

Continual Zero-Shot Learning through Semantically Guided Generative Random Walks

Wenxuan Zhang¹ Paul Janson^{1,2 *} Kai Yi¹ Ivan Skorokhodov¹
Mohamed Elhoseiny¹
KAUST¹ University of Moratuwa²

{wenxuan.zhang, remond.janson, kai.yi, ivan.skorokhodov, mohamed.elhoseiny}@kaust.edu.sa

Abstract

Learning novel concepts, remembering previous knowledge, and adapting it to future tasks occur simultaneously throughout a human’s lifetime. To model such comprehensive abilities, continual zero-shot learning (CZSL) has recently been introduced. However, most existing methods overused unseen semantic information that may not be continually accessible in realistic settings. In this paper, we address the challenge of continual zero-shot learning where unseen information is not provided during training, by leveraging generative modeling. The heart of the generative-based methods is to learn quality representations from seen classes to improve the generative understanding of the unseen visual space. Motivated by this, we introduce generalization-bound tools and provide the first theoretical explanation for the benefits of generative modeling to CZSL tasks. Guided by the theoretical analysis, we then propose our learning algorithm that employs a novel semantically guided Generative Random Walk (GRW) loss. The GRW loss augments the training by continually encouraging the model to generate realistic and characterized samples to represent the unseen space. Our algorithm achieves state-of-the-art performance on AWA1, AWA2, CUB, and SUN datasets, surpassing existing CZSL methods by 3-7%. The code has been made available here <https://github.com/wx-zhang/IGCZSL>.

1. Introduction

Researchers have devoted significant effort to developing AI learners to mimic human cognition. One such endeavor is zero-shot learning (ZSL), which aims to identify unseen classes without accessing any of their images during training. However, human zero-shot learning abilities improve dynamically over time. As individuals acquire more knowledge of seen tasks, they become better at recognizing unseen tasks. To evaluate the zero-shot learning in such a dynamic

seen-unseen distribution, the continual zero-shot learning problem (CZSL) has been proposed [53]. CZSL emulates the continuous learning process of a human’s life, where the model continually sees more classes from the unseen world and is evaluated on both seen and unseen classes. This CZSL skill, may it get maturely developed to the world scale, has the potential to accelerate research in species discovery, for example, as known species grow continually, but close to 90% of the species are not yet discovered[55].

Generative models (e.g., GANs[26]) have made significant progress in producing photorealistic images by learning high-dimensional probability distributions. This ability motivated researchers to adapt GANs to ZSL to generate missing data of unseen classes conditioning on unseen semantic information, known as generative-based ZSL. Training the classifier on synthetic unseen samples can reduce model prediction bias towards seen classes and thus achieves competitive zero-shot learning performance [38, 58, 42]. Some CZSL works directly adopt this framework continually, known as transductive continual zero-shot learning [25, 36]. However, in CZSL, the unseen world changes dynamically and unexpectedly, making it unrealistic to use prior knowledge about unseen classes[53]. When we do not assume access to unseen semantic information in the CZSL setting, which is known as *inductive continual zero-shot learning*, most existing methods struggle to perform well, as we show in our experiments. Furthermore, the theoretical understanding of how zero-shot learning benefits from synthetic data is limited, which poses an obstacle to developing purely inductive continual zero-shot methods. Recent analyses of training generative models with synthetic data [8] provide a possible avenue for developing the desired theoretical explanation. This led us to develop a generalization-bound tool to understand the learning mechanism in generative-based CZSL and further develop inductive methods based on it.

In our analysis, we have identified it is crucial to reduce the distance between the generated and actual visual space of unseen classes. This requires the model to generate realistic samples to represent unseen space to augment the training

*Work done during internship at KAUST

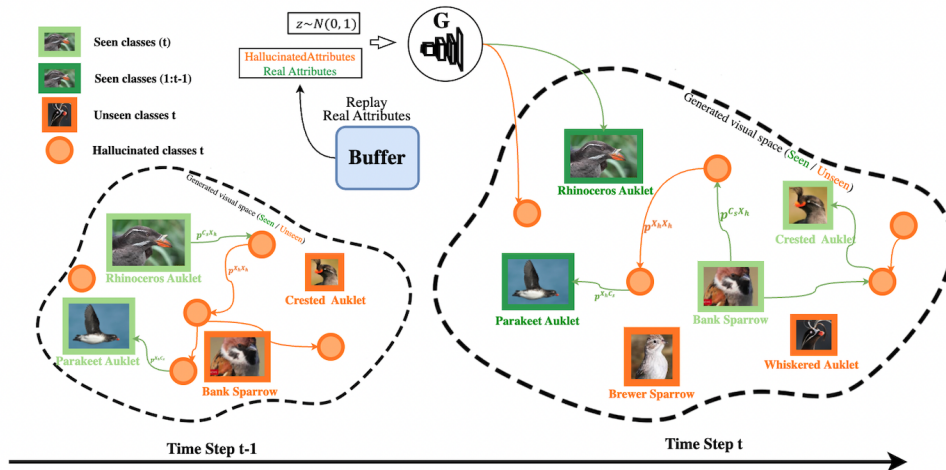


Figure 1. **Semantically guided generative random walk (GRW)**: At each time step, new classes are added to the seen classes space, and the random walk starts from each seen class center (in green) and transitions through generated samples of hallucinated classes (in orange), then the landing probability distribution over the seen classes is predicted. The GRW loss encourages the generated samples from the hallucinated classes to be distinguishable from the seen classes by encouraging the landing probability over seen classes starting from any seen center to be uniformly distributed, and hence hard to classify to any seen class.

of the classifier. However, the lack of ground truth semantic descriptions for unseen classes and the lack of previously seen classes data often leads to the generated samples collapsing to the seen classes. A similar problem has been addressed in generating novel style artworks, where GAN training is augmented to encourage the generated styles to deviate from existing art style classes [15, 51, 27, 30, 28, 31]. Drawing inspiration from the improved feature representation achieved by generative models in producing novel art, and the connection between the ability to generate novel styles in art generation and to generate samples to represent the unseen space in generative-based CZSL, we propose a purely inductive, **Generative Random Walk (GRW)** loss, guided only by semantic descriptions of seen classes.

In each continual learning task, we first hallucinate some classes by interpolating on or sampling from a learnable dictionary based on the current and previous classes, with the belief that the realistic classes, both seen and unseen, should be relatable to each other [16, 17]. We then generate samples from the hallucinated classes. To prevent the generated samples of hallucinated classes from collapsing to the seen classes, we apply the GRW loss, as illustrated in Figure 1. We perform a random walk starting from the **seen class** and moving through generated examples of hallucinated classes for R steps, as described in detail later in Section 5.2.2. The GRW loss encourages high transition probabilities to the **realistic unseen space** by deviating from the visual space of the seen classes and avoiding less realistic areas. The resulting representations are both realistic and distinguishable from seen classes, which enhances the generative understanding of unseen classes. This approach is particularly effective when the model is updated continually,

as it enables the model to use the newly learned knowledge to improve further the generated examples of hallucinated classes. Our contributions lie in

- We provide a theoretical analysis of continual zero-shot learning. This analysis guides us to use proper signals to make up for the missing unseen information. We present these generalization-bound tools for the analysis in Section 4.
- Guided by the analysis, we develop a method for purely inductive continual zero-shot learning; described in detail in Section 5. Our method, ICGZSL, first provides two ways to hallucinate classes, *i.e.* interpolation of two seen classes and learning a dictionary based on the seen classes. Then, we integrate our introduced semantically guided Generative Random Walk (GRW) loss to generate distinguishable and realistic samples to represent unseen classes.
- We performed comprehensive experiments (Section 6) that demonstrate the effectiveness of our approach. Specifically, our model achieves state-of-the-art results on standard continual zero-shot learning benchmarks, AWA1, AWA2, CUB, SUN, and performs often better than transductive methods

2. Related Works

Inductive and Transductive Zero-Shot Learning.

There are varying degrees of accessibility to unseen information in zero-shot learning. Transductive methods use both unlabeled samples and attributes of unseen classes during

training [44, 48]. Semantically transductive methods, on the other hand, only use attributes of unseen classes in training [63, 60]. In the inductive setting, however, no unseen information is allowed to be used (e.g., [69, 16, 41, 64]). This can result in a bias towards seen classes [45]. Generative methods, such as those used by [69, 41, 16], can produce unseen samples using only seen class information during training to solve this issue. For example, [16] relate zero-shot learning to human creativity to generate images that deviate from seen classes during training. [7] used unlabeled samples from out-of-distribution data to gather knowledge about unseen data. [54] utilize two variational autoencoders to generate latent representations for visual and semantic modalities in a shared latent space. In contrast, our approach focuses on investigating the relationship between the generated samples of hallucinated classes and the seen classes, which leads to GRW loss.

Continual Learning. The majority of continual learning works aim to tackle the problem of catastrophic forgetting, where the data representation becomes biased towards the most recent task in sequential learning. Regularization-based methods [39, 4], structure-based methods [49, 13], and replay-based methods [52, 65] have been proposed to resolve this problem. More recently, research has explored forward transfer in continual learning, with the belief that as knowledge accumulates, higher next-task transferability, as measured by zero-shot assessment, should be attained. Their evaluation space either includes the next task [40] or the whole class space [12]. However, compared to our setting, [40] did not evaluate the model in a generalized manner, and [12] only paid attention to the seen accuracy.

Continual Zero-shot learning. [11] introduced A-GEM for continual learning, which was later applied to deal with zero-shot tasks sequentially, laying the foundation for the initial work on CZSL. [53] proposed the inductive CZSL scenario and demonstrated that a class-based normalization approach can improve performance in continual zero-shot learning. Both [21] and [24] explore the CZSL problem, but rely on unseen class descriptions to train a classifier before inference. [36] proposed a generative adversarial approach with a cosine similarity-based classifier that supports the dynamic addition of classes without requiring unseen samples for training. Their approach also relies on unseen class descriptions for seen-unseen deviation, making it a semantically transductive method. This motivated us to explore a purely inductive method for handling seen-unseen deviation and improving the realism of unseen samples.

3. Problem Setup and Notations

3.1. Formulation

We start by defining our problem and notations. A labelled dataset is defined as a tuple $D = \{(\mathbf{x}, \mathbf{a}, y) | y =$

$f(\mathbf{x}), (\mathbf{x}, \mathbf{a}, y) \sim \mathcal{D}\}$, where \mathcal{D} represents the data distribution. Each data point is a tuple of image feature $\mathbf{x} \in \mathbb{R}^{d_x}$, class attribute $\mathbf{a} \in \mathbb{R}^{d_a}$, and a class label y . Here d_x is the dimension of the visual feature space, and d_a is the dimension of the attribute space. Each distribution has a specific labeling function f . Our goal is to learn a model \hat{f} on top of D to estimate f . We study the continual zero-shot learning setting proposed by [53], where we seek to learn the model \hat{f} on a stream of tasks. In each task t , the model is learned on the seen dataset D_s^t , and is evaluated on both the seen distribution \mathcal{D}_s^t and unseen distribution \mathcal{D}_u^t . Moreover, we assume that the set of seen class and unseen class are disjoint, that is $\mathcal{D}_s \cap \mathcal{D}_u = \phi$. This procedure is illustrated in the bottom part of Figure 2.

We use generative models as the backbone. During the training time, the model \hat{f} is trained on the seen dataset D_s as well as the synthesized dataset D_h . D_h is generated by conditioning on hallucinated attributes \mathbf{a}_h and prior $\mathcal{Z} \sim \mathcal{N}(0, 1)$. The labeling function f_h of the generated dataset is a look-up table of the generated features $\mathbf{x} \in X_h$ and the corresponding attribute condition \mathbf{a}_h .

3.2. Notations.

In our theoretical analysis, we use the following notations: 1) We discuss the relationship between the three types of variables, namely, real seen sample, real unseen samples, and generated samples from the hallucinated classes. To specify the variables related to these types of samples, we use subscripts $\cdot_s, \cdot_u, \cdot_h$ respectively, e.g., f_s, f_u, f_h ; 2) We denote the values and model empirically computed by a variable with a hat, e.g., \hat{f} ; 3) We use superscripts \cdot^t or $\cdot^{1:t}$ to indicate that a variable is for task t or for tasks $1 : t$ respectively, e.g., $f_s^t, f_s^{1:t}$; 4) D is used for the empirical sample set, and \mathcal{D} is used for the distribution; 5) We use N_s and N_u to denote the number of seen and unseen classes.

In practice, the unseen information, i.e., \mathbf{a}_u, D_u, N_u , is not available. Therefore, we hallucinate some classes denoted by \mathbf{a}_h and generate samples $D_h = \{(\mathbf{x}_h, \mathbf{a}_h)\}$ by conditioning on these attributes. We use N_h to represent the number of hallucinated classes. Additionally, we do not have access to all the previous data, so $X^{1:t}$ refers to the current samples as well as the previous ones in the buffer. We also use generated seen samples \cdot_{sg} for GAN training.

4. Theoretical Analysis

As mentioned in the introduction, we propose using hallucinated classes to represent the unseen space. By training our model on synthetic samples generated from these classes, we improve the model’s generalization ability to the actual unseen classes during the testing time of continual zero-shot learning. In this section, we quantify the model’s generalization ability by measuring the distance between the synthetic samples that represent the unseen space and the actual un-

seen samples. Additionally, we explain our motivation for using a random walk-based method to reduce this distance when no information about the unseen space is available.

4.1. Generalization Bound Inductive Continual Zero-Shot Learning

In this section, we present a generalization bound for a continual zero-shot learning algorithm. Given the entire training distribution, a learning algorithm can output an optimal hypothesis h that estimates the ground truth labeling function f . However, since the learning algorithm can only be trained on a finite sample from the training set, it outputs an empirical hypothesis \hat{h} to estimate the ground truth labeling function. We define the generalization error [33] for these two types of hypotheses. We define the actual risk,

$$\epsilon(h, f) = \mathbb{E}_{(\mathbf{x}, \mathbf{a}) \sim \mathcal{D}} [\mathbb{1}_{f(\mathbf{x}) \neq h(\mathbf{x}, \mathbf{a})}] , \quad (1)$$

which measures the expected probability of a disagreement between the ground truth and the optimal hypothesis. We also define the empirical risk on the finite sample set \mathcal{D} ,

$$\hat{\epsilon}(\hat{h}, f) = \frac{1}{|\mathcal{D}|} \sum_{(\mathbf{x}, \mathbf{a}) \in \mathcal{D}} \mathbb{1}_{f(\mathbf{x}) \neq \hat{h}(\mathbf{x}, \mathbf{a})} , \quad (2)$$

which measures the probability of a disagreement between the ground truth and the empirical hypothesis.

In a continual zero-shot learning algorithm, given a training set \mathcal{D}_s^t , the algorithm outputs \hat{h} to estimate $f_s^{1:t} \cup f_u$ instead of the ground truth labeling function f_s . To begin our analysis, we propose a distance measure between the generated unseen distribution¹ and the real unseen distribution $\bar{d}_{GDB}(\mathcal{D}_h, \mathcal{D}_u)$ as follows:

Definition 4.1 (Empirical Generative distance). *Given the training set \mathcal{D}_s and the synthetic set \mathcal{D}_h , the ground truth labeling functions f_s, f_h , and f_u , and the optimal hypothesis $\hat{h}^* = \arg \min_{h \in H} \hat{\epsilon}_s(h, f_s) + \hat{\epsilon}_h(h, f_h)$ obtained by training the model on \mathcal{D}_h and \mathcal{D}_s , we can define the distance between \mathcal{D}_h and \mathcal{D}_u as follows:*

$$\bar{d}_{GDB}(\mathcal{D}_h, \mathcal{D}_u) = |\hat{\epsilon}(\hat{h}^*, f_u) - \hat{\epsilon}(\hat{h}^*, f_h)| . \quad (3)$$

Our proposed \bar{d}_{GDB} is a feasible distance measure that satisfies the properties of a pseudo-metric. In the following, we present our generalization bound following [8] for the continual zero-shot learning algorithm, which shows how the generalization ability of the zero-shot learning algorithm is mainly influenced by this distance.

Theorem 4.2 (Generalization bound of the generative-based CZSL). *Given the CZSL procedure described in section 3.1,*

¹In the transductive setting, the unseen distribution is generated by conditioning on the unseen semantic information. In our work, we utilize generated samples from hallucinated classes to represent the generated unseen distribution.

with confidence $1 - \delta$ the risk on the unseen distribution is bounded by

$$\begin{aligned} \epsilon(h, f_u^t) \leq & \hat{\epsilon}(\hat{h}^*, f_s^{1:t}) + \frac{1}{2} d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}_s^{1:t}, \mathcal{D}_u^t) + \bar{\lambda} \\ & + \frac{1}{2} \bar{d}_{GDB}(\mathcal{D}_u^t, \mathcal{D}_h^t) + C\left(\frac{1}{m} + \frac{1}{\delta}\right) \end{aligned} \quad (4)$$

where $\hat{h}^* = \arg \min_{h \in H} \sum_{i=1}^t \hat{\epsilon}(h, f_s^i) + \hat{\epsilon}(h, f_h^t)$, $\bar{\lambda} = \hat{\epsilon}(\hat{h}^*, f_s^{1:t}) + \hat{\epsilon}(\hat{h}^*, f_h^t)$.

In Equation 24, measurement $d_{\mathcal{H}\Delta\mathcal{H}}$ [6] is used to quantify the difference between two distributions for domain adaptation based on the type of model, and is fixed for a specific problem. $\bar{\lambda}$ and $\hat{\epsilon}_s(h, f_s)$ are highly depended on the optimization algorithm. However, if we hallucinate a diverse set of classes, f_u can be compactly supported by f_h . If we further generate realistic samples for each of the hallucinated classes, the optimal solution trained on the synthetic set, $\hat{h}^* = \arg \min_{h \in H} \hat{\epsilon}_s(h, f_s) + \hat{\epsilon}_h(h, f_h)$, should perform well on the real unseen dataset. This can lead to a reduction in $\bar{d}_{GDB}(\mathcal{D}_u, \mathcal{D}_h)$ in Equation 3. We will discuss this further in the following section. The detailed derivation of this theorem can be found in Appendix A.1.

4.2. Reducing the bound using Markov Chain.

To reduce $\bar{d}_{GDB}(\mathcal{D}_h, \mathcal{D}_u)$ in Equation 24, we need to decrease the difference between \mathcal{D}_u and \mathcal{D}_h . One approach proposed by [16] is to hallucinate \mathbf{a}_h as a compact support of \mathbf{a}_u . Once we have achieved this, we can further generate high-quality samples to increase $\mathbb{P}[\mathcal{D}_u \subset \mathcal{D}_h]$, where the probability is taken over all possible generations.

To quantify the probability value $\mathbb{P}[\mathcal{D}_u \subset \mathcal{D}_h]$, we follow the approach of [57] and view the generations as nodes in a Markov chain. We define the transition probability between two states as the probability with which one sample is classified as another. Then, we can bound $\mathbb{P}[\mathcal{D}_u \subset \mathcal{D}_h]$ by the self-transition probability using a generalization bound. When the self-transition probability is the same in two sets of generations, we prefer the one with higher diversity quantified by DDP, as suggested by [32] and [14].

For detailed explanations, please refer to Appendix A.2. Here, we provide an informal statement.

Statement 4.3. *Finding generated samples from hallucinated classes to “carefully” increase the determinant and the diagonal entries of the transition matrix of the above described Markov Chain can reduce \bar{d}_{GDB} .*

We can now design an algorithm that first hallucinates classes and then generates diverse samples from these classes to represent the unseen space that follows Statement 4.3. However, the transition matrix of the Markov chain described above is intractable to compute in practice. To quantify the transition probability, we adapt the random walk framework [29, 5] originally used in semi-supervised few-shot

learning to generative zero-shot learning with a few yet important changes. Please refer to Appendix B.4 the relation between our work and previous work.

We also make the following two adjustments to the statement to encourage the generated samples from hallucinated classes to be consistently realistic like the real samples. Firstly, we represent the transition matrix among hallucinated classes (noted as $\mathbf{P}^{X_h X_h} \in \mathbb{R}^{N_h \times N_h}$) in the seen class space using a congruent transformation $\mathbf{P}^{C_s X_h} \mathbf{P}^{X_h X_h} \mathbf{P}^{X_h C_s}$, where $\mathbf{P}^{C_s X_h} \in \mathbb{R}^{N_s \times N_h}$ is the transition probability matrix from seen prototypes to generated samples from hallucinated classes, and $\mathbf{P}^{X_h C_s}$ is the opposite. Secondly, hallucinating a compact support of unseen class attributes and encouraging the transition matrix to be diagonal requires a huge number of generations. To reduce this number, we encourage the generated samples of hallucinated classes to have a "relatable deviation" to the seen classes. The relationship between the two types of samples is that both should be realistic. This means that the transition matrix $\mathbf{P}^{X_h X_h}$ may not be strictly diagonal, and our goal is to reduce the non-diagonal entries, *i.e.*, to reduce the transition probability between different generated samples of hallucinated classes. We further repeat the transition among generated samples of hallucinated classes to further reduce the non-diagonal entries.

In conclusion, our transitions start from the seen prototypes to generated samples of hallucinated classes for R steps and back to seen prototypes, the transition matrix of which is $\mathbf{P}^{C_s X_h} (\mathbf{P}^{X_h X_h})^R \mathbf{P}^{X_h C_s} \in \mathbb{R}^{N_s \times N_s}$. To encourage "relatable deviation" of the generated samples of hallucinated classes from seen classes, we aim to reduce the non-diagonal entries of the transition matrix, as detailed later in Section 5. This approach intuitively prevents the generations from being attracted by any seen classes, and theoretically can reduce the distance d_{GDB} . Intentionally, this method also transfers knowledge between seen and hallucinated classes, which is useful for generating realistic images.

5. Generative-based Inductive CZSL Approach

Method overview. Generative-based inductive CZSL algorithms adopt generative models as their architecture, where seen samples are used to train the classifier to correctly classify seen classes, and generators are trained to generate realistic samples. At the same time, the generator is encouraged to synthesize samples to represent unseen classes to train the classifier to perform classification on these samples. In our work, we can only hallucinate some classes to represent the actual unseen space and generate samples from the hallucinated classes. As guided by our analysis, the key point of inductive zero-shot learning is to generate realistic and diverse samples from hallucinated classes that are deviated from the real seen space. We introduce the gen-

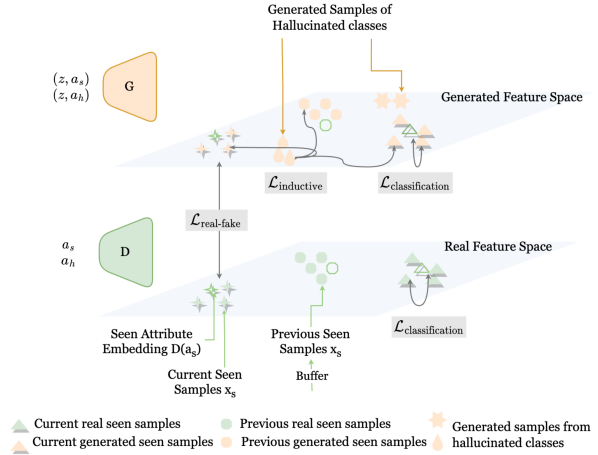


Figure 2. The discriminator embeds attributes $\mathbf{a}_s^{1:t}$ into the real feature space to perform classification with real samples x_s^t , while the generator produces features x_h^t and x_{sg}^t conditioning on the corresponding attributes. The real-fake loss and classification loss encourage the generated sample distribution consistent with the real samples. Then the inductive loss applied to the generated feature space, which encourages the characterization of generated samples from hallucinated classes, can reduce the bias towards the current seen classes of the classifier and improve continual zero-shot learning performance.

erative model backbone in Section 5.1, and how we generate the abovementioned samples in Section 5.2. The overall procedure is shown in Algorithm 1 in Appendix.

5.1. Generative-based CZSL baseline

We follow [36] as our baseline. The model contains a generator $G(\mathbf{a}, \mathbf{z}) : \mathbb{R}^{d_a + d_z} \rightarrow \mathbb{R}^{d_x}$ and a discriminator $D(\mathbf{a}) : \mathbb{R}^{d_a} \rightarrow \mathbb{R}^{d_x}$. The generator takes the semantic information (denoted by \mathbf{a}) and the prior (denoted by \mathbf{z}) sampled from a standard normal distribution \mathcal{Z} as input and outputs visual features. Discriminator projects semantic information \mathbf{a} into visual space. The conditional adversarial training can be illustrated by the discriminator loss and generator loss as:

$$\begin{aligned} \mathcal{L}_D &= -\mathcal{L}_{\text{real-fake}} + \lambda_{\text{cls}} \mathcal{L}_{\text{classification}} + \lambda_{\text{rd}} \mathcal{R}_D, \\ \mathcal{L}_G &= \mathcal{L}_{\text{real-fake}} + \lambda_{\text{cls}} \mathcal{L}_{\text{classification}} + \mathcal{L}_{\text{inductive}} + \lambda_{\text{rg}} \mathcal{R}_G. \end{aligned} \quad (5)$$

As shown in Figure 2, we use $\mathcal{L}_{\text{real-fake}}$ to denote the GAN loss that discriminates between the real and fake samples for the current task, and $\mathcal{L}_{\text{classification}}$ to denote the entropy loss based on cosine similarity that is used to perform the classification of all seen classes up to the current task. The equations for $\mathcal{L}_{\text{real-fake}}$ and $\mathcal{L}_{\text{classification}}$ are shown below

$$\begin{aligned} \mathcal{L}_{\text{real-fake}} &= \mathbb{E}_{(\mathbf{x}, \mathbf{a}) \sim D_s^t} [\log(\langle \mathbf{x}, D(\mathbf{a}) \rangle)] \\ &\quad - \mathbb{E}_{\mathbf{z} \sim \mathcal{Z}, (\mathbf{a}, \mathbf{x}) \sim D_s^t} [\log(\langle G(\mathbf{z}, \mathbf{a}), D(\mathbf{a}) \rangle)] \\ \mathcal{L}_{\text{classification}} &= \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim D_s^{1:t}} [L_e(\langle \mathbf{x}, D(\mathbf{A}_s^{1:t}) \rangle, \mathbf{y})] \\ &\quad + \mathbb{E}_{\mathbf{z} \sim \mathcal{Z}, (\mathbf{x}, \mathbf{a}, \mathbf{y}) \sim D_s^{1:t}} [L_e(\langle G(\mathbf{z}, \mathbf{a}), D(\mathbf{A}_s^{1:t}) \rangle, \mathbf{y})] \end{aligned} \quad (6)$$

where $\langle \cdot, \cdot \rangle$ represents the cosine similarity, $\mathbf{A}_s^{1:t}$ is a matrix of attributes of seen classes up to the current task, and L_e is the cross-entropy loss. In practice, $\mathbf{D}_s^{1:t}$ consists of the current samples and previous samples in the buffer. We follow [36] for regularization terms $\mathcal{R}_D, \mathcal{R}_G$ and $\lambda_{c,rd,rg}$. See Appendix B.1 for more details about the baseline algorithm. $\mathcal{L}_{\text{inductive}}$ with its corresponding λ_i is the main component to improve inductive continual zero-shot learning, which will be described in detail in section 5.2.

5.2. Inductive Loss

5.2.1 Hallucinate Attributes

To begin our method, we first hallucinate classes to represent the unseen space. During this procedure, we aim to generate a diverse and compact set of attributes without using any information from the unseen test set.

Interpolation-based method. When the attributes are distributed uniformly in the attribute space, which can be compactly supported by the seen attribute, we use interpolation-based method. To hallucinate the attributes at every mini-batch, we use an interpolation-based method that was introduced by [16]. Hallucinated attributes are generated using the formula $\mathbf{a}_{ug} = \alpha \mathbf{a}_{s_1} + (1 - \alpha) \mathbf{a}_{s_2}$, where α is drawn from the uniform a distribution $\mathcal{U}(0.2, 0.8)$, and \mathbf{a}_{s_1} and \mathbf{a}_{s_2} are two randomly chosen seen attributes. The sample interval is chosen to be $(0.2, 0.8)$ to ensure that the interpolated attributes are not too close to the seen attributes.

Dictionary-based method. We further propose to learn an attribute dictionary containing N_s^t attribute vectors during training. The use of a learnable dictionary allows the attributes to change more freely in accordance with the loss function. The dictionary is randomly initialized by interpolating seen attributes, and during the computation of GRW loss, we randomly pick attributes from it. This approach is particularly useful for classification at a finer level, where the attributes are more specific.

If the hallucinated class can accurately represent the actual unseen space, which is only accessible during the test time, then the model will have good generalization ability on the test set. We visualize the hallucinated classes to examine if this assumption holds. Please refer to Appendix B.2 for the visualization of our hallucinated attributes.

5.2.2 Improve Generation Quality by Inductive Loss

As we discussed in Section 4.2, we use GRW loss to improve the generation quality such that the generated samples are realistic, diverse and characterized. To encourage diversity of the samples generated from the hallucinated attributes, we firstly generate only one sample for each hallucinated attribute. And then, we perform a random walk to compute the transition probability using generated seen samples \mathbf{X}_{sg} and generated samples from hallucinated samples \mathbf{X}_h . The

random walk starts from each generative seen class center $\mathbf{C}_s \in \mathbb{R}^{N_s^{1:t} \times d_x}$ computed by the mean of generated seen samples from the corresponding class attributes, where $N_s^{1:t}$ are the number of seen classes until step t . Then we take R steps of transitions within generated samples of hallucinated classes \mathbf{X}_h with the final landing probability over seen classes so far. The transition probability matrix from seen class centers to generated samples of hallucinated classes is defined as

$$\mathbf{P}^{C_s X_h} = \sigma(\langle \mathbf{C}_s, \mathbf{X}_h^\top \rangle), \quad (7)$$

where $\langle \cdot, \cdot \rangle$ is a similarity measure, and $\sigma(\cdot)$ is a softmax operator applied on rows. In practice, we use negative Euclidean distance for similarity, that is, suppose \mathbf{x}_h is the row i of \mathbf{X}_h and \mathbf{c} is the class center j ,

$$\langle \mathbf{C}_s, \mathbf{X}_h^\top \rangle_{i,j} = -\|\mathbf{x}_h - \mathbf{c}\|^2. \quad (8)$$

Similarly, the transition probability matrix within generated samples of hallucinated classes and from generated samples of hallucinated classes to seen class centers are defined as

$$\mathbf{P}^{X_h X_h} = \sigma(\langle \mathbf{X}_h, \mathbf{X}_h^\top \rangle), \mathbf{P}^{X_h C_s} = \sigma(\langle \mathbf{X}_h, \mathbf{C}_s^\top \rangle). \quad (9)$$

Then the random walk starting from each seen class center and transiting R steps within generated samples of hallucinated classes and back to seen centers are computed by

$$\mathbf{P}^{C_s X_h C_s}(R) = \mathbf{P}^{C_s X_h} (\mathbf{P}^{X_h X_h})^R \mathbf{P}^{X_h C_s} \quad (10)$$

In practice, we set the diagonal values of $\mathbf{P}^{X_h X_h}$ to small values and hope to reduce the non-diagonal values. This equals to encourage the probability $\mathbf{P}^{C_s X_h C_s}(R)$ to be uniformly distributed over all the seen classes. We further encourage the probability $\mathbf{P}^{C_s X_h} \in \mathbb{R}^{N_s^{1:t} \times N_h}$ to be uniformly distributed over all the generated examples to encourage as many generations to be visited in the random walk, and hence encourage the diversity. Hence, our *Generative Random Walk* (GRW) loss is defined by

$$L_{\text{GRW}} = \sum_{r=0}^R \gamma^r L_e(\mathbf{P}^{C_s X_h C_s}(r), \mathcal{U}) + L_e(\mathbf{P}_v(C_s, X_h), \mathcal{U}_v), \quad (11)$$

where $L_e(\cdot, \cdot)$ is the cross-entropy loss, $\mathcal{U} \in \mathbb{R}^{N_s^{1:t} \times N_s^{1:t}}$ is uniform distribution, R is the transition steps, and γ is exponential decay. We compute the probability that each generated point be visited by any seen class as $P_v(C_s, X_h) = \frac{1}{N_s^{1:t}} \sum_{i=0}^{N_s^{1:t}} \mathbf{P}_i^{C_s X_h}$, where $\mathbf{P}_i^{C_s X_h}$ represents the i^{th} row of the $\mathbf{P}^{C_s X_h}$ matrix. The visit loss is then defined as the cross-entropy between P_v and the uniform distribution $\mathcal{U}_v \in \mathbb{R}^{N_h}$, encouraging all the generated examples to be visited. In addition, we empirically found that the GRW loss can also work as a regularizer to encourage the consistency of generated seen visual space as well, which we defined as

$$R_{\text{GRW}} = \sum_{r=0}^R \gamma^r L_e(\mathbf{P}^{C_s X_{sg} C_s}(r), \mathcal{I}) + L_e(\mathbf{P}_v(C_s, X_{sg}), \mathcal{U}_v), \quad (12)$$

Dataset	AWA1			AWA2			CUB			SUN		
	mSA	mUA	mHA	mSA	mUA	mHA	mSA	mUA	mHA	mSA	mUA	mHA
EWC (<i>cl</i>) [34]	29.4	9.0	13.8	30.8	10.5	15.8	12.2	0.8	1.3	11.6	2.6	4.1
A-GEM (<i>cl</i>) [11]	64.2	3.9	7.2	65.8	6.7	11.9	14.4	0.4	0.8	8.6	3.0	4.2
Tf-GZSL (<i>tr</i>) [22]	70.8	27.4	37.9	78.6	28.7	41.1	46.3	30.8	35.3	15.3	30.7	18.7
DVGR (<i>tr</i>) [25]	65.1	28.5	38.0	73.5	28.8	40.6	44.9	14.6	21.7	22.4	10.7	14.5
A-CGZSL (<i>tr</i>) [24]	71.0	24.3	35.8	70.2	25.9	37.2	34.3	12.4	17.4	17.2	6.3	9.7
BD-CGZSL (<i>tr</i>) [36]	62.9	29.9	39.0	68.1	33.9	42.9	19.8	17.2	17.8	27.5	15.9	20.0
CN-CZSL (<i>in</i>) [53]	-	-	-	33.6	6.4	10.8	44.3	14.8	22.7	22.2	8.2	12.5
BD-CGZSL-in (<i>in</i>) [36]	62.1	31.5	40.5	67.7	32.9	42.3	37.8	9.1	14.4	34.9	14.9	20.8
CARNet (<i>in</i>) [23]	67.6	27.4	37.0	-	-	-	42.4	12.4	18.8	31.5	15.9	20.9
ours + interpolation	67.0	34.2	43.4	71.1	34.9	44.5	42.2	22.7	28.4	36.0	21.6	26.8
ours + dictionary	67.1	33.5	41.6	70.2	35.1	44.6	42.4	23.6	28.8	36.5	21.8	27.1

Table 1. Our proposed method achieves state-of-the-art results when compared with traditional continual learning method (*cl*) recent inductive (*in*) methods and even shows competitive results in mHA with recent semantic transductive methods (*tr*).

where \mathcal{I} is identity distribution, and D_{sg} represents the matrix for generated seen samples.

We numerically show that the random walk-based penalty can reduce \bar{d}_{GDB} (Def 4.1) by the relationship between \bar{d}_{GDB} and L_{GRW} . Details are shown in Appendix B.3.

We also adapt the loss proposed in [16] to directly prevent the generated unseen samples from being classified into seen classes, i.e.,

$$L_{creativity} = \mathbb{E}_{z \sim \mathcal{Z}, a_h \sim D_h} D_{KL}(\langle G(z, a_h), D(A_s^{1:t}) \rangle \| \mathcal{U}), \quad (13)$$

where $D_{KL}(\cdot \| \cdot)$ is the KL divergence, $A_s^{1:t} \in \mathbb{R}^{N_s^{1:t} \times d_a}$ is the matrix of seen classes attributes vectors until task t , a_h is hallucinated attributes according to Section 5.2.1, $\langle G(z, a_{ug}), D(A_s^{1:t}) \rangle \in \mathbb{R}^{N_s^{1:t}}$ are the logits over seen classes so far for a given $G(z, a_h)$, \mathcal{U} is the uniform distribution.

Inductive loss Combining Equation 11, 12 and 13 our final inductive loss is

$$\mathcal{L}_{inductive} = \lambda_c L_{creativity} + \lambda_i L_{GRW} + \lambda_r R_{GRW} \quad (14)$$

where λ_i is the scaling weight for both the GRW loss term and regularization term.

6. Continual Zero-Shot Learning Experiment

6.1. Experiment Setup

Data Stream and Benchmarks: We adopt the continual zero-shot learning framework proposed in [53]. In this setting, a T -split dataset $D^{1:T}$ forms $T - 1$ tasks. At time step t , the split $D^{1:t}$ is defined as a seen set of tasks, and the split $D^{t+1:T}$ is an unseen set of tasks. We conduct experiments on four widely used CGZSL benchmarks for a fair comparison: AWA1 [37], AWA2 [67], Caltech UCSD Birds 200-2011 (CUB)[59], and SUN[43]. We follow [53, 36] for the class split in the continual zero-shot learning setting. More details can be found in Appendix D.

Baselines, backbone, and training: We use the method proposed in [36] as the main baseline and compare it with recent CGZSL methods in the setting we mentioned above, including the transductive method Tf-GZSL [22], DVGR [25], A-CGZSL [24], BD-CGZSL [36], and the inductive method CN-CZSL [53], CARNet [23]. ‘BD-CGZSL-in’ denotes our modified inductive version of [36] by naively removing unseen information. Following [36], we also compare our baseline with the classical continual learning methods EWC [34] and A-GEM [11]. We use vanilla GAN’s Generator and Discriminator, both of which are two-layer linear networks. Image features are extracted by ResNet-101, pre-trained on ImageNet 1k. The are attributes from [62] and extracted features are used as our model input. We use a replay buffer with a fixed size of 5k.

We run all experiments for 50 epochs and 64 batch sizes with the Adam optimizer. We use a learning rate of 0.005 and a weight decay of 0.00001. Results reported in Table 1 are based on one NVIDIA Tesla P100 GPU. We select our random walk steps R , weight decay γ and coefficient of inductive loss terms λ_i according to prior exploratory zero-shot learning experiments shown in Appendix C.

Metrics: We use the mean seen accuracy, mean unseen accuracy and mean harmonic seen/unseen accuracy [53] to measure the zero-shot learning ability. These metrics are defined as follows,

$$\begin{aligned} \text{mSA} &= \frac{1}{T} \sum_{t=1}^T S_t(D^{1:t}), \text{mUA} = \frac{1}{T-1} \sum_{t=1}^{T-1} U_t(D^{t+1:T}) \\ \text{mHA} &= \frac{1}{T-1} \sum_{t=1}^{T-1} H(S_t(D^{1:t}), U_t(D^{t+1:T})), \end{aligned} \quad (15)$$

where $H(\cdot, \cdot)$ is the harmonic mean and S_t, U_t are seen and unseen per-class accuracy using the model trained after time t . We also use the backward transfer [11, 66, 53] to measure the continual learning ability, which is defined in [53]

$$\text{BWT} = \frac{1}{T-1} \sum_{t=1}^{T-1} (S_T(D^{1:t}) - S_t(D^{1:t})) \quad (16)$$

Note that this should only be conducted on seen set, since part of the early unseen set become seen set later. The BWT on unseen set cannot reflect the knowledge retain ability of the model.

6.2. Results

The mean harmonic accuracy of the four benchmarks is shown in Table 1, and the task-wise mHA of the CUB dataset is shown in Figure 3. In coarse-grained datasets AWA1 and AWA2, our proposed learner achieves 43.4% and 44.6% in mHA, respectively, surpassing all the current inductive and transductive methods. In the fine-grained datasets and tasks with long steps (CUB, SUN), our method achieves 28.8% and 27.1%, surpassing all the current CZSL methods. We observe that even though other methods have comparable mSA, they have far lower mUA than ours. We believe that our method achieves this improved knowledge transfer ability from seen visual space to unseen visual space through the proposed inductive learning signals, *i.e.*, $\mathcal{L}_{\text{inductive}}$. Table 2 displays the backward transfer of different continual zero-shot algorithms, where higher results indicate better knowledge retention. Our model exhibits a strong backward transfer capability, particularly on longer task sequences where it is needed the most. We achieved the highest BWT score of 0.19 on CUB. On SUN, negative BWT scores (*i.e.*, forgetting) are observed in most other models, but our method can still retain knowledge from the past. These results suggest that the analysis tools we created allow us to identify the critical factors for zero-shot learning, and the development of tools for continual learning can improve our ability to retain information.

6.3. Ablation Study

To assess the impact of our novel random walk-based penalties, L_{GRW} and \mathcal{R}_{GRW} , we conducted ablation experiments; see Table 3. The results in Table 3 indicate that the improvements are mainly attributed to L_{GRW} , while \mathcal{R}_{GRW} contributes an additional 1%. $L_{\text{creativity}}$ is also part of the inductive loss. Additionally, removing $L_{\text{creativity}}$ while using our GRW losses has little effect on the performance, as shown in Table 3. More details can be found in

Dataset		AWA1	AWA2	CUB	SUN
DVGR [25]	tr	0.09	0.10	-0.07	-0.20
A-CGZSL [24]	tr	0.11	0.05	0.10	0.005
BD-CGZSL [36]	tr	0.18	0.14	0.13	-0.02
CN-ZSL [53]	in	-	-	-0.04	-0.02
BD-CGZSL-in [36]	in	0.18	0.15	0.14	-0.03
ours + interpolation	in	0.12	0.10	0.19	0.01
ours + dictionary	in	0.11	0.11	0.19	0.01

Table 2. Backward transfer of different CGZSL methods, where higher results indicate less forgetting.

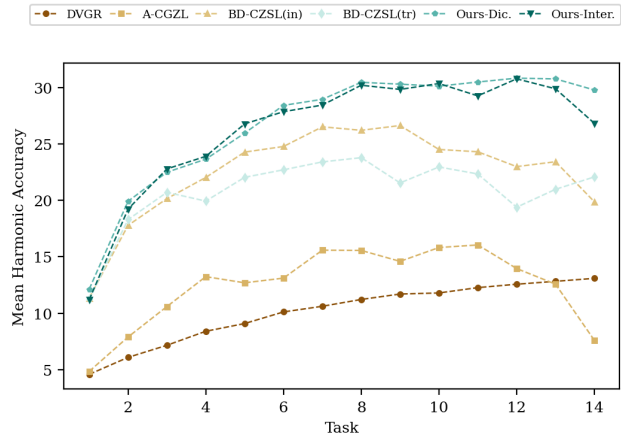


Figure 3. Mean Harmonic accuracy up to each task on SUN dataset. Our method outperforms both transductive and inductive methods

Appendix D.

6.4. More CZSL settings

Our focus lies on assessing performance under varying seen/unseen class ratios during knowledge accumulation, which is proposed in [53] and referred to as static setting in [36]. There are other continual zero-shot learning settings proposed in [36], such as dynamic and online settings. In the dynamic setting, the seen and unseen classes dynamically increase, while in the online setting, certain unseen classes are continually converted to seen classes. We find our explored static setting a more informing benchmark for the inductive CZSL skill, as the evaluation after every task is always performed on all classes in the dataset and hence is more challenging. Despite this, we still provide a comparison between our method and the baseline methods on the dynamic and online settings; see Table 4. The results show that our method is superior to the baseline in the dynamic and online settings in almost all datasets, and gains the most improvements in the most challenging static setting.

6.5. Replay Buffer Analysis

Some existing methods [24, 25, 36] tend to use the generative replay method proposed by [21], where the correctly predicted seen generated features from the previous task are

	interpolation	dictionary
with $R_{GRW} + L_{GRW}$	28.4	28.8
- $L_{\text{creativity}}$	27.72	27.66
w/o R_{GRW}, L_{GRW}	19.07	20.75
- $L_{\text{creativity}}$	14.43	14.43
with L_{GRW}	26.73	27.39

Table 3. Effect of the random walk-based penalty with mH measure on CUB dataset.

	setting	AWA1	AWA2	CUB	SUN
BD-CGZSL	D	56.9/49.1	56.4	16.8	28.0
ours + inter.	D	60.0	58.8	32.8	41.6
ours + dic.	D	59.7	55.5	31.8	40.2
BD-CGZSL	O	56.9/49.1	53.4	28.4	33.7
ours + inter.	O	49.6	48.5	32.3	39.6
ours + dic.	O	46.1	47.3	31.2	39.2

Table 4. mH in dynamic setting (D) and online setting (O)

	Buffer Size	Ours		BD-CGZSL (<i>tr</i>)
		BWT	mHA	mHA
generative	28.5k	0.14	21.06	17.76
real	10k	0.17	28.44	27.79
real	5k	0.19	28.8	26.55
real	2.5k	0.08	26.99	26.77

Table 5. Comparison of generative replay and real replay methods on CUB [59]. Dictionary-based attribute generation is used

stored in buffers. However, the buffer size increases significantly over tasks since a fixed number of samples for each class is stored, and if the model struggles to make accurate predictions for certain classes, samples from these classes are absent in the buffer.

We empirically found that the class-balanced experience replay method proposed by [46] can be extremely helpful. At every task, we save the class attribute in $A^{1:t}$, the class center matrix C , and modify the buffer with current features noted as $D_s^{1:t}$, such that the buffer is balanced across all the seen classes.

In this comparison on the CUB dataset, we observe in Table 5 that the method using real replay can achieve better harmonic accuracy with a smaller buffer size (around 1/10 of the generative replay buffer size) and comparable backward transfer with a slightly larger buffer size (around 1/5 of the generative replay buffer size). Moreover, the real replay-based method is not as sensitive to the buffer size as the generative replay-based methods. It is worth noting that DVGR, A-CGZSL, and BD-CGZSL typically use generative replay, while only CN-CGZSL uses real replay. In addition, the last column in Table 5 shows that our proposed real replay method can also improve the harmonic accuracy of other methods.

To understand the real replay and generative replay, we extend our analysis by visualizing the distribution of buffer features across various classes in task 2 of the SUN dataset, as illustrated in Figure 4. Real replay approach exhibits a balanced allocation of features across all classes. Conversely, the generative replay technique displays an intriguing pattern, wherein certain classes lack a substantial number of stored features, while others exhibit a twofold increase. Notably, the classes with fewer stored features coincide with instances where the model’s performance is suboptimal. This discrepancy can be attributed to the generative replay method’s

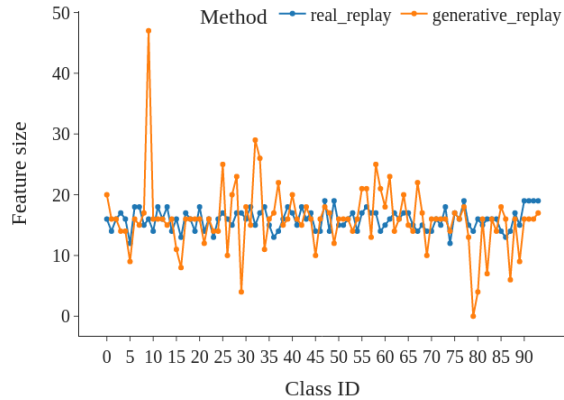


Figure 4. Comparison of replayed number of features per class in different replay method at task 2 in SUN dataset

propensity to store exclusively the accurately classified generated data. Consequently, this uneven distribution of stored features could potentially lead to a compromised performance in these classes during subsequent tasks.

7. Conclusion and Discussion

In this paper, we focus on inductive continual zero-shot learning (CZSL) to eliminate the need of unseen information for more realistic learning systems. To this end, we developed a framework for the theoretical analysis of generative zero-shot learning, introducing a distance metric to measure the ability of generated samples to represent the unseen space when the unseen information is inaccessible during training. We also proposed a continual zero-shot algorithm, ICGZSL, which can reduce the distance without using unseen information during training. We conducted experiments on four popular continual zero-shot learning benchmarks: AWA1, AWA2, CUB, and SUN. Our approach achieved around 3% higher harmonic accuracy in the small dataset and around 7% in the larger dataset compared to previous inductive and transductive methods. These results demonstrate that unseen semantic information is not essential when a well-analyzed seen distribution and method are used.

However, it is important to acknowledge certain limitations in our work. While the developed theoretical bounds and distance measures hold promise for methodical numeric analysis, a more stringent alignment between empirical and anticipated distance measures could substantially enhance algorithmic design. Moreover, the consideration of multi-class classification conditions warrants attention. Additionally, the use of a frozen backbone for image feature extraction, though effective, encourages further exploration into continual learning methods that facilitate viable zero-shot learning capabilities while enabling the backbone to progressively accumulate knowledge.

References

- [1] Yaser S Abu-Mostafa, Malik Magdon-Ismail, and Hsuan-Tien Lin. *Learning from data*, volume 4. AML-Book New York, 2012. 13
- [2] Zeynep Akata, Florent Perronnin, Zaid Harchaoui, and Cordelia Schmid. Label-embedding for image classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(7):1425–1438, 2016. 20
- [3] Zeynep Akata, Scott Reed, Daniel Walter, Honglak Lee, and Bernt Schiele. Evaluation of output embeddings for fine-grained image classification. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015. 20
- [4] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory Aware Synapses: Learning what (not) to forget. In *European Conference on Computer Vision*, 2018. 3
- [5] Ahmed Ayyad, Yuchen Li, Nassir Navab, Shadi Albarqouni, and Mohamed Elhoseiny. Semi-supervised few-shot learning with prototypical random walks. In *AAAI Workshop on Meta-Learning and MetaDL Challenge*, pages 45–57, 2021. 4, 17
- [6] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine Learning*, 79(1-2):151–175, 2010. 4, 13
- [7] Shivam Chandhok, Sanath Narayan, Hisham Cholakkal, Rao Muhammad Anwer, Vineeth N Balasubramanian, Fahad Shahbaz Khan, and Ling Shao. Structured latent embeddings for recognizing unseen classes in unseen domains. In *British Machine Vision Conference*, 2021. 3
- [8] Fu-Chieh Chang, Hao-Jen Wang, Chun-Nan Chou, and Edward Y. Chang. G2R Bound: A Generalization Bound for Supervised Learning from GAN-Synthetic Data. *arXiv:1905.12313*, 2019. 1, 4, 13
- [9] Soravit Changpinyo, Wei-Lun Chao, Boqing Gong, and Fei Sha. Synthesized classifiers for zero-shot learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 5327–5336, 2016. 19, 20
- [10] 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 *European Conference on Computer Vision*, pages 52–68, 2016. 19
- [11] Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient Lifelong Learning with A-GEM. In *International Conference on Learning Representations*, 2019. 3, 7
- [12] Arthur Douillard, Eduardo Valle, Charles Ollion, Thomas Robert, and Matthieu Cord. Insights from the future for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3482–3491, 2021. 3
- [13] Sayna Ebrahimi, Franziska Meier, Roberto Calandra, Trevor Darrell, and Marcus Rohrbach. Adversarial continual learning. In *European Conference on Computer Vision*, pages 386–402, 2020. 3
- [14] Mohamed Elfeki, Camille Couprie, Morgane Riviere, and Mohamed Elhoseiny. Gdpp: Learning diverse generations using determinantal point processes. In *International conference on machine learning*, pages 1774–1783, 2019. 4, 15
- [15] Ahmed Elgammal, Bingchen Liu, Mohamed Elhoseiny, and Marian Mazzone. Can: Creative adversarial networks, generating "art" by learning about styles and deviating from style norms. In *International Conference on Computational Creativity*, 2017. 2
- [16] Mohamed Elhoseiny and Mohamed Elfeki. Creativity inspired zero-shot learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5784–5793, 2019. 2, 3, 4, 6, 7, 19, 20
- [17] Mohamed Elhoseiny, Kai Yi, and Mohamed Elfeki. Cizsl++: Creativity inspired generative zero-shot learning. *arXiv preprint arXiv:2101.00173*, 2021. 2
- [18] Mohamed Elhoseiny, Yizhe Zhu, Han Zhang, and Ahmed Elgammal. Link the head to the "beak": Zero shot learning from noisy text description at part precision. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 19
- [19] Ali Farhadi, Ian Endres, Derek Hoiem, and David Forsyth. Describing objects by their attributes. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1778–1785, 2009. 20
- [20] Rafael Felix, Vijay BG Kumar, Ian Reid, and Gustavo Carneiro. Multi-modal cycle-consistent generalized zero-shot learning. In *European Conference on Computer Vision*, pages 21–37, 2018. 20
- [21] Chandan Gautam, Sethupathy Parameswaran, Ashish Mishra, and Suresh Sundaram. Generative Replay-based Continual Zero-Shot Learning. *arXiv:2101.08894*, 2021. 3, 8
- [22] Chandan Gautam, Sethupathy Parameswaran, Ashish Mishra, and Suresh Sundaram. Tf-gczsl: Task-free generalized continual zero-shot learning. *Neural Networks*, 155:487–497, 2022. 7
- [23] Chandan Gautam, Sethupathy Parameswaran, Vinay Verma, Suresh Sundaram, and Savitha Ramasamy. Refinement matters: Textual description needs to be refined for zero-shot learning. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 6127–6140, 2022. 7

- [24] Subhankar Ghosh. Adversarial training of variational auto-encoders for continual zero-shot learning. In *International Joint Conference on Neural Networks*, 2021. 3, 7, 8
- [25] Subhankar Ghosh. Dynamic VAEs with Generative Replay for Continual Zero-shot Learning. In *CVPR workshop*, 2021. 1, 7, 8
- [26] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, volume 27, 2014. 1
- [27] Aaron Hertzmann. Can computers create art? In *Arts*, volume 7, page 18, 2018. 2
- [28] Aaron Hertzmann. Visual indeterminacy in gan art. *Leonardo*, 53(4):424–428, 2020. 2
- [29] Philip Häusser, Alexander Mordvintsev, and Daniel Cremers. Learning by association — a versatile semi-supervised training method for neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 626–635, 2017. 4, 17
- [30] Divyansh Jha, Hanna H. Chang, and Mohamed Elhoseiny. Wölflin’s affective generative analysis of visual art. *The International Conference on Computational Creativity*, 2021. 2
- [31] Divyansh Jha, Kai Yi, Ivan Skorokhodov, and Mohamed Elhoseiny. Creative walk adversarial networks: Novel art generation with probabilistic random walk deviation from style norms. In *ICCC*, pages 195–204, 2022. 2
- [32] Byungkon Kang. Fast determinantal point process sampling with application to clustering. *Advances in Neural Information Processing Systems*, 26, 2013. 4, 15
- [33] Michael J Kearns and Umesh Vazirani. *An introduction to computational learning theory*. MIT press, 1994. 4
- [34] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13):3521–3526, 2017. 7
- [35] Elyor Kodirov, Tao Xiang, and Shaogang Gong. Semantic autoencoder for zero-shot learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 3174–3183, 2017. 20
- [36] Hari Chandana Kuchibhotla, Sumitra S Malagi, Shivam Chandhok, and Vineeth N Balasubramanian. Unseen classes at a later time? no problem. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 1, 3, 5, 6, 7, 8, 15
- [37] Christoph H Lampert, Hannes Nickisch, and Stefan Harmeling. Learning to detect unseen object classes by between-class attribute transfer. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 951–958, 2009. 7, 20
- [38] Jingjing Li, Mengmeng Jing, Ke Lu, Zhengming Ding, Lei Zhu, and Zi Huang. Leveraging the Invariant Side of Generative Zero-Shot Learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7394–7403, 2019. 1
- [39] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12):2935–2947, 2017. 3
- [40] Zhiqiu Lin, Jia Shi, Deepak Pathak, and Deva Ramanan. The clear benchmark: Continual learning on real-world imagery. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021. 3
- [41] Yang Liu, Jishun Guo, Deng Cai, and Xiaofei He. Attribute attention for semantic disambiguation in zero-shot learning. In *IEEE/CVF International Conference on Computer Vision*, 2019. 3
- [42] Sanath Narayan, Akshita Gupta, Fahad Shahbaz Khan, Cees GM Snoek, and Ling Shao. Latent embedding feedback and discriminative features for zero-shot classification. In *European Conference on Computer Vision*, pages 479–495, 2020. 1
- [43] Genevieve Patterson and James Hays. SUN attribute database: Discovering, annotating, and recognizing scene attributes. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2751–2758, 2012. 7, 20
- [44] Akanksha Paul, Narayanan C Krishnan, and Prateek Munjal. Semantically aligned bias reducing zero shot learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7056–7065, 2019. 3
- [45] Farhad Pourpanah, Moloud Abdar, Yuxuan Luo, Xinlei Zhou, Ran Wang, Chee Lim, and Xi-Zhao Wang. A review of generalized zero-shot learning methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PP, 07 2022. 3
- [46] Ameya Prabhu, Philip HS Torr, and Puneet K Dokania. Gdumb: A simple approach that questions our progress in continual learning. In *European conference on computer vision*, pages 524–540, 2020. 9
- [47] Ruizhi Qiao, Lingqiao Liu, Chunhua Shen, and Anton van den Hengel. Less is more: Zero-shot learning from online textual documents with noise suppression.

- In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 19
- [48] Shafin Rahman, Salman Khan, and Nick Barnes. Transductive learning for zero-shot object detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6082–6091, 2019. 3
- [49] Jathushan Rajasegaran, Munawar Hayat, Salman Khan, Fahad Shahbaz Khan, and Ling Shao. Random path selection for incremental learning. *Advances in Neural Information Processing Systems*, 2019. 3
- [50] Gerard Salton and Christopher Buckley. Term-weighting approaches in automatic text retrieval. *Information processing & management*, 24(5):513–523, 1988. 19
- [51] Othman Sbai, Mohamed Elhoseiny, Antoine Bordes, Yann LeCun, and Camille Couprie. Design: Design inspiration from generative networks. In *ECCV workshop*, 2018. 2
- [52] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. *Advances in Neural Information Processing Systems*, 30, 2017. 3
- [53] Ivan Skorokhodov and Mohamed Elhoseiny. Class normalization for (continual)? generalized zero-shot learning. In *International Conference on Learning Representations*, 2021. 1, 3, 7, 8
- [54] Hongzu Su, Jingjing Li, Zhi Chen, Lei Zhu, and Ke Lu. Distinguishing unseen from seen for generalized zero-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7885–7894, June 2022. 3
- [55] Lee Sweetlove. Number of species on earth tagged at 8.7 million. *Nature*, 23, 2011. 1
- [56] Grant Van Horn, Steve Branson, Ryan Farrell, Scott Haber, Jessie Barry, Panos Ipeirotis, Pietro Perona, and Serge Belongie. Building a bird recognition app and large scale dataset with citizen scientists: The fine print in fine-grained dataset collection. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015. 18
- [57] Twan van Laarhoven and Elena Marchiori. Unsupervised domain adaptation with random walks on target labelings. *arXiv:1706.05335*, 2017. 4, 14, 17
- [58] Maunil Vyas, Hemanth Venkateswara, and Sethuraman Panchanathan. Leveraging seen and unseen semantic relationships for generative zero-shot learning. In *European Conference on Computer Vision*, pages 70–86, 2020. 1, 19, 20
- [59] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011. 7, 9, 18, 21
- [60] Wenlin Wang, Hongteng Xu, Guoyin Wang, Wenqi Wang, and Lawrence Carin. Zero-Shot Recognition via Optimal Transport. In *IEEE Winter Conference on Applications of Computer Vision*, pages 3470–3480, 2021. 3
- [61] Yongqin Xian, Zeynep Akata, Gaurav Sharma, Quynh Nguyen, Matthias Hein, and Bernt Schiele. Latent embeddings for zero-shot classification. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 69–77, 2016. 20
- [62] Yongqin Xian, Christoph H Lampert, Bernt Schiele, and Zeynep Akata. Zero-shot learning—a comprehensive evaluation of the good, the bad and the ugly. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018. 7, 20
- [63] Yongqin Xian, Tobias Lorenz, Bernt Schiele, and Zeynep Akata. Feature Generating Networks for Zero-Shot Learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5542–5551, 2018. 3, 19, 20
- [64] Yongqin Xian, Saurabh Sharma, Bernt Schiele, and Zeynep Akata. F-vaegan-d2: A feature generating framework for any-shot learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019. 3
- [65] Ye Xiang, Ying Fu, Pan Ji, and Hua Huang. Incremental learning using conditional adversarial networks. In *IEEE/CVF International Conference on Computer Vision*, pages 6619–6628, 2019. 3
- [66] Kai Yi and Mohamed Elhoseiny. Domain-Aware Continual Zero-Shot Learning. *arXiv:2112.12989*, 2021. 7
- [67] Lu Yu, Bartłomiej Twardowski, Xialei Liu, Luis Heranz, Kai Wang, Yongmei Cheng, Shangling Jui, and Joost van de Weijer. Semantic Drift Compensation for Class-Incremental Learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6980–6989, 2020. 7
- [68] Li Zhang, Tao Xiang, and Shaogang Gong. Learning a deep embedding model for zero-shot learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 20
- [69] Yizhe Zhu, Mohamed Elhoseiny, Bingchen Liu, Xi Peng, and Ahmed Elgammal. A generative adversarial approach for zero-shot learning from noisy texts. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018. 3, 19, 20