of our assumptions, we first restate their assumptions formally. To this end, we first define the following set:

$$\mathcal{S} = \{d_z\} \cup \{I_S(z) := f \mid \exists 2 \times [d_x] : D_S f|_1(z) \neq 0\} \quad (\text{E.1})$$

The assumption of compositionality in Brady et al. (2023) can now be stated:

Definition E.1 (Compositionality) A differentiable function  $f : Z \rightarrow X$ , is said to be compositional if:

$$\exists z \in Z; k, j \in [K] : I_k(z) \setminus I_j(z) = \emptyset \quad (\text{E.2})$$

We now state the second assumption in Brady et al. (2023), irreducibility.

Definition E.2 (Irreducibility). A differentiable function  $f : Z \rightarrow X$ , is said to be irreducible if for all  $z \in Z$  and  $k \in [K]$  and any partition  $I_k(z) = S_1 \cup S_2$  (i.e.,  $S_1 \setminus S_2 = \emptyset$ ; and  $S_1, S_2 \neq \emptyset$ ), we have:

$$\text{rank } Df|_{S_1}(z) + \text{rank } Df|_{S_2}(z) > \text{rank } Df|_{I_k}(z) \quad (\text{E.3})$$

We now prove that compositionality and irreducibility are equivalent to satisfying interaction asymmetry (3.5) for all equivalent generators (4.1) for  $\theta = 0$ .

Theorem E.3. A  $C^1$  diffeomorphism  $f : Z \rightarrow X$  satisfies compositionality (Def. E.1) and irreducibility (Def. E.2) if and only if  $f$  has at most 0 order interaction across slots (Defn. 3.2) and satisfies interaction asymmetry (Assm. 3.5) for all equivalent generators (4.1).

Proof. We start by proving the forward direction, i.e., that compositionality and irreducibility imply that  $f$  has at most 0 order interaction across slots and satisfies interaction asymmetry for all equivalent generators.

The definitions of compositionality and at most 0 order interaction across slots are precisely equivalent, thus we only need to show that compositionality and irreducibility imply that  $f$  satisfies interaction asymmetry for all equivalent generators. To show this we will prove the following contraposition: that if  $f$  has at most 0 order interaction across slots and does not satisfy interaction asymmetry for all equivalent generators, then  $f$  is not irreducible.



Since  $f$  has at most  $0$ th order interaction across slots and does not satisfy interaction asymmetry for all equivalent generators, this implies that there exists a matrix  $R \in \mathbb{R}^{B_k \times B_k}$  and a partition of  $B_k$ , into  $A; B$  ( $A \cup B = B_k; A \cap B = \emptyset$ ) such that within the function defined as:

$$\forall z \in Z; \quad f(A_1 z_{B_1}; \dots; A_K z_{B_K}) = f(z_{B_1}; \dots; z_{B_K}) \quad (\text{E.4})$$

where  $A_i$  such that  $I_{B_k}$  is the identity matrix, the latents  $z_A; z_B$  have no interaction. This implies that under  $f$ ,  $I_A(z)$  does not intersect with  $I_B(z)$ . Further, because  $\epsilon$  is invertible, we know that  $Df_{B_k}(z)$  is full column rank. Coupling these two properties, we conclude that  $\text{rank}(Df_{B_k}(z)) = \text{rank}(Df_A(z)) + \text{rank}(Df_B(z))$ . Furthermore, the Jacobians  $Df(z)$  and  $Df_S(z)$  will be related by an invertible linear map by construction. Thus,  $Df_S(z)$  and  $Df_S(z)$  have equal rank for any subset  $[d_z]$ . Therefore, we conclude that  $\text{rank}(Df_{B_k}(z)) = \text{rank}(Df_A(z)) + \text{rank}(Df_B(z))$ . Because  $A$  and  $B$  form a partition of  $B_k$  we conclude that  $f$  is not irreducible.

We now prove the reverse direction: if  $f$  has at most  $0$ th order interaction across slots and satisfies interaction asymmetry for all equivalent generators then  $f$  satisfies compositionality and irreducibility. As noted before, the definitions of compositionality and at most  $0$ th order interaction across slots are precisely equivalent. Thus, we only need to show that if  $f$  has at most  $0$ th order interaction across slots and satisfies interaction asymmetry then this implies  $f$  is irreducible. To show this, will prove the following contraposition: that if  $f$  does not satisfy irreducibility, then it does not satisfy interaction asymmetry for all equivalent generators with  $n = 0$ .

Since  $f$  is not irreducible, we know that there exists a slot  $k \in [K]$ , and a partition of  $B_k$  into  $A; B$  such that  $\text{rank}(Df_{B_k}(z)) = \text{rank}(Df_A(z)) + \text{rank}(Df_B(z))$ . Because  $Df_{B_k}(z)$  is full column rank this implies that  $\text{rank}(Df_A(z)) = |A|$  and  $\text{rank}(Df_B(z)) = |B|$ . Now take two matrices  $M_{S_1} \in \mathbb{R}^{d_x \times |A|}$  and  $M_{S_2} \in \mathbb{R}^{d_x \times |B|}$  such that the column space of  $M_{S_1}$  is the same as  $Df_A(z)$  and the column space of  $M_{S_2}$  is the same as  $Df_B(z)$ . Now construct the following matrix  $M \in \mathbb{R}^{d_x \times B_k}$  as follows:

$$M := [M_{S_1}; M_{S_2}] \quad (\text{E.5})$$

Note that by construction this matrix has a block structure such that rows  $M_i$  for  $i \in A$  are never non-zero for the same rows  $M_j$  for  $j \in B$ . Because  $M$  and  $Df_{B_k}(z)$  are both full column rank, then there exist a matrix  $A_k \in \mathbb{R}^{B_k \times B_k}$  such that:

$$M := Df_{B_k}(z) A_k \quad (\text{E.6})$$

Now define the function  $f'$  as follows:

$$\forall z \in Z; \quad f'(A_1^{-1} z_{B_1}; \dots; A_K^{-1} z_{B_K}) = f(z_{B_1}; \dots; z_{B_K}) \quad (\text{E.7})$$

such that  $A_i^{-1}$  is defined as above when  $i = k$ , and otherwise it is the identity matrix.

Writing the derivative of  $Df_{B_k}(z)$  in terms of  $f'$  we get  $Df_{B_k}(z) A_k = M$ . Because  $M$  has a block structure we conclude that there exist a partition of  $B_k$  such that these latents have no interaction within  $f'$  at  $z$ . Because  $\epsilon$  is equivalent to  $f'$  we conclude that the function does not satisfy interaction asymmetry for  $n = 0$ .  $\square$

## E.2 AT MOST 1<sup>ST</sup> ORDER INTERACTION ACROSS SLOTS

We now prove that the assumptions in [Lachapelle et al. \(2023\)](#) are a special case of our assumptions for  $n = 1$ . To this end, we first restate their assumptions. The first assumption in [Lachapelle et al. \(2023\)](#) is that the generator is additive

**Definition E.4 (Additive decoder)** A  $C^2$  diffeomorphism  $f : Z \rightarrow X$  is said to be additive if:

$$f(z) = \prod_{k \in [K]} f^k(z); \quad \text{where } f^k : \mathbb{R}^{B_k} \rightarrow \mathbb{R}^{d_x} \text{ for any } k \in [K] \text{ and } z \in Z \quad (\text{E.8})$$

**Definition E.5 (Sufficient Nonlinearity)** Let  $f : Z \rightarrow X$  be a  $C^2$  diffeomorphism. For all  $k \in [K]$ , let  $B_k^2 := B_k \setminus \{(i_1; i_2) \mid i_1 = i_2\}$ .  $f$  is said to satisfy sufficient nonlinearity if  $\forall z \in Z$  the following matrix has full column-rank:

$$W(z) := [Df(z)]_{i_2 \in B_k} \quad D_{i_1}^2 f(z) \quad (i_1, i_2) \in B_k^2 \quad k \in [K] \quad (\text{E.9})$$

We now state our result.

**Theorem E.6.** Let  $f : Z \rightarrow X$  be a  $C^2$  diffeomorphism. If  $f$  satisfies additivity (Def. E.4) and sufficient nonlinearity (Def. E.5) then  $f$  has at most  $\ell$  order interactions across slots (Defn. 3.3), satisfies sufficient independence (Defn. A.9), and satisfies interaction asymmetry (Asm. 3.5) for all equivalent generators (Defn. 4.1) for  $\ell = 1$ .

**Proof.** We note that  $f$  having at most  $\ell$  order interactions across slots is equivalent to having a block-diagonal Hessian for every observed component. Such functions were proven to be equivalent to additive functions in Lachapelle et al. (2023). Furthermore, sufficient independence is clearly implied by sufficient nonlinearity as if all columns of the matrix  $W^f(z)$  are linearly independent, then blocks  $[D_{i,j} f(z)]_{i,j \in B_k}$   $D_{i,i}^2 f(z)_{(i,i) \in B_k^2}$  will have non-intersecting columns spaces for all  $k \in [K]$  and will thus satisfy sufficient independence (Def. A.9). Consequently, the only thing we need to show is that sufficient nonlinearity (Def. E.5) implies interaction asymmetry (Asm. 3.5) for all equivalent generators (4.1).

Assume for a contradiction that sufficient nonlinearity (Def. E.5) did not imply interaction asymmetry (Asm. 3.5) for all equivalent generators (4.1) with  $\ell = 1$ . This would imply that there exists an equivalent generator  $f$  defined in terms of a slot-wise linear function

$$f = f \circ h \tag{E.10}$$

such that  $f$  has at most  $\ell$  order interaction within some slot  $B_k$ . In other words, leveraging Lemma A.21, there exist  $(j; j^0) \in B_k^2$  and  $z \in Z$  s.t.

$$0 = D_{j,j^0}^2 f(z) = W^f(h(z)) m^h(z; (j; j^0)); \tag{E.11}$$

Because  $W^f(h(z))$  is assumed to be full rank by sufficient nonlinearity (Def. E.5), then in order for this equation to hold  $m^h(z; (j; j^0))$  must be zero. Note, however, that by construction  $z_j, z_{j^0}$  map to the same slot  $B_k$ . By construction, if  $z_j, z_{j^0}$  affect the same slot  $B_k$  then  $m^h(z; (j; j^0))$ , cannot be zero. Thus, we obtain a contradiction.  $\square$

## F TRANSFORMERS FOR INTERACTION REGULARIZATION

Each layer of a Transformer (Vaswani et al., 2017) consist of two main components: an MLP sub-layer and an attention mechanism. Notably, in the MLP sub-layer, MLPs are applied separately to each slot or pixel query and their outputs are then concatenated. Further, additional layer normalization operations (Ba, 2016) are typically used in Transformers but are also separately applied to each slot or pixel query. Thus, the only opportunity for interaction between slots in a Transformer occurs through the attention mechanism. Our focus in this work is on the cross-attention mechanism, opposed to the alternative self-attention. As noted in § 5, cross-attention takes the form:

$$K = W^K [z_{B_k}]_{k \in [K]}; \quad V = W^V [z_{B_k}]_{k \in [K]}; \quad Q = W^Q [o_d]_{d \in [d_x]}; \tag{F.1}$$

$$A_{d;k} = \frac{\exp(Q_{:,d}^> K_{:,k})}{\sum_{l \in [K]} \exp(Q_{:,d}^> K_{:,l})}; \quad x_d = A_{d,:} V^>; \quad \hat{x}_d = (x_d); \tag{F.2}$$

where  $K_{:,k}; V_{:,k} \in \mathbb{R}^{d_q}$ ,  $W^K; W^V \in \mathbb{R}^{d_q \times |B_k|}$  for query dimension  $d_q$ . Further,  $o_d \in \mathbb{R}^{d_o}$ ,  $Q_{:,d} \in \mathbb{R}^{d_q}$ ,  $W^Q \in \mathbb{R}^{d_q \times d_o}$ , where  $d_o$  is the dimension of a pixel coordinate vector, and  $\mathbb{R}^{d_q} \rightarrow \mathbb{R}$ .

**Additional Details.** In Eq. (F.2), we do not include the scaling factor  $\frac{1}{\sqrt{d_q}}$  for  $A_{d;k}$ , that is typically used as it does not affect our arguments below. We do, however, include it in our experiments. Further, when  $x$  is an RGB image  $\hat{x}_d$  will not be a scalar but will instead be a vector  $\mathbb{R}^3$  since each pixel has 3 color channels. Additionally, in our experiments, multi-head attention is used. In this case, slot keys and values and pixel queries are partitioned into sub-vectors. Eqs. (F.1) and (F.2) are then applied separately to each resulting sub-latent, and the resulting outputs are concatenated. When using multiple layers of cross-attention, as we do in our experiments, only applied at the last layer and vectors for a subsequent layer are defined as the vectors from the prior layer. Eqs. (F.1) and (F.2) is then repeated. We discuss how these additional caveats are dealt with empirically when implementing  $\text{interact}$  below in Appx. F.2, however, they do not affect our formal argument regarding regularizing interactions in Appx. F.1.

### F.1 JACOBIAN OF CROSSATTENTION MECHANISM

Our goal is to show that  $\frac{\partial A_{d;k}}{\partial z_{B_k}}$  in equation is 0, then partial derivative of Eqs. (F.1) and (F.2) w.r.t slot  $z_{B_k}$ , i.e,  $\frac{\partial A_{d;k}}{\partial z_{B_k}}$  will also be zero. This would then imply that  $A_{d;:}$  is non-zero for at most one slot  $k$  for every  $d \in [d_x]$ , and every  $\lambda \in \lambda_{\text{supp}}$  then slots do not interact in the sense of Defn. 3.2, since all such derivative products  $\frac{\partial A_{d;k}}{\partial z_{B_k}} \frac{\partial A_{d;l}}{\partial z_{B_l}}$  for  $l \neq k$  are zero. To this end, we are interested in computing the derivative:

$$\frac{\partial A_{d;k}}{\partial z_{B_m}} = \frac{\partial (x)}{\partial z_{B_m}} \frac{\partial (x_d)_i}{\partial z_{B_m}} \quad (\text{F.3})$$

where we here and from now on use the convention that we sum over every index that appears only on one side. To evaluate this we decompose the terms

$$(x_d)_i = A_{d;k} V_{i;k} = A_{d;k} W_{ij}^V (z_{B_k})_j \quad (\text{F.4})$$

We set  $M_{d;:} := \sigma_d^> (W^Q)^> W^K$  so that

$$Q_{:,d}^> K_{:,k} = M_{d;i} (z_{B_k})_i \quad (\text{F.5})$$

This implies that

$$\frac{\partial}{\partial z_{B_m}} \exp(Q_{:,d}^> K_{:,k}) = M_{d;i} \delta_{im} \exp(Q_{:,d}^> K_{:,k}) \quad (\text{F.6})$$

where  $\delta$  is the Kronecker-Delta (and here no summation over  $d$  is done). This implies using the product rule and the chain rule that

$$\frac{\partial A_{d;k}}{\partial z_{B_m}} = M_{d;i} \delta_{i,m} A_{d;k} = M_{d;i} A_{d;k} A_{d;m} \quad (\text{F.7})$$

Plugging this together we get

$$\begin{aligned} \frac{\partial A_{d;k}}{\partial z_{B_m}} &= \frac{\partial (x)}{\partial z_{B_m}} \frac{\partial (x_d)_i}{\partial z_{B_m}} \\ &= A_{d;m} W_{i,r}^V \frac{\partial (x)}{\partial z_{B_m}} + \frac{\partial (x)}{\partial z_{B_m}} W_{ij}^V (z_{B_k})_j \frac{\partial A_{d;k}}{\partial z_{B_m}} \\ &= A_{d;m} W_{i,r}^V \frac{\partial (x)}{\partial z_{B_m}} + \frac{\partial (x)}{\partial z_{B_m}} W_{ij}^V (z_{B_k})_j (M_{d;r} \delta_{r,m} A_{d;k} = M_{d;r} A_{d;k} A_{d,m}) \\ &= A_{d;m} \frac{\partial (x)}{\partial z_{B_m}} (W_{i,r}^V + W_{ij}^V (z_{B_k})_j M_{d;r}) + \frac{\partial (x)}{\partial z_{B_m}} W_{ij}^V (z_{B_k})_j M_{d;r} A_{d;k} A_{d,m} \end{aligned} \quad (\text{F.8})$$

From this, we can see that  $A_{d;m} = 0$ , then the partial derivative  $\frac{\partial A_{d;k}}{\partial z_{B_m}}$ , will indeed be zero as  $A_{d;m}$  scales both terms in the last line of Eq. (F.8).

### F.2 INTERACTION REGULARIZER

Based on Appx. F.1, we propose to regularize the interaction in a Transformer by minimizing the sum of all pairwise products  $A_{i,j} A_{l,k}$ , where  $j \neq k$ . More specifically, we minimize the following loss:

$$L_{\text{interact}} := \sum_{l \in [d_x]} \sum_{j \in [K]} \sum_{k=j+1}^K A_{i,j} (z) A_{l,k} (z) \quad (\text{F.9})$$

where  $A_{i,k}(z)$  is used to indicate the input dependence of attention weights on inputs,  $L_{\text{interact}}$  is a non-negative quantity which will be zero if and only if a matrix has at most one non-zero for each row (Brady et al., 2023).

Code to compute  $L_{\text{interact}}$  for a batch of attention matrices can be seen in Fig. 4. We note that when using multiple attention heads, we first sum the attention matrices over all heads to ensure consistent pixel assignments across different heads.

When using multiple layers, we also sum the attention matrices over each layer, for the same reason.  $L_{\text{interact}}$  is then computed on the resulting attention matrix.

**Regularizing Higher Order Interactions.** We note that while we motivate  $L_{\text{interact}}$  as a regularizer for 1<sup>st</sup> order interactions, we do not explicitly address regularizing for higher order interactions, i.e., for  $k > 2$ . We conjecture there is a relationship between regularizing  $L_{\text{interact}}$  and higher order interactions but that it is less direct than the 1<sup>st</sup> order case. We leave it to future work to explore these connections further, as well as alternative, computationally efficient regularizers which can more directly penalize higher order interactions.

Figure 4: PyTorch code to compute  $L_{\text{interact}}$

**Computational Efficiency.** We note that regularizing with  $L_{\text{interact}}$  adds minimal additional computational overhead since attention weights are already computed at each forward pass through the model, and, moreover can be easily optimized using gradient descent. This is in contrast to [Brady et al. \(2023\)](#) which required computing the Jacobian of the decoder at each forward pass and then optimizing it using gradient descent. This results in second-order optimization which is computationally intractable for high-dimensional data such as images ([Brady et al., 2023](#)).

## G EXTENDED RELATED WORK

### G.1 THEORY

**Relationship Between Principle 3.1 and Other Principles.** The principle of interaction asymmetry, “parts of the same concept have more complex interactions, than parts of different concepts” (3.1), is intuitively similar to several prior principles explored for learning concepts. For example, the prior works of [Baldwin et al. \(2008\)](#); [Reynolds et al. \(2007\)](#); [Schmidhuber \(1990\)](#); [Zacks et al. \(2011\)](#) on disentangling events/sub-task (e.g., “making coffee”, “driving to work”), [Greff et al. \(2015\)](#); [Hyvarinen and Perkio \(2006\)](#) on disentangling objects in an image, and [Schmidhuber \(1992\)](#) are all essentially based on the principle that parts of same concepts are mutually predictable than parts of different concepts. Similarly, [Hochreiter and Schmidhuber \(1999\)](#); [Jiang et al. \(2022\)](#) implicitly use the idea that parts of same concepts are more compressible than different concepts. Research on networks, use the idea that nodes from the same “community” interact more strongly than nodes from different communities ([Fortunato and Hric, 2016](#)), which also resembles ideas from clustering that points from the same cluster have higher mutual information than from different clusters ([Kraskov et al., 2005](#)). This network-based framework was applied by [Schapiro et al. \(2013\)](#) as a model for grouping temporal events. Lastly, [Greff et al. \(2020\)](#) propose that objects do not interact much with their surroundings but internally have a strong structure. While these different ideas are intuitively similar to interaction asymmetry, they take on different formalizations. Moreover, these principles are generally used as high-level heuristics for designing a learning algorithm, and their theoretical implications for disentanglement and compositional generalization are not explored.

**Connection with Information Bottleneck Principle.** Another notable principle for learning representations is the Information Bottleneck principle ([Alemi et al., 2016](#); [Tishby et al., 2000](#)) which has also been applied in the context of learning disentangled representations ([Meo et al., 2024](#)). In the context of disentanglement, this principle suggest learning a representation which tries to balance a trade-off between minimizing the mutual information between a latent vector and an observation  $x$ , and ensuring that  $z$  contains sufficient information to predict, i.e., reconstruct  $x$ . From a theoretical standpoint, the Information Bottleneck principle differs from the principle of interaction asymmetry as defined in [Asm. 3.5](#). Specifically, [Asm. 3.5](#) is an assumption on the generator and does not place assumptions on the latent distribution. Consequently, our theory describes a setting in which disentanglement can be achieved without explicitly enforcing any additional properties on  $z$ . We note, however, that despite this key difference, our theory does yield insights which resemble the Information Bottleneck principle. Specifically, as noted in [§ 5](#), our theory suggest that if a model uses an inferred latent dimension greater than the ground-truth dimension, then it should aim to encode  $x$  using the minimal necessary latent dimension, i.e., the mutual information

between  $x$  and unnecessary latent dimensions should be minimized, while ensuring that  $z$  can be reconstructed from  $x$ .

**On the Relationship Between Disentanglement and Compositional Generalization** A key premise motivating our theoretical study of compositional generalization is that, from a theoretical perspective, disentanglement does not directly imply compositional generalization. Specifically, this would require that equality between  $f$  and  $f^\wedge$  on  $Z_{\text{supp}}$  (disentanglement Defn. 2.1) implies that these functions were also equal on all  $Z$  (compositional generalization Defn. 2.2). As noted by Lachapelle et al. (2023); Wiedemer et al. (2024a), this will not be true for arbitrary functions and necessitates restrictions on the form of  $f^\wedge$  on all of  $Z$ . While several works have provided empirical corroboration of this theoretical statement for concepts of objects (Wiedemer et al., 2024a) and object attributes (Montero et al., 2021; 2022b; Schott et al., 2022), prior works in disentanglement have suggested that disentanglement can in some cases enable compositional generalization (Esmaeili et al., 2018; Higgins et al., 2017; Mahon et al., 2023). We hypothesize that the compositional generalization abilities observed in the latter works are a consequence of only leaving a small number of novel combinations out of the training set, such that compositional generalization becomes much easier compared to the more restricted training domains explored in (Montero et al., 2021; Wiedemer et al., 2024a). With this being said, it is possible that through hidden inductive biases in a model, disentanglement can directly lead to compositional generalization, which would not be at odds with our theoretical observation.

**Polynomial Decoders.** As noted in § 4.2, Asm. 3.5 implies that the cross-partial derivatives of the generator consisting of components from different slots will be finite-degree polynomials. This partially resembles the polynomial constraints from Ahuja et al. (2023) for disentanglement. Importantly, however, Ahuja et al. (2023) assume that cross-partial derivatives of  $g$  are polynomial such that the entire function is a finite-degree polynomial. In contrast, Asm. 3.5 constrains the form of cross-partial derivatives across slots to be polynomial, but does not constrain the form of cross-partial derivatives within the same slot. In other words, Asm. 3.5 only constrains the interactions across slots, while Ahuja et al. (2023) constrains all possible interactions. This is an important distinction since the former gives rise to much more flexible generators than the latter (see Eq. (4.2)).

## G.2 METHOD AND EXPERIMENTS

**VAE Losses in Object-Centric Models.** Prior work in Wang et al. (2023) also apply a VAE loss to an unsupervised object-centric learning setting. However, while we minimize directly on inferred slots  $z$  given by our Transformer encoder, Wang et al. (2023) minimize on an intermediate representation which is then further processed to yield  $z$ . Furthermore, the focus of Wang et al. (2023) is on scene generation and not penalizing the capacity of  $z$ . Additionally, Kori et al. (2024) explore a loss for object-centric learning resembling a VAE loss, though their aim is to enforce a certain probabilistic structure on  $z$  implied by their theoretical disentanglement result, opposed to penalize latent capacity.

**Inductive Bias Through Explicit Supervision.** Recently, many works have shown remarkable empirical success in disentangling (Kirillov et al., 2023; Ravi et al., 2024) and composing (Brooks et al., 2023; Ramesh et al., 2021; 2022; Ruiz et al., 2023; Saharia et al., 2022) visual concepts in images on web-scale data. These works achieve this through explicit supervision via segmentation masks or natural language descriptions of each concept, opposed to constraints on the generative process in Eq. (2.1). Notably, however, many species in human’s evolutionary lineage disentangle and compose concepts in sensory data without using explicit supervision like natural language (Behrens et al., 2018; LeCun, 2022; Summerfield, 2022; Tolman, 1948). This suggests the existence of a self-supervised coding mechanism for disentanglement and compositional generalization that is still lacking in current machine learning models. The present work aims to make theoretical and empirical progress towards such a mechanism.

**Relation Between a Transformer Regularized with  $L_{\text{interact}}$  and Prior Works.** Goyal et al. (2021) proposed RIMs which is a Transformer-style architecture aimed at enforcing a “modular” structure. Contrary to our work, Goyal et al. (2021) do not regularize for modularity, but posit that it may emerge from “competition” induced by an attention mechanism. Similarly, Lamb et al. (2021) propose an alternative Transformer architecture aimed at enforcing modularity, which also tries to enforce competition using a mechanism similar to Goyal et al. (2021). More recently, Vani et al. (2024) propose a novel Transformer component which is aimed at yielding disentanglement by processing a

Transformer embedding into different slots using separate attention heads for each slot. While these works are similar to ours in that they aim to learn disentangled representations of concepts using a Transformer-style architecture, they are based on architectural changes to a Transformer, whereas we use a standard cross-attention Transformer decoder and regularize it explicitly towards having a “modular” structure using  $\mathcal{G}_{\text{interact}}$ .

## H EXTENDED DISCUSSION

### H.1 THEORETICAL ASSUMPTIONS

**Non-Homogeneous Interactions** One potential limitation of our formulation of interaction asymmetry (Asm. 3.5) is that the order of interaction, across slots, must be the same for all latent vectors  $\mathbf{z} \in \mathcal{Z}$  and for any two slots  $\mathbf{z}_{B_i}; \mathbf{z}_{B_j}$ . This assumption will potentially be violated in practice. For example, for concepts of visual objects, it is likely that within each image, only a few objects, i.e., slots, interact at a time (e.g., see Fig. 1), such that different slots will have different orders of interaction within  $\mathcal{Z}$ . We conjecture that our theory can be extended to handle such non-homogeneous interactions, however, we leave this for future work. Furthermore, we note that despite this potential mismatch between theory and practice, our method still achieved robust object-disentanglement on data in which the order of interaction appears to be non-homogeneous, e.g., CLEVR6.

**Requirements on the Observed Dimension** We note that an implication of sufficient independence for  $n > 0$  (where  $n$  is the order of interaction across slots) is that the observed dimension must be greater than the latent dimension. Moreover, the required  $d_x$  will scale as a function of the number of latent slots  $K$ , the slot dimension  $B_k$ , and the order of interaction across slots. For example, for functions with at most  $n$  order interactions across slots, ensuring that the rank condition in sufficient independence (Defn. A.9) is met requires that

$$d_x \geq \sum_{k \in [K]} \frac{j_{B_k} j_{(B_k+1)}}{2} + d_z.$$

Furthermore, for functions with at most  $n$  order interactions, satisfying this condition (Defn. 4.2) requires  $d_x \geq \sum_{k \in [K]} \frac{j_{B_k} j_{(B_k+1)} (j_{B_k+2})}{6} + \frac{d_z (d_z + 1)}{2} + d_z$ . We note that we are interested in modelling high-dimensional sensory data, such as images, in which the observed dimension be much greater than the latent dimension. Thus, for practical cases of interest, we expect these requirements on  $d_x$  to be met.

**Concepts Potentially not Captured by Interaction Asymmetry.** For certain concepts, it is not obvious if interaction asymmetry will always hold. For example, consider object attributes such as the x-y-position of an object, which can be modelled by one-dimensional slots. For such concepts, the interaction within a slot, i.e., the interaction of each latent component w.r.t. itself, should, intuitively, be a simple function. It is thus not obvious if the interaction within each slot will necessarily be more complex than interactions across slots, such that they may not satisfy interaction asymmetry (Asm. 3.5). Additionally, it is not immediately clear how interaction asymmetry can be applied to more abstract concepts which are not directly grounded in sensory data such as the concept of “democracy” or the concept of a “function” in mathematics.

**Restrictiveness of the Aligned-Connected Assumption** Our theoretical results in § 4 leverage the assumption that the latent space  $\mathcal{Z}_{\text{app}}$  is aligned-connected (Defn. A.16). To assess whether the aligned-connectedness assumption is realistic, we believe it is helpful to look at concrete mathematical examples of supports that satisfy it. For example, the whole space is aligned-connected. More generally, any convex set is aligned-connected. This includes the hypercube, any closed ball, and much more. Some aligned-connected sets are not convex. For example, the “L-shaped” set  $[0; 2]^2 \cap [1; 2]^2$  is aligned-connected but not convex. This last example is useful to model concrete settings where some combinations of latent factors are not observed at training time. This corresponds to the running example of Lachapelle et al. (2023) consisting of two balls moving up and down where the configurations where both balls appear in the top half of the image are never observed.

### H.2 METHOD AND EXPERIMENTS

**Self-Attention in Transformer Decoders.** Our Transformer decoder in § 5 resembles the models from Jaegle et al. (2022); Sajjadi et al. (2022a;b) which only rely on a cross-attention mechanism. However, other works in object-centric learning leverage Transformer decoders which also include a self-attention mechanism between queries at each layer (Seitzer et al., 2023; Singh et al., 2022a).

When reconstructing individual pixels, e.g., on Sprites and CLEVR6 in § 6, applying self-attention between queries will not scale to high-dimensional images since it requires computing attention weights, where  $n$  is the number of pixels. However, when reconstructing image patches as was done in our experiments on CLEVRText, using self-attention is scalable since we have a significantly smaller number of queries, e.g. 16 on CLEVRText. While we found that we could achieve strong disentanglement using only a cross-attention mechanism on CLEVRText, it is possible that using self-attention could be advantageous when reconstructing image patches in even more complex settings. For such models, however, it is not immediately obvious how to regularize the decoder to match our theory since adding self-attention between pixels will introduce additional interactions between slots. We leave it for future work to investigate if our current training objective will still yield robust object-disentanglement for such a model and, furthermore, if such a model can be regularized to be in line with our theory.

**Trade-Offs with Slot Attention.** On Sprites and CLEVR6 (§ 6) as well as CLEVRText, we found that our regularized Transformer autoencoder achieved superior disentanglement, based on our metrics, to an unregularized variant with a Slot Attention encoder. Despite this, we emphasize that our goal is not to propose our method as superior to Slot Attention-based methods. Instead, we highlight that both methods offer different trade-offs. For example, training with our proposed loss (Eq. (5.3)) enables using a general Transformer encoder, thus potentially allowing our model to be applied more generally at scale compared to encoders with more explicit object-centric priors such as Slot Attention. This, however, comes at the cost of training with regularizers which require hyperparameter selection. While our experiments did not require extensive hyperparameter tuning, it is possible that certain datasets will exhibit increased sensitivity to these hyperparameters. Additionally, our interaction regularizer is based on decoders which only use a cross-attention mechanism. While this architecture yielded strong disentanglement in our experiments, Slot Attention encoders have been shown to enable disentanglement using more expressive decoders which also use self-attention (Seitzer et al., 2023; Singh et al., 2022a).

**Latent Prediction-Based Disentanglement Metrics** One potential issue with our Jacobian-based disentanglement metrics is that they may fail to measure whether multiple slots actually encode the same object. Specifically, if two slots affect the same object in pixel space, this could be due to both slots encoding the object in latent space, or it could be due to slot interactions modelled by the decoder. Definitely resolving this potential ambiguity would require measuring the information encoded in each slot directly in latent space. Along this line, prior works have considered latent prediction metrics in which the R2 score is computed from the predictions of a model  $t$  between each inferred slot and the best matching ground-truth slot (Dittadi et al., 2022; Jiang et al., 2023; Locatello et al., 2020b). While these metrics indicate if an inferred slot contains all information for a given object, they are insufficient for resolving the possible ambiguity of our current metrics. This is because these metrics do not indicate if an inferred slot contains information about more than one object. This issue with latent prediction metrics was pointed out by Brady et al. (2023) who aimed to address it by measuring the R2 score from an additional predictor  $t$  to the second-best matching ground-truth slot. We found this metric to yield inconsistent results on CLEVR6, which we hypothesize was due to issues when determining the second-best matching ground-truth slot. This lead us to focus on decoder-based metrics which are more straightforward to compute. We leave it for future work to formulate a latent prediction metric which overcomes the aforementioned issues of prior works.

**On Hyperparameter Selection.** One potential limitation of using our regularized loss (Eq. (5.3)) in terms of scalability is that hyperparameter selection is required. In our experiments on Sprites and CLEVR6 (§ 6), extensive hyperparameter tuning was not required. Furthermore, we found the values of  $\alpha = 0.05$ ,  $\beta = 0.05$  to work robustly across both datasets, though the reconstruction loss was weighted by a factor of 5 on Sprites and 1 on CLEVR6. This indicates some level of robustness of these hyperparameters across datasets which contain varying complexity of interaction. On CLEVRText (Appx. I), we found that these exact hyperparameter values did not transfer directly and a small amount of tuning was required to arrive at our values (.1 for all terms in the loss). We hypothesize that this is because, in our current implementation of our loss, the magnitude of the reconstruction loss scales with the dimension of the data. To this end, because the data dimension increased significantly on CLEVRText (256 encoded image patches, each with 768 dimensions), the contribution of the reconstruction term to the loss needed to be slightly diminished. With all this being said, it is possible that more complex dataset could require more extensive hyperparameter tuning, however, we leave this for future work to investigate.

Applying our Method to Other Types of Concepts. One important direction for future work is to apply our method to data consisting of different types of concepts such as object-attributes or temporal events. For object-attributes, our same empirical framework can be applied, but with the additional caveat that the Transformer is permutation invariant, while object-attributes do not possess the same permutation invariance as objects. To this end, methods such as adding a positional encoding to each slot, must be used to address this. Additionally, as noted by Gopalakrishnan et al. (2023); Kipf et al. (2019), the problem of disentangling temporal events in image sequences can also be modelled naturally using a slot-based framework. In this case, the “tokens” that a slot encoder, e.g., a Transformer or Slot Attention, operates on are not pixels processed by a CNN, as in our current model. Instead, they would correspond to individual images in the sequence which are each mapped into representation “tokens”. These tokens can then be mapped into slots by, e.g., a Transformer, and then decoded back to output space, where the queries for the Transformer decoder also would not correspond to individual pixels but instead to images in the temporal sequence.

Limitations of  $L_{\text{disent}}$ . One potential issue with  $L_{\text{rec}}$  is that for real-world data, reconstructing every pixel in an image exactly, may not be necessary and could lead to overly prioritizing tasks irrelevant information in  $z$  such as the background (Seitzer et al., 2023). It would thus be interesting to see if our theory and method could be extended to a self-supervised setting, as in Seitzer et al. (2023), in which exact invertibility is not strictly necessary. Regarding, we first note that in addition to a model having inferred latent dimensionality equal to ground-truth dimensionality, our theory also requires that the inferred slot dimensions equals the ground-truth slot dimensions. While  $L_{\text{KL}}$  explicitly regularizes for the former, it does not directly regularize for the latter. More specifically,  $L_{\text{KL}}$  could, in principle, penalize latent capacity by putting information from all, e.g., objects, in one slot (assuming the slot size is large enough), opposed to distributing this information over components from different slots. Despite this, we found that this failure mode did not occur in our experiments. Another potential issue with  $L_{\text{KL}}$  is that it aims to enforce statistically independent latents which could lead to sub-optimal solutions if the ground-truth latents exhibit strong statistical dependencies. Lastly, regarding  $L_{\text{interact}}$  a shortcoming of this regularizer is that, while it directly regularizes  $k^{\text{th}}$  order interactions (Appx. F.1), its connection to regularizing higher order interactions is not as direct. Future work should thus aim to investigate this point further both theoretically and empirically.

### H.3 ENFORCING THEORETICAL CRITERIA OUT-OF-DOMAIN.

As noted in § 5, enforcing (i) invertibility and (ii) at most  $k^{\text{th}}$  order interactions across slots  $\mathcal{O}^n$  out-of-domain, i.e., globally on all  $\mathcal{O}^n$ , poses distinct practical challenges. We now discuss this in detail. To this end, we first discuss enforcing (ii) globally on

Restricting Interactions Globally. The easiest way to enforce that  $f$  has at most  $k^{\text{th}}$  order interactions across slots  $\mathcal{O}^n$  is to directly parameterize  $f$  to match the form of such functions for some  $n$  (Eq. (4.2)). This is, for example, how at most  $k^{\text{th}}$  order interactions were enforced in Lachapelle et al. (2023), i.e., by defining  $f$  to be an additive function (Defn. E.4) on all  $\mathcal{O}^n$ . We found for higher order interactions, parameterizing  $f$  directly to match the form of Eq. (4.2) leads to training difficulties on toy data. Moreover, even if we could easily train such a model, this explicit form would pose an overly restrictive inductive bias when scaling to more realistic data. This motivated us to consider how to regularize for (ii) opposed to enforce it explicitly. The issue with this approach is that we only regularize the derivatives of  $f$  in-domain on  $Z_{\text{supp}}$ . Yet, enforcing structure on the derivatives of  $f$  on  $Z_{\text{supp}}$  does not imply that same structure will be enforced on all  $\mathcal{O}^n$  as noted in § 4.2, however, by knowing the behavior of the derivatives of  $f$  on  $Z_{\text{supp}}$  we can infer their behavior everywhere on  $\mathcal{O}^n$ . Thus, in principle, it should be possible to propagate the correct derivative structure learned by  $f$  locally on  $Z_{\text{supp}}$  to all of  $Z$ . Practically, however, it is not obvious how this can be done in an effective manner. Thus, properly addressing this challenge would require further methodological and empirical contributions, which are not within the scope of the present work.

Enforcing Invertibility Globally. Additionally, even if  $f$  satisfies (ii) globally, we still must enforce (i) invertibility, globally. As noted in § 5, it is not feasible to define  $f$  such that it is an invertible function from  $Z$  to  $X$  by construction. This necessitated parameterizing the inverse with an encoder  $g$  which was trained to invert the decoder  $f$  via a reconstruction loss. Assuming that a decoder  $f$  satisfies (ii) globally, and is invertible on  $Z_{\text{supp}}$ , it is possible to show that  $f$  will be



invertible on all of  $Z$  and thus generalize compositionally. The issue, however, is that our encoder is only trained to invert  $f^\wedge$  on  $Z_{\text{supp}}$  but not on unseen data from the rest of  $Z$ . Consequently, even if  $f^\wedge$  generalizes compositionally, an encoder will not invert  $f^\wedge$  out-of-domain, and can thus yield an arbitrary representation on such data. This “encoder-decoder inconsistency” was pointed out by Wiedemer et al. (2024a), which studied compositional generalization for decoders with at most 0<sup>th</sup> and 1<sup>st</sup> order interactions. They proposed a loss which addresses this problem by first generating out-of-domain samples using  $g$  and then training the encoder to invert  $f^\wedge$  on this “imagined” data. The implementation of this loss in Wiedemer et al. (2024a), deemed compositional consistency was shown to be ineffective for images consisting of more than 2 objects, however (Wiedemer et al., 2024a). Consequently, scaling this loss, or exploring alternative losses for encoder-decoder consistency, remain open research question that require a deeper investigation to properly address.

For these reasons, the empirical aspects of this work focus on enforcing (i) and (ii) in-domain to achieve disentanglement  $Z_{\text{supp}}$  (Thm. 4.3). As highlighted above, however, our theory elucidates the core problems that need to be solved empirically to also achieve compositional generalization, thus giving a clear direction for future work.

## I EXPERIMENTS ON CLEVRTEX

In this section, we conduct additional experiments on the CLEVRTex dataset (Karazija et al., 2021). This dataset constitutes a significant step up in complexity from CLEVR6 and has been shown to be highly challenging for existing object-centric models (Biza et al., 2023; Karazija et al., 2021). We outline our experimental setup and results below.

### I.1 EXPERIMENTAL SETUP

**Data.** Each image in CLEVRTex consist of between 3 and 10 objects with rich textures, set against complex backgrounds (see Fig. 7 for example images). The dataset consists of 50,000 images. We use 40,000 images for training and 5,000 for validation and testing, respectively.

**Models.** We train 4 models on this data. The first model is our regularized Transformer autoencoder from § 5, for which we weight each term in the loss in Eq. (5.3) by a hyperparameter value of  $\lambda$ . The second model is an unregularized Transformer autoencoder, and the third model is an unregularized autoencoder which uses a Slot Attention encoder with both a Transformer and slot-wise MLP decoder. We train all models using the same setup as in § 6, however, instead of reconstructing the original images, we reconstruct a representation of each image given by a Vision Transformer (ViT) (Dosovitskiy et al., 2021), which is pretrained using the DINO method (Caron et al., 2021). This approach, deemed DINOSAUR (Seitzer et al., 2023), was shown to help object-centric models scale to datasets with increased visual complexity. We thus replace the CNN backbone used in our experiments on Sprites and CLEVR6 with a pretrained ViT which operates on patches of the original images. These patches are mapped to features which are then processed by either a Transformer or Slot Attention encoder. For all models, we use 11 slots with a slot dimension of 64.

**Training and Evaluation Details** We train all models across 3 random seeds using batches of 32. In all cases, we use the Adam optimizer (Kingma and Ba, 2015) with a learning rate of  $10^{-4}$  which we warm-up for the first 10,000 training iterations and then decay by a factor of 10 throughout training. We also warm-up the value of  $\lambda$  for the first 25,000 training iterations. We report the J-ARI and JIS for each model after training for 300,000 iterations. To compute these scores, we bilinearly interpolate our normalized Jacobian maps to match the shape of the original image, since we are reconstructing image patches. When computing J-ARI and JIS for the slot-wise MLP decoder, we rely on the alpha-mask of the decoder opposed to its Jacobian due to computational issues when computing the Jacobian for this model.

### I.2 RESULTS

We report our results in Tab. 2. As we can see, similar to on Sprites and CLEVR6, our regularized Transformer achieves strong object-disentanglement, outperforming both unregularized baseline methods in terms of J-ARI. Our model also achieve superior JIS compared to all baseline models with the exception being the slot-wise MLP decoder. This is not unexpected, however, as this decoder explicitly constrains interactions in the same way as the Spatial Broadcast Decoder used

Table 2: Empirical Results. We show the mean std. dev. for J-ARI and JIS (in %) over 3 seeds for different choices of encoders and weights of the loss terms in Eq. (5.3) on CLEVRTex.

Encoder	Decoder	Loss	J-ARI(")		JIS (")	
Transformer	Transformer	$= 0; = 0$	81:4	3:7	50:9	2:8
Slot Attention	Transformer	$= 0; = 0$	94:2	0:2	54:4	0:3
Slot Attention	Slot-wise MLP	$= 0; = 0$	92:8	0:2	84:3	0:4
Transformer	Transformer	$= 0:1; = 0:1$ (Ours)	95:9	0:06	65:4	0:6

in other experiments. We also visually corroborate these results by plotting normalized slot-wise Jacobians for each model which can be seen in Fig. 7.

## J EXPERIMENTAL DETAILS

### J.1 DATA, MODEL, AND TRAINING DETAILS

**Data.** The Sprites dataset used in § 6 was generated using the Spriteworld renderer (Watters et al., 2019a) and consist of 100,000 images of size  $64 \times 3$  each with between 2 and 4 objects. The CLEVR6 dataset (Johnson et al., 2017; Locatello et al., 2020b) consist of 53,483 images of size  $128 \times 128 \times 3$  each with between 2 and 6 objects. For Sprites, we use 5,000 images for validation, 5,000 for testing, and the rest for training, while for CLEVR6, we use 2,000 images for validation and 2,000 for testing.

**Encoders.** All models use encoders which first process images using the same CNN of Locatello et al. (2020b). When using a Transformer encoder, these CNN features are fed to a 5 layer Transformer which uses both self- and cross-attention with 4 attention heads. When using a Slot Attention encoder, we use 8 Slot Attention iterations, and use the improved implicit differentiation proposed in Chang et al. (2022). Both the Transformer and Slot Attention encoders use learned query vectors opposed to randomly sample queries, which was shown by Biza et al. (2023) to yield improved performance for Slot Attention. On Sprites, all models use 5 slots, each with 32 dimensions, while on CLEVR6, all models use 7 slots, each with 64 dimensions. When using a VAE loss, this slot dimension doubles since we must model the mean and variance of each latent dimension.

**Decoders.** When using a Spatial Broadcast decoder (Watters et al., 2019b), we use the same architecture as (Locatello et al., 2020b) across all experiments, using a channel dimension of 32 for both datasets. When using a Transformer decoder, we first upscale slots to 516 dimensions by processing them separately using a 2 layer MLP, with a hidden dimension of 2064. We then apply a 2 layer cross-attention Transformer to these features which uses 12 attention heads. To obtain the vectors  $o_i$  in Eq. (5.1), we apply a 2D positional encoding to each pixel coordinate. This vector is then mapped by a 2 layer MLP with a hidden dimension of 360 to yield  $o_i$  which has dimension 180. The function  $\psi$  in Eq. (5.2) is implemented by a 3 layer MLP with a hidden dimension of 180, which outputs a 3 dimensional pixel  $a_i$  for each pixel. We additionally note that this architecture does not rely on auto-regressive masking as in Singh et al. (2022a).

**Training Details.** We train all models on Spriteworld across 3 random seeds using batches of 64 for 500,000 iterations. For CLEVR6, we use batches of 32 and train for 400,000 iterations. In all cases, we use the Adam optimizer (Kingma and Ba, 2015) with a learning rate of  $10^{-4}$  which we warm-up for the first 30,000 training iterations and then decay by a factor of 10 throughout training. When training with  $L_{KL}$  and  $L_{interact}$  we use hyperparameter weights of 0.05, which we found to work well across both datasets. We found much larger values could lead to more training instability and, in some cases, insufficient optimization of  $\beta_{ec}$ , while smaller values often did not lead to sufficient optimization of the regularizers. We warm-up the value of  $\beta_{ec}$  for the first 30,000 training iterations. Additionally, when training with  $L_{KL}$  or  $L_{interact}$ , we drop the value of the learning rate after 30,000 training iterations to  $10^{-4}$ , which improved training stability. Lastly, on Sprites, we weight  $L_{KL}$  by a factor of 5, when training with  $L_{KL}$  or  $L_{interact}$ .

## J.2 METRICS AND EVALUATION

**Computing ARI with Attention Scores.** To compute the Adjusted Rand Index (ARI), each pixel must first be assigned to a unique model slot. To this end, prior works typically choose the slot with the largest attention score from either Slot Attention or the alpha mask of a Spatial Broadcast decoder (Locatello et al., 2020b; Seitzer et al., 2023). This approach can be problematic since the attention scores used are model-dependent, making a direct comparison of ARI across models challenging. Further, the relationship between attention scores and the pixels encoded in a model slot is somewhat indirect. As noted in § 6, we consider an alternative and compute the ARI using the Jacobian of a decoder (J-ARI). Specifically, we assign a pixel  $l$  to the slot with the largest  $L_1$  norm for the slot-wise Jacobian  $D_{B_k} \hat{f}_l(\mathbf{z})$ . This can be done for any autoencoder and provides a more principled metric for object disentanglement since a decoder’s Jacobian directly describes the pixels each slot encodes (assuming  $\hat{f}, \hat{g}$  invert each other).

**Evaluation.** We select models for testing which had the highest average values for J-ARI and JIS (each of which take values from 0 to 1) on the validation set. These models were then evaluated on the test set yielding the scores reported in Tab. 1.

## J.3 ADDITIONAL FIGURES

In this subsection, we include 3 additional experimental figures. In Fig. 5, we compare the value of  $L_{\text{interact}}$  throughout training for a model with a Transformer encoder and decoder, trained using our regularized loss Eq. (5.3), the VAE loss and a standard autoencoder loss on both Sprites and CLEVR6. We plot values over 3 random seeds; the shaded regions in the plots indicate one standard deviation from the mean. We find on Sprites (A) and CLEVR6 (B) that the VAE loss achieves much lower  $L_{\text{interact}}$  than the unregularized model. This provides a possible explanation for the solid object disentanglement often achieved by the VAE loss in Tab. 1. We also observe, however, that using  $\alpha > 0$  leads to much lower values for  $L_{\text{interact}}$  compared to the implicit regularization from the VAE loss.

In Fig. 6, we compare slot-wise Jacobians for our model versus baseline models across both Sprites (A) and CLEVR6 (B). To create these plots, we normalize the partial derivatives across slots such that they only take values between 0 and 1. The colors associated with partial derivative values can be interpreted using the color bar at the bottom of (A). We only compute partial derivatives on the foreground pixels and set the derivatives of background pixels w.r.t each slot to 0. We see that when regularizing interactions via our model, slots rarely affect the same pixels (i.e., interact) unnecessarily, while for unregularized models, multiple slots often affect the same pixels even when no interactions should occur, e.g., for images in Sprites (A).

In Fig. 7, we compare slot-wise Jacobians on CLEVRTex as was done in Fig. 6 and also observe here that the regularized Transformer achieves cleaner object decompositions compared to baseline models.

In Fig. 8, we compare decoder attention maps w.r.t. each slot for our model versus baseline models from § 6, which also use a Transformer decoder. These maps, which indicate the slots each pixel attends to, are plotted for both Sprites (A) and CLEVR6 (B). We compute these values by taking the mean attention weight over decoder layers. Similar to Fig. 6, we see that, in our model, pixels rarely unnecessarily attend to multiple slots. On the other hand, for unregularized models, pixels often attend to multiple slots in cases where no interactions between slots should occur.

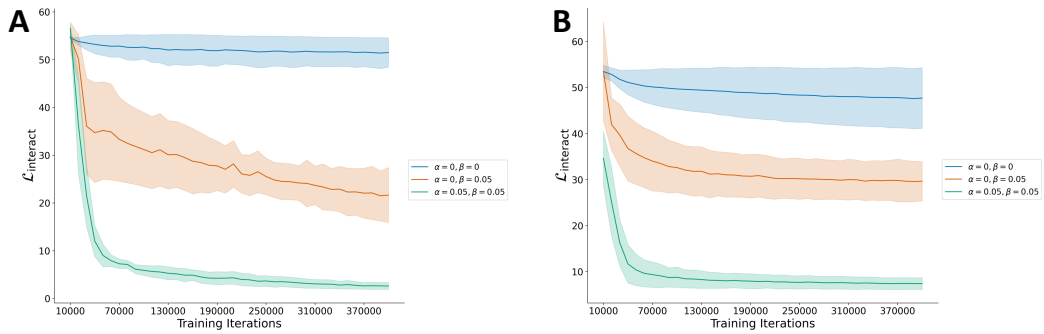


Figure 5: **Analysis of  $\mathcal{L}_{\text{interact}}$  when using a VAE loss.** We plot  $\mathcal{L}_{\text{interact}}$  for the first 400,000 training iterations for a Transformer autoencoder trained without regularization ( $\alpha = 0, \beta = 0$ ), with a VAE loss which does not explicitly optimize  $\mathcal{L}_{\text{interact}}$  ( $\alpha = 0, \beta = 0.05$ ), and with the loss in Eq. (5.3) which regularizes both losses ( $\alpha = 0.05, \beta = 0.05$ ).

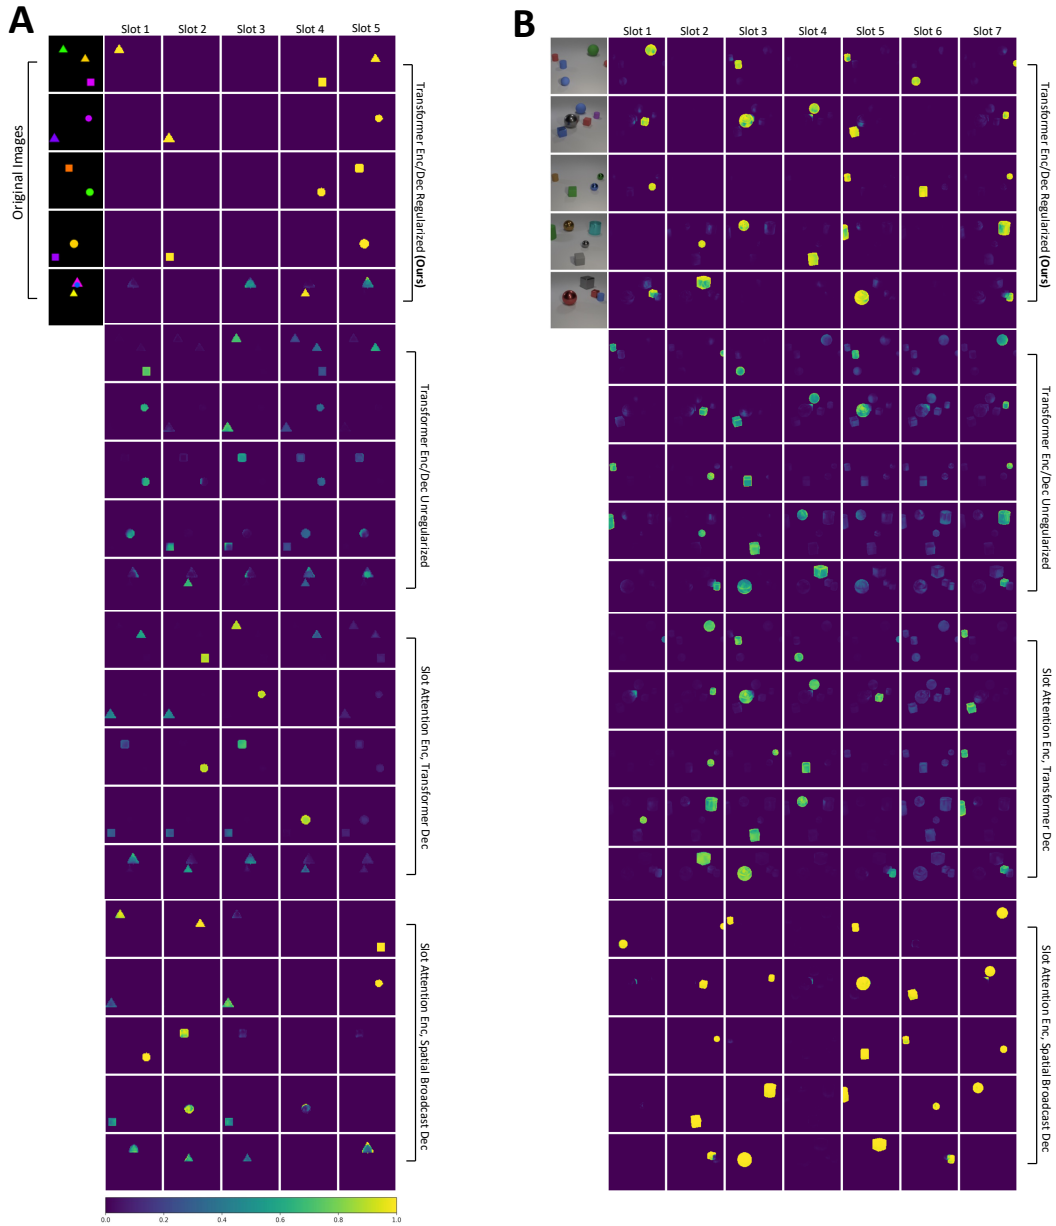


Figure 6: **Normalized Slot-wise Jacobians.** We plot the Jacobians w.r.t. each slot (columns) for 5 random test images (rows) from (A) Sprites and (B) CLEVR6 for our regularized Transformer model and the baseline models used in our experiments in § 6.

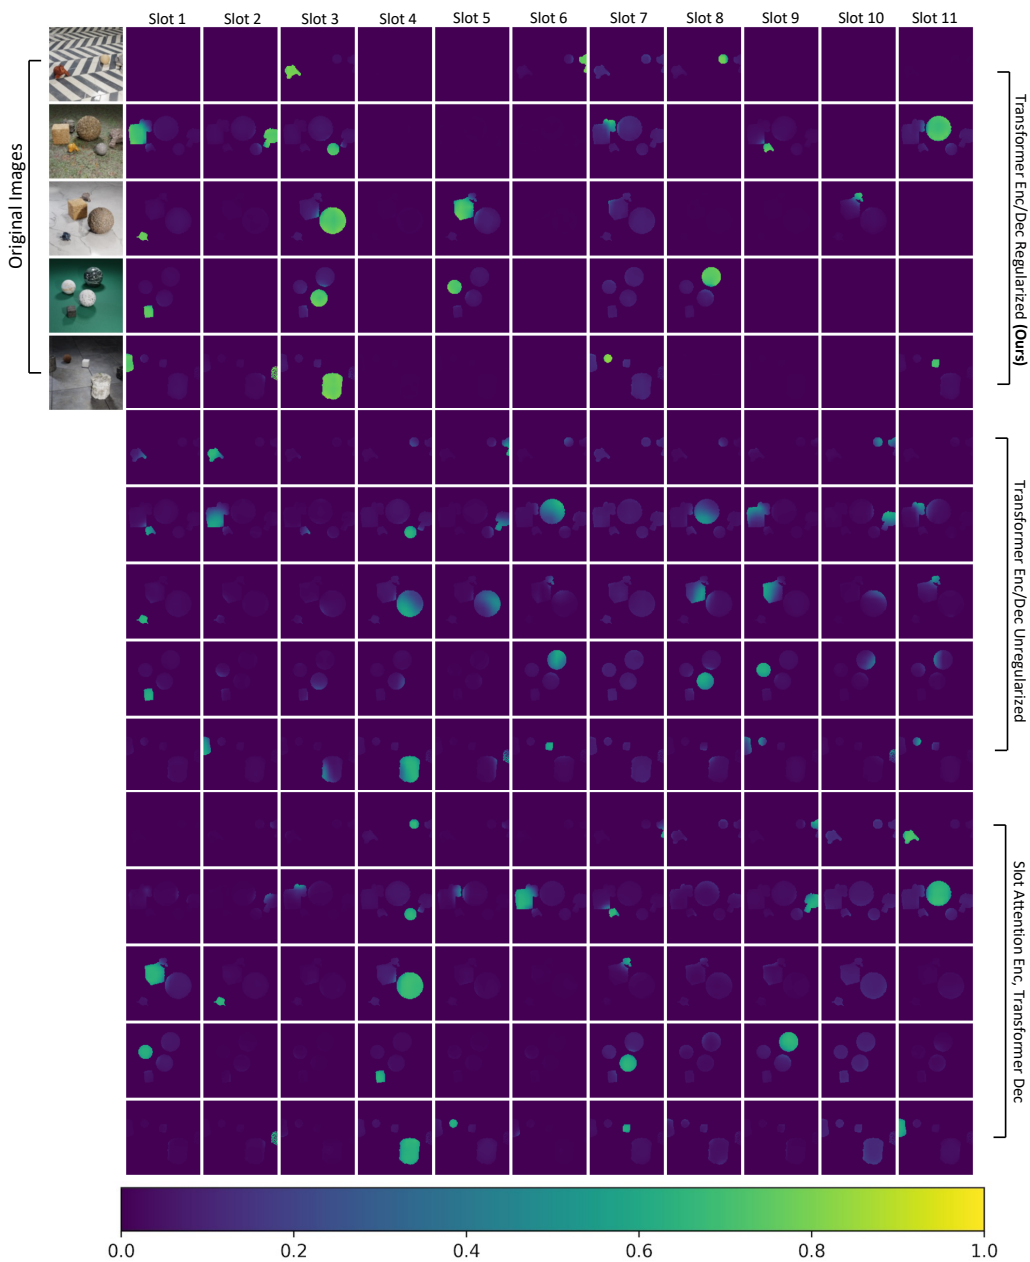


Figure 7: **Normalized Slot-wise Jacobians (CLEVRTex)**. We plot the Jacobians w.r.t. each slot (columns) for 5 random test images (rows) from CLEVRTex for our regularized Transformer model and the baseline models used in our experiments in Appx. I.




Figure 8: **Slot-wise Transformer Decoder Attention Maps.** We plot decoder attention maps w.r.t. each slot (columns) for 5 random test images (rows) from (A) Sprites and (B) CLEVR6 for our regularized Transformer decoder and the baseline models in § 6 which also use a Transformer decoder.