

Meta-Learned Attribute Self-Interaction Network for Continual and Generalized Zero-Shot Learning

Vinay Verma^{2*}, Nikhil Mehta^{2*}, Kevin J Liang³, Aakansha Mishra⁴, Lawrence Carin⁵
²Duke University, ³FAIR, Meta, ⁴IITG, ⁵KAUST
 {¹vverma.vinay, ²nikhilmehta.dce}@gmail.com

Abstract

Zero-shot learning (ZSL) is a promising approach to generalizing a model to categories unseen during training by leveraging class attributes, but challenges remain. Recently, methods using generative models to combat bias towards classes seen during training have pushed state of the art, but these generative models can be slow or computationally expensive to train. Also, these generative models assume that the attribute vector of each unseen class is available a priori at training, which is not always practical. Additionally, while many previous ZSL methods assume a one-time adaptation to unseen classes, in reality, the world is always changing, necessitating a constant adjustment of deployed models. Models unprepared to handle a sequential stream of data are likely to experience catastrophic forgetting. We propose a Meta-learned Attribute self-Interaction Network (MAIN) for continual ZSL. By pairing attribute self-interaction trained using meta-learning with inverse regularization of the attribute encoder, we are able to outperform state-of-the-art results without leveraging the unseen class attributes while also being able to train our models substantially faster ($> 100\times$) than expensive generative-based approaches. We demonstrate this with experiments on five standard ZSL datasets (CUB, aPY, AWA1, AWA2, and SUN) in the generalized zero-shot learning and continual (fixed/dynamic) zero-shot learning settings. Extensive ablations and analyses demonstrate the efficacy of various components proposed.

1. Introduction

Deep learning has demonstrated the ability to learn powerful models given a sufficiently large, labeled dataset of a pre-defined set of classes [17,27]. However, such models often generalize poorly to classes unseen during training. In an ever-evolving world in which new concepts or applications are to be expected, this brittleness can be an undesirable characteristic. In recent years, zero-shot learning (ZSL) [1, 16, 20, 50] has been proposed as an alterna-

tive framework to classify the novel class data. ZSL approaches seek to leverage auxiliary information about these new classes, often in the form of class attributes. This side information allows for reasoning about the relations between classes, enabling the adaptation of the model to recognize samples from one of the novel classes. In the more general setting, a ZSL model should be capable of classifying inputs from both seen and unseen classes; this difficult setting is commonly referred to as generalized zero-shot learning (GZSL) [51, 59].

Some of the strongest results in GZSL [16, 29, 46, 50, 51, 61] have come from utilizing generative models. By learning a generative mapping between the attributes and the data, synthetic samples can be conditionally generated from the unseen class attributes. The model can then be learned in the usual supervised manner on the joint set of seen and (generated) unseen class data, mitigating model bias toward the seen classes. While effective, training the requisite generative model, generating data, and training the model on this combined dataset can be expensive [33, 40, 50]. Furthermore, these generative methods require the unseen class attributes during training; if the unseen classes are not yet known, defining attributes for unseen classes is challenging and error-prone. Thus, the requirement of knowing the unseen class attributes during training makes these methods less practical, hindering their applicability to the real world.

If assuming a one-time adaptation from a pre-determined set of training classes, the cost of such a one-off process may be considered acceptable, but it may present challenges if model adaptations need to be done repeatedly. Most ZSL methods commonly consider only one adaptation, but in reality environments are often dynamic, and new class data may appear sequentially. For example, if it is important for a model to be able to classify a previously unseen class, it is natural that a future data collection effort may later make labeled data from these classes available [20]. Alternatively, changing requirements may require the model to learn from and then generalize to entirely new seen and unseen classes [12]. In such cases, the model should be able to learn from new datasets without catastrophically for-

*Equal contribution.

getting [38] previously seen data, even if that older data is no longer available in its entirety. Thus, it is important that ZSL methods can work in continual learning settings as well.

To address these issues, we propose the Meta-learned Attribute self-Interaction Network (MAIN) for Generalized Zero-Shot Learning (GZSL), which neither requires unseen class attributes *a priori* nor any expensive model adaptations. The MAIN framework learns a unified visual embedding space for the class-specific attributes and the corresponding images belonging to a class. To extract the attribute embeddings, MAIN learns an attribute encoder that maps the semantic information in attributes to the visual embedding space. The attribute encoder in MAIN uses a novel self-interaction module for attributes and meta-learning to generalize the encoder to unseen class attributes. Moreover, MAIN incorporates a theoretically motivated inverse regularization loss, which preserves the semantic attribute information in the visual embedding space. In experiments, we discuss the importance of various components introduced as part of the MAIN framework. We show that training in MAIN is $100\times$ faster than the prior methods that use generative models. We extend MAIN to the continual GZSL setting using a small reservoir of samples and show that MAIN outperforms the recent alternatives proposed for CZSL. Extensive experiments on CUB, aPY, AWA1, AWA2, and SUN datasets demonstrate that MAIN achieves state-of-the-art results in ZSL, GZSL, and continual GZSL settings.

2. Related Work

2.1. Zero-Shot Learning

The ZSL literature is vast, with approaches that can be roughly divided into (i) non-generative and (ii) generative approaches. Initial work [1, 2, 11, 19, 42, 58] mainly focused on non-generative methods, with a primary objective of learning a function from the seen classes that can measure the similarity between the image embeddings (typically features extracted from a pre-trained model) and the attribute embeddings. [31, 42, 58] measure the linear compatibility between the image and attribute embeddings; however, a linear map assumption does not hold for complex relationships between the two spaces. Another set of works [24, 65] focuses on modeling the relation with bi-linear functions. These approaches show promising results for the ZSL setup (where only unseen classes are evaluated), but they perform poorly in the GZSL setting (where during inference, both seen and unseen classes are present).

Of late, generative approaches have been popular for GZSL. Due to rapid progress in generative modeling (*e.g.*, VAEs [23], GANs [15]) generative approaches have been able to synthesize increasingly high-quality samples. For example, [6, 8, 21, 39, 40, 46, 47, 51, 60, 61] have used conditional VAEs or GANs to generate samples for un-

seen classes conditioned on the class attributes, which can then be used for training alongside samples from the seen classes. Given the ability to generate as many samples as needed, these approaches can easily handle model bias towards seen classes, leading to promising results for both ZSL and GZSL [16, 61]. However, these generative models assume all seen, and unseen class attributes are present during training. In many practical settings, this can be a strong assumption, as the model may not know what the unseen classes will be ahead of time. Also, the full pipeline required learning a generative model, synthesizing samples for the unseen classes, and training a classifier on the given and synthesized data makes the generative model-based ZSL framework expensive.

The aforementioned problems with generative models and recent promising results of non-generative models [20, 34] motivate us to revisit non-generative approaches. [20] propose an inexpensive approach for the GZSL and shows that normalization and initialization may play a crucial role for GZSL and outperform the expensive generative model. These normalization techniques can be tricky: high sensitivity to hyperparameters means small changes quickly degrade model performance. We propose a self-gating mechanism, meta-learning-based training, and multi-modal regularization in a non-generative framework. Even without tricky and unstable normalization, we can outperform the recent generative or normalization-based models by a significant margin.

2.2. Zero-Shot Continual Learning

One of the desiderata of continual learning techniques is the forward transfer of previous knowledge to future tasks, which may not be known ahead of time; similarly, GZSL approaches seek to adapt models to the novel, unseen classes while still being able to classify the seen classes. As such, there are clear connections between the two problem settings. Some recent works [12, 13, 20, 28, 37, 57] have drawn increasing attention towards continual zero-shot learning (CZSL). For example, [57] considers a task incremental learning setting, where task ID for each sample is provided during train and test, leading to an easier and perhaps less realistic setting than class incremental learning. A-GEM [3] proposed a regularization-based model to overcome catastrophic forgetting while maximizing the forward transfer. [20] proposed a simple class normalization as an efficient solution to ZSL and extended it to CZSL, where the attributes for the unseen classes of future tasks are known *a priori*. This setting is referred to as Fixed Continual GZSL. Meanwhile, [12, 13, 28] proposed a replay-based approach, showing state-of-the-art results in a more realistic setting for CZSL called Dynamic Continual GZSL, where the attributes for unseen classes of future tasks are not known *a priori*.

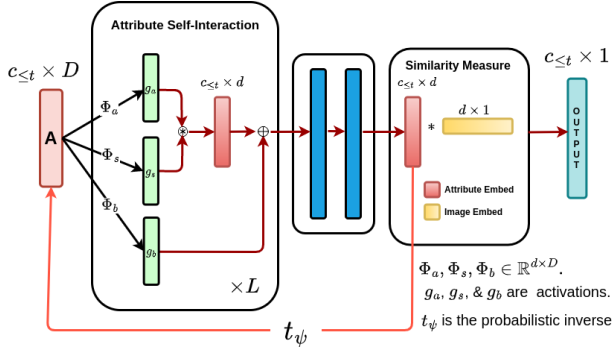


Figure 1. Meta-learned Attribute Self-Interaction Network.

3. Proposed Approach

Our proposed Meta-Learned Attribute self-Interaction Network (MAIN)¹ for continual ZSL can be divided into three major components: (i) an attribute encoder designed using our novel self-interaction module, which maps the structured class-specific attribute to an attribute embedding in the visual feature space, (ii) Inverse Regularization (IR) that preserves the semantic information in the visual feature space, preventing over-fitting to the seen classes, and (iii) a meta-learning [41] framework for training the attribute encoder with reservoir sampling to prevent catastrophic forgetting of previous tasks, while also learning a model that generalizes to novel classes without fine-tuning.

3.1. Background and Notation

The ZSL and GZSL follow the standard setting proposed by [59], we provide more details in the supplementary material. Here we mostly focus on the continual setting. Let $T_t = \{\mathcal{D}_{tr}^t, \mathcal{D}_{ts}^t\}$ be a task arriving at time t , where \mathcal{D}_{tr}^t and \mathcal{D}_{ts}^t are the train and test sets associated with the t^{th} task, respectively. In a continual learning setup, we assume that a set of tasks arrive sequentially, such that the training data for only the current task is made available. Let this sequence of tasks be $\{T_1, T_2, \dots, T_K\}$, where at time t , the training data for only the t^{th} task is available. In continual learning, the goal is to learn a new task while preventing catastrophic forgetting of previously seen tasks. Hence, during testing, the model is evaluated on the current and previous tasks that the model has encountered before. We assume that for a given task t , we have $c_{\le t}^s$ number of seen classes that we can use for training and $c_{\le t}^u$ unseen/novel classes that we are interested in adapting our model to. Analogous to GZSL, during test time, we assume that samples come from any of the seen or unseen classes of the current or previous tasks, *i.e.*, the test samples for task T_t contain $c_{\le t}^s$ ($c_{\le t}^s = \sum_{k=1}^t c_k^s$) number of seen classes and $c_{\le t}^u$ number of unseen classes, where $c_{\le t}^u$ will depend on the chosen evaluation protocol as described in supplementary. We denote $\mathcal{D}_{tr}^t = \{x_i^t, y_i^t, a_{y_i}, t\}_{i=1}^{N_{tr}^s}$ to be the training data, where $x_i^t \in \mathbb{R}^d$ is the visual feature

of sample i of task t , y_i^t is the label of x_i^t , a_{y_i} is the attribute/description vector of the label, t is the task identifier (id) and \mathcal{Y}_c is the label set for c classes. Similarly, testing data for task t is

$$\mathcal{D}_{ts}^t = \left\{ \{x_i^t\}_{i=1}^{N_{ts}^u}, y_{x_i^t}, \mathcal{A}_t : \mathcal{A}_t = \{a_y\}_{y=1}^{c_{\le t}^s + c_{\le t}^u} \right\} \quad (1)$$

where $y_{x_i^t} \in \{k \in \mathbb{N} \mid k \leq c_{\le t}^s + c_{\le t}^u\}$. For each task T_t , the seen and unseen classes are disjoint: $\mathcal{Y}_{c_{\le t}^s} \cap \mathcal{Y}_{c_{\le t}^u} = \emptyset$. In each task t , seen and unseen classes have an associated attribute/description vector $a_y \in \mathcal{A}_t$ that helps to transfer knowledge from seen to unseen classes. We will drop the class-specific indexing for attributes in $\{a_y\}$ and simply denote the attributes as $\{a\}$ in the following sections of the paper. We use N_{tr}^s and N_{ts}^u to denote the number of training and testing samples for the task t , respectively. We assume that there are a total of S seen classes and U unseen classes for each dataset, and the attribute set is $\mathcal{A} \in \mathbb{R}^{S \cup U \times D}$, where D is the attribute dimension. In the continual learning setup, test samples contain both unseen and seen classes up to task t , and the objective of continual ZSL is to predict future task classes while also overcoming catastrophic forgetting of the previous tasks.

3.2. Self-Interaction for Attributes (SIA)

We develop a self-interaction module for the attribute vector that can continually learn from sequential data and provide a robust class-representative vector in the visual embedding space. Let $\mathcal{A}_{\le t}^s \subset \mathbb{R}^D$ be the set of seen class attributes up to the training of task T_t , where $|\mathcal{A}_{\le t}^s| = c_{\le t}^s$. The self-interaction module comprises of three transformations parameterized as Φ_a , Φ_s , and Φ_b followed by activations g_a , g_s and g_b , respectively. Consider an attribute encoder f_{Φ} with L self-interaction modules, where the output of each module is denoted as \mathbf{a}_{ℓ} . Given an attribute vector $\mathbf{a} \in \mathcal{A}_{\le t}^s$, we define $\mathbf{a}_{\ell+1}$ as:

$$\mathbf{a}_{\ell+1} = g_a(\Phi_a(\mathbf{a}_{\ell})) * g_s(\Phi_s(\mathbf{a}_{\ell})) + g_b(\Phi_b(\mathbf{a}_{\ell})), \quad (2)$$

where $\mathbf{a}_0 = \mathbf{a}$. The final output of the attribute encoder $\mathbf{z} = f_{\Phi}(\mathbf{a})$ is computed by passing \mathbf{a}_L to two fully-connected layers as shown in figure 1. Next, we analyze two types of self-interactions, depending on type of activation functions.

3.2.1 Polynomial Kernels

The self-interaction module described above is equivalent to the polynomial kernel when the activations g_a , g_s and g_b are linear. If we stack L self-interaction modules, the output of layer- ℓ approximates the class of polynomials that grow exponentially in degree.

Lemma 1 (Polynomial Approximation) Consider a model with L layers of self-interaction modules as defined in (2) with parameters $\{\Phi_a^{\ell}, \Phi_s^{\ell}, \Phi_b^{\ell}\}_{\ell=1}^L$ and identity activation $g_a^{\ell}(x) = g_s^{\ell}(x) = g_b^{\ell}(x) = x$. Let input to

¹https://anonymous.4open.science/r/main_zsl-C1BD

the model be: $\mathbf{a} = [a_1, a_3, \dots, a_D]$. Then, the output of the model \mathbf{a}_L approximates following class of polynomial functions:

$$\left\{ P_\ell(\mathbf{a}) = \sum_{\beta} w_{\beta} a_1^{\beta_1} a_2^{\beta_2} \dots a_D^{\beta_D} \mid 0 \leq |\beta| \leq 2^\ell, \beta \in \mathbb{N}^D \right\},$$

where the sum is across multiple terms (monomials), $\beta = [\beta_1, \dots, \beta_D]$ is a vector containing the exponents of each attribute in a given term having degree $|\beta| = \sum_{i=1}^D \beta_i$, and w_{β} is the coefficient of the corresponding term that depends on the module parameters. Furthermore, the degree of the polynomial grows exponentially with the model depth.

The proof of Lemma 1 is in the Appendix. Notably, the output polynomial features contain high-order interactions of the input attribute, resulting in an architecture that combines explicit feature engineering of attributes with implicit functions in deep neural networks.

3.2.2 Self-Gating

The second type of self-interaction we propose is self-gating. For self-gating, we set g_a and g_b as the ReLU and g_s as the sigmoid activation:

$$\mathbf{a}_{\ell+1} = \text{ReLU}(\Phi_a(\mathbf{a}_\ell)) * \sigma(\Phi_s(\mathbf{a}_\ell)) + \text{ReLU}(\Phi_b(\mathbf{a}_\ell)) \quad (3)$$

where σ is the sigmoid activation function. The function Φ_s with a sigmoid activation maps the projected attribute to the range $[0, 1]$, acting as a gate for each dimension of the function Φ_a . Values closer to one signify the higher importance of a particular attribute dimension, while values closer to zero imply the opposite. Finally, function Φ_b projects the attribute to the same space as Φ_s and Φ_a , resulting in a bias vector similar to the residual connection. We empirically find that Φ_a and Φ_s help learn a robust and global attribute vector while Φ_b helps stabilize model training. In our experiments, we found self-gating to perform better than polynomial-based self-interactions. Unless stated otherwise, we use self-gating modules in MAIN for the main results in experiments. In section 5.4, we include an ablation comparing the two types of self-interactions.

3.3. Inverse Regularization (IR)

In the previous section, we proposed two types of SIA modules that map the semantic attribute space to the visual feature space. As noted in 3.1, the SIA-based attribute encoder is learned from only the seen class data. Because of the propensity for model bias towards the seen class samples, the unseen class attributes can be projected to the seen class features. To overcome this problem, we propose using inverse regularization that learns to project the resulting visual features from the attribute encoder back to the semantic space. We first provide a theoretical motivation for IR, which shows that IR is closely connected to entropy regularization in the visual feature space.

Lemma 2 (Maximize Entropy with IR) Let $t_\xi(\mathbf{a}|z) = \mathcal{N}(\mathbf{a}; \mathcal{R}_\xi(z), I)$ be the probabilistic inverse map associated with the attribute encoder f_Φ , where $z = f_\Phi(\mathbf{a})$ denotes the attribute embedding and \mathcal{R}_ξ is the inverse regressor network with parameters ξ that projects the attribute embedding in the visual feature space back to the semantic attribute space. Then, the mutual information between the attribute \mathbf{a} and the attribute embedding z is defined as:

$$I(\mathbf{a}; z) = H(z; \Phi) \geq H(\mathbf{a}) + \mathbb{E}_{\mathbf{a} \sim p(\mathbf{a})} [\log t_\xi(\mathbf{a}|f_\Phi(\mathbf{a}))]. \quad (4)$$

The proof is included in the Appendix. Since $H(\mathbf{a})$ is a constant with respect to (Φ, ξ) , maximizing $\mathbb{E}_{\mathbf{a} \sim p(\mathbf{a})} [\log t_\xi(\mathbf{a}|f_\Phi(\mathbf{a}))]$ increases the entropy $H(z; \Phi)$ in the embedding space. Hence, adding IR maximizes the entropy in the induced visual feature space, which acts as a regularization when training the attribute encoder and avoids over-fitting the model to seen classes.

The IR can be easily achieved by minimizing the cyclic-consistency loss between the semantic and visual space. The inverse regularization loss corresponding to Lemma 2 is defined as:

$$\mathcal{L}_c = \sum_{\mathbf{a} \in \mathcal{A}_{\leq t}^s} \|\mathcal{R}_\xi(f_\Phi(\mathbf{a})) - \mathbf{a}\|^2 \quad (5)$$

Few recent works [8, 24, 51] have also explored similar regularization for improved performance. However, a theoretical justification of how IR for the attribute encoder is connected to entropy maximization in the visual feature space is one of the contributions of this work. In our experiments, we show that adding entropy regularization improves the model’s generalization to unseen classes.

Overall training loss: For the classification task, the probabilities are computed as $p(t = y|x) = e^{s_y} / \sum_{y'} e^{s_{y'}}$, where $s_y = f_\Phi(\mathbf{a}_y)^T x$ is the class logit and t is predicted label. We use the standard cross entropy (H) as the loss:

$$\mathcal{L}_s = H(p(t = y|x), e_y) \quad (6)$$

where e_y is the one-hot representation for the ground-truth label y corresponding to the image embedding x in the visual feature space. We optimize the joint loss obtained from Equations 5 and 6:

$$\mathcal{L}_\tau = \mathcal{L}_s + \lambda \mathcal{L}_c \quad (7)$$

where τ is the sampled batch and λ is a hyperparameter tuned on the validation data.

3.4. Reservoir Sampling

In continual learning settings, we assume that tasks arrive sequentially, such that at the i^{th} step, training samples are only available from the task T_i , and samples from previous tasks (T_1, \dots, T_{i-1}) are not accessible. Due to the tendency of neural networks to experience catastrophic forgetting [38], the model is likely to forget previously learned knowledge when learning new tasks. To mitigate

catastrophic forgetting, MAIN incorporates reservoir sampling [54] using a small and constant memory size to store selected samples from previous tasks and replaying these samples while training the attribute encoder on the current task. This has several advantages: (i) each task can be trained with constant time and computational resources, and (ii) the size of the memory does not grow as the tasks increase. Replaying samples from the reservoir effectively mitigates forgetting. We select each sample for replay with probability M/N , where M is the memory budget and N is the number of samples selected so far.

3.5. Training with Meta-Learning

While reservoir sampling is effective for mitigating forgetting, it also has several drawbacks. As the task number grows, a constant memory budget means the number of samples of each class diminishes, as the same amount of memory has to accommodate a large number of classes. Similarly, the current task will have a sizably larger number of samples for each class from the current task than the number of samples in the reservoir for classes of past tasks. Therefore, training can be difficult, resulting in models that generalize unevenly across all tasks/classes. This issue is similar to the problem of few-shot learning, where we have to learn a model using only a few training examples present in the reservoir from the previous tasks. Meta-learning models have shown promising results in the few-shot learning settings [10, 41]. Recent work [44, 48, 53] shows that even without any fine-tuning meta-learned models generalize well to novel classes.

Inspired by the recent success of meta-learning, MAIN leverages Reptile [41], a first-order meta-learning approach to train the attribute encoder. The training does not require storing gradients in memory for the inner loop, resulting in fast training of a generalized model. Given an attribute encoder f_Φ parameterized with Φ , we define U_τ^m as the operator representing m number of gradient updates of the loss function (7) with respect to Φ . With Reptile, the training batch τ is sampled from an augmented dataset containing the training samples from the current task and the stored memory. The encoder parameters are updated as follows:

$$\Phi \leftarrow \Phi - \eta(\Phi - \tilde{\Phi}), \quad (8)$$

where $\tilde{\Phi} = U_\tau^m(\Phi)$. The algorithm 1 shows the Reptile-based training for continual ZSL with an augmented dataset. As we later show in the ablations, meta-learning significantly improves various experimental settings considered for both continual and generalized ZSL. To the best of our knowledge, this is the first work to show the merits of meta-learning in the continual ZSL setting.

4. Experiments

While the primary focus of this work is on continual ZSL settings (*i.e.*, Fixed continual [20] and Dynamic con-

Algorithm 1 Reptile [41] for continual ZSL on a task t .

- 1: Given encoder f_Φ , task \mathcal{D}_t , memory \mathcal{C} , & $k > 1$.
- 2: **for** iteration = 1, 2, ... **do**
- 3: Sample task τ from the augmented dataset $\mathcal{C} \cup \mathcal{D}_t$.
- 4: Compute $\tilde{\Phi} = U_\tau^k(\Phi)$, denoting k steps of SGD starting at Φ .
- 5: Update: $\Phi \leftarrow \Phi + \epsilon(\tilde{\Phi} - \Phi)$.
- 6: **end for**

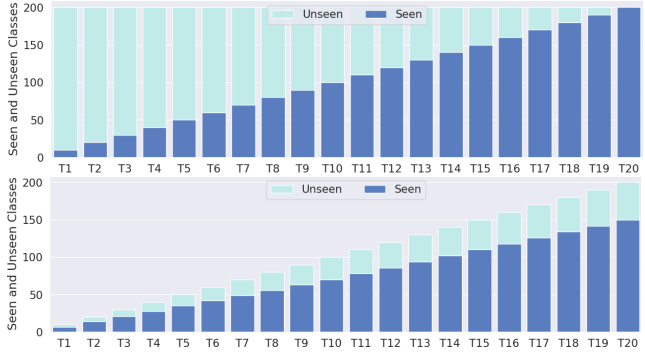


Figure 2. Seen and Unseen class distribution for the *fixed* and *dynamic* GZSL setting for the CUB dataset. In the fixed model as task proceed unseen classes decrease, however in dynamic GZSL as the task increase seen unseen increase. In both with the increase of the task forgetting increases.

tinual [12]), we also analyze and compare MAIN with several baselines on the non-continual and standard ZSL and GZSL settings. In supplementary, we discuss these different settings and the corresponding evaluation protocols in detail. In section 4.2, we show that MAIN outperforms recent continual GZSL baselines while also achieving state-of-the-art results on the standard GZSL and ZSL settings. Furthermore, the training of MAIN is $\sim 100\times$ faster than other alternatives which require training generative models. Additional implementation details are provided in the supplementary.

4.1. Datasets and Baselines

We conduct experiments on five widely used datasets for ZSL. In our experiment we used CUB-200 [56], AWA2 [59], AWA1 [30], SUN [43] and aPY [7] datasets. Further details about the datasets are in the supplementary.

4.1.1 Fixed and Dynamic Continual GZSL Tasks

In the fixed continual generalized zero-shot learning (GZSL) setting, we assume that we have a fixed number of classes that arrive sequentially. Figure-2 (top) illustrates the seen-unseen distribution for the CUB dataset, where the data is divided into 20 tasks. Initially, the problem is highly challenging as only a few classes are visible, while the remaining classes are unseen. As the tasks progress, the number of seen classes grows, and the number of unseen classes decreases, as the latter are added to the seen classes.

Table 1. Mean seen accuracy (mSA), mean unseen accuracy (mUA), and their harmonic mean (mH) for fixed continual GZSL.

	CUB			aPY			AWA1			AWA2			SUN		
	mSA	mUA	mH	mSA	mUA	mH	mSA	mUA	mH	mSA	mUA	mH	mSA	mUA	mH
Seq-CVAE [39]	24.66	8.57	12.18	51.57	11.38	18.33	59.27	18.24	27.14	61.42	19.34	28.67	16.88	11.40	13.38
Seq-CADA [46]	40.82	14.37	21.14	45.25	10.59	16.42	51.57	18.02	27.59	52.30	20.30	30.38	25.94	16.22	20.10
CZSL-CA+res [12]	43.96	32.77	36.06	57.69	20.83	28.84	62.64	38.41	45.38	62.80	39.23	46.22	27.11	21.72	22.92
NM-ZSL [20]	55.45	43.25	47.04	45.26	21.35	27.18	70.90	37.46	48.75	76.33	39.79	51.51	50.01	19.77	28.04
GRCZSL [13]	41.91	14.12	20.48	66.47	12.06	20.12	78.66	21.86	33.56	81.01	23.59	35.82	17.74	11.50	13.73
CARNet [14]	43.41	47.44	44.23	—	—	—	—	—	—	—	—	—	23.64	30.95	26.03
ULCT [28]	42.39	37.09	36.35	60.96	26.50	33.96	66.30	50.63	55.59	67.15	54.19	61.32	21.44	27.57	23.87
MAIN w/o IR	52.27	48.59	49.02	52.02	26.04	34.30	65.86	45.66	52.33	64.43	48.20	54.45	48.53	24.31	31.86
MAIN (Ours)	58.58	47.81	51.43	62.81	25.17	35.42	65.15	47.05	53.34	64.02	49.60	55.36	48.50	25.27	32.72

Table 2. Mean seen accuracy (mSA), mean unseen accuracy (mUA), and harmonic mean (mH) for dynamic continual GZSL.

	CUB			aPY			AWA1			AWA2			SUN		
	mSA	mUA	mH	mSA	mUA	mH	mSA	mUA	mH	mSA	mUA	mH	mSA	mUA	mH
Seq-CVAE [39]	38.95	20.89	26.74	65.87	17.90	25.84	70.24	28.36	39.32	73.71	26.22	36.30	29.06	21.33	24.33
Seq-CADA [45]	55.55	26.96	35.62	61.17	21.13	26.37	78.12	35.93	47.06	79.89	36.64	47.99	42.21	23.47	29.60
CZSL-CV+res [12]	63.16	27.50	37.84	78.15	28.10	40.21	85.01	37.49	51.60	88.36	33.24	47.89	37.50	24.01	29.15
CZSL-CA+res [12]	68.18	42.44	50.68	66.30	36.59	45.08	81.86	61.39	69.92	82.19	55.98	65.95	47.18	30.30	34.88
NM-ZSL [20]	64.91	46.05	53.79	79.60	22.29	32.61	75.59	60.87	67.44	89.22	51.38	63.41	50.56	35.55	41.65
GRCZSL [13]	59.27	26.03	35.67	77.61	22.26	33.01	87.88	30.34	44.32	90.45	36.66	51.74	30.78	22.59	25.54
UCLT [28]	50.51	52.2	51.18	74.92	33.94	46.26	79.51	69.13	73.49	79.19	70.71	74.09	45.43	44.59	44.56
MAIN w/o IR	61.73	71.14	65.94	55.41	39.29	43.51	82.37	64.99	72.11	86.79	60.82	70.52	46.26	50.70	48.25
MAIN (Ours)	64.59	70.40	67.20	57.50	38.59	45.10	80.56	69.05	73.83	85.56	65.59	73.45	47.25	52.01	49.39

The dynamic continual GZSL setting is more practical, where both seen and unseen classes increase as the tasks advance. Figure-2 (bottom) depicts the dynamic setting for the CUB dataset. Here, 150 seen classes and 50 unseen classes are divided into 20 tasks, with $[7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8]$ and $[3, 3, 3, 3, 3, 3, 3, 3, 3, 3, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2]$ classes in each task for seen and unseen data, respectively. In this scenario, both seen and unseen classes grow with each task, but only the data from the current task is available for training. For instance, in the 3rd task, there are 21 seen classes ($7 + 7 + 7$) and 9 unseen classes ($3 + 3 + 3$), yet only the data from the 3rd task’s seven classes are accessible for training. The reservoir memory for each dataset is set to $B \times \#total_classes$, where B represents the reservoir parameter, indicating the number of images per class. Specifically, the reservoir sample B is set to 25 for AWA1, AWA2, and aPY datasets, 10 for CUB dataset, and 5 for the SUN dataset. Further details for other datasets can be found in the supplementary materials.

4.1.2 Continual ZSL Baselines

This section describes the baselines used for the continual ZSL setting. For the non-continual (standard) ZSL and GZSL baselines, please refer directly to tables 3 and 4 for a variety of non-continual baselines we considered. To the best of our knowledge, continual ZSL has mainly been explored in [3, 12, 20, 28, 57]. For comparison, we consider CZSL-CV+res and CZSL-CA+res [12], which use conditional variational auto-encoders (VAE) and CADA [46] with a memory reservoir, respectively. We also consider the recently proposed GRCZSL [13] that uses generative replay to overcome catastrophic forgetting in the continual GZSL setting. NM-ZSL [20] is also a competitive continual ZSL baseline and other continual ZSL baselines include

Seq-CVAE [39] and Seq-CADA [46], which sequentially train a VAE over tasks to augment the training dataset using replay. The recent work UCLT [28] uses a generative model and replay to solve the CZSL problem.

4.2. Results

4.2.1 Fixed Continual GZSL

In the Fixed Continual GZSL [20] setting, as the model is trained on a sequence of tasks, the number of seen classes increases while the number of unseen classes decreases. The results for the CUB, aPY, AWA1, AWA2, and SUN datasets are shown in table 1. We use the same memory buffer size as in [12] for a fair comparison. On CUB, aPY, and SUN, the proposed model shows 4.39%, 1.46%, and 4.68% absolute increase over the best baseline. However, compared to AWA1 and AWA2, we have a marginal difference. We also analyze the performance of various methods across each task in the CUB dataset in Figure 3. We observe that our proposed MAIN consistently outperforms recent baselines. Notably, these significant improvements are achieved without using any costly generative models.




Figure 3. Harmonic mean per task for continual GZSL on the CUB dataset. **Left:** Represents the Fixed Continual GZSL setting. **Right:** Represents the Dynamic Continual GZSL setting.

4.2.2 Dynamic Continual GZSL

An alternate framing of continual GZSL is explored in the Dynamic Continual GZSL [12] setting, where each task consists of its own disjoint set of seen and unseen classes. In

Table 3. Mean seen accuracy (mSA), mean unseen accuracy (mUA), and their harmonic mean (mH) for GZSL. **Top:** Methods similar to the proposed non-generative model where the model does not require the unseen class attributes. **Bottom:** Expensive generative model approaches, which require unseen class attributes during training. w/o IR: no inverse regularization.

	SUN			CUB			AWA1			AWA2			Train Time
	mSA	mUA	mH	mSA	mUA	mH	mSA	mUA	mH	mSA	mUA	mH	
DEM [64]	34.3	20.5	25.6	57.9	19.6	29.2	84.7	32.8	47.3	86.4	30.5	45.1	20 minutes
CRNet [63]	36.5	34.1	35.3	56.8	45.5	50.5	78.8	52.6	63.1	74.7	58.1	65.4	–
LFCAA [36]	34.9	20.8	26.1	79.6	43.4	56.2	–	–	–	90.3	50.0	64.4	–
NM-ZSL [20]	44.7	41.6	43.1	49.9	50.7	50.3	63.1	73.4	67.8	60.2	77.1	67.6	30 seconds
CVC-ZSL [33]	36.3	42.8	39.3	47.4	47.6	47.5	62.7	77.0	69.1	56.4	81.4	66.7	3 hours
SGAL [62]	42.9	31.2	36.1	47.1	44.7	45.9	52.7	75.7	62.2	55.1	81.2	65.6	50 minutes
TF-VAEGAN [40]	45.6	40.7	43.0	52.8	64.7	58.1	–	–	–	59.8	75.1	66.6	1.75 Hours
LsrGAN [55]	44.8	37.7	40.9	48.1	59.1	53.0	–	–	–	54.6	74.6	63.0	1.25 hours
F-VAEGAN-D2 [61]	45.1	38.0	41.3	48.4	60.1	53.6	–	–	–	57.6	70.6	63.5	–
ZSML [50]	45.1	21.7	29.3	60.0	52.1	55.7	57.4	71.1	63.5	58.9	74.6	65.8	3 hours
DAZLE [18]	24.3	52.3	33.2	59.6	56.7	58.1	–	–	–	75.7	60.3	67.1	–
CARNet [14]	49.4	40.5	44.5	65.0	59.6	61.2	69.5	74.7	72.0	65.7	79.7	72.0	–
TDCSS [9]	–	–	–	44.2	62.8	51.9	54.4	69.8	60.9	59.2	74.9	66.1	–
TransZero [5]	52.6	33.4	40.8	69.3	68.3	68.8	60.3	81.1	69.2	61.3	82.3	70.2	–
ICCE [26]	–	–	–	67.3	65.5	66.4	67.4	81.2	73.6	65.3	82.3	72.8	–
MAIN w/o IR	40.3	46.9	43.4	57.2	66.4	61.4	78.9	64.6	71.8	77.9	67.1	72.1	30 seconds
MAIN (Ours)	40.0	51.1	44.8	58.7	65.9	62.1	77.9	71.9	74.8	81.8	72.1	76.7	45 seconds

this setting, both the seen and unseen classes increase with t . As in fixed continual GZSL, we conduct experiments in the dynamic continual GZSL on CUB, aPY, AWA1, AWA2, and SUN, as shown in table 2. We observe that our proposed MAIN performs well, seeing absolute gains of 13.41%, 0.34%, and 4.83% on the CUB, AWA1, and SUN datasets while being relatively competitive on aPY and AWA2. In figure 3, we show the harmonic mean per task for the CUB dataset. This is a harder setting compared to the fixed continual GZSL setting since each task brings more seen and unseen classes, making the classification task challenging due to more classes to distinguish from and increasing the opportunity for catastrophic forgetting of previous tasks; thus, accuracy tends to drop with more tasks. Regardless, we again observe that the proposed model consistently outperforms recent baselines.

4.2.3 Generalized Zero-Shot Learning

We conduct experiments on CUB-200, SUN, AWA1, and AWA2 datasets in the GZSL setting, reporting the accuracy over the seen and unseen classes in table 3. We observe significant improvements using MAIN over previous methods on the SUN, AWA1, and AWA2 datasets, with absolute gains of 0.3%, 1.2%, and 3.9%, respectively. Notably, the training² of MAIN relative to the baseline approaches can be 100-300× faster without requiring the unseen class attributes *a priori*. The generative baselines require learning complex generative models (*e.g.*, VAE or GAN) to synthesize realistic samples, which can be expensive. Moreover, generative baselines require an additional assumption about the knowledge of unseen class attributes *a priori* during training. We also evaluate our model on the standard ZSL setting; the results are shown in table 4. Once again, we observe that MAIN significantly outperforms both non-

²Time to train on the whole dataset with an Nvidia GTX 1080Ti, code is available in supplementary.

Table 4. Results on the standard ZSL setting. **Top:** Non-generative methods (like ours). **Bottom:** Expensive generative methods.

	AWA1	AWA2	CUB	SUN
LATEM [58]	55.1	55.8	49.3	55.3
ALE [1]	59.9	62.5	54.9	58.1
RelationNet [49]	68.2	64.2	55.6	–
SAE [25]	53.0	54.1	33.3	40.3
EF-ZSL [52]	57.0	57.4	44.7	63.3
DCN [35]	65.2	–	56.2	61.8
SE-ZSL [51]	69.5	69.2	59.6	63.4
f-CLSWGAN [60]	68.2	–	57.3	60.8
LisGAN [32]	70.6	–	58.8	61.7
ADA [22]	–	70.4	70.9	63.3
f-VAEGAN-D2 [61]	71.1	–	61.0	65.6
LFCAA [36]	–	68.1	67.6	62.0
SP-AEN [4]	–	58.5	55.4	59.2
OCD-CVAE [21]	–	71.3	60.3	63.5
CARNet [14]	75.0	73.7	73.1	63.1
TransZero [5]	71.9	70.1	76.8	65.6
MAIN w/o IR	76.9	76.5	71.4	64.8
MAIN (Ours)	79.0	76.8	73.7	65.8

generative and generative baselines.

5. Ablation Studies

We conduct extensive ablation studies on the proposed model’s different components, observing that each of the proposed components play a critical role. We show the effects of different components on the AWA1 and CUB datasets in the fixed continual GZSL setting, with more ablation studies for dynamic continual GZSL in the supplementary material.

5.1. Effect of Self-Gating on the Attribute

As proposed in section 3.2, we use the self-gating module to learn the attribute embeddings. We perform an ablation to verify the importance of the self-gating module in figure 4. As shown in the figure 4 (left), the performance (harmonic mean) decreases when we remove the self-gating on all the datasets. In particular, we observed that self-gating increased the harmonic mean by an average of 1.1




Figure 4. The proposed model with and without self-gating (SG) in the Fixed continual GZSL setting (Left) and the standard non-continual GZSL setting (Right).

on all the datasets. We observed similar gains in the standard GZSL setting (figure 4 (right)), where the self-gating led to an average improvement of 1.95. This shows that self-gating is an important component of MAIN.

5.2. Importance of Inverse Regularization (IR)

The IR imposing cyclic consistency loss (5) plays a crucial role in the generalization capability of the attribute encoder to novel classes. In tables 3, 1, 2, 4, we have shown the effect of IR. MAIN w/o IR shows the result without IR. We observe a consistent improvement from the cyclic consistency loss. Notably, even without IR, MAIN outperformed recent approaches by a significant margin.

5.3. Effect of Meta-training




Figure 5. The proposed model with meta-learning and without meta-learning in the Fixed Continual GZSL setting (Left) and the standard non-continual GZSL setting (Right).

While we used meta-learning in our proposed MAIN (See section 3.5), the attribute encoder can also be learned directly by minimizing the loss in (7) using a standard optimizer without using meta-learning. We evaluate the model performance with and without meta-learning-based training on the fixed continual GZSL and the standard non-continual GZSL settings. The result is shown in figure 5. We observed that if we withdraw the meta-learning-based training, the MAIN performance drops significantly. On AWA1, AWA2, CUB, and SUN datasets for fixed continual GZSL, the harmonic mean of MAIN drops from 53.3 to 45.2, 55.4 to 54.2, 51.4 to 48.0 and 32.7 to 30.2, respectively. Similarly, in the GZSL setting, MAIN achieves the harmonic mean of 74.8, 76.7, 62.1, and 44.8 for the AWA1, AWA2, CUB, and SUN datasets, respectively, which drops to 67.9, 64.3, 59.8 and 43.4 in the absence of meta-learning.

5.4. Polynomial Kernel vs. Self-gating

We performed the ablation over the two types of SIA modules proposed in section 3.2 for varying depth $L = \{1, 2, 3\}$. The results are shown in figure-6; here, we can observe that the self-gating module with $L = 1$ has the best results. We found that using polynomial-based self-interaction leads to model overfitting due to an exponen-




Figure 6. Comparison of the two types of Attribute self-Interaction modules as proposed in section 3.2; Polynomial Kernels (PK) vs. Self-Gating (SG). We also show the performance as a function of the depth of the self-interaction module with $L = \{1, 2, 3\}$. The only difference between SG and PK is that SG includes non-linear transformations.

tial rise in the degree with L . The only difference between the polynomial kernel and self-gating is, in the self-gating, we use non-linear transformation, and we observed which leads to better performance. Therefore if the polynomial kernel is equipped with the non-linearity, it can learn more complex functions with a smaller layer, and an increase in the layer may lead to overfitting because of the exponential increase in the complexity. In conclusion, for the $L = 1$, the polynomial kernel with the non-linear projection is the best-performing model.

6. Conclusions

We propose the Meta-learned Attribute self-Interaction Network (MAIN), a method capable of operating in both generalized zero-shot learning and continual zero-shot learning settings. Through first-order meta-learning-based training of the network and a novel self-interaction module for attributes, we obtain state-of-the-art results in GZSL settings across multiple common benchmarks. MAIN also incorporates a theoretically motivated regularization scheme for better generalization of the attribute encoder to the unseen class attributes. Notably, our approach does not rely on expensive generative models; this allows for considerably faster speed during training and relaxes the restrictive constraint requiring unseen class attributes to be known at training. Therefore, unlike methods relying on generative models, the inductive setting assumption in the attribute space makes the proposed model highly practical and easy to use.

Given clear connections between zero-shot and continual learning, we extend our approach to settings where class data arrive sequentially. We adopt a data reservoir to mitigate catastrophic forgetting and combine it with meta-learning-based training, leading to a few-shot-based approach that naturally enables the model to efficiently learn from the few samples that can be saved in a buffer. Our results on the CUB, aPY, AWA1, AWA2, and SUN datasets in GZSL and two different protocols of continual GZSL demonstrate that MAIN outperforms a wide array of strong baselines. Ablation studies demonstrate that each component of our proposed approach is critical for its success.

References

- [1] Zeynep Akata, Florent Perronnin, Zaid Harchaoui, and Cordelia Schmid. Label-embedding for attribute-based classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 819–826, 2013.
- [2] Zeynep Akata, Scott Reed, Daniel Walter, Honglak Lee, and Bernt Schiele. Evaluation of output embeddings for fine-grained image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2927–2936, 2015.
- [3] Arslan Chaudhry, Marc’ Aurelio Ranzato, Marcus Rohrbach, Mohamed Elhoseiny, and Dummy. Efficient lifelong learning with a-gem. In *International Conference on Learning Representations*, 2019.
- [4] Long Chen, Hanwang Zhang, Jun Xiao, Wei Liu, and Shih-Fu Chang. Zero-shot visual recognition using semantics-preserving adversarial embedding networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [5] Shiming Chen, Ziming Hong, Yang Liu, Guo-Sen Xie, Baigui Sun, Hao Li, Qinmu Peng, Ke Lu, and Xinge You. Transzero: Attribute-guided transformer for zero-shot learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 330–338, 2022.
- [6] Yu-Ying Chou, Hsuan-Tien Lin, and Tyng-Luh Liu. Adaptive and generative zero-shot learning. In *International Conference on Learning Representations*, 2021.
- [7] Ali Farhadi, Ian Endres, Derek Hoiem, and David Forsyth. Describing objects by their attributes. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1778–1785. IEEE, 2009.
- [8] Rafael Felix, Ian Reid, Gustavo Carneiro, et al. Multi-modal cycle-consistent generalized zero-shot learning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 21–37, 2018.
- [9] Yaogong Feng, Xiaowen Huang, Pengbo Yang, Jian Yu, and Jitao Sang. Non-generative generalized zero-shot learning via task-correlated disentanglement and controllable samples synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9346–9355, 2022.
- [10] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, pages 1126–1135. PMLR, 2017.
- [11] Zhenyong Fu, Tao Xiang, Elyor Kodirov, and Shaogang Gong. Zero-shot object recognition by semantic manifold distance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2635–2644, 2015.
- [12] Chandan Gautam, Sethupathy Parameswaran, Ashish Mishra, and Suresh Sundaram. Generalized continual zero-shot learning. *arXiv preprint arXiv:2011.08508*, 2020.
- [13] Chandan Gautam, Sethupathy Parameswaran, Ashish Mishra, and Suresh Sundaram. Generative replay-based continual zero-shot learning. *arXiv preprint arXiv:2101.08894*, 2021.
- [14] 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.
- [15] 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*, pages 2672–2680, 2014.
- [16] Zongyan Han, Zhenyong Fu, Jian Yang, and Dummy. Learning the redundancy-free features for generalized zero-shot object recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12865–12874, 2020.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [18] Dat Huynh and Ehsan Elhamifar. Fine-grained generalized zero-shot learning via dense attribute-based attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4483–4493, 2020.
- [19] Sung Ju Hwang and Leonid Sigal. A unified semantic embedding: Relating taxonomies and attributes. *Advances in Neural Information Processing Systems*, 2014.
- [20] Skorokhodov Ivan and Elhoseiny Mohamed. Class normalization for zero-shot learning. In *International Conference on Learning Representations*, 2021.
- [21] Rohit Keshari, Richa Singh, Mayank Vatsa, and dummy dummy. Generalized zero-shot learning via over-complete distribution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13300–13308, 2020.
- [22] Varun Khare, Divyat Mahajan, Homanga Bharadhwaj, Vinay Kumar Verma, and Piyush Rai. A generative framework for zero shot learning with adversarial domain adaptation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3101–3110, 2020.
- [23] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *International Conference on Learning Representations*, 2014.
- [24] Elyor Kodirov, Tao Xiang, Zhenyong Fu, and Shaogang Gong. Unsupervised domain adaptation for zero-shot learning. In *The IEEE International Conference on Computer Vision*, pages 2452–2460, 2015.
- [25] Elyor Kodirov, Tao Xiang, and Shaogang Gong. Semantic autoencoder for zero-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3174–3183, 2017.
- [26] Xia Kong, Zuodong Gao, Xiaofan Li, Ming Hong, Jun Liu, Chengjie Wang, Yuan Xie, and Yanyun Qu. En-compactness: Self-distillation embedding & contrastive generation for generalized zero-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9306–9315, 2022.

- [27] Alex Krizhevsky, Ilya Sutskever, Geoffrey Hinton, and E. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 1097–1105, 2012.
- [28] Hari Chandana Kuchibhotla, Sumitra S Malagi, Shivam Chandhok, and Vineeth N Balasubramanian. Unseen classes at a later time? no problem. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9245–9254, 2022.
- [29] Vinay Kumar Verma, Aakansha Mishra, Ashish Mishra, and Piyush Rai. Generative model for zero-shot sketch-based image retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019.
- [30] Christoph H Lampert, Hannes Nickisch, and Stefan Harmeling. Learning to detect unseen object classes by between-class attribute transfer. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 951–958. IEEE, 2009.
- [31] Christoph H Lampert, Hannes Nickisch, Stefan Harmeling, and dummy dummy. Attribute-based classification for zero-shot visual object categorization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(3):453–465, 2014.
- [32] Jingjing Li, Mengmeng Jing, Ke Lu, Zhengming Ding, Lei Zhu, and Zi Huang. Leveraging the invariant side of generative zero-shot learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [33] Kai Li, Martin Renqiang Min, and Yun Fu. Rethinking zero-shot learning: A conditional visual classification perspective. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3583–3592, 2019.
- [34] Lu Liu, Tianyi Zhou, Guodong Long, Jing Jiang, Xuanyi Dong, and Chengqi Zhang. Isometric propagation network for generalized zero-shot learning. In *International Conference on Learning Representations*, 2021.
- [35] Shichen Liu, Mingsheng Long, Jianmin Wang, and Michael I Jordan. Generalized zero-shot learning with deep calibration network. In *Advances in Neural Information Processing Systems*, pages 2009–2019, 2018.
- [36] Yang Liu, Jishun Guo, Deng Cai, and Xiaofei He. Attribute attention for semantic disambiguation in zero-shot learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6698–6707, 2019.
- [37] David Lopez-Paz, Ranzato, Marc’Aurelio, and D Dummy. Gradient episodic memory for continual learning. In *Advances in Neural Information Processing Systems*, pages 6467–6476, 2017.
- [38] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier, 1989.
- [39] Ashish Mishra, M Reddy, Anurag Mittal, and Hema A Murthy. A generative model for zero shot learning using conditional variational autoencoders. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshop*, 2017.
- [40] Sanath Narayan, Akshita Gupta, Fahad Shahbaz Khan, Cees GM Snoek, and Ling Shao. Latent embedding feedback and discriminative features for zero-shot classification. *arXiv preprint arXiv:2003.07833*, 2020.
- [41] Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018.
- [42] Mohammad Norouzi, Tomas Mikolov, Samy Bengio, Yoram Singer, Jonathon Shlens, Andrea Frome, Greg S Corrado, and Jeffrey Dean. Zero-shot learning by convex combination of semantic embeddings. *arXiv preprint arXiv:1312.5650*, 2013.
- [43] Genevieve Patterson and James Hays. Sun attribute database: Discovering, annotating, and recognizing scene attributes. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2751–2758. IEEE, 2012.
- [44] Aniruddh Raghu, Maithra Raghu, Samy Bengio, and Oriol Vinyals. Rapid learning or feature reuse? towards understanding the effectiveness of maml. *ICLR*, 2020.
- [45] Edgar Schonfeld, Sayna Ebrahimi, Samarth Sinha, Trevor Darrell, and Zeynep Akata. Generalized zero- and few-shot learning via aligned variational autoencoders. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [46] Edgar Schonfeld, Sayna Ebrahimi, Samarth Sinha, Trevor Darrell, and Zeynep Akata. Generalized zero-and few-shot learning via aligned variational autoencoders. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8247–8255, 2019.
- [47] Yuming Shen, Jie Qin, Lei Huang, Li Liu, Fan Zhu, and Ling Shao. Invertible zero-shot recognition flows. In *European Conference on Computer Vision*, pages 614–631. Springer, 2020.
- [48] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, pages 4077–4087, 2017.
- [49] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1199–1208, 2018.
- [50] Vinay Verma, Kumar, Dhanajit Brahma, and Piyush Rai. A meta-learning framework for generalized zero-shot learning. *Association for the Advancement of Artificial Intelligence*, 2020.
- [51] Vinay Kumar Verma, Gundeep Arora, Ashish Mishra, and Piyush Rai. Generalized zero-shot learning via synthesized examples. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018.
- [52] Vinay Kumar Verma and Piyush Rai. A simple exponential family framework for zero-shot learning. In *ECML-PKDD*, pages 792–808. Springer, 2017.
- [53] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *Advances in Neural Information Processing Systems*, pages 3630–3638, 2016.

- [54] Jeffrey S Vitter. Random sampling with a reservoir. *ACM Transactions on Mathematical Software (TOMS)*, pages 37–57, 1985.
- [55] Maunil R Vyas, Hemanth Venkateswara, Sethuraman Panchanathan, and dummy dummy. Leveraging seen and unseen semantic relationships for generative zero-shot learning. In *European Conference on Computer Vision*, pages 70–86. Springer, 2020.
- [56] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011.
- [57] Kun Wei, Cheng Deng, and Xu Yang. Lifelong zero-shot learning. *IJCAI*, 2020.
- [58] Yongqin Xian, Zeynep Akata, Gaurav Sharma, Quynh Nguyen, Matthias Hein, and Bernt Schiele. Latent embeddings for zero-shot classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 69–77, 2016.
- [59] 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.
- [60] Yongqin Xian, Tobias Lorenz, Bernt Schiele, and Zeynep Akata. Feature generating networks for zero-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [61] Yongqin Xian, Saurabh Sharma, Bernt Schiele, and Zeynep Akata. f-vaegan-d2: A feature generating framework for any-shot learning. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.
- [62] Hyeonwoo Yu and Beomhee Lee. Zero-shot learning via simultaneous generating and learning. In *Advances in Neural Information Processing Systems*, pages 46–56, 2019.
- [63] Fei Zhang and Guangming Shi. Co-representation network for generalized zero-shot learning. In *International Conference on Machine Learning*, pages 7434–7443, 2019.
- [64] Li Zhang, Tao Xiang, and Shaogang Gong. Learning a deep embedding model for zero-shot learning. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3010–3019. IEEE, 2017.
- [65] Ziming Zhang and Venkatesh Saligrama. Learning joint feature adaptation for zero-shot recognition. *arXiv preprint arXiv:1611.07593*, 2016.