

Improving Few-Shot Image Classification Using Machine- and User-Generated Natural Language Descriptions

Kosuke Nishida, Kyosuke Nishida, Shuichi Nishioka
NTT Human Informatics Laboratories, NTT Corporation
kosuke.nishida.ap@hco.ntt.co.jp

Abstract

Humans can obtain the knowledge of novel visual concepts from language descriptions, and we thus use the few-shot image classification task to investigate whether a machine learning model can have this capability. Our proposed model, LIDE (Learning from Image and DEscription), has a text decoder to generate the descriptions and a text encoder to obtain the text representations of machine- or user-generated descriptions. We confirmed that LIDE with machine-generated descriptions outperformed baseline models. Moreover, the performance was improved further with high-quality user-generated descriptions. The generated descriptions can be viewed as the explanations of the model’s predictions, and we observed that such explanations were consistent with prediction results. We also investigated why the language description improved the few-shot image classification performance by comparing the image representations and the text representations in the feature spaces.

1 Introduction

Humans can efficiently learn about new concepts from language (Chopra et al., 2019). Hence, in this paper, we focus on the few-shot image classification problem to verify machine learning models’ capability to understand new concepts from language. This problem is a kind of meta-learning problem in which a model first learns from the concepts of classes by training on a few instances and then learns unseen classes in the same way.

In our problem setting, the model can use language descriptions of images as additional information. This setting is similar to teaching a new concept to others by explaining it deductively from small amounts of data, unlike most machine learning models that learn inductively from large amounts of data.

For this setting, we propose a new model, called *LIDE* (*Learning from Image and DEscription*). As




Figure 1: Concept of the LIDE model.

shown in Figure 1, LIDE consists of an image encoder, an image classifier, a text decoder, and a text encoder.

LIDE has the advantage of providing explainability. It passes an image representation encoded by the image encoder to the text decoder, which then generates a language description of the image as an explanation of the model’s prediction. The image classifier then classifies the image in accordance with the text representation, which is encoded by the text encoder, in addition to the image representation.

LIDE also provides high accuracy due to its explainability. It has been difficult to use machine-generated descriptions to improve image classification performance because of their low quality (Mu et al., 2020). Therefore, we design a training algorithm and a text encoding method to obtain robust text representations. In addition, LIDE includes a feature fusion module added to the image classifier to combine the information from both the image representation and the text representation.

Moreover, LIDE can take user-generated descriptions as input instead of machine-generated descriptions. We can use a description that contains textual features captured from an input image by human perception or a post-edited text in the text decoder output. The resulting high-quality descriptions provided by users can improve the image classification accuracy.

Our contributions are summarized as follows:

- We confirmed that LIDE with machine-generated descriptions outperformed previous models, and thus the explanations of the

model’s predictions were helpful to improve the classification accuracy.

- We observed that the performance improved further when gold captions were fed to LIDE as users’ high-quality descriptions.
- We investigated whether the generated explanations were consistent with the image classification predictions, and we found a positive correlation between the quality of the generated captions and the classification accuracy.
- We thoroughly investigated why the text representations explaining the input image contributed to the image classification task, specifically in terms of the distributions of the representations in the feature space, the robustness of the representations for noisy images, and the knowledge of concepts that can be extracted from the representations.

2 Background

2.1 Few-Shot Image Classification

N -way K -shot classification involves three data splits, \mathcal{T}_{train} , \mathcal{T}_{dev} , and \mathcal{T}_{test} , and each split consists of many classes and instances. The classes in the splits are disjoint. Each N -way K -shot classification task is a classification problem with N classes. Each task provides K training instances for each class, called *support* instances. A task entails the evaluation of M instances for each class, called *query* instances. The tasks, which consist of the classes and instances, are randomly sampled from the data splits.

This problem is a meta-learning problem. In the training phase, we use the episodic training (Ravi and Larochelle, 2017), where many mini-batches of size B consisting of B tasks, each of which consists of N classes and $N(K + M)$ instances, are independently sampled from \mathcal{T}_{train} . In the test phase, the model learns new N -way K -shot classification tasks with unseen classes and instances sampled from \mathcal{T}_{dev} or \mathcal{T}_{test} . For each sampled task, a model learns from the supports, and we evaluate the classification performance of the queries.

A major approach for N -way K -shot classification is the prototypical network (ProtoNet) (Snell et al., 2017). In both the training and test phase, instead of updating the model parameters for each sampled task with few support instances, the prototypical network computes the class prototypes. Let

h_k^c be the k -th support feature of class c . Here, the class prototype z^c is

$$z^c = \frac{1}{K} \sum_k W_{proto} h_k^c.$$

In the training phase, the model is trained with the cross-entropy loss of the queries from \mathcal{T}_{train} :

$$L_{class} = -\frac{1}{M} \sum_i \sum_c y_i^c \log \frac{\exp[s(z^c, h_i)]}{\sum_{c'} \exp[s(z^{c'}, h_i)]},$$

where h_i is the i -th query feature, $s(z_c, h_{mm,i}) = z_c^\top W_{proto} h_{mm,i}$ is a score function containing W_{proto} as a trainable parameter, and $y_i^c \in \{0, 1\}$ is the ground-truth label of the i -th query. In the test phase, the class prototypes are obtained from the supports, and the average score over the sampled tasks is reported.

2.2 Few-Shot Image Classification with Language Description

We focus on the few-shot image classification to verify the machine learning models’ capability to learn new concepts from language descriptions in addition to the images. Previous studies have used the language description to improve the few-shot image classification. The classifiers in the studies are based on ProtoNet, which corresponds to the model with the image encoder and prediction module in Figure 2. Mu et al. (2020) proposed LSL by introducing a text decoder to ProtoNet to avoid overfitting by training the image encoder with a language generation loss. They observed that a text representation from a noisy machine-generated description was harmful for image classification. Accordingly, they viewed the text decoder as a regularizer and did not use a text encoder. RS-FSL (Afham et al., 2021) replaced the GRU (Cho et al., 2014) text decoder of LSL with a bi-directional transformer (Vaswani et al., 2017).

Andreas et al. (2018) were interested in describing the hidden states with natural language, but not in improving the image classification performance. They proposed L3 by adding both a text decoder and encoder to ProtoNet. They encoded an image into an image representation and decoded it into an explanation. The input of their image classifier was only the text representation encoded by the text encoder. L3 provided the explainability, but their model performance decreased.

All of the aforementioned papers assumed that the image encoder was not pre-trained in a super-




Figure 2: Overall structure of LIDE. The text decoder generates a caption on the basis of the image representation encoded by the image encoder. The text encoder obtains the text representation from the generated caption or a gold caption. The feature fusion module combines the image and text representation to generate the multi-modal representations. The prediction module outputs the classification probability.

vised or self-supervised fashion with external images. Our motivation is to clarify the benefit of the language description to learn novel visual concepts, and so we follow their setting.

3 Methods

3.1 Model

We show the overall model structure of LIDE in Figure 2. The model components are explained in the following.

Image Encoder We can use any network as the image encoder, and we used a 4-layer CNN as in the previous studies. The output is the image feature h_{img} .

Text Decoder First, we map the image feature to the text feature space as follows:

$$f_{I2T}(h_{img}) = \text{Linear}(\text{LayerNorm}(h_{img})),$$

where LayerNorm is layer normalization (Ba et al., 2016). Then, we pass the text feature vector to the text decoder as an encoder hidden state sequence of length 1. The text decoder autoregressively generates the j -th token t_j . The j -th token generation probability p_j is written as

$$p_j = \text{Pr}(t_j; f_{I2T}(h_{img}), t_{0:j-1}).$$

We used a uni-directional three-layer transformer.

Text Encoder We use BERT (Devlin et al., 2019) for the text encoder, which outputs the last hidden states H_{BERT} . The text feature h_{text} is a weighted average pooling of H_{BERT} :

$$h_{text} = \frac{1}{\sum p_j w_j} \sum p_j w_j h_{BERT,j},$$

where the weight p_j is the token generation probability in the text decoder. If a caption is user-generated, then p_j is 1 for all tokens. The weight

w_j is 1 if the j -th token is not a stop word; otherwise, $w_j = 0$.

The use of weighted average pooling with the text generation probability has two advantages. First, the text encoder can ignore the low-confidence tokens. Second, the image classification loss back-propagates to the text decoder through the weight p_j . Because the discrete operation of text generation breaks the computation graph, we cannot back-propagate the gradient of H_{BERT} to the decoder without the method.

Feature Fusion Module The feature fusion module combines the single-modal features h_{img} and h_{text} into a multi-modal feature h_{mm} . Let $[\cdot; \cdot]$ be the vector concatenations, f_{T2I} be a mapping function from the text feature space to the image feature space, and g be a linear function to \mathbb{R}^2 . f_{T2I} is a three-layer FFNN with ReLU activation. The feature fusion operation is the weighted sum of the two features:

$$[w_{img}; w_{lang}] = \text{softmax}(g([h_{img}; h_{text}])),$$

$$h_{mm} = w_{img}h_{img} + w_{text}f_{T2I}(h_{text}).$$

Prediction Module We use ProtoNet for the prediction module and replace h with h_{mm} .

3.2 Algorithms

Loss Function For image classification, we compute two image classification losses: $L_{class,gold}$ is computed from the image feature and the gold caption, while $L_{class,gen}$ is computed from the image feature and the generated caption. For text generation, we compute L_{text} with teacher-forcing and cross-entropy loss.

To enrich the mapped text feature $f_{T2I}(h_{text})$, we use the contrastive loss (Sohn, 2016; Oord et al., 2018) between the gold and generated captions.

Let v_{gold}^c and v_{gen}^c be the averages of the mapped support text features in class c from the gold and generated captions, respectively. The similarity is $\cos(v_{gold}^c \top v_{gen}^c)$. Then, the contrastive loss L_{cntr} is

$$L_{cntr} = -\frac{1}{2N} \sum_c \log \frac{\exp[\cos(v_{gold}^c \top v_{gen}^c)/\tau]}{\sum_{c'} \exp[\cos(v_{gold}^c \top v_{gen}^{c'})/\tau]} - \frac{1}{2N} \sum_{c'} \log \frac{\exp[\cos(v_{gold}^{c'} \top v_{gen}^{c'})/\tau]}{\sum_c \exp[\cos(v_{gold}^c \top v_{gen}^{c'})/\tau]},$$

where τ is a temperature parameter.

The total loss that we minimize in the training phase to optimize the whole model parameters in an end-to-end manner is

$$L = L_{class,gold} + L_{class,gen} + \lambda_{text} L_{text} + \lambda_{cntr} L_{cntr},$$

where λ_{text} and λ_{cntr} are hyperparameters.

Pre-Training Following Afham et al. (2021); Wang et al. (2019), we pre-train the model with the training data for the downstream task. The pre-training consists of the standard image classification task, and we replace the prediction module with a linear classifier for all training classes. The loss function is $L = L_{class,gen} + \lambda_{text} L_{text}$.

Caption Generation In the training phase, we use a greedy algorithm and random sampling for computational reasons. In each step, we uniformly and randomly choose between the two algorithms. In random sampling, we restrict the candidate tokens to the top 20 tokens at each position.

In the test phase, we input the generated captions to the text encoder in the setting where the user-generated description is not available. The generation algorithm is beam search with a beam width of five and a length penalty of 0.5. Thus, the token sequence $t_{1:l}$ is generated as

$$\operatorname{argmax}_{l, t_{1:l}} \frac{1}{l^{0.5}} \sum_{j=1}^l \log \Pr(t_j; f_{I2T}(h_{img}), t_{1:j-1}),$$

where l is the text length. The length penalty reduces the preference for words consisting of multiple subwords, such as ‘point _y’ (‘pointy’).

4 Evaluation

4.1 Dataset

We used the Caltech-UCSD Birds (CUB) dataset (Wah et al., 2011) for evaluation. It contains 200

bird species (classes) and 40-60 images for each class. The classes are split into 100 training classes, 50 development classes, and 50 test classes. We used this dataset for the $N = 5$ -way $K = 1$ -shot classification problem. The number of query instances M per class was 15.

The CUB dataset has 10 captions for each image (Reed et al., 2016). For each step, we randomly sampled one caption from the 10 gold captions.

4.2 Metrics

For the image classification, we report the average accuracy over 600 tasks, following the previous studies. To evaluate the generated caption quality, we used major metrics for image captioning, BLEU₄ (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005), and ROUGE_L (Lin, 2004).

4.3 Implementation

We pre-processed the images in the same way as Mu et al. (2020). The dimension of the image feature h_{img} was 1600. The text encoder and tokenizer were the pre-trained BERT-base-uncased model. The text encoder output dimension was 768. The configuration of the transformer layers in the text decoder was the same as that of the T5-base decoder model (Raffel et al., 2020), but we did not use the pre-trained parameters for the text decoder because they did not contribute to the performance. The parameter size of W_{proto} was 1600×1600 . The other hyperparameters and optimization details are given in Appendix A.

4.4 Compared Models

ProroNet (Snell et al., 2017) was the baseline, with only the image encoder. As for the other compared models, L3 (Andreas et al., 2018) used a 200-dimensional GRU text encoder and decoder but did not use image representation for classification. LSL (Mu et al., 2020) used a 200-dimensional GRU text decoder for regularization. RS-FSL (Afham et al., 2021) used a 2-layer, 768-dimensional, bi-directional transformer as the text decoder. All models used a 4-layer CNN as the image encoder, along with the ProtoNet-based prediction module.

4.5 Ablated Models

To evaluate the models that use a single-modal representation for image classification, we introduced the models “No Text,” “No Image,” and “No Text Encoder” by removing the feature fusion module

Model	Text Dec.	Text Enc.	Modal	Baseline
No Text			Image	ProtoNet
No Image	✓	✓	Text	L3
No Text Enc.	✓		Image	LSL, RS-FSL

Table 1: Setting of ablated single-modal models.

Model	Accuracy	Img. Enc.	Text Enc.	Fusion	Text Dec.
ProtoNet	57.97 \pm 0.96	✓			
L3	53.96 \pm 1.06		✓		✓
LSL	61.24 \pm 0.96	✓			✓
RS-FSL	65.66 \pm 0.90	✓			✓
LIDE	67.53 \pm 0.91	✓	✓	✓	✓

Table 2: Performance of the compared models.

from LIDE. Each model corresponds to our implementation of the baseline models, as shown in Table 1.

4.6 Evaluation Results for LIDE as Few-Shot Image Classification Model

4.6.1 Main Results

Performance with Machine-generated Description Table 2 summarizes the results. LIDE outperformed ProtoNet, LSL, and RS-FSL, which use an image classifier using image representations only. LIDE improved the image classification performance with the mechanism of encoding text and combining multi-modal representations, while L3, which uses an image classifier with text representations only, harmed their performance.

Performance with User’s Description The advantage of the text encoder is that it enables textual input by users. Specifically, a user can use language as an explanation from humans to machines by feeding the textual features captured from the input image to the model. In addition, if a user objects to the model’s explanations, the user can edit them to correct the model’s misunderstanding.

To evaluate LIDE in this setting, we viewed the gold captions as user descriptions. As the CUB dataset has 10 gold captions per image, we selected one gold caption and fed it to the model to maximize the similarity to a generated caption and thus simulate a user editing the generated caption. The similarity was defined as the bi-gram precision of the gold caption with respect to the generated caption.

Table 3 shows that the performance was improved significantly when a high-quality gold caption was given. However, the performance declined

Random Description	58.89 \pm 0.93
Generated Description	67.53 \pm 0.91
Gold Description	73.08 \pm 0.88

Table 3: Performance with each kind of description.

when a wrong caption was given, which was a randomly sampled caption from all captions in \mathcal{T}_{test} . We conclude that the model’s output depends on the quality of the description.

4.6.2 Ablation Study

Evaluation on modalities The first set of rows in Table 4 lists ablation study results for the different modalities that LIDE uses. We confirmed that LIDE using all the modalities outperformed the compared models using part of the modalities. The image and text representations complemented each other, as will be discussed later.

Evaluation on introduced techniques The second set of rows in Table 4 lists the techniques that were introduced in LIDE: image classification loss with generated captions, contrastive learning, weighted average pooling, and random sampling during caption generation. We confirmed that all of these techniques contributed to the performance of LIDE.

We assume that the loss with the generated captions decreases the discrepancy between the train and test phases. The contrastive loss enriches the text representation and makes the representations of the generated text and the gold caption close. The weighted average pooling reduces the effect of noisy generated text. The random sampling in the training phase contributes to the generation of diverse captions.

Evaluation on pre-training of text encoder We also found that the multi-modal feature was useful even when the text encoder was trained from scratch. The text representation could assist in the image representation without BERT pre-training, because the captions in the CUB dataset are restricted to the descriptions of birds, and the training data thus covers the space of the captions well.

4.7 Evaluation Results for LIDE as Interpretable Machine Learning Model

Evaluation on quality of generated captions

First, generated captions are insufficient as explanations if they are not accurate. The upper bound was the CNN-LSTM model pre-trained with the

	LIDE	67.53 \pm 0.91
No Text (ProtoNet)		62.22 \pm 0.92
No Image (L3)		49.60 \pm 1.03
No Text Encoder (LSL, RS-FSL)		63.10 \pm 0.90
No $L_{class,gen}$		61.96 \pm 0.90
No Contrastive Loss		64.60 \pm 0.88
No Weighted Average Pooling		66.16 \pm 0.93
No Random Sampling		66.40 \pm 0.93
None of the Above		59.68 \pm 0.98
No BERT Pre-Training		66.42 \pm 0.94

Table 4: Ablation study results.

	BLEU ₄	METEOR	ROUGE _L
UB: Caption	59.0	36.1	69.7
No Text Enc.: Caption	50.0	34.6	67.2
Correlation	0.114	0.201	0.217
LIDE: Caption	48.1	34.1	66.4
Correlation	0.309 [†]	0.468 [*]	0.436 ^{**}

Table 5: Captioning scores and correlations to the prediction scores. [†]: $p < 0.1$, ^{*}: $p < 0.05$, ^{**}: $p < 0.01$

MSCOCO (Lin et al., 2014) dataset from Chen et al. (2017). Although we had no training data without the CUB dataset, Table 5 shows that the differences between LIDE and the upper bound were only 2.0 points for METEOR and 3.3 points for ROUGE.

Evaluation on consistency between generated captions and image classification When the generated captions are correct (respectively, incorrect), the classification results should also be correct (incorrect) in terms of the consistency between explanations and classification results.

Accordingly, we calculated the Spearman rank-order correlation coefficient between the captioning scores and the classification scores. We divided 2,953 test-split images into 30 bins in accordance with the ascending order of each captioning score, and we computed the average image classification accuracy in each bin.

Table 5 lists the results. We confirmed a positive correlation between the quality of the generated captions and the prediction accuracy of LIDE. However, LIDE without the text encoder, which corresponds to LSL (Mu et al., 2020) and RS-FSL (Afham et al., 2021), showed a low correlation. This result demonstrated the importance of the text encoder and the feature fusion modules introduced in LIDE.

Qualitative analysis on generated captions Figure 3 shows examples of the generated captions. We found that the captions captured the birds’ char-



	Ours	this bird is yellow with black and has a very short beak.
	Gold	the bird has a small black bill and a yellow crown
	Ours	this bird has wings that are brown and has a white belly.
	Gold	a small bird with a red face black crown white cheeks brown back and black and yellow wings

Figure 3: Examples of generated captions and gold captions.

acteristics. However, the structures of the captions were uniform, and they could not describe a birds’ most distinctive element, such as the red face in the second example. We believe that overfitting to the 5-class classification problems with the small dataset caused this problem.

We restricted the image encoder to a 4-layer CNN for fair comparison to the existing models, and we did not use external training data to validate the ability to learn novel classes from language descriptions. Removal of the limitations would improve the performance of LIDE as an explainable machine learning model.

4.8 Discussion

In this section, we clarify four reasons why text representations are useful in the few-shot image classification task. Specifically, we compared the multi-modal feature space of LIDE with the gold captions to the image feature space of the “No Text” and “No Text Encoder” models.

How are classes distributed in each modal feature space? First, we calculated the inner- and inter-class distances in each feature space. Table 6 lists the results. The inner-class distances in the multi-modal feature space of LIDE were smaller than those in the image feature space; the inter-class distances of LIDE were larger. As a result, the clusters were distributed well in the multi-modal feature space.

We believe that this is because the captions describe the similarities and differences between the images more obviously than the images themselves do. For example, to determine that two birds belong to different species, one piece of evidence is the belly color. The captions can explain this information clearly, e.g., “yellow belly” and “white belly”. From the image, however, the extraction of this information requires multiple steps such as locating the belly and specifying its color.

	Inner-Class Dist.	Inter-Class Dist.	LID
No Text (Image)	0.504	0.592	17.8
No Text Enc. (Image)	0.526	0.609	19.0
LIDE (Fusion)	0.459	0.709	6.73

Table 6: Distribution of feature representations.

What are characteristics of latent feature spaces? Second, we examined the dimensions of the feature spaces. [Lee and Chung \(2021\)](#) observed that the features embedded in a manifold with a smaller latent dimension are more generalized. They evaluated the latent dimension by using the average of the local intrinsic dimensions (LID) of the features. The LID measures the number of dimensions of a feature manifold in the neighbor of x , and it can be estimated as

$$\hat{\text{LID}}(x) = - \left\{ \frac{1}{n_{nn}} \sum_i^{n_{nn}} \log \left(\frac{r_i(x)}{r_{n_{nn}}(x)} \right) \right\}^{-1}$$

by maximum likelihood estimation, where n_{nn} is the number of the nearest neighbors and r_i is the Euclidean distance from x to the i -th nearest neighbor ([Levina and Bickel, 2005](#); [Amsaleg et al., 2015](#)). We set $n_{nn} = 20$ in accordance with the previous studies.

Table 6 lists the estimated LIDs of the features of each model. In addition, we applied principal component analysis (PCA) to the features, and Figure 4 shows the cumulative contribution rates. All features were embedded in \mathbb{R}^{1600} .

The multi-modal features existed in a manifold with a smaller latent dimension than those of the image-only representations. Therefore, the text representation contributed to shrinking the representation manifold to a smaller dimension. We assume that the text controlled the main focus among the many objects in an image. For example, the captions in the CUB dataset describe the characteristics of birds. As a result, the model can extract the important information from captions for the downstream image classification task. In contrast, an image has much information, such as the background, and image features thus require a larger-dimension manifold.

Are multi-modal representations robust for noisy images? Next, we hypothesize that language descriptions can help to classify noisy or obscure images. To verify our hypothesis, we performed experiments in two heuristic noisy settings: One consisted of grayscale images and the

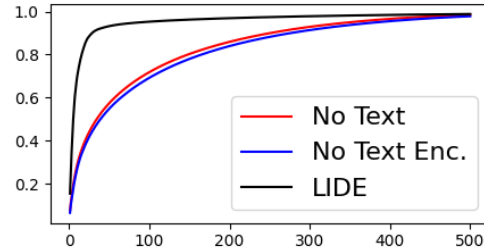


Figure 4: Cumulative contribution rates for PCA of each representations.

	Original	Grayscale	Adversarial
No Text	62.22 \pm 0.92	38.04 \pm 0.78	33.09 \pm 0.69
No Text Enc.	63.10 \pm 0.90	38.55 \pm 0.80	31.30 \pm 0.65
LIDE	78.09 \pm 0.79	59.81 \pm 0.97	56.31 \pm 1.01

Table 7: Performance in noisy settings.

other consisted of images that were adversarially attacked via the fast gradient sign method ([Goodfellow et al., 2015](#)). We compared LIDE with gold captions to the ablated models using the image representations and evaluated their performance with the noisy images in the test phase.

Table 7 lists the result. When a caption was not provided, the classification accuracy dropped greatly in both settings. However, LIDE reduced the decline by virtue of the textual information. These results indicate that the text representation may be useful for classifying certain ill-conditioned images, such as an image in which the bird is extremely small.

What information do text representations have?

Finally, we performed a probing test for each modality of representation, as in previous natural language processing (NLP) studies that discovered linguistic properties ([Conneau et al., 2018](#); [Hupkes and Zuidema, 2018](#)) in the text representations. The CUB dataset has annotations of the birds’ attributions, and we recover the attribution labels from h_{img} and h_{text} . Each image has $\{0, 1\}$ labels for 312 attributions.

We used our trained model to obtain h_{img} and h_{text} for all images and captions. Then, we used linear classifiers $W_{img,attr} \in \mathbb{R}^{1600 \times 312}$ and $W_{lang,attr} \in \mathbb{R}^{768 \times 312}$, which were trained the linear classifiers with binary-cross-entropy loss in the training split. Next, we determined the thresholds in the development split and obtained prediction results in the test split. Finally, we performed a Wilcoxon signed-rank test between the image re-

Image	Text	No Significance
43	17	8

Table 8: Numbers of significant attributions.

has_bill_shape::all-purpose
has_wing_color::white
has_back_color::black
has_breast_color::white
has_throat_color::white
has_eye_color::brown
has_eye_color::white
has_nape_color::black
has_nape_color::white
has_nape_color::red
has_belly_color::white
has_size::small_(5_-_9_in)
has_back_pattern::spotted
has_tail_pattern::spotted
has_belly_pattern::spotted
has_crown_color::grey
has_crown_color::black

Table 9: Attributions that were significantly recovered from text representations.

sults and the text results.

Table 8 lists the numbers of attributions having significance at a p -value of 0.05. Among the 68 attributions, 60 were recovered more easily from one modality than from the other modality. In other words, the image and text representations complemented each other. Most of the attributes favored the image representation, but 17 of them favored the text representation. Table 9 lists those 17 attributes, and we can observe two main characteristics among them. First, the colors black and white were recovered from the text representation. These attributes may be difficult to recover from an image because of light and shadow. The second characteristic was a spotted pattern, which is obscure in an image with 84×84 pixels.

5 Related Work

Image classification with language Several studies have provided gold language information in the test phase. For zero-shot learning or few-shot learning, class-label words are used as additional information (Frome et al., 2013; Socher et al., 2013; Xing et al., 2019). Moreover, He and Peng (2017); Liang et al. (2020) used language descriptions for the standard image classification problem, in which the classes in the training and test phases are the same. The few-shot image classification is a more challenging problem that requires the capability

of learning novel concepts from language descriptions.

Textual explanation of image representation

Explainability of artificial intelligence (XAI) has attracted much attention (Bastings et al., 2021). Papers have proposed methods to generate an image description for XAI in the image classification task and visual question answering (VQA) task (Hendricks et al., 2016, 2018; Li et al., 2018). In contrast to those studies, the motivation of this paper is to decode and encode such descriptions to improve the few-shot image classification performance. We could integrate the findings of XAI studies into LIDE.

Analysis of image and text representations Previous papers have investigated why language information is useful in vision and language tasks. Collell Talleda and Moens (2016) also found that the image and text representations complemented each other; for example, taxonomic attributes are captured well in the language. Li et al. (2020) observed that the attention heads of the multi-modal pre-trained models ground elements of language to image regions.

6 Conclusion

We tackled the few-shot image classification task through learning of novel concepts from language descriptions of images. We observed that machine- and user-generated descriptions improved the few-shot image classification performance. We also found that the generated captions explained the input image and were consistent with the prediction performance.

Our experiments also revealed four reasons why the text representation improved the performance: the inner-class distances of the multi-modal representations are smaller and the inter-class distances are larger than those of image representations; multi-modal representations are embedded in a space with a smaller latent dimension; multi-modal representations are robust for noisy images; and certain types of knowledge are easily recovered from text representations.

Humans can learn concepts from language and explain them with language, but this is still difficult for machine learning models. This study sheds light on the importance of interactivity in explaining with language in machine learning.

References

- Mohamed Afham, Salman Khan, Muhammad Haris Khan, Muzammal Naseer, and Fahad Shahbaz Khan. 2021. Rich semantics improve few-shot learning. *arXiv preprint arXiv:2104.12709*.
- Laurent Amsaleg, Oussama Chelly, Teddy Furon, Stéphane Girard, Michael E. Houle, Ken-ichi Kawarabayashi, and Michael Nett. 2015. [Estimating local intrinsic dimensionality](#). In *KDD*, page 29–38.
- Jacob Andreas, Dan Klein, and Sergey Levine. 2018. Learning with latent language. In *NAACL-HLT*, pages 2166–2179.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Satanjeev Banerjee and Alon Lavie. 2005. [METEOR: An automatic metric for MT evaluation with improved correlation with human judgments](#). In *Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization@ACL*, pages 65–72.
- Jasmijn Bastings, Yonatan Belinkov, Emmanuel Dupoux, Mario Giulianelli, Dieuwke Hupkes, Yuval Pinter, and Hassan Sajjad, editors. 2021. *Proceedings of the Fourth BlackboxNLP@EMNLP*.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural Language Processing with Python*. O'Reilly Media, Inc.
- Tseng-Hung Chen, Yuan-Hong Liao, Ching-Yao Chuang, Wan-Ting Hsu, Jianlong Fu, and Min Sun. 2017. Show, adapt and tell: Adversarial training of cross-domain image captioner. In *ICCV*, pages 521–530.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. [Learning phrase representations using RNN encoder–decoder for statistical machine translation](#). In *EMNLP*, pages 1724–1734.
- Sahil Chopra, Michael Henry Tessler, and Noah D Goodman. 2019. The first crank of the cultural ratchet: Learning and transmitting concepts through language. In *CogSci*, pages 226–232.
- Guillem Collell Talleda and Marie-Francine Moens. 2016. Is an image worth more than a thousand words? on the fine-grain semantic differences between visual and linguistic representations. In *COLING*, pages 2807–2817.
- Alexis Conneau, German Kruszewski, Guillaume Lample, Loïc Barrault, and Marco Baroni. 2018. [What you can cram into a single \$\\$&!#*\$ vector: Probing sentence embeddings for linguistic properties](#). In *ACL*, pages 2126–2136.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, pages 4171–4186.
- Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc' Aurelio Ranzato, and Tomas Mikolov. 2013. [Devise: A deep visual-semantic embedding model](#). In *NIPS*, volume 26.
- Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and harnessing adversarial examples. In *ICLR*.
- Xiangteng He and Yuxin Peng. 2017. Fine-grained image classification via combining vision and language. In *CVPR*, pages 5994–6002.
- Lisa Anne Hendricks, Zeynep Akata, Marcus Rohrbach, Jeff Donahue, Bernt Schiele, and Trevor Darrell. 2016. Generating visual explanations. In *ECCV*, pages 3–19.
- Lisa Anne Hendricks, Ronghang Hu, Trevor Darrell, and Zeynep Akata. 2018. Grounding visual explanations. In *ECCV*, pages 264–279.
- Dieuwke Hupkes and Willem Zuidema. 2018. [Visualisation and 'diagnostic classifiers' reveal how recurrent and recursive neural networks process hierarchical structure \(extended abstract\)](#). In *IJCAI*, pages 5617–5621.
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Dong Hoon Lee and Sae-Young Chung. 2021. [Unsupervised embedding adaptation via early-stage feature reconstruction for few-shot classification](#). In *ICML*, volume 139, pages 6098–6108.
- Elizaveta Levina and Peter Bickel. 2005. [Maximum likelihood estimation of intrinsic dimension](#). In *NIPS*, volume 17.
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2020. [What does BERT with vision look at?](#) In *ACL*, pages 5265–5275.
- Qing Li, Jianlong Fu, Dongfei Yu, Tao Mei, and Jiebo Luo. 2018. Tell-and-answer: Towards explainable visual question answering using attributes and captions. In *EMNLP*, pages 1338–1346.
- Weixin Liang, James Zou, and Zhou Yu. 2020. Alice: Active learning with contrastive natural language explanations. In *EMNLP*, pages 4380–4391.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out@ACL*, pages 74–81.
- TY Lin, M Maire, S Belongie, J Hays, P Perona, D Ramanan, P Dollar, and CL Zitnick. 2014. Microsoft coco: Com-mon objects in context. In *ECCV*, pages 740–755.

- Jesse Mu, Percy Liang, and Noah Goodman. 2020. Shaping visual representations with language for few-shot classification. In *ACL*, pages 4823–4830.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *ACL*, pages 311–318.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in pytorch. In *Autodiff@NIPS*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *JMLR*, 21:1–67.
- Sachin Ravi and Hugo Larochelle. 2017. Optimization as a model for few-shot learning. In *ICLR*.
- Scott Reed, Zeynep Akata, Honglak Lee, and Bernt Schiele. 2016. Learning deep representations of fine-grained visual descriptions. In *CVPR*, pages 49–58.
- Jake Snell, Kevin Swersky, and Richard Zemel. 2017. Prototypical networks for few-shot learning. In *NIPS*, volume 30, pages 4077–4087.
- Richard Socher, Milind Ganjoo, Christopher D Manning, and Andrew Ng. 2013. Zero-shot learning through cross-modal transfer. In *NIPS*, pages 935–943.
- Kihyuk Sohn. 2016. [Improved deep metric learning with multi-class n-pair loss objective](#). In *NIPS*, volume 29.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NIPS*, pages 5998–6008.
- Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. 2011. The caltech-ucsd birds-200-2011 dataset.
- Yan Wang, Wei-Lun Chao, Kilian Q Weinberger, and Laurens van der Maaten. 2019. Simpleshot: Revisiting nearest-neighbor classification for few-shot learning. *arXiv preprint arXiv:1911.04623*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *ACL: System Demonstrations*, pages 38–45.
- Chen Xing, Negar Rostamzadeh, Boris Oreshkin, and Pedro O O Pinheiro. 2019. Adaptive cross-modal few-shot learning. In *NeurIPS*, volume 32, pages 4847–4857.

A Experimental Setup

We trained all the models on an NVIDIA Quadro RTX 8000 (48GB), and each experiment took almost one day. The hyperparameter settings are listed in Table 10. We used the Adam optimizer (Kingma and Ba, 2014), PyTorch (Paszke et al., 2017), and transformers (Wolf et al., 2020). Stop words were implemented with NLTK (Bird et al., 2009), and “bird” was added to the stop words. The training of the “No Image” ablated model was unstable due to the transformer architecture, so we used greedy decoding in the test phase to reduce the discrepancy between the train and test phases.

	Pre-Training	Fine-Tuning
Batch size	128	100
# Epochs	100	1500
Learning rate for main model	1e-3	1e-3
Learning rate for text encoder	1e-3	1e-4
Learning rate for text decoder	1e-5	1e-5
λ_{text}	—	10
λ_{cntr}	—	0.1
τ	—	0.05

Table 10: Hyperparameters.

For the experiments on recovering attributions (Section 4.8), we used sigmoid activation and determined the thresholds for each attribution. We sampled one caption for each image in the development and test splits. When an attribution was not described in a caption even though the label was 1, the example was ignored. We also removed attributions when the number of correct predictions was less than 20 or the number of positive examples was less than 20. Here, a positive example means one for which the label is 1 and the attribution is described in the caption. As a result, 68 attributions remained.