

# Counterfactual Zero-Shot and Open-Set Visual Recognition

Zhongqi Yue<sup>1,3\*</sup>, Tan Wang<sup>1\*</sup>, Qianru Sun<sup>2</sup>, Xian-Sheng Hua<sup>3</sup>, Hanwang Zhang<sup>1</sup>

<sup>1</sup>Nanyang Technological University, <sup>2</sup>Singapore Management University, <sup>3</sup>Damo Academy, Alibaba Group

yuez0003@ntu.edu.sg, TAN317@e.ntu.edu.sg, qianrusun@smu.edu.sg,

xiansheng.hxs@alibaba-inc.com, hanwangzhang@ntu.edu.sg

## Abstract

We present a novel counterfactual framework for both Zero-Shot Learning (ZSL) and Open-Set Recognition (OSR), whose common challenge is generalizing to the unseen-classes by only training on the seen-classes. Our idea stems from the observation that the generated samples for unseen-classes are often out of the true distribution, which causes severe recognition rate imbalance between the seen-class (high) and unseen-class (low). We show that the key reason is that the generation is not **Counterfactual Faithful**, and thus we propose a faithful one, whose generation is from the sample-specific counterfactual question: What would the sample look like, if we set its class attribute to a certain class, while keeping its sample attribute unchanged? Thanks to the faithfulness, we can apply the **Consistency Rule** to perform unseen/seen binary classification, by asking: Would its counterfactual still look like itself? If “yes”, the sample is from a certain class, and “no” otherwise. Through extensive experiments on ZSL and OSR, we demonstrate that our framework effectively mitigates the seen/unseen imbalance and hence significantly improves the overall performance. Note that this framework is orthogonal to existing methods, thus, it can serve as a new baseline to evaluate how ZSL/OSR models generalize. Codes are available at <https://github.com/yue-zhongqi/gcm-cf>.

## 1. Introduction

Generalizing visual recognition to novel classes unseen in training is perhaps the Holy Grail of machine vision [37]. For example, if machines could classify new classes accurately by Zero-Shot Learning (ZSL<sup>1</sup>) [29, 62], we could collect labelled data as many as possible for free; if machines could reject samples of unknown classes by Open-Set Recognition (OSR) [50, 10], any recognition system

\*Equal contribution

<sup>1</sup>Conventional ZSL [29] only evaluates the recognition on the unseen-classes. We refer to ZSL as a more challenging setting: Generalized ZSL [11], which is evaluated on the both seen- and unseen-classes.

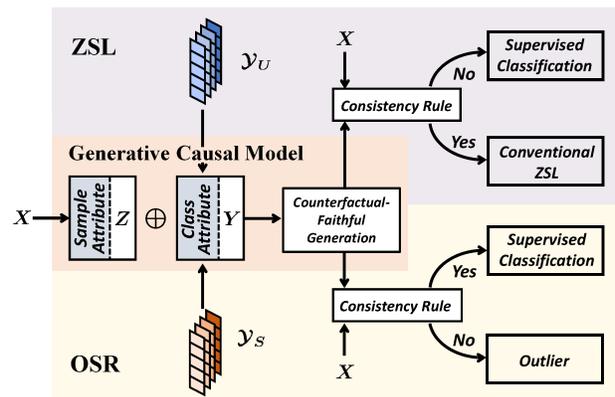


Figure 1: Our counterfactual framework for ZSL and OSR. Here  $\oplus$  denotes vector concatenation,  $\mathcal{Y}_U$  and  $\mathcal{Y}_S$  denote the set of unseen and seen class attributes, respectively.

would be shielded against outliers. Unfortunately, this goal is far from achieved, as it is yet a great challenge for them to “imagine” the unseen world based on the seen one [28, 51].

Over the past decade, all the unseen-class recognition methods stem from the same grand assumption: *attributes (or features) learned from the training seen-classes are transferable to the testing unseen-classes*. Therefore, if we have the ground-truth *class attributes* describing both of the seen- and unseen-classes (or only those of the seen in OSR), ZSL (or OSR) can be accomplished by comparing the predicted class attributes of the test sample and the ground-truth ones [15, 5]; or by training a classifier on the samples generated from the ground-truth attributes [65, 43].

Not surprisingly, the above assumption is hardly valid in practice. As the model only sees the seen-classes in training, it will inevitably cater to the seen idiosyncrasies, and thus result in an unrealistic imagination of the unseen world. Figure 2a illustrates that the samples, generated from the class attribute of an unseen-class, do not lie in the sample domain between the ground-truth seen and unseen, *i.e.*, they resemble neither the seen nor the unseen. As a result, the seen/unseen boundary learned from the generated unseen

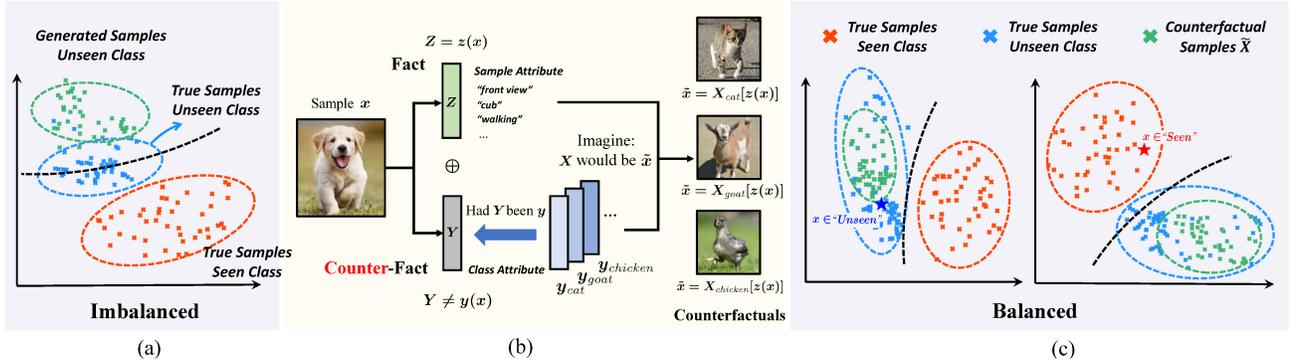


Figure 2: (a) t-SNE [35] plot of the CUB [60] samples in ZSL using a conventional unseen generation method [39], where a single pair of seen- and unseen-class is shown to avoid clutter. Due to the out-of-distribution generation (green), the decision boundary (black dashed lines) is imbalanced between the true seen (red) and unseen (blue) samples. (b) Illustration of our counterfactual generation. (c) t-SNE plot of the CUB using our counterfactual generation. The decision boundary is balanced. See Figure 4 for more examples of OSR.

and the true seen samples is imbalanced. Such imbalanced classification increases the recall of the seen-class by sacrificing that of the unseen. Interestingly, we find that all the existing ZSL methods suffer from this imbalance: the seen accuracy is much higher than the unseen (*cf.* Table 1).

Astute readers may intuitively realize that it is all about *disentanglement*: if every attribute is disentangled, any unseen combination will be sensible. However, it is well-known that learning disentangled features is difficult, or even impossible without proper supervision [34]. In this paper, we propose another way around: **Counterfactual Inference** [47], which does not require the disentanglement for the class attributes. As illustrated in Figure 2b, denote  $X$  as the sample variable, which can be encoded into the *sample* attributes  $Z = z(X = x)$  (*e.g.*, “front view”) and *class* attributes  $Y = y(X = x)$ . The counterfactual [46]:  $\tilde{x} = X_y[z(x)]$  of sample  $x$ , is read as:

$$\boxed{\begin{array}{l} X \text{ would be } \tilde{x}, \text{ had } Y \text{ been } y, \\ \text{given the fact that } Z = z(X = x). \end{array}}$$

Note that the given “fact” is  $Z = z(x)$  and the “counterfact” (what if) is  $Y \neq y(x)$ , indicating that the encoded  $y(x)$  — also an observed fact — clashes with the counterfact  $y$ . The key difference between the conventional unseen generation and the counterfactual is that: the former is purely based on the **sample-agnostic** class attributes  $y$  (or together with Gaussian prior  $z$  [64]), while the latter is also grounded by the **sample-specific** attributes  $z(x)$ .

So, why sample-specific? The imagination of an unseen-class is indeed not necessary to start from scratch, *i.e.*, purely from the class definition  $Y$ , where the effect of some attributes may be lost due to entanglement [12]; instead, it should also start from the observed fact  $Z = z(x)$ , which can make up those entangled attributes. In fact, such grounded generation simulates how we humans imagine an unseen sample [48]: we imagine a “dinosaur” (class at-

tributes) based on its fossil (sample attributes). More formally, by disentangling the two groups: class attribute  $Y$  and sample  $Z$ , we can use the theorem of **Counterfactual Faithfulness** (Section 3.3 & 3.4) to guarantee that the counterfactual distribution is coherent with the ground-truth seen/unseen distribution. As shown in Figure 2c (left), the unseen-class generation grounded by a sample is in the true unseen domain, leading to a more balanced decision boundary. It is worth noting that the group disentanglement between  $Z$  and  $Y$  is much more relaxed and thus more approachable than the full disentanglement [6, 7], *i.e.*, the attributes within each group are not required to be disentangled. In Section 3.4, we design a training procedure to achieve this.

We propose a counterfactual framework for ZSL and OSR, because they are both underpinned by the generalization to unseen-classes. As summarized in Figure 1, the counterfactual is powered by a Generative Causal Model (Section 3.2). The counterfactual faithfulness allows us to use the **Consistency Rule**: if the counter-fact  $y$  is indeed the underlying ground-truth, the counterfactual  $\tilde{x}$  equals to the factual  $x$ . Therefore, we can use the rule as a seen/unseen binary classifier by varying  $y$  across the seen/unseen class attributes (Section 3.3): If a sample is seen, we can apply the conventional supervised learning classifier; otherwise, for ZSL, we apply the conventional ZSL classifier, and for OSR, we reject it.

To the best of our knowledge, the proposed counterfactual framework is the first to provide a theoretical ground for balancing and improving the seen/unseen classification. In particular, we show that the quality of disentangling  $Z$  and  $Y$  is the key bottleneck, so it is a potential future direction for ZSL/OSR [57, 18, 45, 8]. Our method can serve as an unseen/seen binary classifier, which can plug-and-play and boost existing ZSL/OSR methods to achieve new state-of-the-arts (Section 4).

## 2. Related Work

ZSL is usually provided with an auxiliary set of class attributes to describe each seen- and unseen-class [29, 62]. Therefore, ZSL can be approached by either inferring a sample’s attribute and finding the closest match in the attribute space [15, 1, 12, 24], or generating features using the attributes and matching in the feature space [64, 32, 39, 44]. OSR is a more challenging open environment setting with no information on the unseen-classes [49, 50], and the goal is to build a classifier for seen-classes that additionally rejects unseen-class samples as outliers. The inference methods in OSR calibrate a discriminative model by adjusting the classification logits [5, 75], and the generation methods estimate the seen-class density and reject test samples in the low-density area as outliers [43, 54]. **Out-Of-Distribution (OOD)** detection [16, 2, 33] also focuses on unseen detection, where seen and unseen samples are usually from different domains [9]. However, in OOD, the unseen-class information is available. In particular, some works employ OOD techniques in ZSL to distinguish seen- and unseen-classes [36, 13]. Our work is based on **Causal Inference** [46], which has shown promising results in various computer vision tasks [72, 59, 68] including few-shot classification [71], long-tailed classification [58], incremental learning [20], vision-language models [69] and visual question answering [42]. A recent paper [4] uses causal intervention on ZSL. However, it builds on a restrictive setting where class attributes are fully disentangled (*e.g.*, shape and color). Our work uses *counterfactual inference* and relies on the much-relaxed group disentanglement [7], allowing us to outperform on complex benchmark datasets.

## 3. Approach

### 3.1. Problem Definitions

We train a model using a labelled dataset  $\mathcal{D} = \{\mathbf{x}_i, l_i\}_{i=1}^N$  on seen-classes  $\mathcal{S}$ , to recognize samples from both  $\mathcal{S}$  and unseen-classes  $\mathcal{U}$ , where  $\mathbf{x}_i \in \mathcal{X} \subset \mathbb{R}^d$  is the  $d$ -dimensional feature vector of the  $i$ -th sample and  $l_i \in \mathcal{S}$  is its corresponding label. We assume that the samples of  $\mathcal{S}$  and  $\mathcal{U}$  are embedded in the same feature space  $\mathcal{X}$ , whose ambient space is the RGB color space or the network feature space provided by the dataset curator [62].

**Zero-Shot Learning (ZSL)**. It includes two settings: 1) *Conventional ZSL*, where the model is only evaluated on  $\mathcal{U}$ , and 2) *Generalized ZSL*, where the model is evaluated on  $\mathcal{S} \cup \mathcal{U}$ . A common practice [11] is to use an additional set of class attribute vectors  $\mathcal{Y}_S$  and  $\mathcal{Y}_U$  to describe the seen- and unseen-classes, respectively. Compared to the one-hot label embeddings, these attributes can be considered as dense label embeddings [53, 38]. When the context is clear, we refer to ZSL as Generalized ZSL.

**Open-Set Recognition (OSR)**. It is evaluated on both  $\mathcal{S}$

and  $\mathcal{U}$ . Compared to ZSL,  $\mathcal{U}$  in OSR is marked as “unknown”. Instead of using dense labels, each seen-class is described by the one-hot embedding of  $K$  dimensions, where  $K$  is the number of seen-classes. The one-hot embeddings are considered as the seen-class attributes set  $\mathcal{Y}_S$ . Since OSR is evaluated in an open environment, there is no unseen-class attributes set  $\mathcal{Y}_U$ .

### 3.2. Generative Causal Model

Our assumption is that both ZSL and OSR follow a Generative Causal Model (GCM) [46] shown in Figure 3, where the *class attribute*  $Y$  and the *sample attribute*  $Z$  jointly determine the observed image feature  $X$ . In general, the generative causal process  $Z \rightarrow X, Y \rightarrow X$  can be confounded, represented as the dashed links in Figure 3. One can remove the confounders through more elaborate data collection [62] and experimental design [23]. In this paper, we follow the conventions in ZSL and OSR to ignore the confounders [32, 43]. Our GCM entails a generation and inference process. Specifically, given  $Z$  and  $Y$ , we can generate  $X$  by sampling from the conditional distribution  $P_\theta(X|Z, Y)$ . While given  $X$ , we can infer  $Z$  and  $Y$  through the posterior  $Q_\phi(Z|X)$  and  $Q_\psi(Y|X)$ . Therefore, GCM can support both generative and inference-based methods. As the former prevails over the latter in literature, we focus on it in this paper.

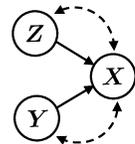


Figure 3: Our GCM for ZSL and OSR.

In ZSL,  $Y$  takes values from the dense labels set  $\mathcal{Y}_S \cup \mathcal{Y}_U$ . Generative ZSL methods [64, 32] aim to learn  $P_\theta(X|Z, Y)$ . In testing, they generate unseen-class samples  $X$  by using a Gaussian prior as  $Z$  and the attributes  $Y$  in  $\mathcal{Y}_U$ . Then, the generated unseen samples and the seen  $\mathcal{D}$  are used to train a classifier to recognize both  $\mathcal{S}$  and  $\mathcal{U}$ . In OSR, the value of  $Y$  is from the one-hot embedding set  $\mathcal{Y}_S$ . Generative OSR methods [43, 54] first infer  $Y$  from a test sample  $\mathbf{x}$  by sampling  $\mathbf{y} \sim Q_\psi(Y|X = \mathbf{x})$ . Then, the inferred  $\mathbf{y}$  and Gaussian noise  $\mathbf{z}$  are used to generate  $\mathbf{x}'$  from  $P_\theta(X|Z = \mathbf{z}, Y = \mathbf{y})$ . Finally, the sample is marked as “unknown” if  $\mathbf{x}$  is dissimilar to  $\mathbf{x}'$  subject to a specified threshold.

Note that all the above generation methods adopt a prior Gaussian noise for  $Z$ , which is not sample-specific. As  $Y$  is inevitably entangled (examples in Appendix), there is no mechanism to make up for the missing sample attribute effect during generation, so, the generated unseen-class samples will be unrealistic and lie outside  $\mathcal{X}$ , rendering the decision rule trained on the seen samples inapplicable.

### 3.3. Counterfactual Generation and Inference

Given a sample  $\mathbf{x}$ , we use the GCM to generate counterfactual samples  $\tilde{\mathbf{x}} = X_{\mathbf{y}}[z(\mathbf{x})]$  following the three steps of computing counterfactual [47]:

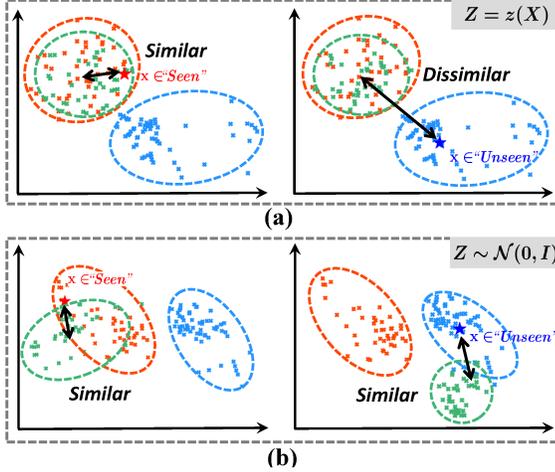


Figure 4: OSR using (a) our counterfactual framework ( $Z=z(X)$ ); (b) existing generative approach ( $Z \sim \mathcal{N}(0, I)$ ) when the test sample is from seen-classes (left) and unseen-classes (right). **Red dashed line** denotes the true samples from the seen-class and **blue dash line** as those from the unseen-class. **Green dash line** denotes generated samples using the seen-class attribute.  $\star$  and  $\star$  denote the samples from seen and unseen class, respectively.

**Abduction**—“given the fact that  $Z = z(\mathbf{x})$ ”. We solve for the endogenous sample attribute  $z(\mathbf{x})$  given the evidence  $X = \mathbf{x}$ . In our GCM, we can sample from the posterior  $z(\mathbf{x}) \sim Q_\phi(Z|X = \mathbf{x})$ .

**Action**—“had  $Y$  been  $y$ ”. Here  $y \in \mathcal{Y}_S \cup \mathcal{Y}_U$  is the intervention target of  $Y$ . Note that the class attribute  $Y$  of this sample can be inferred from  $y(\mathbf{x}) \sim Q_\psi(Y|X = \mathbf{x})$ , but in this step, we intervene on  $Y$  by discarding the inferred value and setting  $Y$  as  $y$ , which may be different from  $y(\mathbf{x})$ .

**Prediction**—“ $X$  would be  $\tilde{\mathbf{x}}$ ”. Conditioning on the inferred  $Z = z(\mathbf{x})$  (fact) and the intervention target  $Y = y$  (counterfact), we can generate the counterfactual sample  $\tilde{\mathbf{x}}$  from  $P_\theta(X|Z = z(\mathbf{x}), Y = y)$ . Notice that the feature imagination from class attribute  $Y = y$  is now grounded by the sample attribute  $Z = z(\mathbf{x})$ .

We want that the above counterfactual sample lies in the true distribution of the seen or unseen samples. Formally, such property is defined as below:

**Definition (Counterfactual Faithfulness)**. Given  $\mathbf{x} \in \mathcal{X}$ , the counterfactual generation  $\tilde{\mathbf{x}}$  using a GCM is faithful whenever  $\tilde{\mathbf{x}} \in \mathcal{X}$ .

In fact, the failure of existing methods, as mentioned in Section 3.2, is because the counterfactual faithfulness does not hold; however, it holds for our generation. We delay the justification to Section 3.4. It assures that any distance metric in  $\mathcal{X}$  is applicable for both  $\mathbf{x}$  and its counterfactual generation  $\tilde{\mathbf{x}}$ . This allows us to build a binary classifier on seen/unseen by applying the **Consistency Rule**:

$$y^*(\mathbf{x}) = y \implies X_y[z(\mathbf{x})] = \mathbf{x}, \quad (1)$$

where  $y^*(\mathbf{x})$  is the (unobserved) ground-truth class attribute of  $\mathbf{x}$  and  $X_y[z(\mathbf{x})]$  is the generated counterfactual with this ground-truth. Our binary classification strategy is based on the equivalent *contraposition* of the rule, which states that *if  $\mathbf{x}$  is dissimilar to  $X_y[z(\mathbf{x})]$ , the ground-truth attribute of  $\mathbf{x}$  cannot be  $y$* :

$$X_y[z(\mathbf{x})] \neq \mathbf{x} \implies y^*(\mathbf{x}) \neq y. \quad (2)$$

Thanks to the counterfactual faithfulness, the dissimilarity can be measured by any distance metric defined in  $\mathcal{X}$  (e.g., Euclidean distance). Figure 4 uses the OSR task to show the benefit of our counterfactual approach grounded by  $Z = z(X)$ . The class-agnostic  $Z$  in the existing methods cannot make up for the attributes that entangled in  $Y$ . This leads to non-faithful generation on unseen-class samples, and the distance can hardly tell apart the seen and unseen-class samples (Figure 4b). In contrast, our approach achieves counterfactual faithfulness and thus the distance remains discriminative (Figure 4a). Next, we detail the binary inference rule for both ZSL and OSR.

**Inference in ZSL**. Since ZSL is evaluated in a closed environment, i.e., the set of unseen-class attributes  $\mathcal{Y}_U$  is known in testing, we use the contraposition—if feature  $\mathbf{x}$  is *dissimilar* to counterfactual generations from the *unseen*-classes,  $\mathbf{x}$  belongs to *seen*. For example in Figure 2c, conditioning on a seen-class sample  $\mathbf{x}$ , the counterfactual generations using the unseen-class attribute are indeed dissimilar to  $\mathbf{x}$  as they lie in the opposite side of the classifier decision boundary. Specifically, we generate a set of counterfactual features  $\tilde{X}$  of the *unseen*-classes by using the inferred  $Z$  from  $Q_\phi(Z|X = \mathbf{x})$  and intervening  $Y$  with the ground-truth attributes from  $\mathcal{Y}^U$ . Using the counterfactual set  $\tilde{X}$  of the *unseen*-classes and  $\mathcal{D}$  of the *seen*-classes, we train a multi-label classifier whose vocabulary is  $\mathcal{S} \cup \mathcal{U}$ . Denote the mean-pooling of the top- $K$  classifier probabilities among seen- and unseen-classes as  $S^K$  and  $U^K$ , respectively. The binary seen/unseen label  $b(\mathbf{x})$  is given by:

$$b(\mathbf{x}) = \begin{cases} \text{seen}, & \text{if } U^K < S^K \\ \text{unseen}, & \text{otherwise} \end{cases} \quad (3)$$

**Inference in OSR**. Since OSR is in an open environment, i.e., there could be infinite number of unseen-classes, it is impossible to generate unseen-class counterfactuals. Therefore, we use the contraposition the other way round—if  $\mathbf{x}$  is *dissimilar* to the counterfactual generations of the *seen*-classes,  $\mathbf{x}$  belongs to the *unseen*. This is shown in Figure 4a, where the unseen-class sample (right), compared to the seen-class sample (left), is much more dissimilar to the seen-class counterfactuals. Hence by thresholding the dissimilarity, we can classify both samples correctly. Specifically, we generate a set of counterfactual features  $\tilde{X}$  on *seen*-classes by setting  $Y$  as the one-hot embeddings in  $\mathcal{Y}_S$ ,

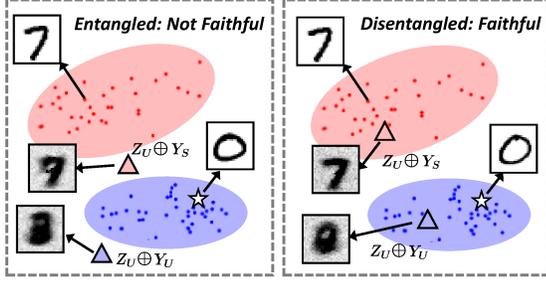


Figure 5: Counterfactual generation conditioned on an unseen sample (denoted as the star).  $Z_U \oplus Y_S$  and  $Z_U \oplus Y_U$  represent counterfactual images conditioned on the unseen sample, using seen and unseen attribute respectively.

and we calculate the minimum Euclidean distance between  $\mathbf{x}$  and each  $\tilde{\mathbf{x}} \in \tilde{X}$ , denoted as  $d_{min}$ . If  $d_{min}$  is larger than a threshold  $\tau$ , the sample is dissimilar to the seen-classes counterfactuals. The binary label  $b(\mathbf{x})$  is given by:

$$b(\mathbf{x}) = \begin{cases} \text{unseen,} & \text{if } d_{min} > \tau \\ \text{seen,} & \text{otherwise} \end{cases} \quad (4)$$

**Two-Stage Inference.** As shown in Figure 1, after the stage-one binary classification, at the second stage, the predicted seen-classes samples can be classified by using the standard supervised classifier, and the predicted unseen-classes samples can be fed into any Conventional ZSL algorithms in ZSL, or rejected as outliers in OSR. The detailed implementation is in Section 4.2.

### 3.4. Counterfactual-Faithful Training

Our framework centers on the counterfactual-faithful generation, which can be guaranteed by the following theorem (proved as a corollary of [6], Appendix):

**Theorem.** *The counterfactual generation  $X_{\mathbf{y}}[z(\mathbf{x})]$  is faithful if and only if the sample attribute  $Z$  and class attribute  $Y$  are group disentangled.*

Figure 5 shows the difference between the generations with and without the group disentanglement. Therefore, to achieve counterfactual faithful generation, the full disentanglement among each dimension of  $Y$  and  $Z$ , which may be impossible in general [34], is relaxed to group disentanglement between  $Y$  and  $Z$ , which can be more easily addressed by the following specially designed training objective:

$$\min_{\theta, \phi} \mathcal{L}_Z + \nu \mathcal{L}_Y + \max_{\omega} \rho \mathcal{L}_F, \quad (5)$$

where  $\nu, \rho$  are a trade-off parameters. All three losses are designed to achieve the counterfactual faithfulness, which assures the distance metric of the consistency rule. Next, we detail each of them.

**Disentangling  $Z$  from  $Y$ .** We minimize the  $\beta$ -VAE [19] loss  $\mathcal{L}_Z$  given by:

$$\mathcal{L}_Z = -\mathbb{E}_{Q_{\phi}(Z|X)} [P_{\theta}(X|Z, Y)] + \beta D_{KL}(Q_{\phi}(Z|X) \| P(Z)), \quad (6)$$

where  $D_{KL}$  denotes KL-divergence,  $P_{\theta}(X|Z, Y)$  and  $Q_{\phi}(Z|X)$  are implemented using the Deep Gaussian family [26], and the prior  $P(Z)$  is set to the isotropic Gaussian distribution. Compared to the standard VAE objective,  $\beta$ -VAE re-weights the KL divergence term by a factor of  $\beta$  ( $\beta > 1$ ). This is shown to be highly effective in learning a disentangled sample attribute  $Z$  [19, 57]. Intuitively, by placing a strict constraint for  $Z$  as a whole to follow the endogenous prior  $P(Z)$ , the value of  $Z$  becomes unaffected by the distribution of  $Y$ , i.e.,  $Z$  is disentangled from  $Y$ .

**Disentangling  $Y$  from  $Z$ .** However, the above regularization cannot guarantee that the GCM correctly uses the class attribute  $Y$  during generation. For example, recent GAN models [74, 25] can generate a large variety of photo-realistic images by only using  $Z$ . In fact, as shown in recent literature [8], it is possible for an over-parameterized model  $P_{\theta}(X|Z, Y)$  to ignore  $Y$  and use purely  $Z$  to generate  $X$ , leading to non-faithful generations. This is because the information in  $Y$  might be *fully contained* in  $Z$ . Therefore, it is necessary to additionally disentangle  $Y$  from  $Z$ . Specifically, given a training sample  $\mathbf{x}$ , its ground-truth attribute  $\mathbf{y}$  and its sample attribute  $z(\mathbf{x})$  sampled from  $Q_{\phi}(Z|X = \mathbf{x})$ , we require  $\mathbf{x}$  to be close to  $\mathbf{x}_{\mathbf{y}} = X_{\mathbf{y}}[z(\mathbf{x})]$ , but far away from the counterfactual generations in the set  $\tilde{X} = \{X_{\mathbf{y}'}[z(\mathbf{x})] \mid \mathbf{y}' \in \mathcal{Y}^S \wedge \mathbf{y}' \neq \mathbf{y}\}$ . We use a contrastive loss as:

$$\mathcal{L}_Y = -\log \frac{\exp(-\text{dist}(\mathbf{x}, \mathbf{x}_{\mathbf{y}}))}{\sum_{\mathbf{x}' \in \tilde{X} \cup \{\mathbf{x}_{\mathbf{y}}\}} \exp(-\text{dist}(\mathbf{x}, \mathbf{x}'))}, \quad (7)$$

where  $\text{dist}$  denotes Euclidean distance. Intuitively, this loss actively intervenes the class attribute  $Y$ , given fixed sample attributes  $Z$ , and therefore maximizes the sample difference (e.g., in terms of contrastiveness) before and after the intervention, which complies with the recent definition of causal disentanglement [18]. Therefore,  $\mathcal{L}_Y$  can disentangles  $Y$  from  $Z$ .

**Further Disentangling by Faithfulness.** Note that the faithfulness  $\tilde{\mathbf{x}} \in \mathcal{X}$  is a necessary condition for disentanglement, so we can also use it as an objective. Note that the VAE objective optimizes a lower bound of the likelihood  $P(X)$ , where the bound looseness undermines the faithfulness. To address this problem, we additionally use the Wasserstein GAN (WGAN) [3] loss. Specifically, we train a discriminator  $D(X, Y)$  parameterized by  $\omega$  that outputs a real value, indicating if feature  $X$  is realistic conditioned on  $Y$  (larger is more realistic). Given a feature  $\mathbf{x}$  with attribute  $\mathbf{y}$ , we can generate the counterfactual  $\mathbf{x}'$  using  $z(\mathbf{x})$  and  $\mathbf{y}$ . The discriminator is trained to mark  $\mathbf{x}$  as real (large  $D(\mathbf{x}, \mathbf{y})$ ) and  $\mathbf{x}'$  as unreal (small  $D(\mathbf{x}', \mathbf{y})$ ). Specifically, the WGAN loss  $\mathcal{L}_F$  is given by:

$$\mathcal{L}_F = \mathbb{E}[D(\mathbf{x}, \mathbf{y})] - \mathbb{E}[D(\mathbf{x}', \mathbf{y})] - \lambda \mathbb{E}[(\|\nabla_{\tilde{\mathbf{x}}} D(\tilde{\mathbf{x}}, \mathbf{y})\|_2 - 1)^2], \quad (8)$$

where  $\lambda$  is a penalty term,  $\hat{\mathbf{x}} = \alpha\mathbf{x} + (1 - \alpha)\mathbf{x}'$  with  $\alpha$  sampled from the uniform distribution defined on  $[0, 1]$ . By training the generation process  $P_\theta(X|Z, Y)$  and the discriminator in an adversarial fashion, the GCM is regularized to generate  $\mathbf{x}'$  that is similar to  $\mathbf{x}$  in the same distribution, *i.e.*, counterfactual-faithful for seen-class samples.

## 4. Experiments

### 4.1. Datasets

**ZSL.** We evaluated our method on standard benchmark datasets: Caltech-UCSD-Birds 200-2011 (*CUB*) [60], *SUN* [66], Animals with Attributes 2 (*AWA2*) [62] and attribute Pascal and Yahoo (*aPY*) [14]. In particular, we followed the unseen/seen split in the *Proposed Split (PS) V2.0* [62], recently released to fix a test-data-leaking bug in the original PS. The granularity, total #images, #seen-classes ( $|\mathcal{S}|$ ) and #unseen-classes ( $|\mathcal{U}|$ ) are given in Table A1 (see Appendix).

**OSR.** We used the standard evaluation datasets: *MNIST* [31], *SVHN* [41], *CIFAR10* [27] and *CIFAR100* [27]. The image size, #classes and # images in train/test split of the datasets are given in Table A2 (see Appendix). Following the standard benchmark [40, 70] in OSR, we split *MNIST*, *SVHN* and *CIFAR10* into 6 seen-classes and 4 unseen-classes, and construct two additional datasets *CIFAR+10 (C+10)* and *CIFAR+50 (C+50)*, where 4 non-animal classes in *CIFAR10* are used as seen-classes, while additional 10 and 50 animal classes from *CIFAR100* are used as unseen-classes.

### 4.2. Evaluation Metrics and Settings

**ZSL Evaluation.** It was conducted in the Generalized ZSL setting. We used two metrics: 1) **ZSL Accuracy.** It consists of 3 numbers ( $U, S, H$ ), where  $U/S$  is the per-class top-1 accuracy of unseen-/seen-classes test samples, and  $H$  is the harmonic mean of  $U, S$ , given by  $H = 2 \times S \times U / (S + U)$ . 2) **CVb.** To measure the balance between unseen/seen classification, we propose to use the Coefficient of Variation of the seen and unseen binary classification accuracy, denoted as CVb. Let  $S_b$  and  $U_b$  be the binary accuracy on seen- and unseen-classes, respectively. CVb is given by  $\sqrt{0.5(S_b - \mu)^2 + 0.5(U_b - \mu)^2} / \mu$ , where  $\mu = (S_b + U_b) / 2$ . Note that the variation between  $S$  and  $U$  of ZSL Accuracy is not a good measure of balance, as they are affected by the number of seen- and unseen-classes, which can be quite different (see *SUN* in Table A1). 3) **AUSUC.** We draw the Seen-Unseen accuracy Curve (SUC) by plotting a series of  $S$  against  $U$  of ZSL Accuracy, where the series is obtained by adjusting a calibration factor  $\omega$  that is subtracted from the classifier logits on the seen-classes. Then we use the Area Under SUC (AUSUC) for evaluation. Compared to a single ZSL Accuracy, SUC and the area pro-

vide a more detailed view of the capability of an algorithm to balance the unseen-seen decision boundary [11, 12].

**OSR Evaluation.** We used the following metrics: 1) Macro-averaged **F1 scores** over seen-classes and “unknown” (for all unseen-classes), which shows how well a method can recognize seen classes while rejecting unseen-classes samples; 2) **Openness-F1 Plot.** We also studied the response of F1 scores under varying *openness* given by  $1 - \sqrt{2N/(N + M)}$ , where  $N$  and  $M$  are number of seen- and unseen-classes, respectively. Compared to a single F1 score where the openness is fixed, this plot shows the robustness of an OSR classifier to the open environment with an unknown number of unseen-classes.

**Implementation Details.** For ZSL, we implemented our GCM based on the network architecture in TF-VAEGAN [39]. Following common protocol [64, 32], we used the ResNet-101 [17] features for  $X$  and attributes provided in [62] for  $\mathcal{Y}_S, \mathcal{Y}_U$ . For OSR, our GCM was implemented using the networks in CGDL [54, 55] and  $X$  represents actual images. Other details are in Appendix.

### 4.3. Results on ZSL

**Mitigate the Imbalance.** As shown in Table 1, our counterfactual approach, denoted as GCM-CF, achieves a **more balanced ZSL Accuracy** and significantly improves the existing state-of-the-art (SOTA) by 2.2% to 4.3%, with a much higher score on  $U$ . For example, compared to LisGAN on *aPY*, our method gains 3.9% on  $U$  while sacrificing only 0.1% on  $S$ . To further show that GCM-CF mitigates the unseen/seen imbalance, we diagnosed the binary classification accuracy using CVb in Table 2. Note that existing methods have very large CVb, which means that there is a large difference between seen and unseen classification accuracy and reveals the imbalance problem. Our method has the **lowest CVb**. This shows that our approach indeed achieves a more balanced binary decision boundary between seen and unseen. However, one may argue that the imbalance problem can be solved by simply adjusting the calibration factor  $\omega$ . Therefore, we plot the SUC using varying  $\omega$  and measured the AUSUC on all datasets. The result is shown in Figure 6, where GCM-CF outperforms other methods in every inch and achieves **the best AUSUC**. This shows that GCM-CF fundamentally improves the unseen/seen classification beyond the reach of simple calibration. Overall, the balanced and much improved ZSL Accuracy, lower CVb for binary classification, and higher accuracy on varying calibrations demonstrate that our method mitigates the unseen/seen imbalance in ZSL. This indicates that our GCM generates faithful counterfactuals (see Figure 2c) and supports the effectiveness of our counterfactual-faithful training in disentangling  $Z$  and  $Y$ .

**Stage-One Binary Classifier.** Our GCM-CF can serve as a stage-one binary unseen/seen classifier and plug into *all*

Method		CUB			AWA2			SUN			aPY		
		<i>U</i>	<i>S</i>	<i>H</i>									
Inf.	ALE <sup>†</sup> [1]	23.7	62.8	34.4	14.0	81.8	23.9	21.8	33.1	26.3	4.6	73.7	8.7
	DEWISE <sup>†</sup> [15]	23.8	53.0	32.8	17.1	74.7	27.8	16.9	27.4	20.9	3.5	78.4	6.7
	LATEM <sup>†</sup> [61]	15.2	57.3	24.0	11.5	77.3	20.0	14.7	28.8	19.5	1.3	71.4	2.6
	RelationNet [56]	36.3	<b>63.8</b>	46.3	22.1	<b>91.4</b>	35.5	15.8	25.5	19.6	11.5	<b>80.7</b>	20.2
Gen.	GDAN [21]	35.0	28.7	31.6	26.0	78.5	39.1	38.2	19.8	26.1	29.0	63.7	39.9
	CADA-VAE [52]	50.3	56.1	53.0	55.4	76.1	64.0	43.6	36.4	39.7	34.0	54.2	41.7
	LisGAN [32]	44.9	59.3	51.1	53.1	68.8	60.0	41.9	37.8	39.8	33.2	56.9	41.9
	TF-VAEGAN [39]	50.7	62.5	56.0	52.5	82.4	64.1	41.0	<b>39.1</b>	40.0	31.7	61.5	41.8
	<b>GCM-CF (Ours)</b>	<b>61.0</b>	59.7	<b>60.3</b>	<b>60.4</b>	75.1	<b>67.0</b>	<b>47.9</b>	37.8	<b>42.2</b>	<b>37.1</b>	56.8	<b>44.9</b>

Table 1: ZSL Accuracy (*U*%, *S*%, *H*%) on the four datasets, where Inf. means inference-based methods and Gen. means generation-based methods. Note that PS V2.0 was released recently in <https://drive.google.com/file/d/1p9gtkuHCCCyjkyezSarCw-1siCSXUyKH/view> to fix a test-data leaking bug. This can have large impacts on the performance of existing methods, such as GDAN [21]. Therefore all our evaluations were conducted on PS V2.0. † indicates that the results are taken from the PS V2.0 report [63], and we reproduced the results on all other methods using the official code. For our method GCM-CF, we used AREN [67] for Conventional ZSL on CUB, and otherwise used TF-VAEGAN [39] for supervised classification and Conventional ZSL.

Method	CUB	AWA2	SUN	aPY
GDAN [60]	15.1	25.1	27.3	14.4
CADA-VAE [66]	6.7	6.0	7.3	11.2
LisGAN [62]	4.9	7.0	4.6	4.0
TF-VAEGAN [14]	8.0	9.4	10.2	5.8
<b>GCM-CF (Ours)</b>	<b>1.5</b>	<b>2.1</b>	<b>1.0</b>	<b>2.3</b>

Table 2: CVb (%) of generative models on all datasets.

Stage 2 \ Stage 1	TF-VAEGAN [39]			GCM-CF (Ours)		
	<i>U</i>	<i>S</i>	<i>H</i>	<i>U</i>	<i>S</i>	<i>H</i>
RelationNet [56]	49.3	81.2	61.3	55.8	75.0	<b>64.0</b>
CADA-VAE [52]	49.5	81.1	61.5	57.6	75.0	<b>65.2</b>
LisGAN [32]	48.8	80.4	60.7	56.1	74.3	<b>63.9</b>
TF-VAEGAN [39]	52.8	83.2	64.6	60.4	75.1	<b>67.0</b>

Table 3: Comparison of the two-stage inference performance on AWA2 [62] using TF-VAEGAN [39] and our GCM-CF as the stage-one binary classifier. The results on other datasets are shown in Appendix, where using GCM-CF as stage-one binary classifier improves all the methods.

ZSL methods for subsequent supervised learning on seen and conventional ZSL on unseen. We performed experiments on 4 representative methods—inference based: RelationNet [56], generation with VAE: CADA-VAE [52], with GAN: LisGAN [32] and with VAE-GAN [30]: TF-VAEGAN [39]. The ZSL Accuracy on AWA2 is shown in Table 3. By comparing with Table 1, we observe that GCM-CF can significantly improve all of them on *H*. For comparison, we used the current SOTA ZSL method TF-VAEGAN as a binary unseen/seen classifier and performed the same experiments. The results on the ZSL Accuracy show that our GCM-CF significantly outperforms TF-VAEGAN. This demonstrates that compared to the class-agnostic  $Z$  in existing methods, the use of sample attribute  $Z$  that is disentangled from class attribute  $Y$  in our GCM-CF is highly effective in improving ZSL performance. Overall, our method as a robust binary classifier can serve as a new strong baseline to evaluate ZSL methods using two-stage inference.

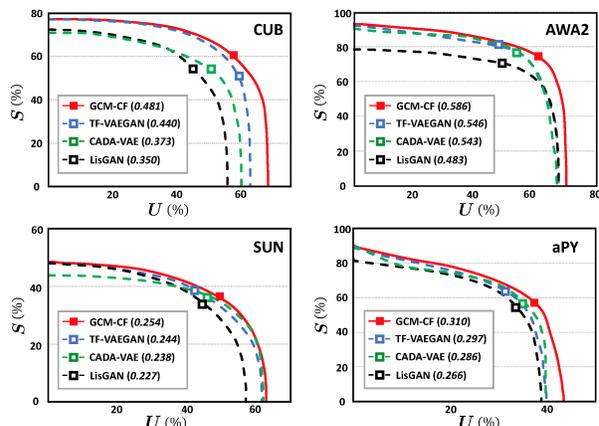


Figure 6: The Seen-Unseen accuracy Curve (SUC) together with the Area Under SUC (AUSUC) on the four datasets. The rectangle denotes the ZSL Accuracy when calibration factor  $\omega=0$ .

#### 4.4. Results on OSR

**Strong Open-Set Classifier.** In Table 4, our GCM-CF achieves SOTA F1 scores in all datasets. We discovered that the common evaluation setting [40, 70] of averaging F1 scores over 5 random splits can result in a large variance in the F1 score. Therefore, we additionally show the results on all 5 splits in the Appendix, where GCM-CF outperforms other methods in every split. This strongly supports that GCM-CF improves F1 scores in general, *i.e.*, not only on some particular splits. So, where does the improvement come from? To rule out the possibility that F1 improves due to a higher classification accuracy on the seen-classes, we measure the Closed-Set Accuracy of a standard CNN trained in a supervised fashion and that of CGDL, which was used in GCM-CF for 2nd stage supervised classification. The results are shown in the Appendix where the performances are similar. Therefore, the increased F1 score indeed comes from improved unseen/seen binary classification. Furthermore, a strong open-set classifier should stay robust regardless of the number of unseen-classes during

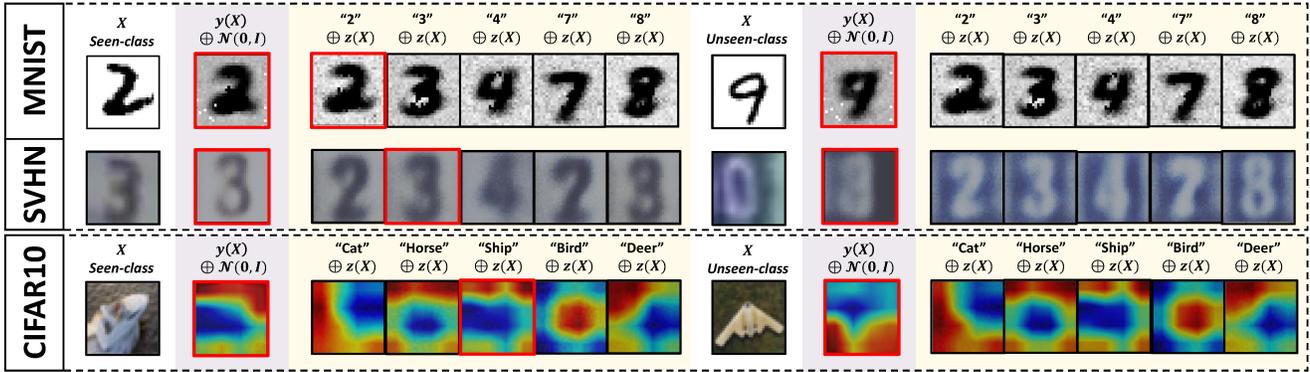


Figure 7: Comparison of the reconstructed images using CGDL [54] ( $y(X) \oplus \mathcal{N}(0, I)$ ) and the counterfactual images generated from our GCM-CF ( $y \oplus z(X)$ ) on the seen- and unseen-class samples. The red box on a generated image denotes that it is similar to the original. For CIFAR10, due to the conflict between visual perception and discriminative training [22], we first train a classifier on seen-classes images and then use CAM [73] to show the sensible yet discriminative features: although the pixel-level generation is not sensible, the pixel-level distance remains discriminative as shown in Table 4.

Method	MNIST	SVHN	CIFAR10	C+10	C+50
Softmax	76.82	76.16	70.39	77.82	65.96
OpenMax [5]	85.93	77.95	71.38	78.68	67.68
CGDL [54]	88.95	76.31	71.03	77.92	70.96
<b>GCM-CF (Ours)</b>	<b>91.37</b>	<b>79.25</b>	<b>72.63</b>	<b>79.38</b>	<b>74.60</b>

Table 4: Comparison of the F1 score averaged over 5 random splits in OSR. We used the official code on CGDL [54] and implemented Softmax and OpenMax [5] for evaluation. For GCM-CF, after binary classification, we used CGDL for supervised classification on the seen-classes.

evaluation. Therefore, we evaluate OSR classifiers with the Openness-F1 Plot. As shown in Figure 8, our GCM-CF achieves the highest F1 score on all openness settings, and our performance is especially competitive in the challenging environment with large openness. Overall, these results clearly demonstrate the effectiveness of our GCM-CF on seen-unseen classification in OSR.

**Qualitative Results.** Figure 7 shows OSR evaluation using existing reconstruction-based approach and our counterfactual approach. Notice the reconstructed image on the unseen-class sample (e.g., “0, 9”) can be similar to the original image. This makes it difficult to tell apart unseen and seen by thresholding reconstruction error. By using our GCM-CF to generate counterfactuals on each seen-class, the generated images are counterfactual-faithful and indeed look like seen-classes. Note that this holds on CIFAR10, where given seen- or unseen-class samples, the CAMs of the generated counterfactuals using the same seen-class attribute look similar (more results in the Appendix). Therefore, by thresholding the dissimilarity between the test sample and its counterfactuals, we can easily distinguish unseen from seen. These results show the effectiveness of our binary inference strategy. Interestingly, we also observe that the generated samples from GCM indeed capture sample attributes better (e.g., text color, background color on SVHN), which validates that we can learn a disentangled  $Z$  using the proposed counterfactual-faithful training.

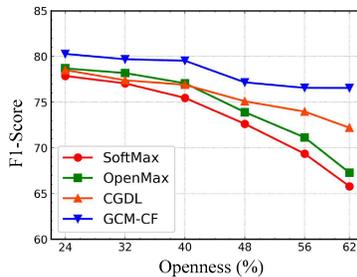


Figure 8: Openness-F1 Plot where 4 non-animal classes from CIFAR10 [27] were used as seen-classes and various classes were drawn from CIFAR100 [27] as unseen-classes.

## 5. Conclusions

We presented a novel counterfactual framework for Zero-Shot Learning (ZSL) and Open-Set Recognition (OSR) to provide a theoretical ground for balancing and improving the seen/unseen classification imbalance. Specifically, we proposed a Generative Causal Model to generate faithful counterfactuals, which allows us to use the Consistency Rule for balanced binary seen/unseen classification. Extensive results in ZSL and OSR show that our method indeed improves the balance and hence achieves the state-of-the-art performance. As future direction, we will seek new definitions on disentanglement [57] and devise practical implementations to achieve improved disentanglement [8].

## 6. Acknowledgements

The authors would like to thank all reviewers and ACs for their constructive suggestions, and specially thank Alibaba City Brain Group for the donations of GPUs. This research is partly supported by the Alibaba-NTU Singapore Joint Research Institute, Nanyang Technological University (NTU), Singapore; A\*STAR under its AME YIRG Grant (Project No. A20E6c0101); and the Singapore Ministry of Education (MOE) Academic Research Fund (AcRF) Tier 1 and Tier 2 grant.

## References

- [1] Zeynep Akata, Florent Perronnin, Zaid Harchaoui, and Cordelia Schmid. Label-embedding for image classification. *TPAMI*, 2015. 3, 7
- [2] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019. 3
- [3] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. *Proceedings of Machine Learning Research*. PMLR, 2017. 5
- [4] Yuval Atzmon, Felix Kreuk, Uri Shalit, and Gal Chechik. A causal view of compositional zero-shot recognition. *Advances in Neural Information Processing Systems*, 2020. 3
- [5] Abhijit Bendale and Terrance E Boult. Towards open set deep networks. In *CVPR*, 2016. 1, 3, 8
- [6] M. Besserve, A. Mehrjou, R. Sun, and B. Schölkopf. Counterfactuals uncover the modular structure of deep generative models. In *ICLR*, 2020. 2, 5
- [7] Michel Besserve, Naji Shajarisales, Bernhard Schölkopf, and Dominik Janzing. Group invariance principles for causal generative models. In *International Conference on Artificial Intelligence and Statistics*, 2018. 2, 3
- [8] Michel Besserve, Rémy Sun, Dominik Janzing, and Bernhard Schölkopf. A theory of independent mechanisms for extrapolation in generative models. *arXiv preprint arXiv:2004.00184*, 2020. 2, 5, 8
- [9] Saikiran Bulusu, Bhavya Kailkhura, Bo Li, Pramod K Varshney, and Dawn Song. Anomalous instance detection in deep learning: A survey. *arXiv preprint arXiv:2003.06979*, 2020. 3
- [10] D. O. Cardoso, F. França, and J. Gama. A bounded neural network for open set recognition. In *2015 International Joint Conference on Neural Networks (IJCNN)*, 2015. 1
- [11] Wei-Lun Chao, Soravit Changpinyo, Boqing Gong, and Fei Sha. An empirical study and analysis of generalized zero-shot learning for object recognition in the wild. In *ECCV*, 2016. 1, 3, 6
- [12] Long Chen, Hanwang Zhang, Jun Xiao, Wei Liu, and Shih-Fu Chang. Zero-shot visual recognition using semantics-preserving adversarial embedding networks. In *CVPR*, 2018. 2, 3, 6
- [13] Xingyu Chen, Xuguang Lan, Fuchun Sun, and Nanning Zheng. A boundary based out-of-distribution classifier for generalized zero-shot learning. In *ECCV*, 2020. 3
- [14] Ali Farhadi, Ian Endres, Derek Hoiem, and David Forsyth. Describing objects by their attributes. In *CVPR*, 2009. 6, 7
- [15] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. In *Advances in neural information processing systems*, pages 2121–2129, 2013. 1, 3, 7
- [16] Will Grathwohl, Kuan-Chieh Wang, Jörn-Henrik Jacobsen, David Duvenaud, Mohammad Norouzi, and Kevin Swersky. Your classifier is secretly an energy based model and you should treat it like one. In *ICLR*, 2020. 3
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 6
- [18] Irina Higgins, David Amos, David Pfau, Sebastien Racaniere, Loic Matthey, Danilo Rezende, and Alexander Lerchner. Towards a definition of disentangled representations. *arXiv preprint arXiv:1812.02230*, 2018. 2, 5
- [19] I. Higgins, Loic Matthey, A. Pal, C. Burgess, Xavier Glorot, M. Botvinick, S. Mohamed, and Alexander Lerchner. beta-vaes: Learning basic visual concepts with a constrained variational framework. In *ICLR*, 2017. 5
- [20] Xinting Hu, Kaihua Tang, Chunyan Miao, Xian-Sheng Hua, and Hanwang Zhang. Distilling causal effect of data in class-incremental learning. In *CVPR*, 2021. 3
- [21] He Huang, Changhu Wang, Philip S Yu, and Chang-Dong Wang. Generative dual adversarial network for generalized zero-shot learning. In *CVPR*, 2019. 7
- [22] Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Logan Engstrom, Brandon Tran, and Aleksander Madry. Adversarial examples are not bugs, they are features. In *Advances in Neural Information Processing Systems*, 2019. 8
- [23] Kosuke Imai, Dustin Tingley, and Teppei Yamamoto. Experimental designs for identifying causal mechanisms. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 2013. 3
- [24] Huajie Jiang, Ruiping Wang, Shiguang Shan, and Xilin Chen. Transferable contrastive network for generalized zero-shot learning. In *ICCV*, 2019. 3
- [25] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019. 5
- [26] Diederik Kingma and Max Welling. Auto-encoding variational bayes. In *ICLR*, 2014. 5
- [27] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 6, 8
- [28] Brenden M Lake, Tomer D Ullman, Joshua B Tenenbaum, and Samuel J Gershman. Building machines that learn and think like people. *Behavioral and brain sciences*, 2017. 1
- [29] C. H. Lampert, H. Nickisch, and S. Harmeling. Learning to detect unseen object classes by between-class attribute transfer. In *CVPR*, 2009. 1, 3
- [30] Anders Boesen Lindbo Larsen, Søren Kaae Sønderby, Hugo Larochelle, and Ole Winther. Autoencoding beyond pixels using a learned similarity metric. In *ICML*, 2016. 7
- [31] Yann LeCun and Corinna Cortes. MNIST handwritten digit database. 2010. 6
- [32] Jingjing Li, Mengmeng Jing, Ke Lu, Zhengming Ding, Lei Zhu, and Zi Huang. Leveraging the invariant side of generative zero-shot learning. In *CVPR*, 2019. 3, 6, 7
- [33] Weitang Liu, Xiaoyun Wang, John Owens, and Sharon Yixuan Li. Energy-based out-of-distribution detection. *Advances in Neural Information Processing Systems*, 2020. 3
- [34] Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Raetsch, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem. Challenging common assumptions in the unsupervised learning of disentangled representations. In *ICML*, 2019. 2, 5

- [35] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 2008. 2
- [36] Devraj Mandal, Sanath Narayan, Sai Kumar Dwivedi, Vikram Gupta, Shuaib Ahmed, Fahad Shahbaz Khan, and Ling Shao. Out-of-distribution detection for generalized zero-shot action recognition. In *CVPR*, 2019. 3
- [37] D. Marr. *Vision: A Computational Investigation Into the Human Representation and Processing of Visual Information*. 2010. 1
- [38] Jonathan Munro and Dima Damen. Multi-modal domain adaptation for fine-grained action recognition. In *CVPR*, 2020. 3
- [39] Sanath Narayan, Akshita Gupta, Fahad Shahbaz Khan, Cees GM Snoek, and Ling Shao. Latent embedding feedback and discriminative features for zero-shot classification. In *ECCV*, 2020. 2, 3, 6, 7
- [40] Lawrence Neal, Matthew Olson, Xiaoli Fern, Weng-Keen Wong, and Fuxin Li. Open set learning with counterfactual images. In *ECCV*, 2018. 6, 7
- [41] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bisaccho, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. In *NeurIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011. 6
- [42] Yulei Niu, Kaihua Tang, Hanwang Zhang, Zhiwu Lu, Xian-Sheng Hua, and Ji-Rong Wen. Counterfactual vqa: A cause-effect look at language bias. In *CVPR*, 2021. 3
- [43] Poojan Oza and Vishal M Patel. C2ae: Class conditioned auto-encoder for open-set recognition. In *CVPR*, 2019. 1, 3
- [44] Ayyappa Pambala, Titir Dutta, and Soma Biswas. Generative model with semantic embedding and integrated classifier for generalized zero-shot learning. In *WACV*, 2020. 3
- [45] Giambattista Parascandolo, Niki Kilbertus, Mateo Rojas-Carulla, and Bernhard Schölkopf. Learning independent causal mechanisms. In *ICML*, 2018. 2
- [46] Judea Pearl. *Causality: Models, Reasoning and Inference*. Cambridge University Press, 2nd edition, 2009. 2, 3
- [47] J. Pearl, M. Glymour, and N.P. Jewell. *Causal Inference in Statistics: A Primer*. Wiley, 2016. 2, 3
- [48] Judea Pearl and Dana Mackenzie. *The Book of Why: The New Science of Cause and Effect*. 2018. 2
- [49] W. J. Scheirer, A. de Rezende Rocha, A. Sapkota, and T. E. Boult. Toward open set recognition. *TPAMI*, 2013. 3
- [50] Walter J Scheirer, Lalit P Jain, and Terrance E Boult. Probability models for open set recognition. *TPAMI*, 2014. 1, 3
- [51] Bernhard Schölkopf. Causality for machine learning. *arXiv preprint arXiv:1911.10500*, 2019. 1
- [52] Edgar Schonfeld, Sayna Ebrahimi, Samarth Sinha, Trevor Darrell, and Zeynep Akata. Generalized zero-and few-shot learning via aligned variational autoencoders. In *CVPR*, 2019. 7
- [53] Eli Schwartz, Leonid Karlinsky, Rogerio Feris, Raja Giryes, and Alex M Bronstein. Baby steps towards few-shot learning with multiple semantics. In *CVPR*, 2019. 3
- [54] Xin Sun, Zhenning Yang, Chi Zhang, Keck-Voon Ling, and Guohao Peng. Conditional gaussian distribution learning for open set recognition. In *CVPR*, 2020. 3, 6, 8
- [55] Xin Sun, Chi Zhang, Guosheng Lin, and Keck-Voon Ling. Open set recognition with conditional probabilistic generative models. *arXiv preprint arXiv:2008.05129v2*, 2020. 6
- [56] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *CVPR*, 2018. 7
- [57] Raphael Suter, Djordje Miladinovic, Bernhard Schölkopf, and Stefan Bauer. Robustly disentangled causal mechanisms: Validating deep representations for interventional robustness. In *ICML*, 2019. 2, 5, 8
- [58] Kaihua Tang, Jianqiang Huang, and Hanwang Zhang. Long-tailed classification by keeping the good and removing the bad momentum causal effect. In *Advances in Neural Information Processing Systems*, 2020. 3
- [59] Tan Wang, Jianqiang Huang, Hanwang Zhang, and Qianru Sun. Visual commonsense r-cnn. In *CVPR*, 2020. 3
- [60] P. Welinder, S. Branson, T. Mita, C. Wah, F. Schroff, S. Belongie, and P. Perona. Caltech-UCSD Birds 200. Technical report, California Institute of Technology, 2010. 2, 6, 7
- [61] Yongqin Xian, Zeynep Akata, Gaurav Sharma, Quynh Nguyen, Matthias Hein, and Bernt Schiele. Latent embeddings for zero-shot classification. In *CVPR*, 2016. 7
- [62] Yongqin Xian, H. Christoph Lampert, Bernt Schiele, and Zeynep Akata. Zero-shot learning - a comprehensive evaluation of the good, the bad and the ugly. *TPAMI*, 2018. 1, 3, 6, 7
- [63] Yongqin Xian, H. Christoph Lampert, Bernt Schiele, and Zeynep Akata. Zero-shot learning - a comprehensive evaluation of the good, the bad and the ugly. 2020. 7
- [64] Yongqin Xian, Tobias Lorenz, Bernt Schiele, and Zeynep Akata. Feature generating networks for zero-shot learning. In *CVPR*, 2018. 2, 3, 6
- [65] Yongqin Xian, Saurabh Sharma, Bernt Schiele, and Zeynep Akata. f-vaegan-d2: A feature generating framework for any-shot learning. In *CVPR*, 2019. 1
- [66] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *CVPR*, 2010. 6, 7
- [67] Guo-Sen Xie, Li Liu, Xiaobo Jin, Fan Zhu, Zheng Zhang, Jie Qin, Yazhou Yao, and Ling Shao. Attentive region embedding network for zero-shot learning. In *CVPR*, 2019. 7
- [68] Xu Yang, Hanwang Zhang, and Jianfei Cai. Deconfounded image captioning: A causal retrospect. *arXiv preprint arXiv:2003.03923*, 2020. 3
- [69] Xu Yang, Hanwang Zhang, Guojun Qi, and Jianfei Cai. Causal attention for vision-language tasks. In *CVPR*, 2021. 3
- [70] Ryota Yoshihashi, Wen Shao, Rei Kawakami, Shaodi You, Makoto Iida, and Takeshi Naemura. Classification-reconstruction learning for open-set recognition. In *CVPR*, 2019. 6, 7
- [71] Zhongqi Yue, Hanwang Zhang, Qianru Sun, and Xian-Sheng Hua. Interventional few-shot learning. In *NeurIPS*, 2020. 3

- [72] Dong Zhang, Hanwang Zhang, Jinhui Tang, Xiansheng Hua, and Qianru Sun. Causal intervention for weakly-supervised semantic segmentation. *NeurIPS*, 2020. 3
- [73] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *CVPR*, 2016. 8
- [74] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *CVPR*, 2017. 5
- [75] Sergey Demyanov Zongyuan Ge and Rahil Garnavi. Generative openmax for multi-class open set classification. In *BMVC*, 2017. 3