

Unseen Classes at a Later Time? No Problem

Hari Chandana Kuchibhotla^{‡*}, Sumitra S Malagi^{‡*}, Shivam Chandhok^{◇ ‡}, Vineeth N Balasubramanian[‡]

[‡] Indian Institute of Technology Hyderabad, India [◇] INRIA, Universite Grenoble Alpes

{ai20resch11006, cs20mtech14006, vineethnb}@iith.ac.in, chandhokshivam@gmail.com

Abstract

Recent progress towards learning from limited supervision has encouraged efforts towards designing models that can recognize novel classes at test time (generalized zero-shot learning or GZSL). GZSL approaches assume knowledge of all classes, with or without labeled data, beforehand. However, practical scenarios demand models that are adaptable and can handle dynamic addition of new seen and unseen classes on the fly (i.e continual generalized zero-shot learning or CGZSL). One solution is to sequentially retrain and reuse conventional GZSL methods, however, such an approach suffers from catastrophic forgetting leading to suboptimal generalization performance. A few recent efforts towards tackling CGZSL have been limited by difference in settings, practicality, data splits and protocols followed – inhibiting fair comparison and a clear direction forward. Motivated from these observations, in this work, we firstly consolidate the different CGZSL setting variants and propose a new Online-CGZSL setting which is more practical and flexible. Secondly, we introduce a unified feature-generative framework for CGZSL that leverages bi-directional incremental alignment to dynamically adapt to addition of new classes, with or without labeled data, that arrive over time in any of these CGZSL settings. Our comprehensive experiments and analysis on five benchmark datasets and comparison with baselines show that our approach consistently outperforms existing methods, especially on the more practical Online setting.

1. Introduction

Deep Neural Networks (DNNs) have shown great promise as predictive models and are increasingly being used for various computer vision applications. However, their reliance on large-scale labeled datasets limit their use in practical scenarios encountered in the real-world. The occurrence of objects in the real world inherently follows long-tailed distributions [16, 35], implying that visual data sampled from the real world may not be readily available

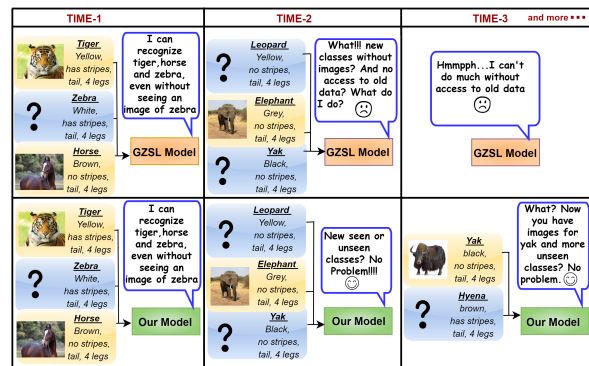


Figure 1. **Illustration of proposed setting.** Applicability of a Generalized Zero-shot Learning (GZSL) model (row 1) and our proposed model (row 2) in a real-world scenario on dynamic addition of new seen (depicted in yellow) and unseen (depicted in blue) categories. The GZSL setting does not allow dynamic addition of newer classes, seen or unseen, that are added over time. Our proposed CGZSL setting is more flexible and can tackle dynamic changes in the initial pool of seen and unseen categories, enhancing model scalability and applicability in practical settings.

for all categories of interest at the same time. Thus, it is required for models to have the ability to generalize and recognize novel objects semantically similar to the ones encountered during training even though visual data for these novel classes is not seen by the model. Existing efforts aim to tackle this problem, by designing generalized zero-shot learning (GZSL) models that are equipped with the ability to generalize to unseen classes at test time.

Another important aspect of object occurrence in the real world is the gradual addition of object categories with time. This can be attributed to the discovery of new objects or the continuous nature of data collection process owing to which a previously rare object may have abundantly available samples at a later stage. However, existing GZSL approaches are not designed to tackle dynamic addition of classes in the initial pool of seen and unseen categories, limiting their scalability and applicability in challenging practical settings. Fig. 1 illustrates the shortcomings of a GZSL model in such practical settings where classes arrive over time. Evidently, the GZSL setting is limiting, and is unable to adapt to dynamic changes to the initial pool of categories brought by the gradual addition of new seen and unseen classes over

*equal contribution

time. This can be attributed to the fact that since the previous data is no longer available, the GZSL model tends to *catastrophically forget* knowledge pertaining to previous tasks when sequentially trained and reused over time. This necessitates concerted effort toward carefully designing problem settings that resemble the occurrence of objects in the real world, and building models that can adapt over time and seamlessly tackle such challenges.

Very recently, sporadic efforts [6, 9, 26] have been made towards designing models that can dynamically adapt and generalize on addition of new seen and unseen classes. The aforementioned works term this setting as *continual generalized zero-shot learning* (CGZSL). However, these efforts are nascent, and vary in the definition of the problem setting, practicality, data splits and protocols followed – thus inhibiting fair comparison and a clear path forward. To address this issue and motivated by the need for progress in this direction, in this work, we firstly consolidate the different CGZSL settings tackled in these recent efforts, and clearly segregate existing methods and settings according to the challenges they tackle. We propose a more flexible and realistic *Online-CGZSL* setting which more closely resembles scenarios encountered in the real world. In addition, we propose a unified framework that employs a bi-directional incremental alignment-based replay strategy to seamlessly adapt and generalize to new seen and unseen classes that arrive over time. Our replay strategy is based on a feature-generative architecture, and hence does not require storing samples from previous tasks. We also use a static architecture for incremental learning (as against using a model-growing one) in order to facilitate scalability and efficiency. In summary, the key contributions of our work are below:

- We identify the different challenges of the relatively new CGZSL setting, and consolidate its different variants based on the challenges they tackle and restrictions they impose. We hope that this will enable fair comparison among such approaches and further progress in the field.
- We establish a practical, but more challenging, *Online-CGZSL* setting which more closely resembles real-world scenarios encountered in practice.
- We propose a novel feature-generative framework to address the different CGZSL setting variants, which avoids catastrophic forgetting through bi-directional incremental alignment thereby allowing forward semantic knowledge transfer from previous tasks and enabling generalization.
- We perform extensive experiments and analysis on three different CGZSL settings on well-known benchmark datasets: AWA1, AWA2, Attribute PASCAL and Yahoo (aPY), Caltech-UCSD-Birds (CUB) and SUN, demonstrating the promise of our approach. We observe that our model consistently improves over baselines and existing

approaches, especially on the more challenging *Online* setting.

2. Related Work

Generalized Zero-Shot Learning (GZSL). Existing GZSL approaches can be broadly classified into *embedding-based methods* and *generative methods*. Traditional embedding-based approaches [4, 12, 23, 34] aim to project visual and/or semantic features onto a common embedding space and use a nearest neighbor based classifier to classify visual samples. On the other hand, a relatively recent and more effective set of approaches [1, 5, 13, 20, 32] propose to use generative models to synthesize unseen visual features, thus converting the GZSL problem into a supervised classification problem. Although, the aforementioned approaches aim to generalize to novel unseen classes at test-time, they are not designed to tackle dynamic addition of categories (seen or unseen) over time.

Continual Learning (CL). Existing approaches that aim to learn continually and tackle catastrophic forgetting [19] can broadly be categorized into: parameter isolation methods [17], regularization-based methods [10, 14, 33], model-growing architectures [24], and rehearsal-based methods [2]. Parameter isolation methods [17] aim to identify the parameters important for a task while regularization-based methods [10, 14, 33] constrain parameters to avoid deviation of weights important to previous tasks. On the other hand, model growing architectures [18] dynamically increase model capacity and rehearsal-based methods store or generate images of previous tasks to avoid forgetting [15]. In addition, various approaches use knowledge distillation technique [15, 22] to transfer knowledge from previous tasks to current task. However, these approaches cannot generalize to unseen classes for which visual samples have not been encountered by the model during training.

Continual Generalized Zero-Shot Learning (CGZSL). While CL aims to learn incrementally by *transferring learned knowledge* to future tasks, GZSL aims to *transfer semantic knowledge* from seen classes to recognize unseen classes. Considering the common objective of transferring learned knowledge, there have been a few recent efforts to unify these paradigms. Lifelong ZSL [30] marked the first attempt to address the problem of CGZSL, where a multi-head architecture was used to accumulate knowledge while training from multiple datasets. However, the approach requires task level supervision in the form of task-ids at test time limiting its applicability in realistic scenarios. On the other hand, [26] suggested a class normalization-based approach to solve the CGZSL problem and [9] learned a new VAE for every task. However, [9, 26, 29] specifically consider previously encountered tasks as seen and future tasks as unseen classes, and thus only tackle the *static-CGZSL*

Different Settings	Base Task						New Tasks (Continual Setting)						Base Task		New Tasks (Continual)									
	Task-1		Task-2		Task-3		Training		Testing		Training		Testing		S	U	S	U	U->S	S _{new}	U _{new}	S	U	
	Seen(S)	Unseen(U)	Seen(S)	Unseen(U)	Seen(S)	Unseen(U)	Seen(S)	Unseen(U)	Seen(S)	Unseen(U)	Seen(S)	Unseen(U)	Seen(S)	Unseen(U)	Seen(S)	Unseen(U)	Seen(S)	Unseen(U)	Seen(S)	Unseen(U)	Seen(S)	Unseen(U)	Seen(S)	Unseen(U)
Semi-Supervised	Seen(S)	Unseen(U)	Seen(S)	Unseen(U)	Seen(S)	Unseen(U)	✓	✓																
Open-set Learning	S	(Unk)	S	(Unk)	S	(Unk)	✓	✓	●															
Zero shot Learning (ZSL)	S	Unseen(U)	S	Unseen(U)	S	Unseen(U)	✓	✓	✓															
Generalized Zero shot Learning (GZSL)	S	S U	S	S U	S	S U	✓	✓	✓															
Continual Learning (CL)	Seen1(S1)	S1	Seen2(S2)	S1 S2	Seen3(S3)	S1 S2 S3	✓	✓			✓	✓												
Open World Recognition	S (Unk)	S	S	S	S	S	✓		●	✓														
Static Continual GZSL*	S1	S1 U	S2	S1 S2 U	S3	S1 S2 S3 U	✓	✓	✓	✓														
Dynamic Continual GZSL*	S1	S1 U1	S2	S1 U1 S2 U2	S3	S1 U1 S2 U3 U3	✓	✓	✓		✓	✓												
Online Continual GZSL*	S1	S1 U1	S2	S1 U1 S2 U2	S3	S1 U1 S2 U3 U3	✓	✓	✓	✓	✓	✓	✓											

Figure 2. Comparison of proposed settings with other known settings. In ZSL, model is trained on seen classes and evaluated on unseen classes. In GZSL, the model is tested on both seen and unseen classes. CL models are trained on classes that arrive sequentially, but do not have unseen classes either during training or testing. Continual GZSL (CGZSL) settings (highlighted with a *) are proposed in this work. In static-CGZSL, classes that arrive in future are considered unseen. In Dynamic-CGZSL, each task has disjoint set of seen and unseen classes. Online-CGZSL allows for conversion of previously unseen classes to seen (based on availability of data) in addition to handling new seen and unseen classes at each task.

setting which is restrictive in practice (Sec.3). [6, 7] recently formulated CGZSL as a problem where each task has its own set of seen and unseen classes and work in the *dynamic-CGZSL* setting. Although this setting is an improvement over *static-CGZSL* setting, it is still restrictive as it does not allow the dynamic conversion of unseen classes to seen as data becomes available over time. Furthermore, learning a new VAE for every task [9] or storing exemplar samples of previous tasks [7] results in progressively increasing memory requirements, which is inefficient. On the other hand, [6] employs a rehearsal based strategy but does not take into account the changing semantic structure of visual spaces due to addition of new categories over time.

3. CGZSL: Settings and Formulation

In this section, we consolidate and provide a holistic overview of the various CGZSL settings tackled by recent efforts and discuss in detail the major differences in their formulations. We also describe our proposed *Online-CGZSL* setting which is more flexible and resembles real world scenarios more closely. Fig. 2 illustrates a pictorial representation of the various CGZSL settings and strives to correlate and appropriately position them w.r.t. other related limited-supervision and continual learning settings. We now describe each of the CGZSL setting variants below.

Static-CGZSL. In this setting [8, 9, 26, 29], the dataset is divided into T subsets, and the model encounters each of these subsets in an incremental fashion over time. The setting assumes all previously encountered tasks as seen and future tasks as consisting of unseen classes. Formally, for a

given task \mathcal{T}_t at a given time step t , the first t subsets i.e data belonging to the current and previous tasks are considered as seen classes while the future tasks are considered unseen.

This setting differs from traditional GZSL in the fact that during evaluation of the t^{th} task, previous training data is unavailable. Thus the model should be capable of retaining previously learned knowledge while adapting to the newly encountered seen classes. However, *static-CGZSL* presents a constrained setting which requires the total number of classes or tasks to be known beforehand (hence the name *static*). Furthermore, the setting mandates that all tasks till current time step t are considered as seen classes, with only the future tasks contain the unseen classes. Thus the dynamic addition of classes is restricted to conversion of an unseen class to seen after a particular task is encountered while learning continually. While it is reasonable to assume that visual features of unseen classes may become available in future, it may be non-viable to assume that novel seen or unseen categories, unknown at the beginning of training, will not be added in future. This fundamentally limits the notion of continual learning where a model should be adaptable to any number of tasks or addition of new classes.

Dynamic-CGZSL. Considering the limitations of *static-CGZSL* setting, [6, 7] proposed another setting where each task has an exclusive set of seen and unseen classes and the model can accommodate any number of tasks over time. We categorize this setting as *dynamic-CGZSL*, which is less restrictive than *static-CGZSL* setting as it allows for addition of both seen and unseen classes in a continual manner. However, this setting imposes a constraint that previously

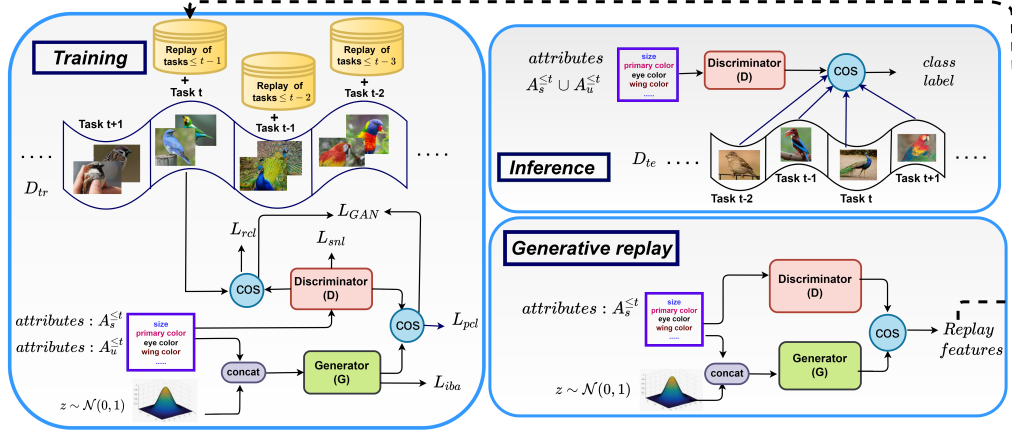


Figure 3. Our approach consists of three phases: (1) Training; (2) Generative replay; and (3) Inference. (1) During training, the discriminator D is trained to map input attribute to the visual space, and generator G is trained to generate pseudo-visual features. The GAN is trained adversarially using similarity scores. L_{rcl} , L_{pcl} and L_{snl} loss terms aid in generating discriminative features for seen and unseen classes. L_{sal} and $L_{nuclear}$ are used for incremental bi-directional alignment; (2) We test on three different settings at a given time step t , as described in Sec. 3. Cosine similarity is used to classify the target visual feature; (3) Trained G is used to replay/generate features of seen classes encountered until task t .

unseen classes may never become seen in future and is unable to tackle the scenario where a rare class may have abundantly available samples at a later time stage. This is limiting in practice as owing to the continuous nature of data collection, it is plausible that some of the unseen classes may become seen when its data becomes available over time.

Online-CGZSL. In order to better align with the scenarios commonly encountered in the real-world, we introduce a new *Online-CGZSL* setting which can handle a variety of dynamic changes in the pool of seen and unseen classes. Specifically, each task has a disjoint set of seen and unseen classes, and an arbitrary number of such tasks/categories can be dynamically incorporated on the fly. Importantly, this setting allows the conversion of previously unseen classes into seen if the corresponding visual features become available depending on changes in data availability in the future. Note that our setting is more flexible as it does not require the knowledge of entire pool of categories beforehand, impose any restrictions on the conversion of unseen class to seen class or require task level supervision at test-time. Next, we formally describe and formulate our proposed *Online-CGZSL* problem setting.

Problem Formulation. Let subscripts s and u denote seen and unseen classes respectively. Each task t consists of train and test data. Let A^t be the union of attributes of seen (A_s^t) and unseen (A_u^t) encountered so far. The training data D_{tr}^t at task t is given by $D_{tr}^t = \{(training\ visual\ samples\ of\ current\ task\ t\ (X_{str}^t),\ their\ class\ labels\ (Y_{str}^t),\ class\ attributes(A^t))\}$. The test data at task t be $D_{te}^t = \{(test\ visual\ samples\ of\ seen\ and\ unseen\ classes\ encountered\ so\ far\ (X_{ste}^t\ and\ X_{ute}^t),\ their\ class\ labels\ (Y_{ste}^t\ and\ Y_{ute}^t),\ class\ attributes(A^t))\}$. Let R^t be the replayed seen visual features of previous tasks. During the training phase of each task, data pertaining to current task and R^t is used as training

data. At the end of each task, the model’s performance is evaluated by testing it on current as well as previous tasks’ seen and unseen data. We operate in semantically transductive setting and do not use unseen visual features during training.

4. CGZSL: Proposed Methodology

Overall Framework. The overall framework of our approach is shown in Fig. 3. Given a time step t , we first train a generator G which employs a cosine similarity based formulation enabling it to dynamically incorporate any number of categories or tasks over time (Sec. 4.1). In order to ensure that generated visual features are discriminative w.r.t class distribution at current time step, we impose normalized discriminative loss (Sec. 4.2). In addition, we propose to utilise incremental bi-directional alignment in order to adapt and ensure knowledge transfer from previous $(t - 1)$ tasks and strengthen semantic relationship among classes encountered till time t reducing catastrophic forgetting (Sec. 4.3). At time step $(t + 1)$ we use generative replay strategy to generate seen class visual features (R^t) for all previous tasks and combine them with the seen class samples from the current task (Sec. 4.4). This procedure is repeated for the next time step.

4.1. Cosine Similarity-based GAN Classifier

We learn a generative model which comprises of a generator, $G_\theta : \mathcal{Z} \times \mathcal{A} \rightarrow \mathcal{X}$ and a discriminator $D : \mathcal{A} \rightarrow \mathcal{X}$. The generator takes as input random noise, $z \in \mathcal{Z}^d$ and class attributes $a \in A^t$ and outputs a generated visual feature belonging to the same class. On the other hand, the discriminator D_ϕ takes class attributes $a \in A^t$ as input and outputs identifier projection of the attribute. *Identifier projection is a projection of the attribute in the visual space.* The G_θ and D_ϕ are trained adversarially where discrimina-

tor tries to minimize the cosine similarity between identifier projection and generated seen features belonging to same class, while the generator tries to maximize this cosine similarity. In addition to minimizing the aforementioned cosine similarity, the discriminator tries to maximize cosine similarity between real visual features and the corresponding identifier projection. Through this adversarial training, G_θ learns to generate visual features similar to the real visual features and D_ϕ learns a better mapping for the attributes. Training the G_θ and D_ϕ adversarially with the cosine similarities between generated seen features (X'), identifier projection ($D(a)$) and the real visual features ($X_{str}^t \cup R^t$), is formulated as:

$$L_{GAN} = \mathbb{E}_{x \sim p_{data}(X_{str}^t \cup R^t)} [\log [\cos(x, D(a))]] + \mathbb{E}_{x' \sim p_\theta(X'|a)} [\log [1 - \cos(x', D)]] \quad (1)$$

where the distributions of real visual feature and generated visual feature are denoted by $p_{data}(X_{str}^t \cup R^t)$ and $p_\theta(X'|a)$. During testing, we compute the cosine similarity between test sample and identifier projections of all the class attributes encountered so far. The test sample is assigned the class label of identifier projection with which it shares maximum similarity. Thus we are able to achieve continual generalized zero-shot classification using a single GAN model without the need for training a linear classifier during each task. The proposed architecture is simple and allows easy adaptation to increasing number of classes, as the classification is performed by merely computing cosine similarities between the identifier projections of the classes encountered so far and the test sample.

4.2. Real and Pseudo Normalized Loss

As new tasks or classes gradually arrive over time, the specific features that allow us to best discriminate among the pool of classes dynamically changes. Thus, in order to ensure that the generated visual features are discriminative w.r.t the class distribution at any given time step t , we impose a set of losses on the generated visual features from current and all previous timesteps. Specifically, we calculate softmax score of the cosine similarity between all identifier projections and visual features. The class label corresponding to identifier projection with highest softmax score is the predicted label. We use three classification losses- (i) Real classification loss (L_{rcl}) is the classification loss corresponding to real visual features (current task + replayed samples). L_{rcl} enforces the discriminator to consider inter-class distances while mapping the attributes. (ii) The classification loss corresponding to generated seen visual features called pseudo-visual classification loss (L_{pcl}) is added to the generator. L_{pcl} encourages generator to generate more discriminative visual features. (iii) Since visual features of unseen classes are unavailable, a classification loss of generated unseen visual features called seen-normalized loss L_{snl} is added to the discriminator. L_{snl} serves as reference

for finding appropriate mapping for unseen attributes.

The classification losses L_{rcl} , L_{pcl} and L_{snl} are defined as follows:

$$L_{rcl}, L_{pcl}, L_{snl} = c.e(\log \frac{\exp(\cos(x, D(a_i)))}{\sum_{i \in A^t} \exp(\cos(x, D(a_i)))}, y_i) \quad (2)$$

x corresponds to the real visual features in L_{rcl} , generated seen visual features in L_{pcl} and generated unseen visual feature in L_{snl} . $c.e$ stands for cross entropy and y_i is the true class label of x . t covers only seen classes in L_{rcl} and L_{pcl} ; but sums over seen and unseen classes for L_{snl} . During testing, classification spans all classes encountered till t .

4.3. Incremental Bi-directional Alignment Loss

Owing to the nature of the CGZSL setting, the visual feature space dynamically changes as new seen and unseen classes are added over time. Furthermore, the visual features for seen classes belonging to previous time steps and entire pool of unseen classes are not available during the current time step. Thus there is a need for a mechanism that can forward transfer knowledge from previous tasks to avoid catastrophic forgetting and exploit the current semantic structure to generate better visual features (especially unseen). To this end, we propose to use an incremental bi-directional alignment loss (L_{iba}) consisting of nuclear loss and semantic alignment loss. Semantic alignment loss helps in using semantic information [27] as a reference for generating unseen visual features, while the nuclear loss aids in transferring visual information from seen classes to identifier projections (projection of attributes in visual space).

The visual similarity between two classes c_i and c_j is the cosine similarity between their class means. It is represented as $X_{sim}(\mu_{c_i}; \mu_{c_j})$ where μ_{c_i} stands for the mean visual feature of class i . Let semantic similarity between the classes be represented as $\tau_{sim}(a_{c_i}, a_{c_j})$ where a_{c_i} is the attribute of class i . Semantic alignment loss constraints $X_{sim}(\mu_{c_i}; \mu_{c_j})$ to lie in a range $\tau_{sim}(a_{c_i}, a_{c_j})$ plus or minus ϵ (hyper-parameter), thus transferring the semantic structure to generated features. With the addition of new classes, the semantically similar classes of c_i may change. Hence we incrementally calculate the semantic alignment loss (L_{sal}) for all classes encountered so far. Nuclear loss is the L2 norm between μ_{c_i} and corresponding identifier projection.

The incremental bi-directional semantic alignment loss is given by:

$$L_{sal} = \min_{\theta_g} \frac{1}{N} \sum_{i=1}^N \sum_{j \in I_{c_i}} \|\max(0, X_{sim}(\mu_{c_j}, \mu'_{c_i}) - (\tau_{sim}(a_{c_j}, a_{c_i}) + \epsilon))\|^2 + \|\max(0, (\tau_{sim}(a_{c_j}, a_{c_i}) - \epsilon) - X_{sim}(\mu_{c_j}, \mu'_{c_i}))\|^2 \quad (3)$$

$$L_{nuclear} = \|\mu_{c_i} - S_{c_i}\|^2 \quad (4)$$

Algorithm 1 Proposed CGZSL Method

Input : D_{tr}^t, D_{te}^t, G, D
Output : Predicted labels y_{pred}
Parameters : θ_G, ϕ_D
if new task arrived **then**
 if task number greater than 1 **then**
 $R^t = \text{replay data till task } t-1; X_{str}^t = X_{str}^{t-1} \cup R^t$
 end if
 for epochs = 1 to N **do**
 $X'_s = G(z, A_s^t)$ // seen pseudo-visual features
 $X'_u = G(z, A_u^t)$ // unseen pseudo-visual features

 $L_{GAN} \leftarrow \text{eqn (1)}$ // using X_{str}^t, X'_s and $D(A_s^t)$
 $L_{rcl} \leftarrow \text{eqn (2)}$ // using X_{str}^t and $D(A_s^t)$
 $L_{snl} \leftarrow \text{eqn (2)}$ // using $X'_u, D(A^t)$
 // Overall D loss
 $L_D^t = \lambda_1 L_{GAN} + \lambda_2 L_{rcl} + \lambda_3 L_{snl}$
 $\phi_D = \phi_D - \eta_1 \times \nabla L_D^t$ // update D
 $L_{pcl} \leftarrow \text{eqn (2)}$ // using X'_s and $D(A_s^t)$
 $L_{iba} \leftarrow \text{eqn (3) and eqn (4)}$ // using X'_s
 and X'_u
 // Overall G loss
 $L_G^t = \lambda_1 L_{GAN} + \lambda_2 L_{pcl} + \lambda_4 L_{iba}$
 $\theta_G = \theta_G - \eta_2 \times \nabla L_G^t$ // Update G
 end for
 Inference: model is evaluated on D_{te}^t
end if

where N is the number of classes encountered so far. For a given class c_i , S_{c_i} is the identifier projection, μ_{c_i} is the mean of real visual features and μ'_{c_i} is mean of generated visual features. Mean of real visual features is available only for seen classes. Let I_{c_i} represent the set of n_c nearest neighbours of c_i .

4.4. Generative Replay

We work in a setting where data arrives incrementally, and samples from previous tasks are not accessible during the current task. This results in catastrophic forgetting. In order to retain previously learned knowledge and adopt new knowledge, we use generative replay. Visual features of previously seen classes are generated by passing a concatenation of attribute and noise vectors to the generator network. A combination of generated features of previous tasks and real features from the current task act as input data to train the model. To ensure that we are replaying credible and good quality visual features at a time step t , we classify the generated visual features and replay only features that are classified correctly.

$$R^t = G(z, A_s^{\leq t-1}) \quad (5)$$

where $z \sim \mathcal{N}(0, 1)$ and $A_s^{\leq t-1}$ denotes attributes of all the seen classes encountered so far.

4.5. Training and Inference

For a given time step $t (> 1)$, we concatenate the replayed visual features and current task’s visual features from seen classes along with their corresponding labels.

The concatenated data acts as input for training the generative model. During training, the discriminator and generator are trained sequentially on different tasks. After the training process, the discriminator learns a mapping function to map attributes to visual space. The generator learns to generate synthetic visual features conditioned on attributes. During testing, we classify the test sample using cosine similarity as described in Sec. 4.2. Based on the setting we are working in, accuracies are computed as described in Appendix.

5. Experiments and Results

For each of the three CGZSL problem settings – static, dynamic and online, we evaluate the proposed approach on five different benchmark datasets: Animals with Attributes (AWA1 and AWA2) [11, 31], Attribute Pascal and Yahoo (aPY) [3], Caltech-UCSD-Birds 200-2011 (CUB) [28] and SUN [21]. We follow the data splits mentioned in [6, 9] for fair comparison. We follow [6, 7, 26] and evaluate our model on mean seen accuracy (mSA), mean unseen accuracy (mUA) and mean harmonic value (mH) over all tasks. Further details on the metrics are provided in the Appendix.

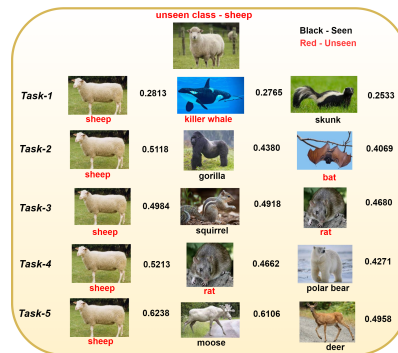


Figure 4. Cosine similarity scores of unseen class ‘sheep’ of AWA2 dataset w.r.t. identifier projection. Top three cosine similarity scores shared with ‘sheep’ during inference at every task is shown. Unseen classes are in red, seen classes are in black. For detailed explanation, refer appendix.

Baselines. In order to analyze the need for a CGZSL model for tackling dynamic addition of new classes, we

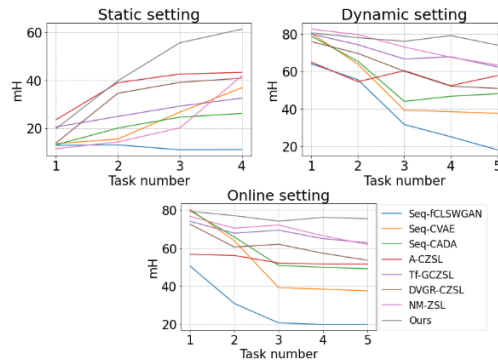


Figure 5. Task-wise mH values for AWA2 dataset in static (top left), dynamic (top right) and online (below) setting.

	aPY			AWA1			AWA2			CUB			SUN		
	mSA	mUA	mH	mSA	mUA	mH	mSA	mUA	mH	mSA	mUA	mH	mSA	mUA	mH
Online-CGZSL															
EWC [10]	40.44	-	-	53.62	-	-	55.18	-	-	38.64	-	-	25.79	-	-
AGEM [2]	-	-	-	-	-	-	57.25	-	-	41.93	-	-	28.61	-	-
Seq-fCLSWGAN [32]	26.92	16.34	20.34	24.83	15.29	18.93	24.56	16.84	19.98	14.26	10.0	11.76	14	10.38	11.92
Seq-CVAE [20]	63.58	16.24	25.87	72.41	26.39	38.68	71.84	25.33	37.45	36.23	21.70	27.14	30.87	22.69	26.16
Seq-CADA [25]	62.54	22.89	33.51	75.45	23.28	45.58	78.59	35.81	49.20	53.18	27.22	36.00	40.67	24.13	30.29
CV +mem [20]	77.74	27.62	40.76	81.35	39.63	53.29	85.29	34.64	49.27	64.27	26.39	37.42	36.11	22.39	27.64
CA +mem [25]	67.18	35.06	46.07	79.23	58.28	67.15	78.49	55.32	64.90	64.83	40.99	50.22	45.25	28.19	34.74
A-CZSL [8]	59.44	18.94	28.51	78.56	37.41	52.23	79.75	38.11	51.57	41.77	24.34	30.39	16.33	9.46	11.98
Tf-GCZSL [7]	69.01	24.81	37.59	65.01	50.43	56.80	65.59	60.17	62.76	65.36	44.19	52.80	40.18	41.05	37.81
DVGR-CZSL [9]	57.97	28.46	35.27	73.55	44.21	55.22	78.36	40.75	53.62	45.68	16.65	23.94	23.02	14.00	16.92
NM-ZSL [26]	-	-	-	-	-	-	84.25	48.94	61.91	61.40	49.64	54.9	53.57	27.51	36.35
Ours	75.79	37.72	49.11	78.99	69.33	73.33	80.06	72.62	75.44	44.09	53.57	48.03	41.87	45.72	43.20

Table 1. Mean seen accuracy (mSA), mean unseen accuracy (mUA), and their harmonic mean (mH) for Online-CGZSL setting.

	Dynamic-CGZSL		Online-CGZSL	
	AWA2	CUB	AWA2	CUB
A-CZSL	0.09	0.14	0.07	0.13
Tf-GCZSL	0.09	0.07	0.09	0.09
DVGR-CZSL	0.15	0.14	0.12	0.12
NM-ZSL	0.11	0.08	0.10	0.08
Ours	0.09	0.13	0.09	0.11

Table 2. Forgetting measure for dynamic & online settings on AWA2 & CUB datasets.

compare our work with sequential versions of well-known traditional GZSL baselines: f-CLSWGAN [32], CVAE [20] and CADA [25] (which we call Seq-fCLSWGAN, Seq-CVAE and Seq-CADA respectively). Existing continual learning methods have shown acceptable performance in retaining previously learned knowledge, but fail to handle unseen classes. We compare our method with notable continual learning frameworks EWC [10] and A-GEM [2]. Considering the importance of memory based replay for overcoming catastrophic forgetting, we assess our model with sequential versions of CVAE and CADA after adding a memory based replay strategy (which we call CV+mem and CA+mem respectively). We further evaluate the performance of our model against recent efforts that try to solve the CGZSL problem: (i) A-CZSL [8]; (ii) Tf-GCZSL [7]; (iii) DVGR-CZSL [9]; (iv) NM-ZSL [26]. We compare these baselines with our model on the three settings described in Sec. 3.

Results. Table 1 shows the performance comparison of our model in online setting with multiple baselines. Table 3 shows the performance of our model on dynamic and static settings. As a first effort in addressing the Online-CGZSL setting, we convert an unseen class from the $t - 1^{th}$ task to a seen class in the t^{th} task, in addition to introducing new seen and unseen classes, and use this for evaluating our method as well as the baselines. We introduce benchmark results on all datasets and baselines for this setting in Table 1. We observe that our model outperforms the baselines by a significant margin of 3.04%, 6.18%, 10.54% and 5.39%

on aPY, AWA1, AWA2 and SUN datasets respectively on mean harmonic value(mH). In Fig. 5 (bottom), we show mH per task for the AWA2 dataset in online setting.

On Dynamic-CGZSL (Table 3), we observe that our model yielded 1.18%, 3.57%, 8.14% and 2.91% increase in mH on aPY, AWA1, AWA2 and SUN datasets and competitive performance on CUB dataset with respect to all baselines. Our model outperforms 4 out of 5 state-of-the-art unseen accuracies on CUB, AWA1, AWA2 and SUN datasets. Considering unseen class accuracy poses a greater challenge than seen classes, the consistent performance of the proposed approach on unseen accuracy corroborates our claim. We believe that adding other continual learning strategies (such as regularizers) can further improve the performance of our method on seen classes. In Fig. 5 (top right), we evaluate mH per task for the AWA2 dataset in dynamic setting.

On Static-CGZSL (Table 3), we observe that our model obtained state-of-the-art performance with noticeable gains of 2.26%, 6.84% and 9.81% over the baselines on aPY, AWA1 and AWA2 datasets for mH. The results support our claim that our model performs exceedingly well on unseen accuracies across all settings, thanks to our bi-directional alignment loss which greatly helped in boosting performance. Fig. 5 (top left) shows mH of the model versus the task number for the AWA2 dataset in static setting.

We report forgetting measures of AWA2 and CUB datasets in dynamic and online settings in Table 2, which shows that addition of generative replay helps to reduce forgetting. Fig. 4 demonstrates how cosine similarity of an unseen class changes across different tasks during inference.

6. Ablation Studies and Analysis

To understand the importance of different components, Table 4 shows ablation studies of our approach on coarse and fine-grained datasets i.e AWA1 and CUB.

Role of replay data. In order to retain previous task information, we replay samples from seen classes using the generator network. To confirm that replaying data helps in overcoming catastrophic forgetting we execute the model without replaying visual features of seen class and observe that

	aPY			AWA1			AWA2			CUB			SUN		
	mSA	mUA	mH	mSA	mUA	mH	mSA	mUA	mH	mSA	mUA	mH	mSA	mUA	mH
Dynamic-CGZSL															
EWC [10]	42.68	-	-	50.41	-	-	53.29	-	-	40.95	-	-	24.46	-	-
AGEM [2]	-	-	-	-	-	-	57.0	-	-	43.1	-	-	27.15	-	-
Seq-fCLSWGAN [32]	25.73	15.27	19.17	23.39	12.57	16.35	22.88	14.92	18.06	13.14	7.47	9.53	12.70	8.63	10.27
Seq-CVAE [20]	65.87	17.90	25.84	70.24	28.36	39.32	73.71	26.22	36.30	38.95	20.89	27.19	29.06	21.33	24.33
Seq-CADA [25]	61.17	21.13	26.37	78.12	25.93	47.06	79.89	36.64	47.99	55.55	26.96	35.62	42.21	23.47	29.60
CV +mem [20]	78.15	28.10	40.21	85.01	37.49	51.60	88.36	33.24	47.89	63.16	27.50	37.84	37.50	24.01	29.15
CA +mem [25]	66.30	36.59	45.08	81.86	61.39	69.92	82.19	55.98	65.95	68.18	42.44	50.68	47.18	30.30	34.88
A-CZSL [8]	64.06	16.82	24.46	78.03	35.38	52.51	82.91	42.19	57.74	47.34	27.67	32.77	15.26	8.93	11.27
Tf-GCZSL [7]	72.12	19.66	30.90	61.79	57.77	59.72	67.42	58.08	62.41	44.52	43.21	43.85	27.76	39.09	32.46
DVGR-CZSL [9]	69.67	31.58	43.44	76.9	42.04	53.38	78.17	40.44	50.89	47.28	22.58	29.77	23.37	16.33	18.65
NM-ZSL [26]	79.60	22.29	32.61	75.59	60.87	67.44	89.22	51.38	63.41	64.91	46.05	53.79	50.56	35.55	41.65
Ours	74.92	33.94	46.26	79.51	69.13	73.49	79.19	70.71	74.09	50.51	52.2	51.18	45.43	44.59	44.56
Static-CGZSL															
EWC [10]	37.65	-	-	42.6	-	-	43.71	-	-	35.89	-	-	21.36	-	-
AGEM [2]	-	-	-	-	-	-	52.61	-	-	39.54	-	-	24.22	-	-
Seq-fCLSWGAN [32]	21.18	8.96	12.59	20.32	9.25	12.71	18.97	8.45	11.69	11.16	6.54	8.25	10.02	6.89	8.16
Seq-CVAE [20]	51.57	11.38	18.33	59.27	18.24	27.14	61.42	19.34	28.67	24.66	8.57	12.18	16.88	11.40	13.38
Seq-CADA [25]	45.25	10.59	16.42	51.57	18.02	27.59	52.30	20.30	30.38	40.82	14.37	21.14	25.94	16.22	20.10
CV +mem [20]	64.88	15.24	23.90	78.56	23.65	35.51	80.97	25.75	38.34	44.89	13.45	20.15	23.99	14.10	17.63
CA +mem [25]	57.69	20.83	28.84	62.64	38.41	45.38	62.80	39.23	46.22	43.96	32.77	36.06	27.11	21.72	22.92
A-CZSL [8]	58.14	15.91	23.05	71.00	24.26	35.75	70.16	25.93	37.19	34.47	12.00	17.15	-	-	-
Tf-GCZSL [7]	57.92	21.22	29.55	64.00	38.34	46.16	64.89	40.23	48.33	46.63	32.42	36.31	28.09	24.70	24.79
DVGR-CZSL [9]	62.5	22.7	31.7	65.1	28.5	38	73.5	28.8	40.6	44.87	14.55	21.66	22.36	10.67	14.54
NM-ZSL [26]	45.26	21.35	27.18	70.90	37.46	48.75	76.33	39.79	51.51	55.45	43.25	47.04	50.01	19.77	28.04
Ours	60.96	26.5	33.96	66.30	50.63	55.59	67.15	54.19	61.32	42.39	37.09	36.35	21.44	27.57	23.87

Table 3. Mean seen accuracy (mSA), mean unseen accuracy (mUA), and their harmonic mean (mH) for dynamic & static-CGZSL settings.

	AWA1			CUB		
	mSA	mUA	mH	mSA	mUA	mH
Ours	78.99	69.33	73.33	44.09	53.57	48.03
Ours w/o replay	21.27	33.64	20.24	1.89	2.74	2.21
Ours w/o L_{sal}	78.16	64.38	69.10	33.83	49.62	40.10
Ours w/o $L_{nuclear}$	72.33	67.62	69.31	37.66	42.36	39.53
Ours w/o L_{sal} and $L_{nuclear}$	57.39	60.22	56.38	14.14	36.52	19.88
Ours w/o L_{rcl} and L_{pcl}	78.65	33.4	45.65	13.79	9.50	10.72
Ours w/o L_{snt}	77.16	61.87	67.74	37.53	45.66	40.92
Ours w/o replay, L_{rcl} and L_{pcl}	3.36	6.29	4.29	0.53	0.81	0.58

Table 4. Ablation study on AWA1 & CUB datasets for online-CGZSL setting.

mean harmonic accuracy value drops by 53.09% on AWA1 and 45.82% on CUB datasets respectively.

Effect of bi-directional alignment. We hypothesized that the bi-directional alignment loss helps the model accommodate and align new classes; we empirically show its significance by validating our model with and without the loss. Having only one of the alignment (one-directional loss) i.e either L_{sal} or $L_{nuclear}$ resulted in the drop of performance by 4% on AWA1 and 8% on CUB datasets. The performance of the model without bi-directional loss significantly drops by 16.95% and 28.15% on AWA1 and CUB datasets which supports the need for it.

Effect of classification loss. The absence of the classification losses which catalyzes the bi-directional alignment clearly indicates a drop of 27.68% on AWA1 and 37.31% on CUB datasets.

7. Conclusions and Future Work

Existing literature has witnessed accelerated efforts towards solving the GZSL problem, but the static nature of the setting does not allow gradual addition of categories

(seen or unseen) over time which hampers applicability in the real-world. In this work, we address the aforementioned issue, and push the existing continuous generalized zero-shot learning (CGZSL) boundaries towards a more pragmatic *Online-CGZSL* setting. To this end, we propose a bi-directional alignment-based generative framework that can tackle changes in pool of categories and data availability, over time. Extensive experiments conducted on benchmark datasets verify that our approach demonstrates superior performance compared to existing baselines. Our future efforts will include exploring dynamic feature-based attention mechanisms for designing adaptable models to better address the CGZSL setting.

Limitation and Broader Impacts. The continual generalized zero-shot setting (CGZSL) is a step towards designing AI systems that are adaptable to changes in data availability and dynamic additions of object categories, which inevitably occur in the real-world over time. However, getting annotated attributes for all datasets may be non-trivial. Our work has no known social detrimental impacts other than the ones generally associated with the development of a new AI system.

Acknowledgements. This work has been partly supported by the funding received from DST through the IMPRINT and ICPS programs. Hari Chandana Kuchibhotla would like to thank MoE for the PMRF fellowship support. We thank the anonymous reviewers for their valuable feedback that improved the presentation of this paper.

References

- [1] Shivam Chandhok and Vineeth N Balasubramanian. Two-level adversarial visual-semantic coupling for generalized zero-shot learning. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 3100–3108, 2021. [2](#)
- [2] Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with a-GEM. In *International Conference on Learning Representations (ICLR)*, 2019. [2](#), [7](#), [8](#)
- [3] Ali Farhadi, Ian Endres, Derek Hoiem, and David Forsyth. Describing objects by their attributes. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1778–1785. IEEE, 2009. [6](#)
- [4] Andrea Frome, Greg S. Corrado, Jonathon Shlens, Samy Bengio, Jeffrey Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. In *Advances in Neural Information Processing Systems (NIPS)*, page 2121–2129, 2013. [2](#)
- [5] Rui Gao, Xingsong Hou, Jie Qin, Li Liu, Fan Zhu, and Ling Shao. “Zero-VAE-GAN: Generating unseen features for generalized and transductive zero-shot learning. In *IEEE Transactions on Image Processing (TIP)*, pages 3665–3680, 2020. [2](#)
- [6] Chandan Gautam, Sethupathy Parameswaran, Ashish Mishra, and Suresh Sundaram. Generalized continual zero-shot learning, 2021. [2](#), [3](#), [6](#)
- [7] Chandan Gautam, Sethupathy Parameswaran, Ashish Mishra, and Suresh Sundaram. Tf-gcysl: Task-free generalized continual zero-shot learning. In *arXiv preprint arXiv:2103.10741*, 2021. [3](#), [6](#), [7](#), [8](#)
- [8] Subhankar Ghosh. Adversarial training of variational autoencoders for continual zero-shot learning(a-czsl). In *International Joint Conference on Neural Networks (IJCNN)*, pages 1–8, 2021. [3](#), [7](#), [8](#)
- [9] Subhankar Ghosh. Dynamic vaes with generative replay for continual zero-shot learning, 2021. [2](#), [3](#), [6](#), [7](#), [8](#)
- [10] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, ..., and Raia Hadsell. Overcoming catastrophic forgetting in neural networks. In *Proceedings of the national academy of sciences*, pages 3521–3526, 2017. [2](#), [7](#), [8](#)
- [11] 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 (CVPR)*, pages 951–958. IEEE, 2009. [6](#)
- [12] Christoph H. Lampert, Hannes Nickisch, and Stefan Harmeling. Attribute based classification for zero-shot visual object categorization. In *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, pages 453–465, 2013. [2](#)
- [13] Jingjing Li, Mengmeng Jing, Ke Lu, Zhengming Ding, Lei Zhu, and Zi Huang. Leveraging the invariant side of generative zero-shot learning.. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7402–7411, 2019. [2](#)
- [14] Zhizhong Li and Derek Hoiem. Learning without forgetting. In *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, pages 2935–2947, 2017. [2](#)
- [15] Xialei Liu, Chenshen Wu, Mikel Menta, Luis Herranz, Bogdan Raducanu, ..., and Joost van de Weijer. Generative feature replay for class-incremental learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR) workshops*, pages 226–227, 2020. [2](#)
- [16] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X. Yu. Large-scale long-tailed recognition in an open world. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [1](#)
- [17] Arun Mallya and Svetlana Lazebnik. PackNet: Adding Multiple Tasks to a Single Network by Iterative Pruning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7765–7773, 2018. [2](#)
- [18] Nikhil Mehta, Kevin Liang, Vinay Kumar Verma, and Lawrence Carin. Continual Learning using a Bayesian Non-parametric Dictionary of Weight Factors. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 100–108, 2021. [2](#)
- [19] MichaelMcCloskey and Neal J.Cohen. Catastrophic interference in connectionist networks. In *Psychology of learning and motivation, vol. 24 academic press*, pages 109–165, 1989. [2](#)
- [20] Ashish Mishra, Shiva Krishna Reddy, Anurag Mittal, and Hema A. Murthy. A generative model for zero shot learning using conditional variational autoencoders. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR) workshops*, pages 2188–2196, 2018. [2](#), [7](#), [8](#)
- [21] Genevieve Patterson and James Hays. Sun attribute database: Discovering, annotating, and recognizing scene attributes. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2751–2758. IEEE, 2012. [6](#)
- [22] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, G. Sperl, and Christoph H. Lampert. icarl: Incremental classifier and representation learning. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5533–5542, 2017. [2](#)
- [23] Bernardino Romera-Paredes and Philip Torr. An embarrassingly simple approach to zero-shot learning. In *International Conference on Machine Learning (ICML)*, pages 2152–2161, 2015. [2](#)
- [24] Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016. [2](#)
- [25] Edgar Schonfeld, Sayna Ebrahimi, Samarth Sinha, Trevor Darrell, and Zeynep Akata. Generalized zero-and few-shot learning via aligned variational autoencoders. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8247–8255, 2019. [7](#), [8](#)
- [26] Ivan Skorokhodov and Mohamed Elhoseiny. Class normalization for (continual)? generalized zero-shot learning. In *International Conference on Learning Representations (ICLR)*, 2021. [2](#), [3](#), [6](#), [7](#), [8](#)

- [27] Maunil R. Vyas, Hemanth Venkateswara, and Sethuraman Panchanathan. Leveraging Seen and Unseen Semantic Relationships for Generative Zero-Shot Learning. In *European Conference on Computer Vision (ECCV)*, pages 70–86, 2020. [5](#)
- [28] C. Wah, Steve Branson, P. Welinder, P. Perona, and Serge J. Belongie. The caltech-ucsd birds-200-2011 dataset. In . California Institute of Technology, 2011. [6](#)
- [29] Kai Wang, Luis Herranz, Anjan Dutta, and Joost van de Weijer. Bookworm continual learning: beyond zero-shot learning and continual learning, 2020. [2](#), [3](#)
- [30] Kun Wei, Cheng Deng, and Xu Yang. Lifelong zero-shot learning. In *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 551–557, 2020. [2](#)
- [31] 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 (TPAMI)*, 41(9):2251–2265, 2018. [6](#)
- [32] Yongqin Xian, Tobias Lorenz, Bernt Schiele, and Zeynep Akata. Feature generating networks for zero shot learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5542–5551, 2018. [2](#), [7](#), [8](#)
- [33] Lu Yu, Bartłomiej Twardowski, Xialei Liu, Luis Herranz, Kai Wang, Yongmei Cheng, and Shangling Jui. Semantic drift compensation for class-incremental learning.. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6982–6991, 2020. [2](#)
- [34] 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 (CVPR)*, pages 2021–2030, 2017. [2](#)
- [35] Yifan Zhang, Bingyi Kang, Bryan Hooi, Shuicheng Yan, and Jiashi Feng. Deep long-tailed learning: A survey, 2021. [1](#)