

---

# A causal view of compositional zero-shot recognition

---

Yuval Atzmon<sup>1</sup> Felix Kreuk<sup>1,2</sup> Uri Shalit<sup>3</sup> Gal Chechik<sup>1,2</sup>

<sup>1</sup>NVIDIA Research, Tel Aviv, Israel

<sup>2</sup>Bar-Ilan University, Ramat Gan, Israel

<sup>3</sup>Technion - Israel Institute of Technology

yatzmon@nvidia.com, gchechik@nvidia.com,

## Abstract

People easily recognize new visual categories that are new combinations of known components. This *compositional generalization* capacity is critical for learning in real-world domains like vision and language because the long tail of new combinations dominates the distribution. Unfortunately, learning systems struggle with compositional generalization because they often build on features that are *correlated* with class labels even if they are not “essential” for the class. This leads to consistent misclassification of samples from a new distribution, like new combinations of known components.

Here we describe an approach for compositional generalization that builds on causal ideas. First, we describe compositional zero-shot learning from a causal perspective, and propose to view zero-shot inference as finding “*which intervention caused the image?*”. Second, we present a causal-inspired embedding model that learns disentangled representations of elementary components of visual objects from correlated (confounded) training data. We evaluate this approach on two datasets for predicting new combinations of attribute-object pairs: A well-controlled synthesized images dataset and a real-world dataset which consists of fine-grained types of shoes. We show improvements compared to strong baselines. Code and data are provided in [https://github.com/nv-research-israel/causal\\_comp](https://github.com/nv-research-israel/causal_comp)

## 1 Introduction

Compositional zero-shot recognition is the problem of learning to recognize new combinations of known components. People seamlessly recognize and generate new compositions from known elements and *Compositional Reasoning* is considered a hallmark of human intelligence [33, 34, 6, 4]. As a simple example, people can recognize a purple cauliflower even if they have never seen one, based on their familiarity with cauliflowers and with other purple objects (Figure 1b). Unfortunately, although *feature compositionality* is a key design consideration of deep networks, current deep models struggle when required to generalize to new *label compositions*. This limitation has grave implications for machine learning because the heavy tail of unfamiliar compositions dominates the distribution of labels in perception, language, and decision-making problems.

Models trained from data tend to fail with compositional generalization for two fundamental reasons: **distribution-shift** and **entanglement**. First, recognizing new combinations is an extreme case of distribution-shift inference, where label combinations at test time were never observed during training (zero-shot learning). As a result, models learn correlations during training that hurt inference at test time. For instance, if all cauliflowers in the training set are white, the correlation between the color and the class label is predictive and useful. A correlation-based model like (most) deep networks will learn to associate cauliflowers with the color white during training, and may fail when presented with a purple cauliflower at test time. For the current scope, we put aside the fundamental

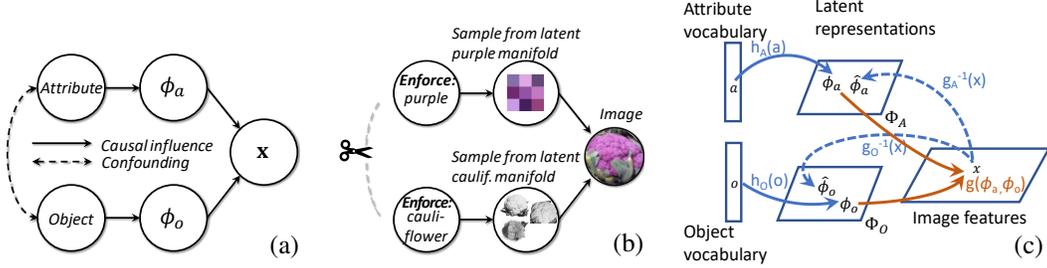


Figure 1: **(a)** The causal graph that generates an image. The solid arrows represent the real-world processes by which the two categorical variables “Object” and “Attribute” each generate “core features” [21, 17]  $\phi_o$  and  $\phi_a$ . The core features then jointly generate an image feature vector  $\mathbf{x}$ . The core features are assumed to be stable for unseen combinations of objects and attributes. The dotted double-edged arrows between the Object and Attribute nodes indicates that there is a process “confounding” the two: they are not independent of each other. **(b)** An intervention that generates a test image of a purple-cauliflower, by enforcing  $a = \textit{purple}$  and  $o = \textit{cauliflower}$ . It cuts the confounding link between the two nodes [49] and changes the joint distribution of the nodes to the “*interventional distribution*”. **(c)** Illustration of the learned mappings, detailed in Section 4

semantic question about what defines the class of an object (cauliflower), and assume that it is given or determined by human observers.

The second challenge is that the training samples themselves are often labeled in a compositional way, and disentangling their “elementary” components from examples is often an ill-defined problem [38]. For example, for an image labeled as white cauliflower, it is hard to tell which visual features capture being a cauliflower, and which, being white. In models that learn from data the representation of these terms may be inherently entangled, and it would be hard to separate which visual features represent white and which represent a cauliflower.

These two challenges are encountered when learning deep *discriminative* models from data. For example, consider a simple model that learns the concept “cauliflower”, by training a deep model over all cauliflower images (VisProd [44]), and the same for the concept “white”. At inference time, simply select the most likely attribute  $\hat{a} = \arg \max_a p(a|\mathbf{x})$  and, independently, the most likely object  $\hat{o} = \arg \max_o p(o|\mathbf{x})$ . Unfortunately, this model, while quite powerful, tends to be sensitive to training-specific correlations in its input.

Here we propose to address compositional recognition by modelling images as being generated, or caused, by real-world entities (labels) (Figure 1). This model recognizes that the distribution  $p(\textit{Image}=\mathbf{x}|\textit{Attr}=a, \textit{Obj}=o)$  is more likely to be stable across the train and test environments ( $p_{\textit{test}}(\mathbf{x}|a, o) = p_{\textit{trn}}(\mathbf{x}|a, o)$ ) [54, 49, 57]: it means that unlike objects or attributes by themselves, *combinations* of objects and attributes generate the same distribution over images in train and test. We propose to consider images of unseen combinations as generated by interventions on the attribute and object labels. In causal inference, intervention means that the value of a random variable is forced to some value, without affecting its causes (but affecting other variables that depend on it, Figure 1b). We cast zero-shot inference as the problem of finding which intervention caused a given image.

In the general case, the conditional distribution  $p(\mathbf{x}|a, o)$  can have arbitrary a complex structure and may be hard to learn. We explain how treating labels as causes, rather than as effects of the image, reveals an independence structure that makes it easier to learn  $p(\mathbf{x}|a, o)$ . We propose conditional independence constraints applied to the structure of this distribution and show how the model can be learned effectively from data.

The paper makes the following novel contributions: First, we provide a new formulation of compositional zero-shot recognition using a causal perspective. Specifically, we formalize inference as a problem of finding the most likely intervention. Second, we describe a new embedding-based architecture that infers causally stable representations for compositional recognition. Finally, we demonstrate empirically that in two challenging datasets, our architecture better recognizes new unseen attribute-object compositions compared to previous methods.

## 2 Related work

**Attribute - object compositionality:** [42] studied decomposing attribute-object combinations. They embedded attributes and object classes using deep networks pre-trained on other large datasets. [45] proposed to view attributes as linear operators in the embedding space of object word-embeddings. Operators are trained to keep transformed objects similar to the corresponding image representation. [46] proposed a method similar to [42] with an additional decoding loss. [60] used GANs to generate a feature vector from the label embedding. [50] trained a set of network modules, jointly with a gating network that rewires the modules according to embeddings of attribute and object labels. [37] is a very recent framework inspired by group theory that incorporates symmetries in label-space.

**Compositional generalizations:** Several papers devised datasets to directly evaluate compositional generalization for vision problems by creating a test set with new combinations of train-set components. [25] introduced a synthetic dataset inherently built with compositionality splits. [1, 9] introduced new compositional splits of VQA datasets [2] and show that the performance of existing models degrades under their new setting. [26] used a knowledge graph for composing classifiers for verb-noun pairs. [6] proposed an experimental framework for measuring compositional generalization and showed that structured prediction models outperform image-captioning models.

**Zero-shot Learning (ZSL):** Compositional generalization can be viewed as a special case of zero-shot learning [62, 35, 8], where a classifier is trained to recognize (new) unseen classes based on their semantic description, which can include a natural-language textual description [48] or predefined attributes [35, 7]. To discourage attributes that belong to different groups from sharing low-level features, [24] proposed a group-sparse regularization term.

**Causal inference for domain adaptation:** Several recent papers take a causal approach to describe and address domain adaptation problems. This includes early work by [55] and [64]. Adapting these ideas to computer vision, [17] were one of the first papers to propose a causal DAG describing the generative process of an image as being generated by a “domain”, which generates a label and an image. They use this graph for learning invariant components that transfer across domains. [36, 3] extended [17] with adversarial training [15]. It learned a single discriminative classifier,  $p(o|\mathbf{x})$ , that is robust against domain shifts and accounts for the dependency of classes on domains. When viewing attributes as “domains”, one of the independence terms in their model corresponds to one term (c) in Eq. (6). [28, 55] discusses image recognition as an “anti-causal” problem, inferring causes from effects. [10, 27] studied learning causal structures under sparse distributional shifts. [40] proposed to learn causal relationships in images by detecting the causal direction of two random variables. [14] used targeted interventions to address distribution shifts in imitation learning. [32] learned a conditional-GAN model jointly with a causal-model of label distribution. [5] proposed a regularization term to improve robustness against distributional changes. Their view is complementary to this paper in that they model the labeling processes, where images cause labels, while our work focuses on the data generating process (labels cause images). [21] proposed a similar causal DAG for images, while adding *auxiliary* information such as that some images are of the *same instance* with a different “style”. This allowed the model to identify core features. The approach described in this paper does not use such auxiliary information. It also views the “object” and “attribute” as playing mutual roles, making their inferred representations invariant to each-other.

**Unsupervised disentanglement of representations:** Several works use a VAE [31] approach for unsupervised disentanglement of representations [39, 22, 11, 13, 52, 41]. This paper focuses on a different problem setup: (1) Our goal is to infer a joint attribute-object pair, disentangling the representation is a useful byproduct. (2) In our setup, attribute-object combinations are dependent in the training data, and new combinations may be observed at test time. (3) We do not use unsupervised learning. (4) We take a simpler embedding based approach.

## 3 Method overview

We start with a descriptive overview of our approach. For simplicity, we skip here the causal motivation and describe the model in an informal way from an embedding viewpoint.

Our model is designed to estimate  $p(\mathbf{x}|a, o)$ , the likelihood of an image feature vector  $\mathbf{x}$ , conditioned on a tuple  $(a, o)$  of attribute-object labels. For inference, we iterate over all combinations of labels and select  $(\hat{a}, \hat{o}) = \operatorname{argmax}_{a, o} p(\mathbf{x}|a, o)$ .

**To estimate** the distribution  $p(\mathbf{x}|a, o)$ , our model learns two embedding spaces:  $\Phi_A$  for attributes, and  $\Phi_O$  for objects (see Figure 1a). These spaces can be thought of as semantic embedding spaces, where an attribute  $a$  (say, “white”) has some dense prototypical representation  $\phi_a \in \Phi_A$ , and an object  $o$  (say, a cauliflower) has a dense representation  $\phi_o \in \Phi_O$ . Given a new image  $\mathbf{x}$ , we learn a mapping to three spaces. First, an inferred attribute embedding  $\hat{\phi}_a \in \Phi_A$  represents the attribute seen in the image (say, how white is the object). Second, an inferred object embedding  $\hat{\phi}_o \in \Phi_O$  represents the object seen in the image (say, how “cauliflowered” it is). Finally, we also represent the image in a general space of image features.

**Learning** the parameters of the model involves learning the three mappings above. In addition, we learn the representation of the attribute prototype (“White”)  $\phi_a \in \Phi_A$  and the representation of the object prototype (“Cauliflower”)  $\phi_o \in \Phi_O$ . Very naturally, we want that a perceived attribute  $\hat{\phi}_a$  would be embedded close to its attribute prototype  $\phi_a$ . Our loss captures this intuition. Finally, we also aim to have the representation spaces of attributes and objects statistically independent. The intuition is that we want to keep the representation of an object (cauliflower) independent of the attribute (white), so we can recognize that object when seen with new attributes (purple cauliflower).

At this point, the challenge remains to build a principled model that can be learned efficiently from data. We now turn to the formal and detailed description of the approach.

## 4 A causal formulation of compositional zero-shot recognition

We put forward a causal perspective that treats labels as causes of an image, rather than its effects. This direction of dependencies is consistent with the mechanism underlying natural image generation and, as we show below, allows us to recognize unseen label combinations.

Figure 1 presents our causal generative model. We consider two “elementary factors” which are categorical variables called “Attribute”  $a \in \mathcal{A}$  and “Object”  $o \in \mathcal{O}$ , and are dependent (confounded) in the training data. As one example, European swans (Cygnus) are white but Australian ones are black. The data collection process may make ‘white’ and ‘swan’ confounded if collected in Europe, even-though this dependency does not apply in Australia.. The model also has two semantic representation spaces: one for attributes  $\Phi_A = \mathbb{R}^{d_A}$  and another for objects  $\Phi_O = \mathbb{R}^{d_O}$ . An attribute  $a$  induces a distribution  $p(\phi_a|a)$  over the representation space, which we model as a Gaussian distribution. We denote by  $h_a$  a function that maps a categorical attributes to the center of this distribution in the semantic space  $h_a : \mathcal{A} \rightarrow \Phi_A$  (Figure 1c). The conditional distribution is therefore  $\phi_a \sim \mathcal{N}(h_a, \sigma_a^2 I)$ . We have a similar setup for  $p(\phi_o|o) \sim \mathcal{N}(h_o, \sigma_o^2 I)$ .

Given the semantic embedding of the attribute and object, the probability of an image feature vector  $\mathbf{x} \in \mathcal{X}$  is determined by the representations  $p(\mathbf{x}|\phi_a, \phi_o)$ , which we model as Gaussian, w.r.t a mapping  $g$ ,  $\mathbf{x} \sim \mathcal{N}(g(\phi_a, \phi_o), \sigma_x^2 I)$ .  $\phi_a$  and  $\phi_o$  can be viewed as an encoding of “core features”, namely encoding a representation of attribute and object that is “stable” in the training set and test set, as proposed by [21, 17]. Namely, the conditional distributions  $p(\phi_a|a)$  and  $p(\phi_o|o)$  do not substantially change for unseen combinations of attributes and objects.

We emphasize that our causal graph is premised on the belief that what we use as objects and attributes are truly distinct aspects of the world, giving rise to different core features. For attributes that have no physical meaning in the world, it may not be possible to postulate two distinct processes giving rise to separate core features.

### 4.1 Interventions on elementary factors

Causal inference provides a formal mechanism to address the confounding effect through a “do-intervention”<sup>1</sup>. A “Do-intervention” *overrides* the joint distribution  $p_{trn}(a, o)$ , enforcing  $a, o$  to specific values and propagates them through the causal graph. With this propagation, an intervention *changes the joint distribution* of nodes in the graph. Therefore, a test image is generated according to a new joint-distribution, denoted in causal language as the *interventional distribution*  $p^{do(A=a, O=o)}(\mathbf{x})$ . Thus, for zero-shot learning, we postulate that inference about a test image is equivalent to asking: **Which intervention on attributes and objects caused the image?**

<sup>1</sup>Our formalism and model can be extended to include other types of intervention on the joint distribution of attributes and objects. For simplicity, we focus here on the most-common “do-intervention”.

## 5 Inference

We propose to infer the attribute and object by choosing the most likely interventional distribution:

$$(\hat{a}, \hat{o}) = \operatorname{argmax}_{a, o \in \mathcal{A} \times \mathcal{O}} p^{do(A=a, O=o)}(\mathbf{x}). \quad (1)$$

This inference procedure is more stable than the discriminative zero-shot approach, since the generative conditional distribution is equivalent to the interventional distribution [55].

$$p^{do(A=a, O=o)}(\mathbf{x}) = p(\mathbf{x}|a, o). \quad (2)$$

This holds both for training and test, so we simply write  $p(\mathbf{x}|a, o)$ . This likelihood depends on the core features  $\phi_A, \phi_O$  which are latent variables; Computing the likelihood exactly requires to marginalize (integrate) over the latent variables. Since this integral is very hard to compute, we take a ‘‘hard’’ approach instead, evaluating the integrand at its most likely value. Since  $\phi_A, \phi_O$  are not known, we estimate them from the image  $x$ , by learning a mapping function  $\hat{\phi}_a = g_A^{-1}(x)$  (see Figure 1c). The supplemental describes these approximation steps in details. It shows that the negative log-likelihood  $-\log p(\mathbf{x}|a, o)$  can be approximated by

$$\hat{L}(a, o) = \frac{1}{\sigma_a^2} \|\hat{\phi}_a - h_a\|^2 + \frac{1}{\sigma_o^2} \|\hat{\phi}_o - h_o\|^2 + \frac{1}{\sigma_x^2} \|\mathbf{x} - g(h_a, h_o)\|^2. \quad (3)$$

Here,  $h_a, h_o$  and  $g(h_a, h_o)$  are the parameters of the Gaussian distributions of  $\phi_a, \phi_o$  and  $\mathbf{x}$ . The factors  $a$  and  $o$  are inferred by taking the  $\operatorname{argmin}_{a, o} \hat{L}(a, o)$  of Eq. (3). Note that in the causal graph,  $\phi_a, \phi_o$  are parent nodes of the image  $\mathbf{x}$ , but  $\hat{\phi}_a, \hat{\phi}_o$  are estimated from  $x$  and are therefore child nodes of  $\mathbf{x}$ , and therefore do not immediately follow the conditional independence relations that  $\phi_a, \phi_o$  obey. We elaborate on this point in section 6.

## 6 Learning

Our model consists of five learned mappings:  $h_A, h_O, g, g_A^{-1}$  and  $g_O^{-1}$ , illustrated in Figure 1c. All mappings are modelled using MLPs. We aim to learn the parameters of these mappings such that the (approximated) negative log-likelihood of Eq. (3) is minimized. In addition, we also include in the objective several regularization terms designed to encourage properties that we want to induce on these mappings. Specifically, the model is trained with a linear combination of three losses.

$$\mathcal{L} = \mathcal{L}_{data} + \lambda_{indep} \mathcal{L}_{indep} + \lambda_{invert} \mathcal{L}_{invert}, \quad (4)$$

where  $\lambda_{indep} \geq 0$  and  $\lambda_{invert} \geq 0$  are hyperparameters. We now discuss these losses in detail.

**(1) Data Likelihood loss.** The first component of the loss,  $\mathcal{L}_{data}$ , corresponds to the (approximate) negative log likelihood of the model, as described by Eq. (3)

$$\mathcal{L}_{data} = \|h_a - g_A^{-1}(\mathbf{x})\|^2 + \|h_o - g_O^{-1}(\mathbf{x})\|^2 + \lambda_{ao} \mathcal{L}_{triplet}(\mathbf{x}, (a, o), (a, o)_{neg}). \quad (5)$$

For easier comparisons with [44], we replaced the rightmost term in Eq. (3) with the standard triplet loss  $\mathcal{L}_{triplet}$  with Euclidean distance  $\|\mathbf{x} - g(h_a, h_o)\|^2$ .  $\lambda_{ao} \geq 0$  is a hyperparameter.

**(2) Independence loss.** The second component of the loss  $\mathcal{L}_{indep}$  is designed to capture conditional-independence relations and apply them to the reconstructed core factors  $\hat{\phi}_a, \hat{\phi}_o$ . By that, the following property is encouraged:  $p^{do(O=o)}(\hat{\phi}_o) \approx p^{do(A=a, O=o)}(\hat{\phi}_o)$  and  $p^{do(A=a)}(\hat{\phi}_a) \approx p^{do(A=a, O=o)}(\hat{\phi}_a)$ . Namely, learning a representation of objects that is robust to attribute interventions, and vice versa.

In more detail, the causal graph (Figure 1a) dictates conditional-independence relations for the latent core factors  $\phi_a, \phi_o$ :

$$\begin{aligned} (a) \quad & \phi_a \perp\!\!\!\perp O | A = a & (b) \quad & \phi_a \perp\!\!\!\perp \phi_o | A = a, \\ (c) \quad & \phi_o \perp\!\!\!\perp A | O = o & (d) \quad & \phi_a \perp\!\!\!\perp \phi_o | O = o. \end{aligned} \quad (6)$$

These relations reflect the independent mechanisms that generate the training data.

Since the core factors  $\phi_a, \phi_o$ , are latent and not observed, we wish that their reconstructions  $\hat{\phi}_a$  and  $\hat{\phi}_o$  maintain approximately the same independence relations. For example, we encourage  $(\hat{\phi}_o \perp\!\!\!\perp A | O = o)$  to capture the independence in Eq. (6)c.

To learn mappings that adhere to these statistical independences over  $\hat{\phi}_a$  and  $\hat{\phi}_o$ , we regularize the learned mappings using a differentiable measure of statistical dependence. Specifically, we use the Hilbert-Schmidt Information Criterion (HSIC) [19, 20]. HSIC is a non-parametric method for estimating the statistical dependence between samples of two random variables, based on an implicit embedding into a universal reproducing kernel Hilbert space. In the infinite-sample limit, the HSIC between two random variables is 0 if and only if they are independent [20]. HSIC also has a simple finite-sample estimator which is easily calculated and is differentiable w.r.t. the input variables. In supplemental Section B, we describe the details of  $\mathcal{L}_{indep}$  and how it is optimized with HSIC, and why minimizing  $\mathcal{L}_{indep}$  indeed encourages the property  $p^{do(O=o)}(\hat{\phi}_o) \approx p^{do(A=a, O=o)}(\hat{\phi}_o)$ . This minimization can be viewed as minimizing the ‘‘Post Interventional Disagreement’’ (PIDA) metric of [58], a recently proposed measure of disentanglement of representations. We explain this perspective in more detail in the supplemental (B.2).

There exist alternative measures for encouraging statistical independence, such as adversarial training [15, 36, 3] or techniques based on mutual information [16, 53, 29]. HSIC has the benefit that it is non-parametric and therefore does not require training an additional network. It was easy to optimize, and was provide useful in previous literature [59, 56, 43, 18].

**(3) Invertible embedding loss.** The third component of the loss,  $\mathcal{L}_{invert}$ , encourages the label-embedding mappings  $h_a$ ,  $h_o$ , and  $g(h_a, h_o)$  to preserve information about their source labels when minimizing  $\mathcal{L}_{data}$ . Without this term, minimizing  $\|\hat{\phi}_a - h_a\|^2$  may reach trivial solutions because the model has no access to ground-truth values for  $\phi_a$  or  $h_a$  (same for  $\phi_o$ ,  $h_o$ ). Similar to [44], we use a cross-entropy (CE) loss with a linear layer that classifies the labels that generate each embedding, and a hyperparameter  $\lambda_g$ :

$$\mathcal{L}_{invert} = CE(a, f_a(h_a)) + CE(o, f_o(h_o)) + \lambda_g [CE(a, f_{ga}(g(h_a, h_o))) + CE(o, f_{go}(g(h_a, h_o)))].$$

## 7 Experiments

### 7.1 Data

Despite several studies of compositionality, current datasets used for evaluations are quite limited. Two main benchmarks were used in evaluations of previous literature: *MIT states* [23] and *UT-Zappos50K* [63].

The MIT-states dataset was labeled automatically using early technology of image search engine based on text surrounding images. As a result, labels are often incorrect. We quantified label quality using human raters, and found that they are too noisy to be useful for proper evaluations. In more detail, we presented images to human raters, along with candidate attribute labels from the dataset. Raters were asked to select the best and second-best attributes that describe the noun (multiple-choice setup). Only 32% of raters selected the correct attribute for their first choice (top-1 accuracy), and only 47% of raters had the correct attribute in one of their choices (top-2 accuracy). The top-2 accuracy was only slightly higher than adding a random label on top of the top-1 label (yielding 42%). To verify that raters were attentive, we also injected 30 ‘‘sanity’’ questions that had two ‘‘easy’’ attributes, yielding top-2 accuracy of 100%. See supplemental (H) for further details. We conclude that this level of  $\sim 70\%$  label noise is too noisy for evaluating noun-attribute compositionality.

**Zappos:** We evaluate our approach on the Zappos dataset, which consists of fine-grained types of shoes, like ‘‘leather sandal’’ or ‘‘rubber sneaker’’. It has 33K images, 16 attribute classes, and 12 object classes. We use the split of [50] and the provided ResNet18 pretrained features. The split contains both seen pairs and unseen pairs for validation and test. It uses 23K images for training of 83 seen pairs, a validation set with 3K images from 15 seen and 15 unseen pairs, and a test set with 3K images from 18 seen and 18 unseen pairs. All the metrics we report for our approach and compared baselines are averaged over 5 random initializations of the model.

**AO-CLEVR:** To evaluate compositional methods on a well-controlled clean dataset, we generated a synthetic-images dataset containing images of ‘‘easy’’ Attribute-Object categories. We used the CLEVR framework [25], hence we name the dataset *AO-CLEVR*. AO-CLEVR has attribute-object pairs created from 8 attributes: { red, purple, yellow, blue, green, cyan, gray, brown } and 3 objects { sphere, cube, cylinder }, yielding 24 attribute-object pairs. Each pair consists of 7500 images. Each image has a single object that consists of the attribute-object pair. The object is randomly assigned

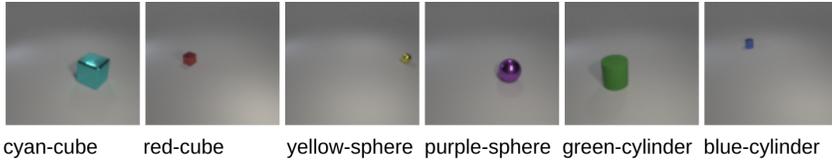


Figure 2:  
Example images  
of AO-CLEVR  
dataset and their  
labels.

one of two sizes (small/large), one of two materials (rubber/metallic), a random position, and random lightning according to CLEVR defaults. See Figure 2 for examples.

For cross-validation, we used two types of splits. The first uses the same unseen pairs for validation and test. This split allows us to quantify the potential generalization capability of each method. In the second split, which is harder, unseen validation pairs are not overlapping with the unseen test pairs. Importantly, we vary the ratio of unseen:seen pairs on a range of (2:8, 3:7, ..., 7:3), and for each ratio we draw 3 random seen-unseen splits. We report the average and the standard error of the mean (S.E.M.) over the three random splits and three random model initialization for each split. We provide more details about the splits in the suppl.

## 7.2 Compared methods

**(1) Causal.** Our approach as described in Section 5. For Zappos it also *learns* a single layer network to project the pretrained image features to the feature space  $\mathcal{X}$ . We also evaluate a variant named **Causal**  $\lambda_{indep}=0$ , which nulls the loss terms that encourage the conditional independence relations

**(2) VisProd:** A common discriminatively-trained baseline [45, 42]. It uses two classifiers over image features to predict the attribute and object independently, and approximates  $p(a, o|\mathbf{x}) \sim p(a|\mathbf{x})p(o|\mathbf{x})$ .

**(3) VisProd&CI:** A discriminatively-trained variant of our model. We use VisProd as a vanilla model, regularized by the conditional independence loss  $\mathcal{L}_{indep}$ , where we use the top network layer activations of attributes and objects as proxies for  $\hat{\phi}_a, \hat{\phi}_o$ .

**(4) LE: Label embedding** [45] trains a neural network to embed images and attribute-object labels to a joint feature space. LE is a vanilla baseline because it approximately models  $p(\mathbf{x}|a, o)$ , but without modelling the core-features.

**(5) ATTOP: Attributes-as-operators** [45] views attributes as operators over the embedding space of object label-embeddings. We use the code of [45] to train ATTOP and LE.

**(6) TMN: Task-modular-networks** [50] trains a set of network modules jointly with a gating network. The gate rewires the modules according to embeddings of attributes and objects. We used the implementation provided by the authors and followed their grid for hyperparameter search (details in suppl.). Our results differ on Zappos because we report an average over 5 random initializations rather than a single initialization as reported in [50]. Some random initializations reproduce well their reported AUC metric.

We explicitly avoid using prior knowledge in the form of *pretrained* label embeddings, because we are interested to quantify the effectiveness of our approach to naturally avoid unreliable correlations. Yet, most of the methods we compare with, rely on pretrained embeddings. Thus, we provide additional results using random initialization for the compared methods, denoted by an asterisk (e.g. LE\*).

**Implementation details** of our approach and reproduced baselines are given in the supplemental.

## 7.3 Evaluation

In zero-shot (ZS) attribute-object recognition, a training set  $\mathcal{D}$  has  $N$  labeled samples of images:  $\mathcal{D} = \{(\mathbf{x}_i, a_i, o_i), i = 1 \dots N\}$ , where each  $\mathbf{x}_i$  is a feature vector,  $a_i \in \mathcal{A}$  is an attribute label,  $o_i \in \mathcal{O}$  is an object label and each pair  $(a_i, o_i)$  is from a set of (source) *seen* pairs  $\mathcal{S} \subset \mathcal{A} \times \mathcal{O}$ . At test time, a new set of samples  $\mathcal{D}' = \{\mathbf{x}_i, i = N + 1 \dots N + M\}$  is given from a set of target pairs. The target set is a union of the set of seen pairs  $\mathcal{S}$  with new *unseen* pairs  $\mathcal{U} \subset \mathcal{A} \times \mathcal{O}$ ,  $\mathcal{U} \cap \mathcal{S} = \emptyset$ . Our goal is to predict the correct pair of each sample.

**Evaluation metrics:** We evaluate methods by the accuracy of their top-1 prediction for recognizing seen and unseen attribute-object pairs. In AO-CLEVR, we compute the balanced accuracy across pairs, namely, the average of per-class accuracy. This is the common metric to evaluate zero-shot

benchmarks [61, 62]. Yet, in Zappos, we used the standard (imbalanced) accuracy, to be consistent with the protocol of [50],

We compute metrics in two main evaluations setups, which differ in their *output* label space, namely, which classes can be predicted. (1) **Closed**: Predictions can be from unseen class-pairs *only*. (2) **Open**: Predictions can be from all pairs in the dataset, seen or unseen. This setup is also known as the *generalized* zero-shot learning setup [62, 12]. Specifically, we compute: **Seen**: Accuracy is computed on *samples from* seen class-pairs. **Unseen**: Accuracy is computed on *samples from* unseen class-pairs. **Harmonic mean**: A metric that quantifies the overall performance of both Open-Seen and Open-Unseen accuracy. It is defined as:  $H = 2(Acc_{seen} * Acc_{unseen}) / (Acc_{seen} + Acc_{unseen})$ . For the harmonic metric, we follow the protocol of [61, 62], which does not take an additional post-processing step. We note that some papers [50, 37, 12] used a different protocol averaging between seen and unseen. Finally, we report the **Area Under Seen-Unseen Curve (AUSUC)**, which uses a post-processing step [12] to balance the seen-unseen accuracy. Inspired by the area-under-the-curve procedure, it adjusts the level of confidence of seen pairs by adding (or subtracting) a constant (see [12] for further details). To compute AUSUC, we sweep over a range of constants and compute the area under the seen-unseen curve.

For early stopping, we use (i) The Harmonic for the open setup. (ii) The closed accuracy for the closed setup. In Zappos, we followed [50] and use the AUSUC for early stopping at the closed setup.

All experiments were performed on a cluster of DGX-V100 machines. Training a single model for 1000 epochs on the 5 : 5 AO-CLEVR split (with ~80K samples) takes 2-3 hours.

## 8 Results

We describe here the results obtained on AO-CLEVR (overlapping-split) and Zappos. Additional results are reported in the supplemental, including the full set of metrics and numeric values; using random initialization; results with the non-overlapping split (showing a similar trend to the overlapping split); studying our approach in greater depth through ablation experiments; and an error analysis.

**AO-CLEVR**: Figure 3 (right) shows the Harmonic metric for AO-CLEVR for the whole range of unseen:seen ratios. Unsurprisingly, the more seen pairs are available, the better all models perform for unseen combinations of attributes and objects. Our approach “Causal”, performs better than or equivalent to all the compared methods. VisProd easily confuses the Unseen classes. ATTOP, is better than LE on the open unseen pairs but performs substantially worse than all methods on the seen pairs. TMN performs equally well as our approach for splits with mostly seen pairs but degrades when the fraction of seen pairs is below 4:6.

**The seen-unseen plane**: By definition, our task aims to perform well in two different metrics (multi-objective): accuracy on seen pairs and unseen pairs. It is therefore natural to compare approaches by their performance on the seen-unseen plane. This is important, because different approaches may select different operating points to trade-off accuracy on unseen for accuracy on seen. Figure 3 (left) shows how the compared approaches trade-off the metrics for the 5:5 split. ATTOP tends to favor unseen-pairs accuracy over the accuracy of seen pairs, vanilla models like VisProd, LE tend to favor seen classes. **Importantly**, it reveals that modelling the core-features largely improves the unseen accuracy, without hurting much the seen accuracy. Specifically, comparing *Causal* to vanilla baseline *LE*, improves the unseen acc. from 26% to 47% and reduces the seen acc. from 86% to 84%. Comparing *VisPros&CI* to *VisProd* improves the unseen acc. from 19% to 38% and reduces the seen Acc. from 85% to 82%.

	UNSEEN	SEEN	HARMONIC	CLOSED	AUSUC
WITH PRIOR EMBEDDINGS					
LE	10.7 ± 0.8	52.9 ± 1.3	17.8 ± 1.1	55.1 ± 2.3	19.4 ± 0.3
ATTOP	22.6 ± 2.9	35.2 ± 2.7	26.5 ± 1.4	52.2 ± 1.8	20.3 ± 1.8
TMN	9.7 ± 0.6	51.9 ± 2.4	16.4 ± 1.0	<b>60.9 ± 1.1</b>	<b>24.6 ± 0.8</b>
NO PRIOR EMBEDDINGS					
LE*	15.6 ± 0.6	52.0 ± 1.0	24.0 ± 0.7	58.1 ± 1.2	22.0 ± 0.9
ATTOP*	16.5 ± 1.5	15.8 ± 1.9	15.8 ± 1.4	42.3 ± 1.5	16.7 ± 1.1
TMN*	6.3 ± 1.4	<b>55.3 ± 1.6</b>	11.1 ± 2.3	58.4 ± 1.5	24.5 ± 0.8
CAUSAL $\lambda_{indep}=0$	22.5 ± 2.0	45.5 ± 3.7	29.4 ± 1.5	55.3 ± 1.1	22.2 ± 0.9
CAUSAL	<b>26.6 ± 1.6</b>	39.7 ± 2.2	<b>31.8 ± 1.7</b>	55.4 ± 0.8	23.3 ± 0.3

Table 1: Results for Zappos. ± denotes the Standard Error of the Mean (S.E.M.) over 5 random model initializations.

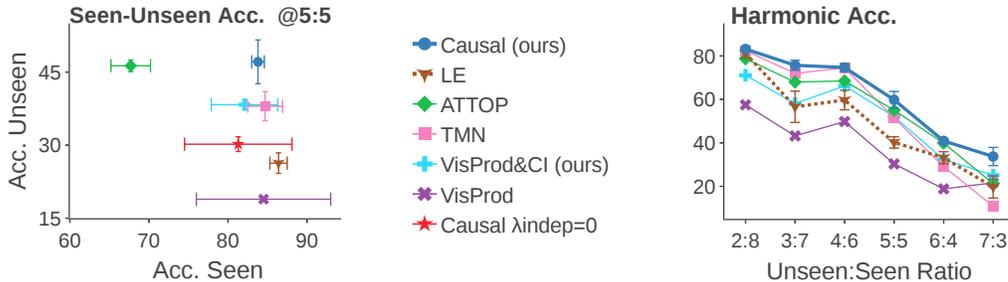


Figure 3: **Left:** The seen-unseen plane for the 5:5 split. Modelling the core features largely improves the unseen accuracy: Compare *Causal* to *LE* or to  $\lambda_{indep}=0$  and compare *VisProd&CI* to *VisProd*. Error bars denote Standard Error of the Mean (S.E.M.) over 3 random splits and three random seeds. **Right:** Harmonic mean of seen-unseen accuracy for AO-CLEVR on a range of 20% to 70% unseen classes. To reduce visual clutter, error bars are shown only for our *Causal* method and for a vanilla baseline (*LE*).

**Zappos:** Our approach improves the Unseen and Harmonic metrics. For the “Closed” and “AUSUC” metrics it loses compared to TMN. We note that results on the Closed metric are less interesting from a causal point of view: A model cannot easily rely on the knowledge of which attribute-object pairs tend to appear in the test set.

For both AO-CLEVR and Zappos, the independence loss improves recognition on unseen pairs but hurts recognition of seen pairs. This is a known and important trade-off when learning models that are robust for interventions [51]. The independence loss discourages certain types of correlations, hence models *do not* benefit from them when the test and train distributions are identical (seen-pairs). However, the loss is constructed in such a way that these are exactly the correlations that fail to hold once the test distribution changes (unseen-pairs). Thus, ignoring these correlations improves performance for samples of unseen-pairs.

## 9 Discussion

We present a new causal perspective on the problem of recognizing new attribute-object combinations in images. We propose to view inference in this setup as answering the question “which intervention on attribute and object caused the image”. This perspective gives rise to a causal-inspired embedding model. The model learns disentangled representations of attributes and objects although they are dependent in the training data. It provides better accuracy on two benchmark datasets.

The trade-off between seen accuracy and unseen accuracy reflects the fact that prior information about co-occurrence of attributes and objects in the training set is useful and predictive. A related problem has been studied in the setting of (non-compositional) generalized zero-shot learning [8]. We suggest that some of these techniques could be adapted to the compositional setting.

Several aspects of the model can be further extended by relaxing its assumptions. First, the assumption that image features are normally distributed may be limiting, and alternative ways to model this distribution may improve the model accuracy. Second, the model is premised on the prior knowledge that the attributes and objects have distinct and stable generation processes. However, this prior knowledge may not always be available, or some attribute may not have an obvious physical meaning. E.g. “cute”, “comfortable” or “sporty”, and in a multi-label setting [47] it is challenging to reveal what are the independent generation mechanisms themselves from confounded training data.

This paper focused on the case where attributes and objects are fully disentangled. Clearly, in natural language, many attributes and object are used in a way are that makes them dependent. For example, white wine is actually yellow, and the attribute a small bear is larger than a large bird. It remains an interesting question to extend the fully disentangled case to learn specific dependencies while leveraging the power of disentangled representations.

## Broader Impact

Compositional generalization, the key challenge of this work, is critical for learning in real-world domains where the long tail of new combinations dominates the distribution, like in vision-and-language tasks or for the perception modules of autonomous driving.

A causality-based approach, like the one we propose, may allow vision systems to make more robust inference, and debias correlations that naturally exist in the training data, allowing to use vision systems in complex environments where the distribution of labels and their combinations is varying. It has been shown in the past that vision systems may emphasize biases in the data, and the ideas put forward in the paper may help make systems more robust to such biases.

Such approach may be useful for improving fairness in various applications, for example by providing a more balanced visual recognition of individuals from minority groups.

## Acknowledgements

We thank Daniel Greenfeld, Idan Schwartz, Eli Meirom, Haggai Maron, Lior Bracha and Ohad Amosy for their helpful feedback on the early version. Uri Shalit was partially supported by the Israel Science Foundation (grant No. 1950/19).

## Funding Disclosure

Uri Shalit was partially supported by the Israel Science Foundation. Yuval Atzmon was supported by the Israel Science Foundation and Bar-Ilan University during his Ph.D. studies.

## References

- [1] Aishwarya Agrawal, Aniruddha Kembhavi, Dhruv Batra, and Devi Parikh. C-vqa: A compositional split of the visual question answering (vqa) v1.0 dataset. *arXiv preprint arXiv:1704.08243*, 2017.
- [2] Aishwarya Agrawal, Jiasen Lu, Stanislaw Antol, Margaret Mitchell, C Lawrence Zitnick, Devi Parikh, and Dhruv Batra. Vqa: Visual question answering. *International Journal of Computer Vision*, 123(1):4–31, 2017.
- [3] Kei Akuzawa, Yusuke Iwasawa, and Yutaka Matsuo. Adversarial invariant feature learning with accuracy constraint for domain generalization. *Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD)*, 2019.
- [4] J. Andreas, M. Rohrbach, T. Darrell, and Klein D. Learning to compose neural networks for question answering. In *NAACL*, 2016.
- [5] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.
- [6] Yuval Atzmon, Jonathan Berant, Vahid Kezami, Amir Globerson, and Gal Chechik. Learning to generalize to new compositions in image understanding. *arXiv preprint arXiv:1608.07639*, 2016.
- [7] Yuval Atzmon and Gal Chechik. Probabilistic and-or attribute grouping for zero-shot learning. In *Proceedings of the Thirty-Forth Conference on Uncertainty in Artificial Intelligence*, 2018.
- [8] Yuval Atzmon and Gal Chechik. Adaptive confidence smoothing for generalized zero-shot learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [9] Dzmitry Bahdanau, Harm de Vries, Timothy J. O’Donnell, Shikhar Murty, Philippe Beaudoin, Yoshua Bengio, and Aaron Courville. Closure: Assessing systematic generalization of clevr models, 2019.
- [10] Yoshua Bengio, Tristan Deleu, Nasim Rahaman, Nan Rosemary Ke, Sebastien Lachapelle, Olexa Bilaniuk, Anirudh Goyal, and Christopher Pal. A meta-transfer objective for learning to disentangle causal mechanisms. In *International Conference on Learning Representations*, 2020.
- [11] Christopher P Burgess, Irina Higgins, Arka Pal, Loic Matthey, Nick Watters, Guillaume Desjardins, and Alexander Lerchner. Understanding disentangling in *beta*-vae. *arXiv preprint arXiv:1804.03599*, 2018.
- [12] R. Chao, S. Changpinyo, B. Gong, and Sha F. An empirical study and analysis of generalized zero-shot learning for object recognition in the wild. In *ICCV*, 2016.
- [13] Tian Qi Chen, Xuechen Li, Roger B Grosse, and David K Duvenaud. Isolating sources of disentanglement in variational autoencoders. In *Advances in Neural Information Processing Systems*, pages 2610–2620, 2018.
- [14] Pim de Haan, Dinesh Jayaraman, and Sergey Levine. Causal confusion in imitation learning. In *Advances in Neural Information Processing Systems 32*, pages 11698–11709. 2019.

- [15] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1):2096–2030, 2016.
- [16] Weihao Gao, Sreeram Kannan, Sewoong Oh, and Pramod Viswanath. Conditional dependence via shannon capacity: Axioms, estimators and applications. *arXiv preprint arXiv:1602.03476*, 2016.
- [17] Mingming Gong, Kun Zhang, Tongliang Liu, Dacheng Tao, Clark Glymour, and Bernhard Schölkopf. Domain adaptation with conditional transferable components. In *International conference on machine learning*, pages 2839–2848, 2016.
- [18] Daniel Greenfeld and Uri Shalit. Robust learning with the hilbert-schmidt independence criterion. *arXiv preprint arXiv:1910.00270*, 2019.
- [19] Arthur Gretton, Olivier Bousquet, Alex Smola, and Bernhard Schölkopf. Measuring statistical dependence with hilbert-schmidt norms. In *International conference on algorithmic learning theory*, pages 63–77. Springer, 2005.
- [20] Arthur Gretton, Kenji Fukumizu, Choon H Teo, Le Song, Bernhard Schölkopf, and Alex J Smola. A kernel statistical test of independence. In *Advances in neural information processing systems*, pages 585–592, 2008.
- [21] Christina Heinze-Deml and Nicolai Meinshausen. Conditional variance penalties and domain shift robustness. *arXiv preprint arXiv:1710.11469*, 2017.
- [22] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. *ICLR*, 2017.
- [23] Phillip Isola, Joseph J Lim, and Edward H Adelson. Discovering states and transformations in image collections. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1383–1391, 2015.
- [24] D. Jayaraman, F. Sha, and K. Grauman. Decorrelating semantic visual attributes by resisting the urge to share. In *CVPR, CVPR '14*, 2014.
- [25] Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *CVPR*, 2017.
- [26] Keizo Kato, Yin Li, and Abhinav Gupta. Compositional learning for human object interaction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 234–251, 2018.
- [27] Nan Rosemary Ke, Olexa Bilaniuk, Anirudh Goyal, Stefan Bauer, Hugo Larochelle, Chris Pal, and Yoshua Bengio. Learning neural causal models from unknown interventions. *arXiv preprint arXiv:1910.01075*, 2019.
- [28] Niki Kilbertus, Giambattista Parascandolo, and Bernhard Schölkopf. Generalization in anti-causal learning. *arXiv preprint arXiv:1812.00524*, 2018.
- [29] Byungju Kim, Hyunwoo Kim, Kyungsu Kim, Sungjin Kim, and Junmo Kim. Learning not to learn: Training deep neural networks with biased data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9012–9020, 2019.
- [30] D.P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- [31] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *ICLR*, 2014.
- [32] Murat Kocaoglu, Christopher Snyder, Alexandros G. Dimakis, and Sriram Vishwanath. CausalGAN: Learning causal implicit generative models with adversarial training. In *International Conference on Learning Representations*, 2018.
- [33] B. Lake. *Towards more human-like concept learning in machines: compositionality, causality, and learning-to-learn*. Massachusetts Institute of Technology, 2014.
- [34] B. Lake, T. Ullman, J. Tenenbaum, and Gershman S. Building machines that learn and think like people. In *Behavioral and Brain Sciences*, 2017.
- [35] C.H. Lampert, H. Nickisch, and S. Harmeling. Learning to detect unseen object classes by between-class attribute transfer. In *CVPR*, 2009.
- [36] Ya Li, Xinmei Tian, Mingming Gong, Yajing Liu, Tongliang Liu, Kun Zhang, and Dacheng Tao. Deep domain generalization via conditional invariant adversarial networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 624–639, 2018.
- [37] Yong-Lu Li, Yue Xu, Xiaohan Mao, and Cewu Lu. Symmetry and group in attribute-object compositions. *CVPR*, 2020.
- [38] Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Rätsch, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem. Challenging common assumptions in the unsupervised learning of disentangled representations. *arXiv preprint arXiv:1811.12359*, 2018.
- [39] Francesco Locatello, Michael Tschannen, Stefan Bauer, Gunnar Rätsch, Bernhard Schölkopf, and Olivier Bachem. Disentangling factors of variation using few labels. *ICLR*, 2020.
- [40] David Lopez-Paz, Robert Nishihara, Soumith Chintala, Bernhard Schölkopf, and Léon Bottou. Discovering causal signals in images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6979–6987, 2017.

- [41] Emile Mathieu, Tom Rainforth, N Siddharth, and Yee Whye Teh. Disentangling disentanglement in variational autoencoders. *ICML*, 2019.
- [42] Ishan Misra, Abhinav Gupta, and Martial Hebert. From red wine to red tomato: Composition with context. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1792–1801, 2017.
- [43] Tanmoy Mukherjee, Makoto Yamada, and Timothy M Hospedales. Deep matching autoencoders. *arXiv preprint arXiv:1711.06047*, 2017.
- [44] Tushar Nagarajan and Kristen Grauman. Attributes as operators: factorizing unseen attribute-object compositions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 169–185, 2018.
- [45] Tushar Nagarajan and Kristen Grauman. Attributes as operators: factorizing unseen attribute-object compositions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 169–185, 2018.
- [46] Zhixiong Nan, Yang Liu, Nanning Zheng, and Song-Chun Zhu. Recognizing unseen attribute-object pair with generative model. In *AAAI*, volume 33, 2019.
- [47] Genevieve Patterson and James Hays. Coco attributes: Attributes for people, animals, and objects. *European Conference on Computer Vision*, 2016.
- [48] Tzaf Paz-Argaman, Yuval Atzmon, Gal Chechik, and Reut Tsarfaty. Zest: Zero-shot learning from text descriptions using textual similarity and visual summarization, 2020.
- [49] Judea Pearl. *Causality: models, reasoning and inference*, volume 29. Springer, 2000.
- [50] Senthil Purushwalkam, Maximilian Nickel, Abhinav Gupta, and Marc’Aurelio Ranzato. Task-driven modular networks for zero-shot compositional learning. In *ICCV*, 2019.
- [51] Dominik Rothenhäusler, Nicolai Meinshausen, Peter Bühlmann, and Jonas Peters. Anchor regression: heterogeneous data meets causality. *arXiv preprint arXiv:1801.06229*, 2018.
- [52] Paul K Rubenstein, Bernhard Schölkopf, and Ilya Tolstikhin. Learning disentangled representations with wasserstein auto-encoders. *ICLR workshop*, 2018.
- [53] Eduardo Hugo Sanchez, Mathieu Serrurier, and Mathias Ortner. Learning disentangled representations via mutual information estimation. *arXiv preprint arXiv:1912.03915*, 2019.
- [54] Bernhard Schölkopf. Causality for machine learning. *arXiv preprint arXiv:1911.10500*, 2019.
- [55] Bernhard Schölkopf, Dominik Janzing, Jonas Peters, Eleni Sgouritsa, Kun Zhang, and Joris Mooij. On causal and anticausal learning. In *Proceedings of the 29th International Conference on Machine Learning*, pages 459–466. Omnipress, 2012.
- [56] Le Song, Alex Smola, Arthur Gretton, Justin Bedo, and Karsten Borgwardt. Feature selection via dependence maximization. *Journal of Machine Learning Research*, 13(May):1393–1434, 2012.
- [57] Raphael Suter, Đorđe Miladinović, Bernhard Schölkopf, and Stefan Bauer. Robustly disentangled causal mechanisms: Validating deep representations for interventional robustness. *arXiv preprint arXiv:1811.00007*, 2018.
- [58] Raphael Suter, Djordje Miladinovic, Bernhard Schölkopf, and Stefan Bauer. Robustly disentangled causal mechanisms: Validating deep representations for interventional robustness. In *ICML*, 2019.
- [59] Yao-Hung Hubert Tsai and Ruslan Salakhutdinov. Improving one-shot learning through fusing side information. *arXiv preprint arXiv:1710.08347*, 2017.
- [60] Kun Wei, Muli Yang, Hao Wang, Cheng Deng, and Xianglong Liu. Adversarial fine-grained composition learning for unseen attribute-object recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3741–3749, 2019.
- [61] Y. Xian, T. Lorenz, B. Schiele, and Z. Akata. Feature generating networks for zero-shot learning. In *CVPR*, 2018.
- [62] Y. Xian, B. Schiele, and Z. Akata. Zero-shot learning - the good, the bad and the ugly. In *CVPR*, 2017.
- [63] A. Yu and K. Grauman. Fine-grained visual comparisons with local learning. In *Computer Vision and Pattern Recognition (CVPR)*, Jun 2014.
- [64] Kun Zhang, Bernhard Schölkopf, Krikamol Muandet, and Zhikun Wang. Domain adaptation under target and conditional shift. In *International Conference on Machine Learning*, pages 819–827, 2013.

# Supplementary Information: A causal view of compositional zero-shot recognition

## A Approximating $\operatorname{argmax}_{a,o} p(\mathbf{x}|a, o)$

The conditional likelihood  $p(\mathbf{x}|a, o)$  can be written by marginalizing over the latent factors  $\phi_a$  and  $\phi_o$

$$p(\mathbf{x}|a, o) = \iint_{\phi_a, \phi_o} p(\mathbf{x}, \phi_a, \phi_o|a, o) = \iint_{\phi_a, \phi_o} p(\mathbf{x}|\phi_a, \phi_o)p(\phi_a|a)p(\phi_o|o)d\phi_o d\phi_a. \quad (\text{S.7})$$

Computing this integral is hard in the general case. Variational autoencoders (VAEs) [31] approximate a similar integral by learning a smaller support using an auxiliary probability function  $Q$  over the latent space. Here we make another approximation, taking a "hard" approach and find the single most likely integrand.

$$\operatorname{argmax}_{(a,o) \in \mathcal{A} \times \mathcal{O}} p(\mathbf{x}|\phi_a, \phi_o)p(\phi_a|a)p(\phi_o|o) \quad (\text{S.8})$$

Based on Section 4, the three factors of the distributions are Gaussians. Therefore, maximizing Eq. (S.8) is equivalent to minimizing the negative log likelihood

$$\operatorname{argmin}_{(a,o) \in \mathcal{A} \times \mathcal{O}} \frac{1}{\sigma_x^2} \|\mathbf{x} - g(\phi_a, \phi_o)\|^2 + \frac{1}{\sigma_a^2} \|\phi_a - h_a\|^2 + \frac{1}{\sigma_o^2} \|\phi_o - h_o\|^2. \quad (\text{S.9})$$

This expression is composed of three components. The components allow to infer  $(a, o)$  by evaluating distances in three representation spaces  $\mathcal{X}$ ,  $\Phi_A$  and  $\Phi_O$  (Section 4).

However, we cannot apply Eq. (S.9) to infer  $(a, o)$  because the core features  $\phi_a, \phi_o$  are latent. Next, we introduce two additional approximations we use to apply Eq. (S.9). The first, approximates  $\|\mathbf{x} - g(\phi_a, \phi_o)\|^2$  by  $\|\mathbf{x} - g(h_a, h_o)\|^2$ , using a Taylor expansion at the means  $(\phi_a, \phi_o) = (h_a, h_o)$ . The second, recovers (infers) the core features  $\phi_a, \phi_o$  from the image, and substitute the recovered features in  $\|\phi_a - h_a\|^2, \|\phi_o - h_o\|^2$ .

### A.1 Approximating $\|\mathbf{x} - g(\phi_a, \phi_o)\|^2$

A causal model can be equivalently represented using a "Structural Causal Model" (SCM) [49]. An SCM matches a set of assignments to a causal graph. Each node in the graph is assigned a deterministic function of its parent nodes and an independent noise term. Specifically, based on the Gaussian assumptions in Section 4, the SCM of our causal graph (Figure 1a) is

$$\phi_a = n_a + h_a \quad (\text{S.10})$$

$$\phi_o = n_o + h_o \quad (\text{S.11})$$

$$\mathbf{x} = n_x + g(\phi_a, \phi_o), \quad (\text{S.12})$$

where  $n_a, n_o, n_x$  are jointly independent Gaussian random variables  $n_a \sim \mathcal{N}(0, \sigma_a^2 I)$ ,  $n_o \sim \mathcal{N}(0, \sigma_o^2 I)$ ,  $n_x \sim \mathcal{N}(0, \sigma_x^2 I)$ .  $n_a, n_o, n_x$  represent sampling from the manifold of attributes, objects and images near their prototypes  $h_a$  and  $h_o$ .

We use a zeroth-order Taylor expansion of  $g(\phi_a, \phi_o)$  at  $(\phi_a, \phi_o) = (h_a, h_o)$  and make the following approximation

$$(\mathbf{x} - g(\phi_a, \phi_o)) \approx (\mathbf{x} - g(h_a, h_o)) \quad (\text{S.13})$$

Next, we discuss the first-order *approximation error*. For brevity, we denote with  $\phi_{ao}$  the concatenation of the elements in  $(\phi_a, \phi_o)$  into a single vector. Similarly  $(h_a, h_o)$  into  $h_{ao}$  and  $(n_a, n_o)$  into  $n_{ao}$ .

We approximate  $g(\phi_{ao})$  by a first-order Taylor expansion at  $\phi_{ao} = h_{ao}$  to

$$g(\phi_{ao}) \approx g(h_{ao}) + [Jg](h_{ao}) \cdot (\phi_{ao} - h_{ao}) = g(h_{ao}) + [Jg](h_{ao}) \cdot n_{ao}, \quad (\text{S.14})$$

where  $[Jg]$  is the Jacobian of  $g$  and the last equality stems from the SCM (Eqs. S.10, S.11).

Using Eq. (S.14) and the Cauchy-Schwarz inequality, the first-order squared error approximation of Eq. (S.13) is:

$$\begin{aligned} \mathbb{E} \|\mathbf{x} - g(\phi_{ao}) - (\mathbf{x} - g(h_{ao}))\|^2 &= \mathbb{E} \|g(\phi_{ao}) - g(h_{ao})\|^2 \\ &\approx \mathbb{E} \|[Jg](h_{ao})n_{ao}\|^2 \leq \|[Jg](h_{ao})\|_F^2 \mathbb{E} \|n_{ao}\|^2 \end{aligned} \quad (\text{S.15})$$

This implies that the error of the approximation Eq. (S.13) is mainly dominated by the gradients of  $g$  at  $h_{ao}$ , and the variance of  $n_{ao}$ . If the gradients and the variance of  $n_{ao}$  are too large, then this approximation may be too coarse, and one may resort to more complex models like variational methods [31]. Empirically, we observe that this approximation is useful.

## A.2 Recovering the core features

To apply Eq. (S.9), we approximate  $\phi_a, \phi_o$ , by reconstructing (inferring) them from the image.

The main assumption we make is that images preserve the information about the core features they were generated from, at least to a level that allows us to differentiate what were the semantic prototypes ( $h_a$  and  $h_o$ ) of the core feature. Otherwise, inference on test data and labeling the training data by human raters will render to random guessing.

Therefore, we assume that there exists an inverse mappings that can infer  $\phi_a$  and  $\phi_o$  from the image up to a reasonably small error:

$$\hat{\phi}_a \equiv g_A^{-1}(\mathbf{x}) = \phi_a + \epsilon_A(x) \quad (\text{S.16})$$

$$\hat{\phi}_o \equiv g_O^{-1}(\mathbf{x}) = \phi_o + \epsilon_O(x), \quad (\text{S.17})$$

where  $\epsilon_A(x), \epsilon_O(x)$  denote the image-based error of inferring the attribute and object core features.

Specifically, we assume that the error in substituting  $\phi_a$  by  $\hat{\phi}_a$  and  $\phi_o$  by  $\hat{\phi}_o$  in Eq. (S.9) is small enough to keep them close to their prototypes ( $h_a$  and  $h_o$ ). Namely, we make the following approximation

$$\begin{aligned} \|\phi_a - h_a\|^2 &= \|\hat{\phi}_a - \epsilon_A(x) - h_a\|^2 \leq \|\hat{\phi}_a - h_a\|^2 + \|\epsilon_A(x)\|^2 \approx \|\hat{\phi}_a - h_a\|^2 \\ \|\phi_o - h_o\|^2 &= \|\hat{\phi}_o - \epsilon_O(x) - h_o\|^2 \leq \|\hat{\phi}_o - h_o\|^2 + \|\epsilon_O(x)\|^2 \approx \|\hat{\phi}_o - h_o\|^2. \end{aligned} \quad (\text{S.18})$$

To conclude, we use Eq. (S.13) and Eq. (S.18) to approximate Eq. (S.9) by :

$$\operatorname{argmin}_{(a,o) \in \mathcal{A} \times \mathcal{O}} \frac{1}{\sigma_x^2} \|\mathbf{x} - g(h_a, h_o)\|^2 + \frac{1}{\sigma_a^2} \|\hat{\phi}_a - h_a\|^2 + \frac{1}{\sigma_o^2} \|\hat{\phi}_o - h_o\|^2,$$

which is the expression for Eq. (3) in the main paper.

## B Independence Loss

Our loss includes a component  $\mathcal{L}_{indep}$ , which is designed to capture the conditional-independence relations that the causal graph dictates (Eq. (6)). We now describe  $\mathcal{L}_{indep}$  in detail.

Since we do not have the actual values of the latent core features  $\phi_a, \phi_o$ , we wish that their reconstructions  $\hat{\phi}_a$  and  $\hat{\phi}_o$  maintain approximately the same independence relations as Eq. (6). To learn mappings that adhere to these statistical independences over  $\hat{\phi}_a$  and  $\hat{\phi}_o$ , we regularize the learned mappings using a differentiable measure of statistical dependence.

Specifically, we use a positive differentiable measure of the statistical dependence, denoted by  $\mathcal{I}$ . For two variables  $(u, v)$ , conditioned on a categorical variable  $Y$ , we denote by  $\mathcal{I}(u, v|Y)$  the positive differentiable measure of the statistical conditional dependence of  $(u, v|Y)$ . For example we encourage approaching the equality in Eq. 6b by minimizing  $\mathcal{I}(\hat{\phi}_a, \hat{\phi}_o|A)$ .  $\mathcal{I}$  is based on measuring the Hilbert-Schmidt Information Criterion (HSIC) [19, 20], which is a non-parametric method for estimating the statistical dependence between samples of two variables. We adapt the HSIC criterion to measure *conditional* dependencies.  $\mathcal{I}$  penalizes conditional dependencies in a batch of samples  $B = \{(u_i, v_i, y_i)\}_{i=1}^{|B|}$ , by summing over groups of samples that have the same label:

$$\mathcal{I}(u, v|Y) = \frac{1}{|Y|} \sum_{y \in Y} \text{HSIC}(U|Y = y, V|Y = y) \quad (\text{S.19})$$

where  $(U|Y = y, V|Y = y) = \{(u_i, v_i) \in B \mid y_i = y\}$ .

Finally, we have  $\mathcal{L}_{indep}$ , a loss term that encourages the four conditional independence relations of 6:

$$\mathcal{L}_{indep} = \mathcal{L}_{oh} + \lambda_{rep} \mathcal{L}_{rep} \quad (\text{S.20})$$

$$\mathcal{L}_{oh} = \mathcal{I}(\hat{\phi}_a, O|A) + \mathcal{I}(\hat{\phi}_o, A|O) \quad (\text{S.21})$$

$$\mathcal{L}_{rep} = \mathcal{I}(\hat{\phi}_a, \hat{\phi}_o|A) + \mathcal{I}(\hat{\phi}_a, \hat{\phi}_o|O), \quad (\text{S.22})$$

where  $\lambda_{rep}$  is a hyper parameter.

Minimizing Eq. (S.21) encourages the representation of the inferred attribute  $\hat{\phi}_a$ , to be invariant to the categorical (“one-hot”) representation of an object  $O$ . Minimizing Eq. (S.22) encourages  $\hat{\phi}_a$  to be invariant to the “appearance” of an object ( $\hat{\phi}_o$ ).

### B.1 An expression for Hilbert-Schmidt Information Criterion (HSIC) with linear kernel

The linear-kernel HSIC between two batches of vectors  $\mathbf{U}, \mathbf{V}$  is calculated in the following manner [20]: It uses two linear kernel matrices  $K_{i,j} = \mathbf{u}_{[i,:]} \mathbf{u}_{[j,:]}^T$ ,  $L_{i,j} = \mathbf{v}_{[i,:]} \mathbf{v}_{[j,:]}^T$ . Then the HSIC is calculated as the scaled Hilbert-Schmidt norm of their cross-covariance matrix:

$$\text{HSIC}(U, V) = \frac{1}{(n-1)^2} \cdot \text{tr}(KHLH),$$

where  $H_{ij} = \delta_{i,j} - \frac{1}{n}$  is a centering matrix, and  $n$  is the batch size

### B.2 Minimizing $\mathcal{L}_{indep}$ encourages robustness of the reconstructed core features $\hat{\phi}_a, \hat{\phi}_o$ .

The conditional-independence term of  $\mathcal{L}_{oh}$  within  $\mathcal{L}_{indep}$  is related to a metric named “Post Interventional Disagreement” (PIDA), recently introduced by [58]. PIDA measures disentanglement of representations for models that are trained from unsupervised data. This section explains their relation in more detail, by showing that minimizing  $\mathcal{L}_{indep}$  encourages the following properties:  $p^{do(O=o)}(\hat{\phi}_o) \approx p^{do(A=a, O=o)}(\hat{\phi}_o)$  and  $p^{do(A=a)}(\hat{\phi}_a) \approx p^{do(A=a, O=o)}(\hat{\phi}_a)$ .

The PIDA metric for attributes is measured by

$$\text{PIDA}(a'|a, o) := d(\mathbb{E}^{do(a)}[\hat{\phi}_{a'}], \mathbb{E}^{do(a,o)}[\hat{\phi}_{a'}]), \quad (\text{S.23})$$

where  $d$  is loosely described as “a suitable” positive distance function (like the  $L_2$  distance).  $\text{PIDA}(a'|a, o)$  quantifies the shifts in the inferred features  $\hat{\phi}_{a'}$  when the object is enforced to  $o$ . Similarly, the PIDA term for objects is  $\text{PIDA}(o'|o, a)$ .

Below we show that encouraging the conditional independence ( $\hat{\phi}_a \perp\!\!\!\perp O|A$ ), is equivalent to minimizing  $d(p^{do(a)}(\hat{\phi}_{a'}), p^{do(a,o)}(\hat{\phi}_{a'}))$  and therefore minimizes  $\text{PIDA}(a'|a, o)$ .

First, in our causal graph (Figure 1a) a do-intervention on both  $a$  and  $o$  is equivalent to conditioning on  $(a, o)$

$$p^{do(a,o)}(\hat{\phi}_{a'}) = p(\hat{\phi}_{a'}|a, o). \quad (\text{S.24})$$

Second, minimizing Eq. (S.21) encourages the conditional independence ( $\hat{\phi}_{a'} \perp\!\!\!\perp O|A$ ), which is equivalent to minimization of  $d(p(\hat{\phi}_{a'}|a), p(\hat{\phi}_{a'}|a, o))$ <sup>2</sup>. Third, when  $d(p(\hat{\phi}_a|a), p(\hat{\phi}_a|a, o))$  approaches zero, then  $p^{do(a)}(\hat{\phi}_a)$  approaches  $p(\hat{\phi}_a|a)$ . This stems from the adjustment formula,  $p^{do(a)}(\hat{\phi}_{a'}) = \sum_o p(\hat{\phi}_{a'}|a, o)p(o)$ : When  $d(p(\hat{\phi}_{a'}|a), p(\hat{\phi}_{a'}|a, o))$  approaches zero, then

<sup>2</sup>When ( $\hat{\phi}_{a'} \perp\!\!\!\perp O|A$ ), then by definition  $p(\hat{\phi}_{a'}|a, o) = p(\hat{\phi}_{a'}|a)$ . Therefore encouraging ( $\hat{\phi}_{a'} \perp\!\!\!\perp O|A$ ), makes  $p(\hat{\phi}_{a'}|a, o)$  approach  $p(\hat{\phi}_{a'}|a)$

$\sum_o p(\hat{\phi}_{a'}|a, o)p(o)$  approaches  $\sum_o p(\hat{\phi}_{a'}|a)p(o) = p(\hat{\phi}_{a'}|a)$ . Therefore  $p^{do(a)}(\hat{\phi}_{a'})$  approaches  $p(\hat{\phi}_{a'}|a)$ .

As a result, we have the following: Minimizing Eq. (S.21) leads to  $p^{do(a,o)}(\hat{\phi}_{a'})$  approaching  $p(\hat{\phi}_{a'}|a)$ , which as we have just shown, leads to  $p(\hat{\phi}_{a'}|a)$  approaching  $p^{do(a)}(\hat{\phi}_{a'})$ . Therefore,  $d(p^{do(a)}(\hat{\phi}_{a'}), p^{do(a,o)}(\hat{\phi}_{a'}))$  is minimized and accordingly,  $PIDA(a'|a, o)$  (Eq. S.23) is minimized.

Similarly, for objects, encouraging the conditional independence  $(\hat{\phi}_{o'} \perp\!\!\!\perp A|O)$  minimizes  $PIDA(o'|a, o)$ . Therefore, minimizing  $\mathcal{L}_{oh}$  (Eq. S.21), optimizes both  $PIDA(a'|a, o)$  and  $PIDA(o'|a, o)$ .

## C Implementation details

### C.1 Architecture

Similar to LE [44], we implemented  $g$ ,  $h_a$  and  $h_o$  by MLPs with ReLU activation. For  $g_A^{-1}$  and  $g_O^{-1}$ , every layer used a batch-norm and leaky-relu activation. All the MLPs share the same size of hidden layer, denoted by  $d_h$ .

For experiments on Zappos, we also learned a single layer network  $f$  to project pretrained image features to the feature space  $\mathcal{X}$ . This strategy was inspired by the baseline models LE and ATTOP. Learning a projection  $f$  finds better solutions on the validation set than using the pretrained features as  $\mathcal{X}$ .

For HSIC we used the implementation of [18] and applied a linear kernel as it does not require tuning additional hyper-parameters.

### C.2 Optimization

**AO-CLEVR:** We optimized AO-CLEVR in an alternating fashion: First we trained  $h_a$ ,  $g_A^{-1}$  and  $g$ , keeping  $h_o$ ,  $g_O^{-1}$  frozen. Then we froze  $h_a$ ,  $g_A^{-1}$  and trained  $h_o$ ,  $g_O^{-1}$  and  $g$ . This optimization strategy allows to stabilize the attribute representation when minimizing Eq. (S.22). The strategy was developed during the early experiments with a low-dimensional ( $\mathcal{X} \subset \mathbb{R}^2$ ) synthetic dataset. In the ablation study (Section F below), we show that a standard (non-alternating) optimization strategy achieves comparable results, but with somewhat higher bias toward seen accuracy.

We used SGD with Nesterov Momentum to train with AO-CLEVR. Empirically, we found that SGD allowed finer control over  $\mathcal{L}_{indep}$  than Adam [30].

In practice, we weighed the loss of  $\|\hat{\phi}_a - h_a\|^2$  and  $\|\hat{\phi}_o - h_o\|^2$  according to the respective attribute and object frequencies in the training set. This detail has a relatively small effect on performance. Without weighing the loss, the Harmonic decreases by 1.2% (from 68.8% to 67.6%) and Unseen accuracy decreases by 1.3% (from 57.5% to 56.2%).

**Zappos:** We optimized Zappos in a standard (non-alternating) fashion. We couldn't use the *alternating* optimization strategy, because in Zappos we also learn a mapping  $f$  that projects pretrained image features to the feature space  $\mathcal{X}$ . Thus, updating the parameters of  $f$  changes the mapping to  $\mathcal{X}$  and we cannot keep  $h_a$ ,  $g_A^{-1}$  frozen once  $\mathcal{X}$  changes.

As in [44], we used Adam [30] to train with Zappos.

### C.3 Early Stopping and Hyper-parameter selection

We trained each model instantiation for a maximum of 1000 epochs and early stopped on the validation set.

We used two metrics for early-stopping and hyper-parameter selection on the validation set: (i) the Harmonic metric for testing the unseen-accuracy, the seen-accuracy and the Harmonic; and (ii) accuracy of the Closed metric for testing the Closed accuracy. In Zappos, we followed [50] and used the AUSUC for testing both the AUSUC metric and the Closed accuracy.

For our approach and all the compared methods, we tuned the hyper-params by first taking a coarse random search, and then further searching around the best performing values on the validation set. As a rule of thumb, we first stabilized the hyper-parameters of the learning-rate, weight-decay, and architecture. Then we searched in finer detail over the hyper-parameters of the loss functions. At the most fine-grained iteration of the random search, each combination of hyper-parameters was evaluated with 3 different random weight initializations, and metrics were averaged over 3 runs. We chose the set of hyper-parameters that maximized the average metric of interest.

Since AO-CLEVR has  $6 \cdot 3 = 18$  different splits, we searched the hyper-parameters over a single split of each of the ratios  $\{2:8, 5:5, 6:4, 7:3\}$ . For  $\{3:7, 4:6\}$  ratios, we used the hyper-parameters chosen for the  $\{5:5\}$  split.

#### C.4 Hyper-parameters for loss function

In practice, to weigh the terms of the loss function, we use an equivalent but different formulation. Specifically, we set  $\lambda_{indep}=1$  and  $\lambda_{invert}=1$ , and use the following expressions for  $\mathcal{L}_{indep}$  and  $\mathcal{L}_{invert}$

$$\mathcal{L}_{indep} = \lambda_{oh}\mathcal{L}_{oh} + \lambda_{rep}\mathcal{L}_{rep} \quad (\text{S.25})$$

$$\begin{aligned} \mathcal{L}_{invert} = & \lambda_{icore}(CE(a, f_a(h_a)) + CE(o, f_o(h_o))) + \\ & \lambda_{ig}(CE(a, f_{ga}(g(h_a, h_o))) + CE(o, f_{go}(g(h_a, h_o))))). \end{aligned} \quad (\text{S.26})$$

where  $\lambda_{oh}$ ,  $\lambda_{rep}$ ,  $\lambda_{icore}$  and  $\lambda_{ig}$  weigh the respective loss terms.

#### C.5 Grid-search ranges

For *Causal* with AO-CLEVR, we started the random grid-search over the following ranges: **Architecture:** (1) Number of layers for  $h_a$  and  $h_o \in \{0, 1, 2\}$ , (2) Number of layers for  $g_A^{-1}$  and  $g_O^{-1} \in \{1, 2, 3\}$ , (3) Number of layers for  $g \in \{2, 4\}$ , (4) Common size of hidden layers  $d_h \in \{10, 30, 150, 300, 1000\}^3$ . **Optimization:** (1) learning rate  $\in \{1e-5, 3e-5, 1e-4, 3e-4, 1e-3\}$ , when using alternate training, we used different learning rates for each alternation. (2) weight-decay  $\in \{1e-5, 1e-4, 1e-3, 0.01, 0.1, 1\}$  (3)  $\lambda_{rep} \in \{0, 0.001, 0.003, 0.01, 0.03, 0.1, 0.3, 1, 3, 10, 30, 100, 300\}$  (4)  $\lambda_{oh} \in \{0, 0.001, 0.003, 0.01, 0.03, 0.1, 0.3, 1, 3, 10, 30, 100, 300\}$  (5)  $\lambda_{ao} \in \{0, 0.1, 0.3, 1, 3, 10\}$  (6)  $\lambda_{icore} \in \{0, 0.01, 0.03, 0.1, 0.3, 1, 3, 10, 30, 100\}$  (7)  $\lambda_{ig} \in \{0, 0.01, 0.03, 0.1, 0.3, 1, 3, 10, 30, 100\}$ . We didn't tune batch-size, we set it to 2048.

For *Causal* with Zappos, we used some hyper-parameters found with Causal&AO-CLEVR and some already known for LE with Zappos. Specifically, for the architecture, we set (1) Number of hidden layers for  $h_a$  and  $h_o = 0$  (linear embedding), (2) Number of layers for  $g_A^{-1}$  and  $g_O^{-1} = 2$ , (3) Number of layers for  $g = 2$  (4) Common size of hidden layers  $d_h = 300$ . For the optimization, in order to find a solution around the solution used by LE, we selected  $\lambda_{ao} = 1000$ ,  $\lambda_{ig} = 0$ , weight-decay =  $5e-5$ , batch-size = 2048; and  $\lambda_{icore} = 100$  as with Causal&AO-CLEVR. We applied the random search protocol over  $\lambda_{rep} \in 15 \cdot \{0, 0.001, 0.003, 0.01, 0.03, 0.1, 0.3, 1, 3, 10, 30, \}$ ,  $\lambda_{oh} \in 15 \cdot \{0, 0.001, 0.003, 0.01, 0.03, 0.1, 0.3, 1, 3, 10, 30, \}$ , and learning rate  $\in \{1e-4, 3e-4\}$

For *TMN*, we applied the random grid-search according to the ranges defined in the supplemental of [50]. Specifically: lr  $\in \{0.0001, 0.001, 0.01, 0.1\}$ , lrg  $\in \{0.0001, 0.001, 0.01, 0.1\}$ , batch-size  $\in \{64, 128, 256, 512\}$ , concept-dropout  $\in \{0, 0.05, 0.1, 0.2\}$ , nmod  $\in \{12, 18, 24, 30\}$ , output-dimension  $\in \{8, 16\}$ , number-of-layers  $\in \{1, 2, 3, 5\}$ . Additionally, we trained TMN for a maximum of 30 epochs, which is  $\times 6$  longer than the recommended length (4-5 epochs). As instructed, we chose the number of negatives to be "all negatives". With Zappos, our grid-search found a hyper-parameters' combination with better performance than the one reported by the authors.

For *ATTOP* with AO-CLEVR, we used the following ranges:  $\lambda_{aux} \in \{0, 0.001, 0.003, 0.01, 0.03, 0.1, 0.3, 1, 3, 10, 30, 100\}$ ,  $\lambda_{comm} \in \{0, 0.001, 0.003, 0.01, 0.03, 0.1, 0.3, 1, 3, 10, 30, 100\}$ ,

<sup>3</sup>We used 150 and not 100, in order to have total size of  $2 \cdot 150 = 300$  for the concatenated representation of  $[\hat{\phi}_a, \hat{\phi}_o]$ . 300 is comparable to the default value for emb-dim in the LE baseline.

$\text{emb-dim} \in \{10, 30, 150, 300, 1000\}$ ,  $\text{weight-decay} \in \{1e-5, 1e-4, 1e-3, 1e-2\}$ . We used the default learning rate  $1e-4$ , and used a batch size of 2048.

For *ATTOP* with Zappos, we used the hyper-parameters combinations recommended by [44, 50], and also searched for  $\text{emb-dim} \in \{300, 1000\}$

For *LE*, we used the following ranges:  $\text{weight-decay} \in \{1e-4, 1e-3, 1e-2\}$ ,  $\text{emb-dim} \in \{10, 30, 150, 300, 1000\}$ . With Zappos we also followed the guideline of [50] and searched  $\text{lr} \in \{0.0001, 0.001\}$ .

Note that for *ATTOP* and *LE*, when Glove embedding is enabled then  $\text{emb-dim}$  is fixed to 300.

For *VisProd* with AO-CLEVR, we used the following ranges: We used the same number of layers, weight-decay and learning rates as used to train “Causal”. We also followed the alternate-training protocol. With Zappos, we searched for  $\text{emb-dim} \in \{100, 300, 1000\}$ ,  $\text{lr} \in \{1e-4, 1e-3\}$ .

### C.6 AO-CLEVR dataset

**Pretrained features:** Similar to Zappos, we extracted pretrained features for AO-CLEVR using a pretrained ResNet18 CNN.

**Cross validation splits:** For cross-validation, we used two types of splits. The first uses the same unseen pairs for validation and test. We call this split the “overlapping” split. The split allows us to quantify the potential generalization capability of each method. The second split, is harder, where unseen validation pairs are not overlapping with the unseen test pairs. We call this split the “non-overlapping” split.

For the overlapping split, we varied the ratio of unseen:seen pairs on a range of (2:8, 3:7, ... 7:3), and for each ratio we drew 3 random seen-unseen splits. For the non-overlapping split, we varied the ratio of unseen:seen pairs on a range of (2:6, 3:5, ... 5:3), and for each ratio we drew 3 random seen-unseen splits. In addition, we always draw 20% of the pairs for validation.

### D AO-CLEVR with a non-overlapping validation set.

Here we present results for AO-CLEVR with the “non-overlapping” split. For this split, the unseen validation pairs are not overlapping with the unseen test pairs. It is harder than the “overlapping” split, which uses the same unseen pairs for validation and test.

The non-overlapping split is important because, in practice, we cannot rely on having labeled samples of the unseen pairs for validation purposes.

Figure S.1 shows the measured metrics when comparing “Causal”, LE, TMN, and ATTOP and varying the ratio of seen:unseen pairs between 2:6 to 5:3.

For the main zero-shot metrics (Open-Unseen, Harmonic and Closed), our approach “Causal”, performs substantially better than the compared methods. ATTOP performs substantially worse on “seen” pairs.

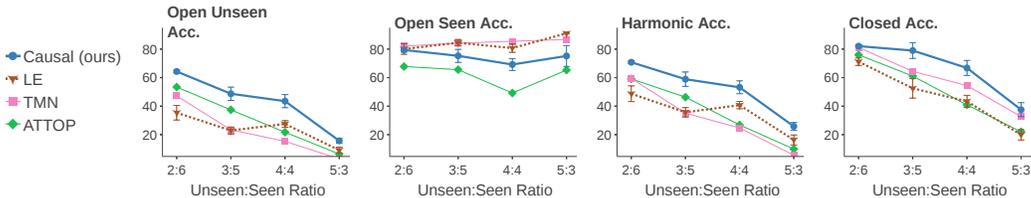


Figure S.1: AO clevr with a *non*-overlapping validation set.

### E Complete results for AO clevr with overlapping split

Figure S.2 shows the accuracy metrics for compared approaches with AO-CLEVR when varying the fraction of seen:unseen classes (between 2:8 to 7:3). The top row in the figure shows the measured

metrics. The bottom row shows the difference (subtraction) from LE. We selected LE as the main reference baseline because its embedding loss approximately models  $p(\mathbf{x}|a, o)$ , but without modeling the core-features.

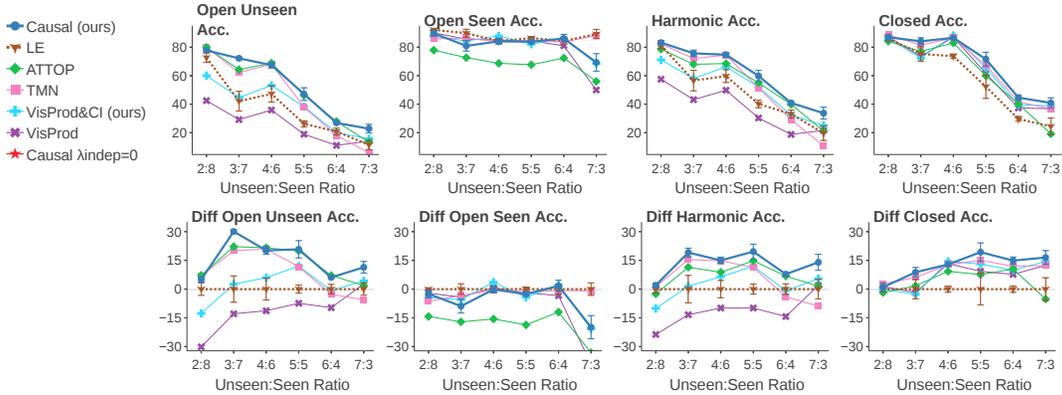


Figure S.2: Accuracy metrics for AO-CLEVR on a sweep of 20% unseen classes up to 70% unseen classes. The top row show the measured metrics. The bottom row show the difference (*subtraction*) of measured metrics from the LE baseline method. Error bars denote Standard Error of the Mean (S.E.M) over 3 random splits. To reduce visual clutter, error bars are shown only for our Causal method and for the reference baseline (LE).

Across the full sweep of unseen:seen ratios, our approach “Causal”, performs better than or equivalent to all the compared methods for the main zero-shot metrics (Open-Unseen, Harmonic and Closed). VisProd, which approximates  $p(a, o|\mathbf{x})$ , has a relatively low Unseen accuracy. VisProd&CI, the discriminatively-trained variant of our model, improves the Unseen performance by a large margin, while not hurting VisProd Seen accuracy. ATTOP is better than LE on open unseen pairs but performs substantially worse than all methods on the seen pairs. TMN performs equally well as our approach for splits with mostly seen pairs (unseen:seen @ 2:8, 3:7, 4:6), but degrades when the fraction of seen pairs reduces below 60%.

### E.1 Baseline models without language embeddings

Figure S.3 compares LE, ATTOP, and TMN with and without initialization by Glove word embedding. It demonstrates that for AO-CLEVR, Glove initialization somewhat hurts LE, improves ATTOP, and is mostly equivalent to TMN.

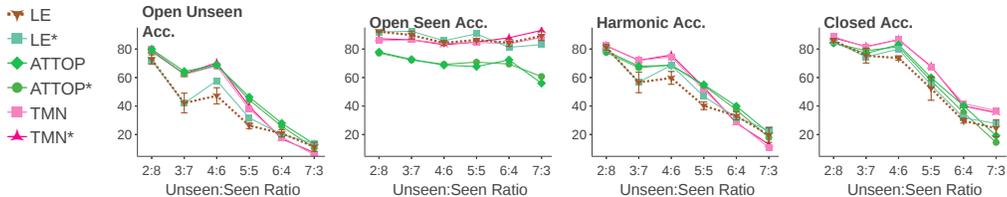


Figure S.3: Comparing LE, ATTOP, and TMN with and without initialization by Glove word embedding. Glove initialization somewhat hurts LE, improves ATTOP, and is mostly equivalent for TMN.

## F Ablation study

To understand the contribution of the different components of our approach, we conducted an ablation study to quantify the effect of the components. We report test metrics for one of the 5:5 “overlapping” splits of AO-CLEVR. Specifically, the split used for hyper-parameter search.

	UNSEEN	SEEN	HARMONIC	CLOSED		UNSEEN	SEEN	HARMONIC	CLOSED
<b>CAUSAL</b>					<b>WITH PRIOR EMBEDDINGS</b>				
$\lambda_{indep} = 0$	57.5 ± 2.3	85.7 ± 3.4	<b>68.8 ± 2.1</b>	73.8 ± 1.0	LE	21.4 ± 1.1	84.1 ± 1.8	34.0 ± 1.3	34.2 ± 2.4
$\lambda_{ao} = 0$	44.9 ± 1.9	83.5 ± 5.6	58.1 ± 0.3	72.9 ± 3.0	ATTOP	48.7 ± 0.5	73.5 ± 0.8	58.5 ± 0.1	58.2 ± 0.5
$\lambda_{invert} = 0$	19.5 ± 2.2	46.5 ± 2.9	27.3 ± 2.3	28.5 ± 4.6	TMN	32.3 ± 2.8	87.3 ± 4.1	47.0 ± 3.3	65.1 ± 3.0
<b>NON-ALTERNATE TRAINING</b>					<b>NO PRIOR EMBEDDINGS</b>				
<b>CAUSAL</b>	53.3 ± 2.0	90.6 ± 2.3	67.0 ± 1.5	<b>74.5 ± 1.2</b>	VISPROD	19.1 ± 1.3	94.3 ± 1.1	31.7 ± 1.8	60.1 ± 0.2
$\lambda_{rep} = 0$	53.0 ± 4.7	91.3 ± 1.4	66.7 ± 3.3	70.7 ± 1.3	LE*	28.2 ± 1.7	87.5 ± 0.5	42.5 ± 1.9	43.4 ± 2.7
$\lambda_{oh} = 0$	52.8 ± 4.0	90.5 ± 0.6	66.5 ± 3.0	71.9 ± 2.0	ATTOP*	45.6 ± 0.5	76.3 ± 1.0	57.0 ± 0.1	54.6 ± 1.3
$\lambda_{indep} = 0$	38.0 ± 2.3	94.6 ± 0.4	54.1 ± 2.4	68.9 ± 0.6	TMN*	36.6 ± 5.2	89.1 ± 3.5	51.6 ± 5.6	66.5 ± 4.3
					VisP&CI	40.5 ± 2.7	84.7 ± 4.3	54.4 ± 1.8	59.9 ± 0.3

Table S.1: **Left:** Ablation study on a 5:5 split of AO-CLEVR. We use the split used for hyper-param search.  $\pm$  denote S.E.M on 3 random initializations. **Right:** Reference metrics for baselines.

Table S.1 reports the test metrics when ablating different components of our approach: We first compared the different components of the model while using alternate-training (see implementation details). Next, we compared the alternate-training strategy to standard (non-alternate) training. Finally, we compared the different components of the conditional-independence loss.

We compared the following components:

1. “Causal” is our approach described in Section 4. We tested it with both alternate training, and standard (non-alternate) training.
2.  $\lambda_{indep} = 0$  indicates nullifying the loss terms that encourage the conditional independence relations.
3.  $\lambda_{ao} = 0$  indicates nullifying the embedding to the image space  $\mathcal{X}$ .
4.  $\lambda_{invert} = 0$  indicates nullifying the term that preserves information about source labels of the attribute and object embeddings.
5.  $\lambda_{rep} = 0$  indicates nullifying  $\mathcal{L}_{rep}$ : The term that encourages invariance between  $\hat{\phi}_a$  to  $\hat{\phi}_o$ .
6.  $\lambda_{oh} = 0$  indicates nullifying  $\mathcal{L}_{oh}$ : The term that encourages invariance of  $\hat{\phi}_a$  to a categorical representation of an object (and similarly for  $\hat{\phi}_o$ ).

First, we find that both alternate-training and standard non-alternate training result in a comparable Harmonic metric. However, alternate training has a better Unseen accuracy (57.5% vs 53.3%), but lower Seen accuracy (85.7% vs 90.6%)

Second, nullifying each of the major components of the loss has a substantial impact on the performance of the model. Specifically, (1) nullifying  $\lambda_{indep}$  reduces the Harmonic from 68.8% to 45.7%, (2) nullifying  $\lambda_{ao}$  reduces the Harmonic to 58.1%, (3) nullifying  $\lambda_{invert}$  reduces the Harmonic to 27.3%.

Finally,  $\mathcal{L}_{oh}$  and  $\mathcal{L}_{rep}$  have a synergistic effect on the performance of AO-CLEVR. Their individual performance metrics are comparable, but jointly they improve the Closed accuracy from ~71.5% to 74.5%.

**Accuracy versus hidden layer size:** Figure S.4 shows the different accuracy metrics when changing the hidden layer size for the 5:5 split of AO-CLEVR. It shows that the seen accuracy increases with the layer size, presumably because it can capture more subtleties of the seen pairs. The unseen accuracy appears to be bi-modal, with peaks at 15-units and 150 units. Indeed the cross-validation procedure selected a layer size of 150, because it maximized the harmonic-mean of the seen and unseen accuracy.

## G Error analysis

### G.1 Zappos

We analyzed the *errors* that *Causal* makes when recognizing unseen pairs in Zappos (open). In 53% of cases the object was predicted correctly, in 14% the attribute was predicted correctly and in

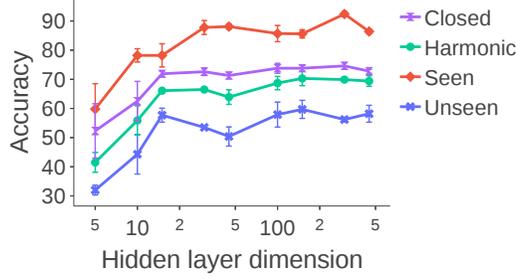


Figure S.4: Ablation: Accuracy versus hidden layer size (in a logarithmic scale)

33% neither. It appears that in Zappos, recognizing the object transferred better to the unseen pairs, presumably because recognizing the attributes is harder in this dataset.

To gain further intuition, we compared the errors that *Causal* makes to those of  $LE^*$ , the strongest no-prior baseline. With *Causal*, 39% of unseen pairs (U) are confused for seen pairs (S), and 36% of unseen pairs are confused for incorrect unseen-pairs. This yields a balanced rate of  $\frac{U \rightarrow S}{U \rightarrow U} = \frac{39\%}{36\%} \approx 1.1$ . For comparison  $LE^*$  errors are largely unbalanced:  $\frac{U \rightarrow S}{U \rightarrow U} = \frac{67\%}{19\%} \approx 3.5$ .



Figure S.5: Qualitative examples of *Causal* success cases for Zappos: Every image was predicted correctly by *Causal*, but erroneously by  $LE^*$ . Above every image we denote the erroneous prediction of  $LE^*$ . (a) An unseen pair: the error rate was 40% for *Causal*, while it was 82% for  $LE^*$ . (b) An unseen pair: the error rate was 39% for *Causal*, while it was 58% for  $LE^*$ . (c) A seen pair: the error rate was 56% for *Causal*, while it was 83% for  $LE^*$ .

We further analyzed wins and losses of *Causal* compared with  $LE^*$ , which are illustrated in Figure S.5. *Causal* succeeded to overcome common failures of  $LE^*$ , sometimes overcoming domain-shifts that exist within the seen pairs. The main weakness of *Causal* is that the error rate for seen pairs is higher compared to  $LE^*$ . This effect is strongest for (1) ( $A=Leather, O=Boots.Ankle$ ), which was mostly confused for either  $A=Full.grain.leather$  or  $A=Suede$ , while the object-class was correctly classified. (2) ( $A=Leather, O=Boots.Mid-Calf$ ), which was mostly confused for  $A=Faux.Leather$  while the object-class was correctly classified. This result shows that *Causal* is less biased toward predicting *Leather*, which is the most common attribute in the training set.

## G.2 AO-CLEVR

We analyzed the errors that *Causal* makes when recognizing unseen pairs at the 5:5 split<sup>4</sup> of AO-CLEVR (open). In 15% of cases the object was predicted correctly, in 82% the attribute was predicted correctly and in 3% neither. It appears that in this case, recognizing the attribute transferred better to the unseen pairs, because in this dataset colors can be easily recognized.

Comparing the errors that *Causal* makes to those of  $TMN^*$ , the strongest no-prior baseline show that with *Causal*, 19% of unseen pairs (U) are confused for seen pairs (S), and 20% of unseen-pairs are confused for incorrect unseen-pairs, a balanced rate of  $\frac{U \rightarrow S}{U \rightarrow U} = \frac{19\%}{20\%} \approx 0.9$ . However, with  $TMN^*$ , confusion is largely unbalanced:  $\frac{U \rightarrow S}{U \rightarrow U} = \frac{41\%}{13\%} \approx 3$ .

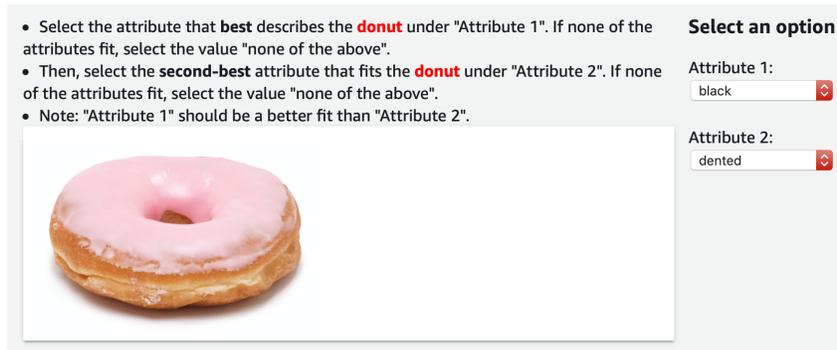


Figure S.6: An example of a task on Amazon Mechanical Turk.

## H Label-quality evaluation: Human-Rater Experiments

To better understand the level of label noise in *MIT-states* dataset we conducted an experiment with human raters.

Each rater was presented with an image of an object and was asked to select the best and second-best attributes that describe this object from a pre-defined list. The list was comprised of attributes that co-occur with the given object in the dataset. For example, the object “apple” had candidate attributes “green”, “yellow” and “red”, but not “black”. Raters were also presented with an option “none of the above” and an option “I don’t know”, see Figure S.6. We sampled 500 instances of objects and attributes, one from each attribute-object pair. As a label-quality score, we computed the balanced accuracy of rater responses compared with the label provided by the dataset across the 500 tasks. To verify that raters were attentive we also introduced a “sanity” set of 30 instances of objects for which there were two clear attributes as answers. We also recorded the rate at which raters chose the “none of the above” and “I don’t know” answers as a proxy for the difficulty of assigning labels to the dataset. Balanced accuracy was computed by averaging the accuracy per attribute.

The average rater top-1 and top-2 accuracies were 31.79% and 47% respectively, indicating a label noise level of  $\sim 70\%$ . The fraction of tasks that raters selected “none of the above” or “I don’t know” was 5%, indicating that raters were confident in about  $\sim 95\%$  of their rating. The top-1 accuracy on the “sanity” set was 88% and the top-2 accuracy was 100%, indicating that the raters were attentive and capable of solving the task at hand.

<sup>4</sup>The split used for ablation study

Finally, Fig. S.7 shows qualitative examples for the label quality of MIT-States. For each label of 5 attribute labels, selected by random, we show 5 images, selected by random. Under each image, we show the choice of the amazon-turker in the label-quality experiment and the provided attribute label.

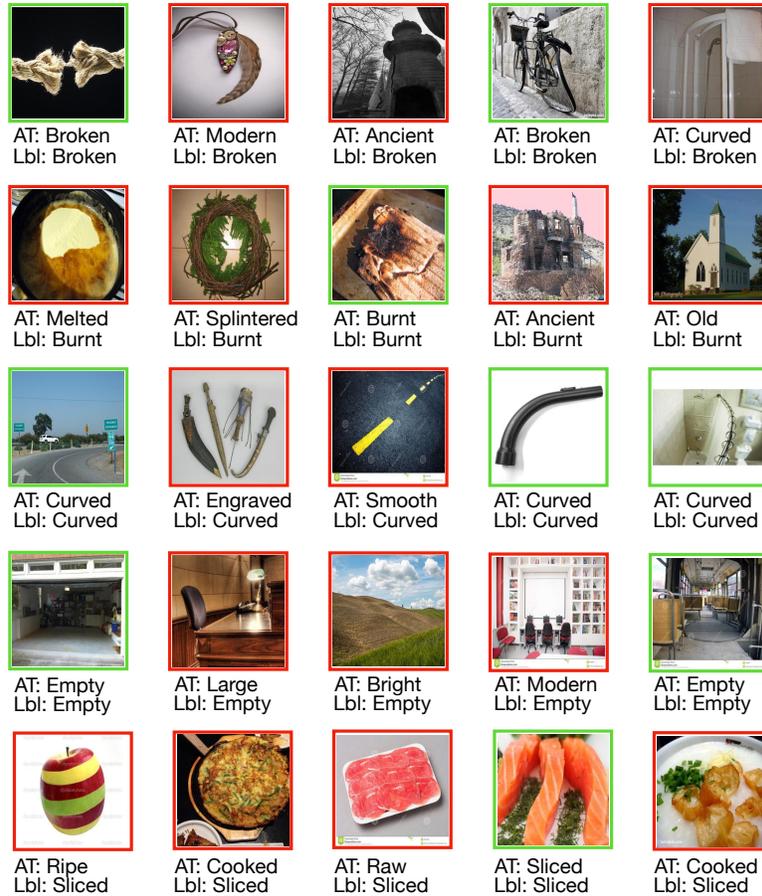


Figure S.7: Label quality of MIT-States. Showing 5 attribute labels, selected by random. For each label, we show 5 images, selected by random. For each image, we show the choice of the amazon-turker (AT) and the provided attribute label (Lbl). Green image margins indicate that the turker choice agrees with the label. Red margins indicate that the turker choice disagrees with the label.

# I Numeric values for the metrics

## I.1 Overlapping split

U:S	CAUSAL				VisPROD&CI				VisPROD			
	UNSEEN	SEEN	HARMONIC	CLOSED	UNSEEN	SEEN	HARMONIC	CLOSED	UNSEEN	SEEN	HARMONIC	CLOSED
2:8	77.7 ± 1.4	89.7 ± 1.9	83.2 ± 1.2	87.0 ± 2.1	60.0 ± 2.7	87.8 ± 2.3	71.1 ± 1.7	85.9 ± 1.0	42.5 ± 3.0	90.1 ± 1.4	57.5 ± 2.9	87.3 ± 1.3
3:7	72.2 ± 1.0	80.9 ± 3.6	75.7 ± 2.3	84.1 ± 2.5	44.7 ± 5.0	84.3 ± 3.8	58.1 ± 4.9	72.6 ± 7.5	29.2 ± 3.7	85.6 ± 3.4	43.2 ± 4.4	73.5 ± 6.9
4:6	67.4 ± 2.0	84.1 ± 1.8	74.7 ± 1.7	86.6 ± 0.7	53.2 ± 1.5	87.9 ± 1.1	66.2 ± 0.9	88.2 ± 1.0	35.9 ± 3.6	85.0 ± 2.2	49.8 ± 2.9	87.0 ± 1.6
5:5	47.1 ± 4.5	83.8 ± 0.8	59.8 ± 3.9	71.6 ± 4.9	38.3 ± 1.1	82.1 ± 4.2	52.1 ± 1.8	65.4 ± 3.6	18.9 ± 0.2	84.5 ± 8.5	30.3 ± 0.8	61.6 ± 2.3
6:4	26.9 ± 0.5	86.1 ± 2.9	40.9 ± 1.0	44.6 ± 1.9	20.0 ± 1.8	86.4 ± 1.0	32.3 ± 2.4	39.6 ± 2.6	11.1 ± 1.4	81.0 ± 9.5	18.8 ± 1.5	37.4 ± 3.5
7:3	22.8 ± 3.0	69.3 ± 6.1	33.7 ± 4.2	40.9 ± 3.6	15.7 ± 2.5	68.5 ± 6.6	25.1 ± 3.4	38.9 ± 1.6	14.1 ± 1.5	50.0 ± 8.6	21.6 ± 2.4	36.8 ± 2.7

U:S	LE*				LE				ATTOP*			
	UNSEEN	SEEN	HARMONIC	CLOSED	UNSEEN	SEEN	HARMONIC	CLOSED	UNSEEN	SEEN	HARMONIC	CLOSED
2:8	72.1 ± 2.8	91.9 ± 0.2	80.7 ± 1.7	85.5 ± 1.8	72.7 ± 3.3	92.2 ± 0.3	81.2 ± 2.0	86.0 ± 1.5	78.4 ± 3.3	77.2 ± 1.0	77.6 ± 1.4	84.1 ± 4.0
3:7	41.9 ± 6.3	92.5 ± 0.3	56.8 ± 6.4	74.5 ± 4.5	42.1 ± 6.9	89.8 ± 2.8	56.6 ± 7.2	75.3 ± 5.1	62.5 ± 2.3	72.3 ± 1.4	66.9 ± 1.2	78.9 ± 4.1
4:6	57.8 ± 4.5	85.8 ± 1.5	68.6 ± 2.7	79.9 ± 2.7	47.2 ± 5.7	84.3 ± 1.6	59.7 ± 4.5	73.7 ± 1.6	68.8 ± 1.2	69.3 ± 2.8	68.8 ± 0.9	81.6 ± 2.1
5:5	31.7 ± 1.6	90.7 ± 1.4	46.8 ± 1.9	54.7 ± 5.6	26.3 ± 2.1	86.4 ± 1.1	40.2 ± 2.6	52.3 ± 8.3	44.3 ± 1.1	70.8 ± 2.3	54.4 ± 1.2	57.8 ± 2.5
6:4	20.1 ± 1.0	81.2 ± 4.7	31.9 ± 0.9	31.8 ± 1.3	20.8 ± 2.4	84.4 ± 2.6	33.1 ± 2.8	29.7 ± 1.7	25.8 ± 1.9	69.5 ± 2.2	37.4 ± 1.7	35.4 ± 0.3
7:3	13.7 ± 2.2	83.1 ± 3.8	23.1 ± 3.3	27.7 ± 3.4	11.4 ± 3.2	89.2 ± 3.3	19.7 ± 5.1	24.4 ± 6.0	10.6 ± 1.7	60.7 ± 7.7	17.9 ± 2.7	14.5 ± 1.6

U:S	ATTOP				TMN*				TMN			
	UNSEEN	SEEN	HARMONIC	CLOSED	UNSEEN	SEEN	HARMONIC	CLOSED	UNSEEN	SEEN	HARMONIC	CLOSED
2:8	79.9 ± 3.5	77.9 ± 1.1	78.7 ± 1.3	84.2 ± 3.9	78.4 ± 5.2	87.2 ± 0.8	82.2 ± 2.7	87.9 ± 1.5	79.7 ± 4.4	85.8 ± 0.9	82.4 ± 2.1	88.7 ± 1.6
3:7	64.3 ± 3.3	72.7 ± 1.8	68.0 ± 1.4	77.0 ± 3.9	62.7 ± 6.1	86.5 ± 0.4	72.1 ± 4.2	81.7 ± 5.0	62.2 ± 4.9	86.5 ± 0.3	72.0 ± 3.4	81.4 ± 5.0
4:6	68.8 ± 1.2	68.7 ± 3.1	68.5 ± 1.4	83.0 ± 2.0	70.4 ± 3.0	83.0 ± 3.2	75.8 ± 0.5	86.6 ± 1.3	68.1 ± 4.1	83.8 ± 2.6	74.5 ± 1.6	86.5 ± 2.5
5:5	46.3 ± 1.2	67.7 ± 2.5	55.0 ± 1.5	59.9 ± 2.8	39.7 ± 3.7	84.9 ± 3.7	53.2 ± 2.4	67.8 ± 5.1	38.0 ± 3.0	84.7 ± 2.2	51.5 ± 2.2	67.3 ± 4.7
6:4	27.9 ± 2.3	72.4 ± 2.7	39.9 ± 2.0	40.2 ± 0.3	17.0 ± 2.6	87.9 ± 1.2	28.1 ± 3.6	40.0 ± 0.9	18.1 ± 2.9	83.7 ± 0.4	29.1 ± 3.7	41.7 ± 0.7
7:3	13.3 ± 1.5	56.0 ± 5.2	21.2 ± 2.2	19.2 ± 2.1	7.3 ± 1.6	93.1 ± 3.3	13.1 ± 2.7	35.4 ± 2.2	5.8 ± 0.9	88.1 ± 2.5	10.8 ± 1.6	36.7 ± 1.0

Table S.2: Numeric values for results of Figure S.2 (top row)

## I.2 Non-overlapping split

U:S	CAUSAL				LE				ATTOP			
	UNSEEN	SEEN	HARMONIC	CLOSED	UNSEEN	SEEN	HARMONIC	CLOSED	UNSEEN	SEEN	HARMONIC	CLOSED
2:8	64.3 ± 1.0	79.4 ± 1.5	70.8 ± 1.0	82.1 ± 0.7	35.4 ± 5.1	80.1 ± 3.7	48.7 ± 5.5	71.2 ± 2.6	53.4 ± 3.7	67.8 ± 3.2	59.1 ± 1.1	76.0 ± 5.4
3:7	48.7 ± 4.7	75.3 ± 4.6	58.9 ± 5.0	79.0 ± 5.5	22.9 ± 2.5	84.5 ± 2.3	35.7 ± 3.3	52.5 ± 6.8	37.4 ± 5.5	65.6 ± 2.4	46.3 ± 4.2	61.1 ± 10.9
4:6	43.5 ± 4.6	69.2 ± 4.2	53.2 ± 4.5	66.8 ± 5.3	27.5 ± 2.3	80.7 ± 2.9	40.7 ± 2.6	43.4 ± 4.2	21.7 ± 9.5	49.2 ± 5.3	26.8 ± 10.6	40.9 ± 10.1
5:5	15.7 ± 1.7	75.2 ± 7.3	25.8 ± 2.7	37.5 ± 4.9	9.1 ± 2.1	91.1 ± 0.5	16.3 ± 3.5	19.9 ± 3.6	6.4 ± 1.9	65.3 ± 7.5	9.9 ± 2.0	22.0 ± 4.2

U:S	TMN			
	UNSEEN	SEEN	HARMONIC	CLOSED
2:8	47.2 ± 2.7	82.1 ± 2.9	59.5 ± 1.6	81.2 ± 2.2
3:7	23.4 ± 4.2	84.2 ± 0.2	35.1 ± 5.4	64.4 ± 5.6
4:6	15.3 ± 4.4	85.5 ± 3.1	24.7 ± 6.3	54.4 ± 4.7
5:5	3.0 ± 1.1	86.8 ± 3.4	5.6 ± 2.0	32.7 ± 1.6

Table S.3: Numeric values for results of Figure S.1