CORRECTION NETWORKS: META-LEARNING FOR ZERO-SHOT LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

We propose a model that learns to perform zero-shot classification using a meta-learner that is trained to anticipate and correct errors based on the learner's training data. The model consists of two modules: a task module that supplies an initial prediction, and a correction module that updates the prediction. The task module is the learner and the correction module is the meta-learner. The correction module takes as input the task module's training data, input, and initial prediction, and updates the initial prediction to be closer to the target value. We demonstrate that this approach leads to state-of-the-art performance on fine-grained zero-shot classification on natural language class descriptions on the CUB and NAB datasets. Correction Networks are independent of the architectures of the modules and can be used for problems for which an updated prediction can be obtained by applying a correction.

1 INTRODUCTION

The ability to solve a task without receiving training examples – zero-shot learning – is desirable. This ability is particularly valuable when training data are limited, difficult to obtain, and/or expensive. We as humans can learn new tasks from descriptions of the tasks, as we learn from reading encyclopedia entries, manuals, handbooks, textbooks, etc. Our artificial agents should be able to do the same. However, generalizing to novel tasks unseen during training is challenging. The agent must integrate its previous training experience to solve the new task, without individual samples from it, while avoiding over-fitting on its previous training experience.

We propose a model that learns a correction on predictions in the zero-shot setting, based on the training data set used to generate the initial prediction. Hence, our model is called Correction Networks. The meta-learner corrects for the zero-shot predictions of the learner.

The intuition for our model is that a zero-shot query sample that is different from samples in the training data will require a different correction than a zero-shot query sample that is similar to samples in the training data. Consider a simple classification model trained to identify different color classes. The training dataset for this simple classification problem consists of the classes of {yellow, green, indigo}. Let there be two zero-shot classes: pink and blue. The deviation of predicted values from true values of pink are different, possibly larger, than deviations of predicted values from true values of blue. These deviations would again be different if the training dataset consisted of the classes of {white, red, purple}.

Correction Networks update the predictions based on the training data. This updated prediction is closer to the target value than the original prediction. Correction Networks consist of two modules: a *task module* that supplies an initial prediction, and a *correction module* that provides a correction to the initial prediction. The task module is the learner and the correction module is the meta-learner. The final prediction is the task module's initial prediction combined with the correction module's correction. This method is illustrated in Figure 1.

The prediction of the meta-learner is used to modify the output of the learner, unlike other meta-learning algorithms that learn to update, initialize, or optimize the learner's weights (Finn et al., 2017) (Ravi & Larochelle, 2017). Our approach is for zero-shot learning while previous meta-learning approaches focus on few shot learn and require a few samples. Another novelty is that the correction module, the meta-learner, takes as input the dataset used to train the task module.

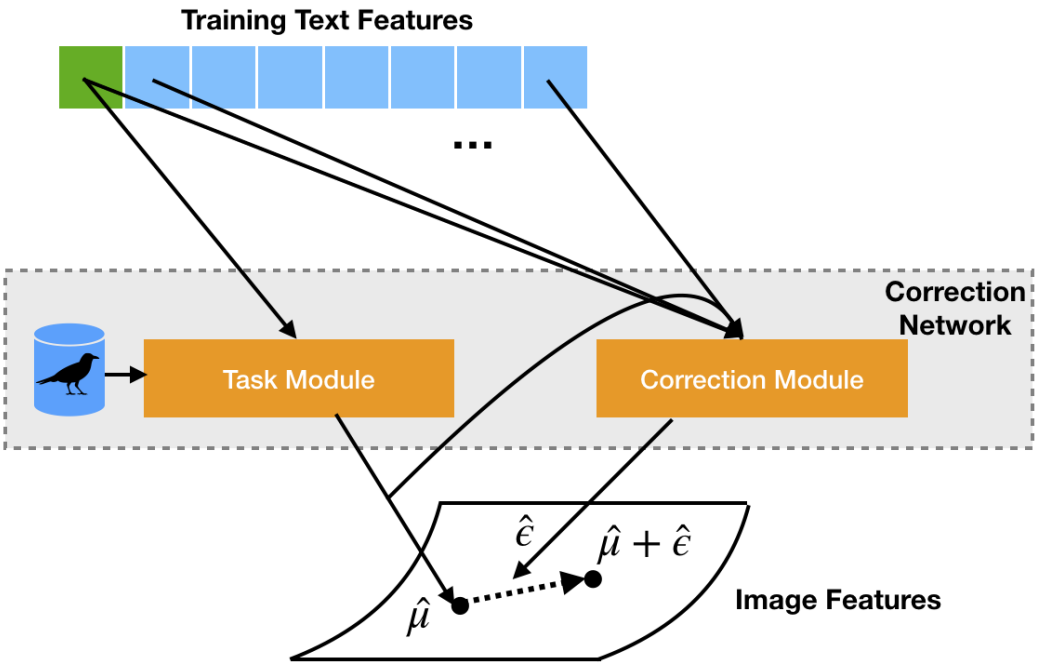

Figure 1: The task module produces an initial prediction. The correction module provides a correction such that when combined with the initial prediction, produces a better prediction.

Correction Networks can be used for problems for which an updated prediction can be obtained by applying a correction. Correction Networks are independent of the representation of the task module. Existing models that provide predictions can be treated as task modules. The correction module only need the inputs to the task module, the predictions of the task module, and the training data of the task module. While our focus is on neural network architectures, both the task module and correction modules can be, but needs not be, neural networks.

Training proceeds by randomly partitioning the training data $\mathcal{D}$ to create disjoint sets $\mathcal{D}_S$ and $\mathcal{D}_U$. In the case of zero-shot classification, the partition is done by class so that the classes in $\mathcal{D}_S$ and $\mathcal{D}_U$ are disjoint. The task module is trained on $\mathcal{D}_S$. After the task module is trained, $\mathcal{D}_U$ is used as the 'zero-shot' queries for the task module. The task module supplies initial predictions on $\mathcal{D}_U$. The correction module is trained to predict the corrections that need to be applied to the task module's initial predictions to obtain the true values associated with $\mathcal{D}_U$. One novelty is that the correction module, the meta-learner, also takes as input the dataset used to train the task module. The prediction of the meta-learner is used to modify the output of the learner, unlike other meta-learning algorithms that learn to update or initialize the learner's weights (Finn et al., 2017) (Ravi & Larochelle, 2017).

We demonstrate Correction Networks on the problem of zero-shot image classification based on natural language class descriptions. Zero-shot image image classification involves classifying images into classes that were unseen in the training data, given a natural language description of the new class.

The main contribution of this paper is a zero-shot learning model that corrects zero-shot predictions based on training data. We demonstrate our model on zero-shot image classification with class text descriptions. Our evaluation shows that Correction Networks compares favorably to state-of-the-art zero-shot learning methods.

The rest of this paper is as follows: Section 2 covers related work, Section 3 describes our Correction Network model, Section 4 summarizes the experiments and results, Section 4.5 outlines future directions, and Section 5 finishes with conclusions.

2 RELATED WORKS

For zero-shot learning for images, the majority of state-of-the-art methods are embedding-based methods. This often involves learning a mapping from the visual space to the semantic space of class labels or vice versa. Alternatively, the embedding function between the visual and semantic spaces is jointly learned through a latent space.

To represent the new, novel classes, hand crafted attributes are often used. The idea is that given a defined attribute ontology, each class name can be converted to an attribute vector against which image features are compared (Lampert et al., 2009) (Palatucci et al., 2009) (Farhadi et al., 2009) Akata et al. (2013). Attribute based embedding spaces more recently gave way to text based embeddings. These text-based embeddings learn from a large external text corpus e.g. Wikipedia or WordNet on a natural language task to learn word embeddings into which class names are then projected (Oquab et al., 2014) (Socher et al., 2013) (Frome et al., 2013) (Fu et al., 2014).

Semantic representations for zero-shot classes have also been created from text documents of the classes eg. Wikipedia articles corresponding to the class. This includes learning attribute representations from a web source by ranking the visualness of attribute candidates (Berg et al., 2010). Internet sources have also been mined for semantic relatedness of attribute-class association (Rohrbach et al., 2013). Directly measuring the the compatibility between documents describing the class and visual features have also been modeled using deep learning (Zhu et al., 2018) (Ba et al., 2015) and non-deep learning models (Elhoseiny et al., 2013) (Romera-Paredes & Torr, 2015) (Fu et al., 2015) (Qiao et al., 2016) (Elhoseiny et al., 2017).

A more recent strategy frames zero-shot recognition as a conventional supervised classification problem by hallucinating samples for unseen classes. Classification performance thus depends on the quality of the hallucinated samples. Guo et al. (2017a) assumes the visual features of each class is distributed as a Gaussian and estimates the distribution of unseen classes by linearly combining the distributions of seen classes. Long et al. (2017) synthesizes visual data from attributes of classes in a one-to-one mapping. Guo et al. (2017b) assigns psuedo labels to samples from seen classes. Zhu et al. (2018) uses generative adversarial training with the generator learning to generate image features from text and query samples are classified based on k-nearest neighbors to labeled or generated samples. In contrast, our approach avoids the data generation phase and instead, directly predicts the center of classes in image feature space. Classification regions are thus defined based on the class of the nearest class center.

Our work is similar to gradient boosting which produces a prediction model using an ensemble of weak models, where each additional model adds an estimator to provide a better model. Unlike gradient boosting, our approach takes in the training data used to create the initial estimate so that the model learns how the initial estimate deviates due to the training data. While the initial learner in gradient boosting is weak, we use a strong learner.

3 CORRECTION NETWORKS

3.1 NOTATION

Let $\mathcal{D}_S$ denote our training data and $\mathcal{D}_U$ our testing data. $\mathcal{D}_S$ is subdivided into disjoint sets $\mathcal{D}_S^s$ and $\mathcal{D}_S^u$. For classification, the classes in $\mathcal{D}_S^s$ are disjoint from the classes in $\mathcal{D}_S^u$.

3.2 MODEL

Correction Networks M consists of two modules: a task module M_T and a correction module M_C. The model components and their inputs and outputs are illustrated in Figure 2.

The task module M_T is so-called because it is task specific and related to the application. The task module is trained on $\mathcal{D}_S^s$. The output of M_T is an estimate $\hat{\mu}$ of a target μ. The predictions of the task module on its training data $\mathcal{D}_S^s$ is $\hat{\mu}_S^s$. Training the task module proceeds by minimizing the distance between $\hat{\mu}_S^s$ and the ground truth μ_S^s. A sparsity regularization term is added to the input layer. The loss is:

$$L_{M_T} = \mathbb{E}[d(M_T(T), \mu)] + \alpha ||w_{M_{T,text}}||^2 \tag{1}$$

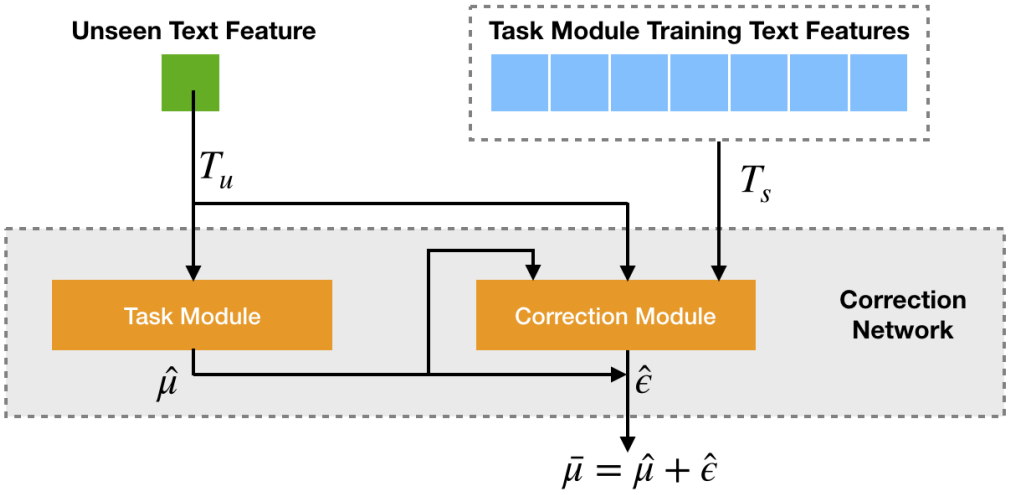

Figure 2: Correction Networks consists of two modules: a task module and a correction module, demonstrated on zero-shot image classification based on a natural language description of the zero-shot class. The task module takes as input the natural language description of the zero-shot class and makes an initial prediction of the cluster center of the class in image feature space. The correction module improves this initial prediction by applying a correction, taking as input the task module's initial prediction, training data, and the zero-shot class description.

where T is the class text description, μ is the empirical mean of samples that belong to the class, and d is a distance function such as the L2 norm.

The task module M_T has not trained on $\mathcal{D}_S^u$ nor $\mathcal{D}_U$. The task module's predictions of $\mathcal{D}_S^u$ are $\hat{\mu}_S^u$. Likewise, the task module's predictions on $\mathcal{D}_U$ are $\hat{\mu}_U$. The task module is trained to minimize. The correction module M_C computes a correction $\hat{\epsilon}$ that is applied to the prediction $\hat{\mu}_U$ $\hat{\mu}_U$ of the task module M_T, where $\hat{\epsilon}$ is calculated based on the data used to train M_T, such that the corrected prediction $(\hat{\mu}_U + \hat{\epsilon})$ is closer than $\hat{\mu}_U$ to the ground truth μ_U. Training the correction module proceeds by minimizing the distance between μ_S^u and $(\hat{\mu}_S^u + \hat{\epsilon}_S^u)$. The training data for the task module D_S^u is inputted into the correction module by representing the training data D_S^u as an unordered collection of data by using a pooling function. The objective function of the correction module is to minimize:

$$L_{M_C} = \mathbb{E}[d(M_C(T), \mu - M_T(T))] \tag{2}$$

We adopt the meta-learning sampling strategy during training as in Snell et al. (2017). Training data for Correction Networks are formed by randomly selecting a subset $\mathcal{D}_S^s$ from the training data $\mathcal{D}_S$. Then, the task module M_T, is trained on $\mathcal{D}_S^s$. The remaining tasks that the task module M_T does not train on are treated as $\mathcal{D}_S^u$ for M_T. The algorithm is detailed in Algorithm 1.

To use Correction Networks for evaluation, the task module M_T is trained on all the available training data $\mathcal{D}_S$. The task module M_T outputs $\hat{\mu}_U$ and the correction network supplies $\hat{\epsilon}_U$. The output of the Correction Networks is $\bar{\mu}_U = \hat{\mu}_U + \hat{\epsilon}_U$. The evaluation algorithm is outlined in Algorithm 2.

Correction Networks can be used for zero-shot problems for which an updated prediction can be obtained by applying a correction. Appropriate problems are those with continuous valued outputs. FOr classification, samples can be represented as continuous values in a feature space and then classified based on the nearest class center. In our experiments, both the base network and the Correction Network are deep neural networks.

3.3 Correction Networks for Zero-Shot Classification

In zero-shot classification, there are seen classes S available during training and unseen classes U during evaluation, where S and U are disjoint.

Algorithm 1 Correction Networks training algorithm. K is the number of classes in the training set S. S is partitioned into S^s and S^u by class. T_S denotes the class descriptions of classes in S. $\mathcal{D}_S$ is the set of individual samples of $(x_i y_i)$ for the classes in the training set S. $\mathcal{D}_S^s$ denotes the subset of $\mathcal{D}_S$ containing all samples from classes in S^s. $\mathcal{D}_k$ is the subset of $\mathcal{D}_S$ containing the individual elements $(x_i y_i)$ such that $y_i = k$. MEAN is a function that returns the mean of it's inputs. OPTIMIZER is an optimizer. d is a distance function, for example, L2 distance. M_T is the task module. M_C is the correction module.

for k to K **do**
  $\mu_k \leftarrow$ MEAN($\mathcal{D}_k$)
end for
while not done **do**
  $S^s \leftarrow$ RANDOMSAMPLE($\{1, ..., K\}, s$)
  $S^u \leftarrow \{1, ..., K\} \backslash S^s$
  while M_T has not finished training **do**
    $\hat{\mu}_S^s \leftarrow M_T(T_S^s)$
    $J_{M_T} \leftarrow d(\hat{\mu}_S^s, \mu_S^s) + \alpha ||w_{M_T}||^2$
    $w_{M_T} \leftarrow$ OPTIMIZER($\nabla_{w_{M_T}} J_{M_T}, w_{M_T}$)
  end while
  $\hat{\mu}_S^u \leftarrow M_T(T_S^u)$
  $\hat{\epsilon}_S^u \leftarrow M_C(T_S^u, \hat{\mu}_S^u, T_S)$
  $\epsilon_S^u \leftarrow \mu_S^u - \hat{\mu}_S^u$
  $J_{M_C} \leftarrow d(\hat{\epsilon}_S^u, \epsilon_S^u)$
  $w_{M_C} \leftarrow$ OPTIMIZER($\nabla_{w_{M_C}} J_{M_C}, w_{M_C}$)
end while

Algorithm 2 Correction Networks evaluation algorithm. S is the set of classes in the training set. U is the set of zero-shot classes. T_U denotes the class descriptions of zero-shot classes in U.

Require Trained $M_T(T_S)$
Output: Prediction $\bar{\mu}_U$

$\hat{\mu}_U \leftarrow M_T(T_U)$
$\hat{\epsilon} \leftarrow M_C(T_U, \hat{\mu}_U, T_S)$
$\bar{\mu}_U = \hat{\mu}_U + \hat{\epsilon}_U$

We are given data $\mathcal{D}_S = \{(x_n, y_n), t_m\})$ containing samples $\{x_n, n = 1, ..., N\} \in X$ and class labels $\{y_n, i = 1, ..., N, y_n \in Y\}$ such that $y_n \in S$, along with class text descriptions $\{t_m, m = 1, ..., M\} \in T_S$ where there exists exactly one t_m for each class in S. For image classification, each sample x_n is an image's features extracted from a pre-trained model. There is no textual data per image. The only text data is t_m with one text description per class. The model does not assume a particular text representation.

At evaluation, we are given new class descriptions for classes in U, denoted by T_U. We seek to learn a function f to minimize the 0-1 loss between the predicted $f(x_i)$ and true class labels y_i for each sample x_i where the samples and class labels can be from the new classes in U.

Correction Networks will map the class text descriptions T_U to their corresponding class centers in image feature space μ_U. Classification of a single image is done by assigning it the class label of the closest class center, measured by L2 distance. We subdivide the training classes S into S^s and S^u. Then, $\mathcal{D}_S^s$ corresponds to the data from $\mathcal{D}_S$ that correspond to the classes in S^s. Similarly, $\mathcal{D}_S^u$ corresponds to the data from $\mathcal{D}_S$ that correspond to the classes in S^u. Training then proceeds as described in the previous section. Algorithm 2 also lays out the training procedure. Correction Networks predict class cluster centers for zero-shot classes. Individual samples are classified to the nearest class cluster center using L2 distance. For classes in the training data, the class cluster center is the empirical mean of samples from the respective class.

In terms of implementation, both the task module and the correction module are feed-forward neural networks consisting of linear layers, activation functions, and dropout. We found the task module performance improves slightly when the output of the task module is feed into a classifier with a single hidden layer that is also trained to classify samples from the task model's training dataset.

4 EXPERIMENTS

4.1 TASK SETUP

We demonstrate Correction Networks on fine-grained zero-shot classification based on natural language text descriptions of the class. Fine-grained image classification is image classification when classes are very similar. This requires the ability to distinguish between minute details and subtle differences between classes. An example of fine-grained classification is distinguishing between different birds that may be very similar.

Features of a class described in a text description may not be visible in all images belonging to that class. This loss of visibility can be due to image cropping, posture, camera angle, age of the bird, and gender of the bird, among other reasons. For example, an image shown from the front results in obfuscation of the back. An image may only be a head shot. Other objects such as branches or leaves may also block the object of interest.

4.2 DATASETS

Image Data We evaluate our method on Caltech UCSD Birds 2011 (CUB) (Wah et al., 2011) and North America Birds (NAB) (Van Horn et al., 2015). The CUB dataset contains 200 classes of birds with a total of 11,788 images. We only use the images and their associated labels. We use a newer version of the NAB dataset that contains 404 classes with a total of 48,562 images. Two splits were proposed for each dataset (Mohamed Elhoseiny & Elgammal, 2017). The splits are named *Super-Category-Shared* (SCS) and *Super-Category-Exclusive* (SCE). In the SCS split, unseen classes and seen classes share the same parent category. In the SCE split, the parent categories are disjoint between seen and unseen classes. The SCE splits are more difficult than the SCS splits because the unseen classes are more different than the seen classes. Where published validation sets are not available, we create our own validation sets by holding out part of the training set.

Visual Features The visual features are activations extracted from the FC layer of a pre-trained parts detector (Mohamed Elhoseiny & Elgammal, 2017) and are the same as used in (Zhu et al., 2018).

Text Features Each class is associated with a Wikipedia article. The articles are tokenized into words, stop words are removed, and words are reduced to their word stems. Then, the processed text is represented by TF-IDF features. We use the same text features as (Zhu et al., 2018).

Table 1: Zero-shot learning classification results accuracy @ 1 on the CUB-200-2011 dataset and the NAB dataset using class descriptions from Wikipedia on the Super-Category-Shared (SCS) and Super-Category-Exclusive (SCE) zero-shot splits

METHOD	**CUB**		**NAB**	
	SCS	SCE	SCS	SCE
MCZSL (Akata et al., 2016)	34.7	-	-	-
WAC-Linear (Elhoseiny et al., 2013)	27.0	5.0	-	-
WAC-Kernel (Elhoseiny et al., 2017)	33.5	7.7	11.4	6.0
ESZSL (Romera-Paredes & Torr, 2015)	28.5	7.4	24.3	6.3
SJE (Akata et al., 2015)	29.9	-	-	-
ZSLNS (Qiao et al., 2016)	29.1	7.3	24.5	6.8
$\text{SynC}_{\text{fast}}$ (Changpinyo et al., 2016)	28.0	8.6	18.4	3.8
SynC_{OVO} (Changpinyo et al., 2016)	12.5	5.9	-	-
ZSLPP (Mohamed Elhoseiny & Elgammal, 2017)	37.2	9.7	30.3	8.1
GAZSL (Zhu et al., 2018)	43.7	10.3	35.6	8.6
Correction Networks	**45.8**	10.0	**37.0**	**9.5**

Examples of the text and images are shown in Figure 4.2.

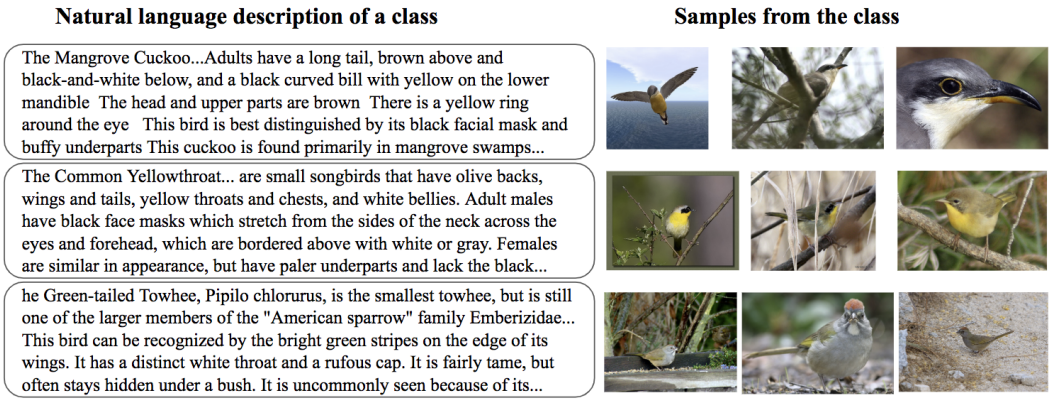

Figure 3: Examples from the dataset where there is one natural language description of a class and the goal is to classify images for new classes without training sample images.

4.3 Conventional Zero-Shot Recognition

The top-1 accuracy of our method and eight state-of-the-art algorithms for the CUB and NAB datasets for both the SCS and SCE splits are tabulated in Table 1. The eight comparison models are MCZSL (Akata et al., 2016), ZSLNS (Qiao et al., 2016), SJE (Akata et al., 2015), WAC (Elhoseiny et al., 2017), SynC (Changpinyo et al., 2016), ZSLPP (Mohamed Elhoseiny & Elgammal, 2017), and GAZSL (Zhu et al., 2018) and the performance numbers are copied from (Zhu et al., 2018). MCZSL directly uses manual part annotations to extract visual representations. On the otherhand our approach, GAZSL, and ZSLPP merely use detected parts for both training and testing. Thus, performance using methods that use detected parts are expected to be poorer than performance on MCZSL that uses manual parts annotations.

Our model increases conventional zero-shot classification accuracy by more than 1% on the SCS-splits.

Table 2: Effects of different components on zero-shot classification accuracy (%) on CUB SCS

METHOD	CUB
Task module only	43.8
Correction module without task module's training dataset	43.4
Correction Networks	45.8

4.4 Ablation Studies

To examine the contributions of separate components of our model, we conduct ablation studies by removing selected components. Then, the model is retrained and evaluated on the test set. The resulting performance of the ablated models are reported in Table 2.

Removing the correction module degrades the performance but the task module itself already achieves the state of the art. The correction module improves the performance by about 1.4%. This suggests that the correction modules makes a slight change to the output of the task module.

When we remove the task module's training data as input into the correction module, the zero-shot accuracy decreases by 1.7%. This demonstrates that the task module's training data is important to the correction module's prediction.

4.5 Extensions

Additional investigations include applying Correction Networks to other tasks with outputs that can be updated, including for example, estimates for regression problems, classification probabilities, or probability distributions for reinforcement learning policies. The type of update can also be multiplicative, weighted sum, or another function, in addition to the summation presented here.

The magnitude of the correction vector predicted by the Correction Network can be interpreted as the size of correction needed. The size of the correction can be compared against the number of training samples given to the base network, defined as the number of seen classes. One would expect that the size of correction decreases as the number of seen classes increases. Also, one would expect that the accuracy stays constant or else improves with the number of seen classes. It would be interesting to see if correction networks can be used to correct for bias in the training data.

5 Conclusion

We propose a zero-shot learning model that consists of a task module and a correction module that is trained to extrapolate from training data to unseen data. The training data is partitioned into a set of data used to train the task module and a disjoint set of data used to train the correction module. This later data is zero-shot with respect to the task module and the correction module is trained to update the task module's zero-shot predictions to create a better prediction.

We demonstrate our model on zero-shot fine-grained classification using only a single natural language description per zero-shot class. Our model performs favorably against the state-of-the-art on zero-shot classification. This framework is flexible to different representations of the task and correction modules and can be extended to other types of model predictions that can be updated to improve predictions.

References

Zeynep Akata, Florent Perronnin, Zaid Harchaoui, and Cordelia Schmid. Label-embedding for attribute-based classification. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pp. 819–826. IEEE, 2013.

Zeynep Akata, Scott Reed, Daniel Walter, Honglak Lee, and Bernt Schiele. Evaluation of output embeddings for fine-grained image classification. In *Computer Vision and Pattern Recognition (CVPR), 2015 IEEE Conference on*, pp. 2927–2936. IEEE, 2015.

Zeynep Akata, Mateusz Malinowski, Mario Fritz, and Bernt Schiele. Multi-cue zero-shot learning with strong supervision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 59–68, 2016.

Lei Jimmy Ba, Kevin Swersky, Sanja Fidler, and Ruslan Salakhutdinov. Predicting deep zero-shot convolutional neural networks using textual descriptions. In *ICCV*, pp. 4247–4255, 2015.

Tamara L Berg, Alexander C Berg, and Jonathan Shih. Automatic attribute discovery and characterization from noisy web data. In *European Conference on Computer Vision*, pp. 663–676. Springer, 2010.

Soravit Changpinyo, Wei-Lun Chao, Boqing Gong, and Fei Sha. Synthesized classifiers for zero-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5327–5336, 2016.

Mohamed Elhoseiny, Babak Saleh, and Ahmed Elgammal. Write a classifier: Zero-shot learning using purely textual descriptions. In *Computer Vision (ICCV), 2013 IEEE International Conference on*, pp. 2584–2591. IEEE, 2013.

Mohamed Elhoseiny, Ahmed Elgammal, and Babak Saleh. Write a classifier: Predicting visual classifiers from unstructured text. *IEEE transactions on pattern analysis and machine intelligence*, 39(12):2539–2553, 2017.

Ali Farhadi, Ian Endres, Derek Hoiem, and David Forsyth. Describing objects by their attributes. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pp. 1778–1785. IEEE, 2009.

Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, pp. 1126–1135, 2017. URL `http://proceedings.mlr.press/v70/finn17a.html`.

Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Tomas Mikolov, et al. Devise: A deep visual-semantic embedding model. In *Advances in neural information processing systems*, pp. 2121–2129, 2013.

Yanwei Fu, Timothy M Hospedales, Tao Xiang, Zhenyong Fu, and Shaogang Gong. Transductive multi-view embedding for zero-shot recognition and annotation. In *European Conference on Computer Vision*, pp. 584–599. Springer, 2014.

Zhenyong Fu, Tao Xiang, Elyor Kodirov, and Shaogang Gong. Zero-shot object recognition by semantic manifold distance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2635–2644, 2015.

Yuchen Guo, Guiguang Ding, Jungong Han, and Yue Gao. Synthesizing samples fro zero-shot learning. IJCAI, 2017a.

Yuchen Guo, Guiguang Ding, Jungong Han, and Yue Gao. Zero-shot learning with transferred samples. *IEEE Transactions on Image Processing*, 26(7):3277–3290, 2017b.

Christoph H Lampert, Hannes Nickisch, and Stefan Harmeling. Learning to detect unseen object classes by between-class attribute transfer. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pp. 951–958. IEEE, 2009.

Yang Long, Li Liu, Ling Shao, Fumin Shen, Guiguang Ding, and Jungong Han. From zero-shot learning to conventional supervised classification: Unseen visual data synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

Han Zhang Mohamed Elhoseiny, Yizhe Zhu and Ahmed Elgammal. Link the head to the" beak": Zero shot learning from noisy text description at part precision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

Maxime Oquab, Leon Bottou, Ivan Laptev, and Josef Sivic. Learning and transferring mid-level image representations using convolutional neural networks. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pp. 1717–1724. IEEE, 2014.

Mark Palatucci, Dean Pomerleau, Geoffrey E Hinton, and Tom M Mitchell. Zero-shot learning with semantic output codes. In *Advances in neural information processing systems*, pp. 1410–1418, 2009.

Ruizhi Qiao, Lingqiao Liu, Chunhua Shen, and Anton van den Hengel. Less is more: zero-shot learning from online textual documents with noise suppression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2249–2257, 2016.

Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *International Conference on Learning Representations*, 2017.

Marcus Rohrbach, Sandra Ebert, and Bernt Schiele. Transfer learning in a transductive setting. In *Advances in neural information processing systems*, pp. 46–54, 2013.

Bernardino Romera-Paredes and Philip Torr. An embarrassingly simple approach to zero-shot learning. In *International Conference on Machine Learning*, pp. 2152–2161, 2015.

Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, pp. 4077–4087, 2017.

Richard Socher, Milind Ganjoo, Christopher D Manning, and Andrew Ng. Zero-shot learning through cross-modal transfer. In *Advances in neural information processing systems*, pp. 935–943, 2013.

Grant Van Horn, Steve Branson, Ryan Farrell, Scott Haber, Jessie Barry, Panos Ipeirotis, Pietro Perona, and Serge Belongie. Building a bird recognition app and large scale dataset with citizen scientists: The fine print in fine-grained dataset collection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 595–604, 2015.

C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011.

Yizhe Zhu, Mohamed Elhoseiny, Bingchen Liu, Xi Peng, and Ahmed Elgammal. A generative adversarial approach for zero-shot learning from noisy texts. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.