

SpatialFormer: Semantic and Target Aware Attentions for Few-Shot Learning

Jinxiang Lai¹, Siqian Yang¹, Wenlong Wu¹, Tao Wu¹, Guannan Jiang², Xi Wang², Jun Liu¹,
Bin-Bin Gao¹, Wei Zhang², Yuan Xie^{*3}, Chengjie Wang^{*1,4}

¹ Tencent YouTu Lab, China

² CATL, China

³ East China Normal University, China

⁴ Shanghai Jiao Tong University, China

{jinxianglai, seasonsyang, ezrealwu, tobinwu}@tencent.com, {jianggn, wangx30, zhangwei}@catl.com
{junsenselee, csqaobb}@gmail.com, yxie@cs.ecnu.edu.cn, jasoncjwang@tencent.com

Abstract

Recent Few-Shot Learning (FSL) methods put emphasis on generating a discriminative embedding features to precisely measure the similarity between support and query sets. Current CNN-based cross-attention approaches generate discriminative representations via enhancing the mutually similar regions of support and query pairs. However, it suffers from two problems: CNN structure produces inaccurate attention map based on local features, and mutually similar backgrounds cause distraction. To alleviate these problems, we design a novel SpatialFormer structure to generate more accurate attention regions based on global features. Different from the traditional Transformer modeling intrinsic instance-level similarity which causes accuracy degradation in FSL, our SpatialFormer explores the semantic-level similarity between pair inputs to boost the performance. Then we derive two specific attention modules, named SpatialFormer Semantic Attention (SFSA) and SpatialFormer Target Attention (SFTA), to enhance the target object regions while reduce the background distraction. Particularly, SFSA highlights the regions with same semantic information between pair features, and SFTA finds potential foreground object regions of novel feature that are similar to base categories. Extensive experiments show that our methods are effective and achieve new state-of-the-art results on few-shot classification benchmarks.

Introduction

Few-Shot Learning (FSL) problem becomes one of the research hotspots due to its special ability that could classify novel (unseen) classes samples relying on few labeled images. In this field, recent works (Hou et al. 2019; Tian et al. 2020; Rizve et al. 2021; Zhengyu et al. 2021; Zhiqiang et al. 2021; Liu et al. 2021; Xu et al. 2021) focus on designing a fine-grained embedding network and improving its robustness to recognize the novel instances.

The criterion of a fine-grained embedding is the ability of distinguish background and the target object. As presented in 1st column of Fig. 1, to generate discriminative features, current Convolutional Neural Network (CNN) based cross-attention methods (Hou et al. 2019; Xu et al. 2021) highlight the mutually similar regions of support and query pairs,

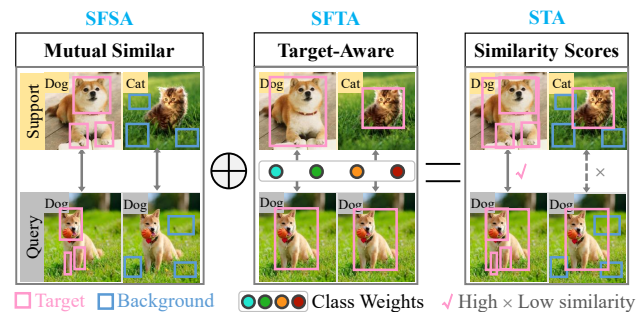


Figure 1: As shown in 1st column, the limitation of cross-attention based approaches is that they only recognize the mutual similar object, but can not solve the distraction problem causing by similar backgrounds. In this paper, we propose two effective modules called SFSA and SFTA, which are utilized to find the mutual similar object and obtain target-aware information, respectively. With their combination, the derived STA module is able to enhance target object while alleviate background distraction.

which assume these similar regions are the target objects. However, these methods suffer from two problems:

- *Problem ①: The CNN-based cross-attention produces inaccurate attention map due to modeling correlation relation based on local features.*
- *Problem ②: The mutually similar backgrounds cause distraction.* For example, in Fig. 1, the query contains dog and grass, while the support image belonging to cat category also has grass. So the meta classifier may regard that this query (dog + grass) is more similar to the support image of cat with grass.

To deal with *Problem ①*, we manage to employ Transformer (Vaswani et al. 2017; Dosovitskiy et al. 2021; Wang et al. 2021), which is able to obtain global information, and has the potential ability to generate a more accurate attention region than CNN-based cross-attention module. However, it raises *Problem ③: The traditional Transformer recommends the inputs are instance-level similar, it against the requirement of FSL which inputs semantic-level similar query and support pairs.* In detailed, SuperGlue (Sarlin et al. 2020) used Transformer to implement cross-attention

*Corresponding Author

between feature pairs for image matching. But our results in Tab. 2 indicates that SuperGlue causes obvious performance degradation in FSL. According to the current applications, Transformer recommends that the input features are instance-level similar. Specifically, the self-attention based Transformer inputs single feature, and utilizes convolution layers to generate perturbed different features. The cross-attention based Transformer, e.g. SuperGlue, inputs pair features that are generated from same instance objects. Obviously, in few-shot classification, the relationship of input support and query pairs violates the recommendation of instance-level similar, i.e. the pairs from same category are semantic-level similar.

To solve *Problem ③*, as illustrated in Fig. 2, we design a novel variant Transformer structure called SpatialFormer to perform spatial-wise cross-attention, which allows the inputs to be semantic-level similar. Our SpatialFormer enhances the spatial response of the input features via the similar regions with *reference object*. The *reference object* can be an instance feature or other forms that has semantic relations with the input features. In practice, given input pair features, SpatialFormer can highlight the mutually similar regions of them with the same semantic information, which is called as SpatialFormer Semantic Attention (SFSA).

To alleviate *Problem ②*, we propose two effective modules called SpatialFormer Target Attention (SFTA) and Novel Task Attention (NTA) to enhance target object while reduce the background distraction. Specifically, as shown in Fig. 1, given input feature, SFTA assigns the class weights (i.e. the weights of linear classifier learned on base dataset) as the *reference object*, to highlight the potential foreground object regions. Particularly, the semantic information of base dataset is embedded into the class weights, which represents base target-aware information. Benefiting from SpatialFormer, SFTA utilizes the class weights to enhance the input novel feature. Besides, NTA involves task-aware information to increase the inter-class feature distance for distinguishing different categories.

In general, our contributions are listed as follows:

- A novel SpatialFormer structure is designed for few-shot learning, which explores semantic-level similar between inputs. It gets rid of intrinsic instance-level similarity in Transformer, which causes accuracy degradation in FSL.
- A Semantic and Target Attentions (STA) module is proposed to enhance the target object helping the subsequent metric classifier to precisely measure similarity. The STA consists of SpatialFormer Semantic Attention (SFSA) and SpatialFormer Target Attention (SFTA), which are utilized to find the mutual similar object and obtain target-aware information, respectively. Besides, for the first time, our SFTA proves that the class weights learned from the base dataset are useful in few-shot learning.
- A Novel Task Attention (NTA) module is introduced to increase the inter-class feature distance via involving task-aware information, which is helpful to generate a more discriminative embedding features.
- Based on STA and NTA modules, a novel STANet framewrok is derived for few-shot classification. Extensive

experiments show that our STANet achieves new state-of-the-art results on few-shot classification benchmarks such as miniImageNet and tieredImageNet.

Related Work

Few-Shot Learning FSL algorithms aim to classify novel classes samples using only few labeled images, and a base dataset containing sufficient labeled images is provided for model pre-training. Four main kinds of FSL algorithms are briefly summarized in below. *Optimization-based* approaches achieve fast model adaptation with few novel images (Ravi and Larochelle 2017; Rusu et al. 2019). *Parameter-generating* approaches (Bertinetto et al. 2016; Gidaris and Komodakis 2019) design a parameter generating network. *Embedding-based* approaches (Tian et al. 2020; Rizve et al. 2021) pay more attention to improve the generalization ability of embedding. *Metric-based* approaches perform classification via calculating similarity between a given input image and the labeled classes samples (Vinyals et al. 2016; Snell, Swersky, and Zemel 2017). Besides, some metric-based approaches (e.g. CAN (Hou et al. 2019) and DANet (Xu et al. 2021)) integrate attention modules to generate discriminative embedding. However, these CNN-based approaches produce the correlation maps based on local features, without global information embedding.

Auxiliary Task for FSL The auxiliary tasks performing supervised (Hou et al. 2019) or self-supervised (Liu et al. 2021; Jinxiang et al. 2022b,a) learning are widely used to improve the generalization ability of few-shot models.

Transformer for FSL The FEAT (Ye et al. 2020a) and CTX (Doersch, Gupta, and Zisserman 2020) apply standard Transformer (Vaswani et al. 2017) structure in feature embedding for FSL. These two methods do not highlight the mutually similar regions between support and query features at semantic level, i.e. they are not able to handle *Problem ①*.

SpatialFormer

Traditional Transformer

Transformer Self-Attention As shown in Fig. 2(a), a standard Transformer (Vaswani et al. 2017) contains an Attention Module and follows a Feed-Forward Network (FFN). Let $f \in \mathbb{R}^{c \times h \times w}$ and $f' \in \mathbb{R}^{c \times h \times w}$ denote the input and output features of a traditional Transformer, respectively. Formally,

$$f' = \text{Transformer}(f) = \text{FFN}(f + A),$$

$$\text{where, } A = \text{Attention}(Q, K, V) = \sigma(QK^T)V, \quad (1)$$

$$\text{where, } Q = W_Q f, \quad K = W_K f, \quad V = W_V f,$$

where, $\{W_Q, W_K, W_V\}$ are convolution layers, and σ is softmax function. $\{Q, K, V\} \in \mathbb{R}^{hw \times c}$ represent $\{query, key, value\}$ as well-known in Transformer. For few-shot classification task, Feat (Ye et al. 2020b) performs embedding adaptation via using the Transformer as a set-to-set function, which makes the modifications as: $Q=W_Q f_{set}$, $K=W_K f_{set}$, $V=W_V f_{set}$, where f_{set} is a feature set of all support samples. Briefly, the vanilla Transformer and Feat

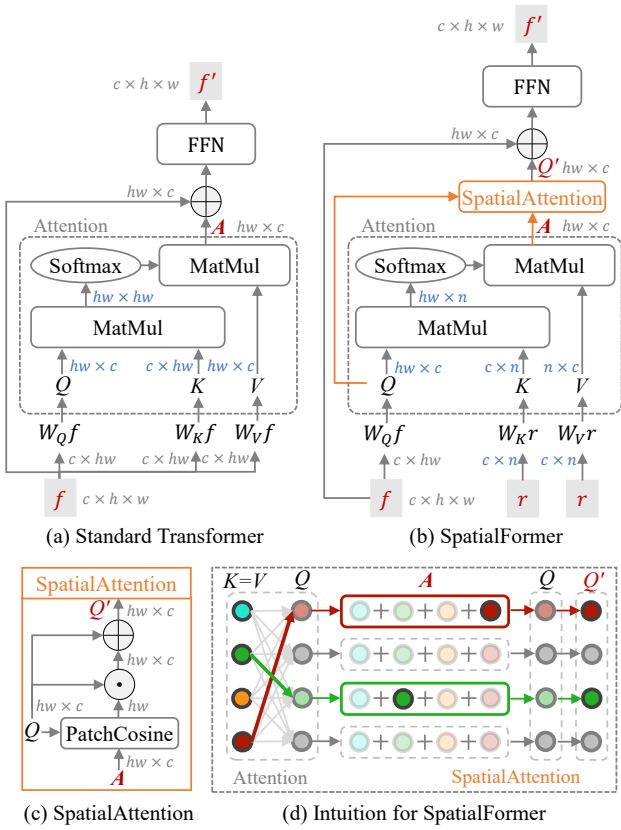


Figure 2: (a) shows a standard Transformer structure. (b) is the proposed SpatialFormer, where (f, r) are input feature and *reference object*, and SpatialAttention operation is presented in (c). (d) gives the intuition for SpatialFormer, which uses A (i.e. Q -aligned V) to strength the spatial response of Q . And the dimension of each node (i.e. the circle in (d)) is \mathbb{R}^c , which denotes each spatial position of features. The gray arrow represents low similarity relationship, and the colorful arrow is high similarity.

learn the correlation of instance features via using self-attention mechanism.

Transformer Cross-Attention Instead of performing self-attention with Transformer, SuperGlue (Sarlin et al. 2020) performs Transformer cross-attention on a pair of input features. Given a pair of input features (f_q, f_s) , it first re-weights f_s with the affinity matrix between (f_q, f_s) , and then the weighted f_s is used to enhance the input feature f_q . Formally, the $\{Q, K, V\}$ of *Transformer cross-attention* are expressed as: $Q=W_Q f_q, K=W_K f_s, V=W_V f_s$.

Transformer Alignment For few-shot classification task, CTX (Doersch, Gupta, and Zisserman 2020) uses the Transformer structure to produce query-aligned prototype f_p . The formulation is listed as follows:

$$f_p = \sum_{i=1}^M \sigma(QK_i^T)V_i, \quad (2)$$

where, $Q = W_Q f_q, K_i = W_K f_s^i, V_i = W_V f_s^i$,

where, M is the number of support instances per category, f_q donates the query feature, and f_s^i is the support feature of i^{th} image. CTX aggregates M support features into alignment with query f_q . Thus CTX achieves feature alignment via implementing cross-attention, while does not enhance the target regions of input feature f_q .

SpatialFormer Methodology

As illustrated in Fig. 2(b), our SpatialFormer applies *Transformer Spatial-Attention*, which enhances the spatial response of the input features f via the similar regions with *reference object* r . Given a pair of feature $f \in \mathbb{R}^{c \times h \times w}$ and *reference object* $r \in \mathbb{R}^{c \times n}$ as input, SpatialFormer outputs the enhanced feature $f' \in \mathbb{R}^{c \times h \times w}$. Formally, SpatialFormer is expressed as:

$$f' = \text{SpatialFormer}(f, r) = \text{FFN}(f + Q'),$$

where, $Q' = \text{SpatialAttention}(Q, A)$,

$$\text{where, } A = \text{Attention}(Q, K, V) = \sigma(QK^T)V,$$

where, $Q = W_Q f, K = W_K r, V = W_V r$.

SpatialAttention is defined as:

$$Q' = \text{SpatialAttention}(Q, A) = Q + Q \odot \text{PatchCosine}(Q, A), \quad (4)$$

where \odot is Broadcasting Element-wise Product to implement spatial-wise attention. Particularly, m^{th} spatial position of $(Q, A) \in \mathbb{R}^{h \times w \times c}$ are denoted as the patch features $(Q_m, A_m) \in \mathbb{R}^c$, then *PatchCosine* operation calculates the cosine similarity between (Q_m, A_m) for each spatial position of (Q, A) , which obtains the result $\text{PatchCosine}(Q, A) \in \mathbb{R}^{h \times w}$.

Next, we discuss the meanings of some key variables in SpatialFormer: (I) (Q, K, V) , are corresponding to inputs (f, r, r) processed by convolution layers for feature adjustment, which are denoted as $(Q, K, V) \approx (f, r, r)$. (II) A , is Q -aligned V . According to the analysis of CTX approach and Fig. 2(d), at m^{th} spatial position, A_m is aggregated from V with the affinity vector between (Q_m, K) . Intuitively, A_m collects all the patch features from V that are semantically similar to the reference Q_m . (III) Q' , is A -enhanced Q . *SpatialAttention* operation uses the cosine similarity score between (Q_m, A_m) to re-weight Q_m to generate Q'_m . Therefore, Q'_m will be enhanced if Q_m is highly similar to A_m (i.e. Q -aligned V), i.e. $Q_m \approx f_m$ can be enhanced to produce Q'_m by the similar regions with $V \approx r$. (IV) f' , is Q' -enhanced f . Thus SpatialFormer can enhance the spatial response of the input features f by the similar regions with *reference object* r . The *reference object* r can be an instance feature or other forms that has semantic relations with the input features.

According to the definition of SpatialFormer in Eq. 3, there is $f' = \text{FFN}(f + Q')$. Let f and r represent two different instance-level objects. Q' is produced via re-weighting $Q \approx f$, thus Q' and f are belong to the same instance-level, which is denoted as $Q' \sim f$. Similarly, according to the analysis of (II), there is $A \sim r$. Comparing to our SpatialFormer, traditional Transformer is defined in

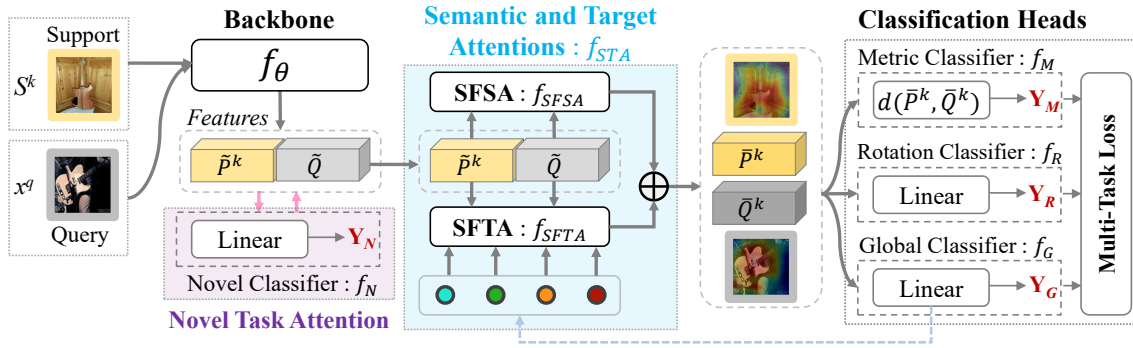


Figure 3: The STANet framework. (a) The designed Novel Task Attention (NTA) module obtains novel task information via utilizing the weights of Novel Classifier, which is fine-tuned on support data of novel set X^{novel} . (b) The proposed Semantic and Target Attentions (STA) module enhances target object regions. Specifically, SpatialFormer Semantic Attention (SFSA) module highlights mutually similar regions between support and query features. Meanwhile, SpatialFormer Target Attention (SFTA) module highlights potential foreground object regions via utilizing base target information embedded in Global Classifier.

Eq. 1 as $f' = FFN(f + A)$. Obviously, the main difference between SpatialFormer and Transformer is from $Q' \sim f$ and $A \sim r$. Transformer updates feature with $FFN(f + A)|A \sim r$, hence the fusion of different instance-level objects may cause serious disturbance on f . Therefore, in current applications such as supervised classification and image matching, the inputs of Transformer are instance-level similar. Differently, our SpatialFormer enhances feature with $FFN(f + Q')|Q' \sim f$, thus the combination of same instance-level objects does not affect the intrinsic semantic essence of f . Therefore, SpatialFormer only requires that the inputs are semantic-level similar, which meets the requirement of few-shot classification task (i.e. the pairs from same category are semantic-level similar).

Semantic and Target Attentions Network

Problem Definition

A N -way M -shot few-shot classification task, aims to train a classification model to recognize N novel categories based on M labeled images per class. Usually, two category-disjoint datasets are provided: a base dataset X^{base} with C^{base} categories for training, and a novel dataset X^{novel} with C^{novel} categories for testing, where $C^{base} \cap C^{novel} = \emptyset$. Following (Vinyals et al. 2016; Hou et al. 2019; Xu et al. 2021), we adopt episodic training paradigm to mimic FSL setting. The episodic training paradigm samples an episode to mimic an individual FSL task. An episode consists of a support set $\mathcal{S} = \{(x_i^s, y_i^s)\}_{i=1}^{m_s}$ ($m_s = N \times M$) and a query set $\mathcal{Q} = \{(x_i^q, y_i^q)\}_{i=1}^{m_q}$, where m_s and m_q are the amount of samples of support set and query set, respectively. Moreover, we denote the k^{th} class support subset as S^k .

Method Framework

As illustrated in Fig. 3, the proposed Semantic and Target Attentions Network (STANet) contains six parts: embedding backbone f_θ , Semantic and Target Attentions (STA)

f_{STA} module, Metric f_M and Novel f_N few-shot classifiers, and auxiliary Rotation f_R and Global f_G classifiers. The proposed STA generates more discriminative representations by highlighting the target object regions of semantic features which benefits the subsequent classifiers. The Novel Task Attention (NTA) operation implemented in Novel Classifier f_N obtains novel task information to increase the inter-class feature distance. our STANet can be divided into two procedures: **Step 1**, train the base model $f_{base} = [f_\theta, f_{STA}, f_M, f_G, f_R]$ on the base set X^{base} ; **Step 2**, fine-tune the Novel Classifier f_N and make predictions with $f_{MN} = [f_\theta, f_{STA}, f_M, f_N]$ on the novel set X^{novel} .

In **Step 1**, the input image x^q in query set $\mathcal{Q} = \{(x_i^q, y_i^q)\}_{i=1}^{m_q}$ is rotated with $[0^\circ, 90^\circ, 180^\circ, 270^\circ]$ and outputs a rotated $\tilde{\mathcal{Q}} = \{(\tilde{x}_i^q, \tilde{y}_i^q)\}_{i=1}^{m_q \times 4}$. The rotated query \tilde{x}^q and support subset S^k are processed by the embedding backbone f_θ , which generates the query feature $\tilde{Q} = f_\theta(\tilde{x}^q) \in \mathbb{R}^{c \times h \times w}$ and support prototype $\tilde{P}^k = \frac{1}{|S^k|} \sum_{x_i^s \in S^k} f_\theta(x_i^s)$, respectively. Then each pair-features (\tilde{P}^k, \tilde{Q}) are processed by STA to enhance the target object regions and generates more discriminative features (\tilde{P}^k, \hat{Q}^k) for the subsequent classification. Finally, f_{base} is optimized by Multi-Task Loss defined in Eq. 9. In **Step 2**, based on NTA operation, firstly the Novel Classifier f_N is fine-tuned on support data of X^{novel} via optimizing $\mathcal{L}_N = CE(Y_N, N)$, where CE is cross-entropy function and Y_N is the prediction result of Novel Classifier. Then the output embedding features of f_θ is updated by NTA for all input support and query images. Finally the fusion model f_{MN} makes predictions on the updated embedding features. Specifically, in inductive inference phase, the overall prediction of STANet is $Y = Y_M + Y_N$, where Y_M is the result of Metric classifier.

In the following context, we first describe two novel components STA and NTA, then the adopted Multi-Task Loss is introduced. The detailed algorithm of STANet is presented in APPENDIX.

Semantic and Target Attentions

The STA module f_{STA} is able to enhance the target object, and consists of SpatialFormer Semantic Attention (SFSA) f_{SFSA} and SpatialFormer Target Attention (SFTA) f_{SFTA} , which are utilized to find the mutual similar object and obtain target-aware information, respectively. The corresponding intuitions of them are presented in Fig. 1. Formally, STA module is expressed as:

$$\begin{aligned} (\tilde{P}^k, \tilde{Q}^k) &= f_{STA}(\tilde{P}^k, \tilde{Q}) \\ &= f_{SFSA}(\tilde{P}^k, \tilde{Q}) + f_{SFTA}(\tilde{P}^k, \tilde{Q}). \end{aligned} \quad (5)$$

SpatialFormer Semantic Attention f_{SFSA} is expressed as:

$$\begin{aligned} (\tilde{P}^{k'}, \tilde{Q}^{k'}) &= f_{SFSA}(\tilde{P}^k, \tilde{Q}) \\ &= SpatialFormer(\tilde{P}^k, \tilde{Q}), SpatialFormer(\tilde{Q}, \tilde{P}^k) \end{aligned} \quad (6)$$

Based on SpatialFormer, SFSA module is able to highlight mutually similar regions between support and query features (\tilde{P}^k, \tilde{Q}) . For example, $\tilde{Q}^{k'} = SpatialFormer(\tilde{Q}, \tilde{P}^k)$: as defined in Eq. 3 and shown in Fig. 2(b), SpatialFormer enhances the spatial response of \tilde{Q} via the similar regions with *reference object* \tilde{P}^k .

Comaring to CNN-based cross-attention, our SFSA produces more complete attention maps by *SpatialAttention* operation. Additionally, inheriting Transformer's property of globally exploring correlations among feature instances, SpatialFormer based SFSA can generate more accurate attention maps via modeling whole data set rather than a meta-task done by CNN-based attention.

SpatialFormer Target Attention f_{SFTA} is expressed as:

$$\begin{aligned} (\tilde{P}^{k'}, \tilde{Q}^{k'}) &= f_{SFTA}(\tilde{P}^k, \tilde{Q}) \\ &= SpatialFormer(\tilde{P}^k, W_G), SpatialFormer(\tilde{Q}, W_G) \end{aligned} \quad (7)$$

where, $W_G \in \mathbb{R}^{C^{base} \times c}$ is the weights of linear Global Classifier which classifies the input into C^{base} classes of base dataset X^{base} . The proposed SFTA module, utilizes the base class weights W_G as the *reference object*, to highlight the potential foreground regions of the input features. Intuitively, the semantic information of base dataset X^{base} is embedded into the class weights W_G , which represents base target-aware information (i.e. corresponding to C^{base} categories). Benefiting from SpatialFormer, SFTA utilizes W_G to enhance the input novel feature. For the first time, our SFTA shows that the semantic information embedded in the base class weights W_G is useful in few-shot learning.

Formally, let's assume the C^{base} contains two subsets C_{sim}^{base} and C_{diff}^{base} , which are semantic similar and different to C^{novel} , respectively. After learning on base dataset X^{base} , $W_G \in \mathbb{R}^{C^{base} \times c}$ gathers the semantic info of C^{base} , i.e. W_G also contains two subsets W_G^{sim} and W_G^{diff} that are semantic similar and different to C^{novel} , respectively. Therefore, in testing on X^{novel} , our SFTA module can enhance the novel feature regions that are similar to W_G^{sim} .

Novel Task Attention

The novel task is sampled from novel set X^{novel} for few-shot classification evaluation, which contains a specific task-aware information for each novel task. As illustrated in

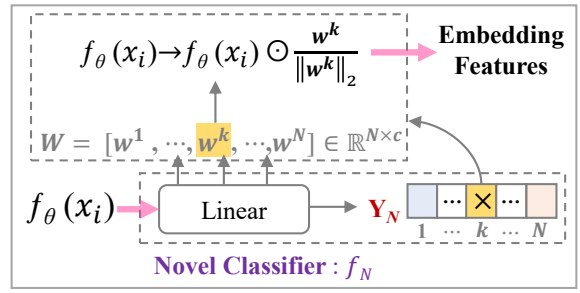


Figure 4: The designed Novel Task Attention (NTA) module obtains novel task information via utilizing the weights of Novel Classifier.

Fig. 4, the Novel Task Attention (NTA) shifts the input feature with the weights of Novel Classifier to generate task-aware feature.

Novel Classifier makes prediction as $Y_N = W f_\theta(x_i)$, where $W \in \mathbb{R}^{N \times c}$ are the weights of Linear layer. After fine-tuning on novel support data, the learned weights W of Novel Classifier is utilized to shift the input features. Specifically, the Novel Classifier recognizes the input x_i as k^{th} category which indicates the strongest response weight is $w^k \subseteq W$, thus $w^k \in \mathbb{R}^c$ is selected as the rectify vector for shifting the embedding feature $f_\theta(x_i)$. Formally, the updated feature is derived by:

$$f_\theta(x_i) \leftarrow f_\theta(x_i) \odot \frac{w^k}{\|w^k\|_2}, \quad (8)$$

where \odot is Broadcasting Element-wise Product to implement channel-wise attention. The learned W represents the response for all classes of novel task, which is considered as task-aware class response, therefore NTA is able to filter out the background noise via the involved task-aware information. Moreover, the features $f_\theta(y = k|x_i)$ of same category inputs have the same corresponding response weight w_k , while the response weights of different categories are not the same, thus NTA can increase the inter-class feature distance based on different rectify vectors.

Multi-Task Loss

The three $\{f_M, f_G, f_R\}$ classifiers category the query image into the corresponding $\{N, C^{base}, B\}$ classes, where B is the four kinds of rotation angle. Their losses are denoted as $\{\mathcal{L}_M, \mathcal{L}_G, \mathcal{L}_R\}$, which are computed by the cross-entropy function between the corresponding predictions $\{Y_M, Y_G, Y_R\}$ and ground truth categories $\{N, C^{base}, B\}$. Then, inspired by (Jinxiang et al. 2022a), the overall multi-task loss is:

$$\mathcal{L} = \frac{1}{2}\mathcal{L}_M + \sum_{j=G,R} \left((\lambda + w_j)\mathcal{L}_j + \log \frac{1}{(\lambda + w_j)} \right), \quad (9)$$

where, $w = \frac{1}{2\alpha^2}$, α is a learnable weight, λ is a hyper-parameter for the balance of few-shot and auxiliary tasks. The ablation study results of λ is presented in Tab. 4.

model	Backbone	miniImageNet		tieredImageNet	
		1-shot	5-shot	1-shot	5-shot
ProtoNet (Snell, Swersky, and Zemel 2017)	Conv4	49.42 ± 0.78	68.20 ± 0.66	53.31 ± 0.89	72.69 ± 0.74
RelationNet (Sung et al. 2018)	Conv4	50.44 ± 0.82	65.32 ± 0.70	54.48 ± 0.93	71.32 ± 0.78
Our STANet	Conv4	57.32 ± 0.47	73.00 ± 0.37	59.65 ± 0.51	74.45 ± 0.41
CAN (Hou et al. 2019)	ResNet-12	63.85 ± 0.48	79.44 ± 0.34	69.89 ± 0.51	84.23 ± 0.37
P-Transfer (Zhiqiang et al. 2021)	ResNet-12	64.21 ± 0.77	80.38 ± 0.59	-	-
MetaOpt+ArL (Hongguang et al. 2021)	ResNet-12	65.21 ± 0.58	80.41 ± 0.49	-	-
DeepEMD (Zhang et al. 2020)	ResNet-12	65.91 ± 0.82	82.41 ± 0.56	71.16 ± 0.87	86.03 ± 0.58
IENet (Rizve et al. 2021)	ResNet-12	66.82 ± 0.80	84.35 ± 0.51	71.87 ± 0.89	86.82 ± 0.58
DANet (Xu et al. 2021)	ResNet-12	67.76 ± 0.46	82.71 ± 0.31	71.89 ± 0.52	85.96 ± 0.35
MCL (Yang et al. 2022)	ResNet-12	67.36 ± 0.20	83.63 ± 0.20	71.76 ± 0.20	86.01 ± 0.20
Meta-DeepBDC (Jiangtao et al. 2022)	ResNet-12	67.34 ± 0.43	84.46 ± 0.28	72.34 ± 0.49	87.31 ± 0.32
Our STANet	ResNet-12	69.84 ± 0.47	84.88 ± 0.30	73.08 ± 0.49	86.80 ± 0.34
PSST (Zhengyu et al. 2021)	WRN-28	64.16 ± 0.44	80.64 ± 0.32	-	-
FEAT (Ye et al. 2020b)	WRN-28	65.10 ± 0.20	81.11 ± 0.14	70.41 ± 0.23	84.38 ± 0.16
DANet (Xu et al. 2021)	WRN-28	67.84 ± 0.46	82.74 ± 0.31	72.18 ± 0.52	86.26 ± 0.35
Our STANet	WRN-28	69.86 ± 0.46	85.16 ± 0.29	74.41 ± 0.50	87.64 ± 0.33

Table 1: Comparing to related methods on 5-way few-shot classification on miniImageNet and tieredImageNet.

Experiment

Datasets and Setting

Datasets and Evaluation Following (Hou et al. 2019; Xu et al. 2021), two popular benchmark datasets *miniImageNet* and *tieredImageNet* are selected, both of which are sampled from ImageNet (Krizhevsky, Sutskever, and Hinton 2012), and the image size is 84×84 pixels. *miniImageNet* dataset contains 100 categories with 600 images per class, which is separated into $\{64, 16, 20\}$ categories for $\{\text{train, validation, test}\}$ respectively. *tieredImageNet* dataset has 608 categories with an average of 1281 images per class, which is divided into $\{351, 97, 160\}$ categories for $\{\text{train, validation, test}\}$ respectively. We evaluate our approaches on 5-way 1-shot and 5-shot classification settings, and report *average accuracy* and *95% confidence interval* on 2000 episodes randomly sampled from the novel test dataset.

Implementation The hyper-parameter λ in Eq. 9 is set to 1.0 according to the results in Tab. 4. More implementation details of STANet for *miniImageNet* and *tieredImageNet* are referred to our released code, such as data augmentation, training epochs, optimizer and learning rate.

Comparing to State-of-the-art Methods

Tab. 1 shows the comparison results between our STANet and the related FSL methods on miniImageNet and tieredImageNet, and more comparisons on CIFAR-FS is presented in APPENDIX. Our STANet outperforms the existing SOTAs, which demonstrates the strength of our method. Different from existing metric-based methods (Zhang et al. 2020; Yang et al. 2022; Jiangtao et al. 2022) extracting support and query features independently, our STANet can enhance the target object regions and obtains more discriminative representations. Comparing to the state-of-the-art metric-based Meta-DeepBDC (Jiangtao et al. 2022), STANet achieves 2.50% higher accuracy on 1-shot task on miniImageNet showing the effectiveness of our STA and NTA modules.

Some metric-based methods (Xu et al. 2021; Hou et al. 2019) integrating cross-attention module, and the proposed STANet still outperforms the competitive DANet (Xu et al. 2021) with a large improvement up to 2.08%, which demonstrates the strength of our Semantic and Target Attentions.

Model Analysis

Influence of STA The results in Tab. 2 shows that: (I) Comparing to ProtoG, our STANet (STA) achieves consistent improvements on 1/5-shot classification tasks, because our method is able to enhance the target object regions and produce more discriminative representations via using the proposed Semantic and Target Attentions. (II) Comparing to CAN (Hou et al. 2019) adopting the CNN-based cross attention, our SpatialFormer based STANet (SFSA, SFTA and STA) achieve obvious improvements. (III) Comparing to ProtoG and CAN, LoFTR and SuperGlue approaches applying Transformer based cross attention, show performance degradation, which indicates that Transformer fuses different instance-level objects may cause serious disturbance on feature updating in few-shot classification task. Benefiting from the designed SpatialFormer integrating features of same instance-level objects, STANet (SFSA) obtains large accuracy improvements than LoFTR and SuperGlue. (V) The proposed SFSA and SFTA are able to find the mutual similar object and obtain target-aware information, respectively. With their combination, the derived STA module achieves further improvements via enhancing the target object while alleviating the background distraction.

Influence of SpatialFormer As shown in Tab. 3, in the setting of assigning (f, r) with (Q, P) and (Q, W_G) , the proposed SpatialFormer achieves consistent improvements on 1/5-shot classification tasks on mini/tieredImageNet, which shows the superiority of our SpatialFormer than Transformer. Besides, comparing to the setting of $(f, r) = (Q, P)$, applying $(f, r) = (Q, W_G)$ obtains higher performance, which indicates that the class weights W_G learned

Model	Attention Module	miniImageNet	
		1-shot	5-shot
ProtoG	None	61.80 ± 0.46	78.59 ± 0.34
CAN	CAM	63.85 ± 0.48	79.44 ± 0.34
STANet	LoFTR	58.30 ± 0.47	72.06 ± 0.36
	SuperGlue	61.99 ± 0.47	73.08 ± 0.36
	SFSA	67.78 ± 0.47	82.05 ± 0.32
	SFTA	67.54 ± 0.47	82.45 ± 0.32
	STA	68.80 ± 0.46	83.12 ± 0.31

Table 2: The 5-way classification results studying the influence of STA module with ResNet-12 backbone. In line with the setting of CAN (Hou et al. 2019), STANet only applies Metric Classifier and Global Classifier (i.e. the Rotation Classifier and Novel Classifier are not adopted). Based on ProtoNet (Snell, Swersky, and Zemel 2017), ProtoG adds Global Classifier to co-train the model.

(f, r)	Module	miniImageNet		tieredImageNet	
		1-shot	5-shot	1-shot	5-shot
(Q, P)	Transformer	60.60	74.13	66.17	79.79
(Q, P)	SpatialFormer	67.78	82.05	71.45	84.73
(Q, W_G)	Transformer	67.39	82.02	71.09	84.54
(Q, W_G)	SpatialFormer	67.54	82.45	71.71	85.24

Table 3: The 5-way classification results studying the influence of SpatialFormer with ResNet-12 backbone. The setting is the same as Tab. 2.

from the base dataset are useful in few-shot learning.

Influence of Multi-Task Loss As shown in Tab. 4, our STANet achieves the optimal results with $\lambda=1.0$, which obtains a large accuracy improvements on 1/5-shot tasks than without Rotation Classifier. These results also indicate that the rotation task is effective for performance improvement due to learning a more representative embedding model.

Influence of NTA As illustrated in Tab. 5, the results under different backbone and dataset show that our NTA module consistently obtains accuracy improvements on 1/5/10-shot tasks, which indicates the proposed NTA module is effective. With the increase of shots, the NTA module obtains larger accuracy improvement, which is due to the increase of support data which leads to learned a more representative weights W of Novel Classifier.

Visualizations

Fig. 5 shows the class activation maps (Bolei et al. 2016) of CAN and our STANet. By comparing STANet to *Embedding*, it shows that STANet is able to highlight the target object which is unseen in the pre-training stage. Comparing to CAN, our STANet is more accurate and has larger receptive field. Obviously, CAN suffers the background distraction problem according to its class activation maps which focus on the background. With the proposed STA and NTA modules, our STANet obtains global attention preventing the target losing in the background, and decreases the background noise to help the model focus on the target objects. More visualizations are presented in APPENDIX.

λ	Loss weights			ResNet-12		WRN-28	
	Metric	Global	Rotation	1-shot	5-shot	1-shot	5-shot
-	0.5	1.0	-	68.80	83.12	67.23	82.27
-	0.5	1.0	1.0	68.97	83.70	68.93	84.14
0.5	0.5	w_G	w_R	69.42	83.74	69.33	83.87
1.0	0.5	w_G	w_R	69.48	84.12	69.55	84.45
1.5	0.5	w_G	w_R	69.29	83.29	69.55	83.94

Table 4: The 5-way classification results of STANet on miniImageNet studying the influence of multi-task loss.

NTA	B.b.	miniImageNet			tieredImageNet		
		1-shot	5-shot	10-shot	1-shot	5-shot	10-shot
-	R.12	69.48	84.12	86.91	72.70	86.28	88.89
✓	R.12	69.84	84.88	87.98	73.08	86.80	89.60
-	W.28	69.55	84.45	86.97	74.01	87.04	89.41
✓	W.28	69.86	85.16	88.09	74.41	87.64	90.10

Table 5: The 5-way classification results of STANet studying the influence of NTA, with $\lambda=1.0$. The B.b., R.12 and W.28 represent Backbone, ResNet-12 and WRN-28, respectively.

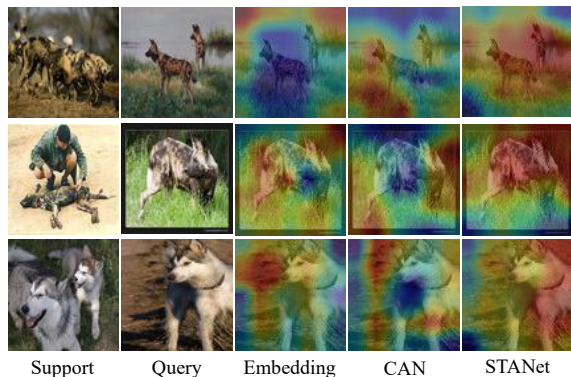


Figure 5: The class activation maps generated from 5-way 5-shot task, where *Embedding* means the features extracted after the backbone in STANet.

Conclusion

In this paper, we proposed a Semantic and Target Attentions Network (STANet) for few-shot classification. Firstly, a novel SpatialFormer structure is designed for few-shot learning, which explores semantic-level similar between inputs. Secondly, two effective modules named STA and NTA are introduced to produce discriminative representations via enhancing target object while filtering out background noise. The NTA module shifts input feature with class weights learned from novel task to increase inter-class feature distance and involve task-aware information. Then, the SpatialFormer based STA module highlights the target object regions to generate more discriminative representations. Extensive experiments show that the proposed approaches are effective, and our STANet achieves new state-of-the-arts on multiple few-shot classification benchmark datasets.

Acknowledgments

This work was supported by the National Key Research and Development Program of China (2021ZD0111000), National Natural Science Foundation of China No. 62222602 and No. 62176092, Shanghai Science and Technology Commission No.21511100700, Natural Science Foundation of Shanghai (20ZR1417700).

References

- Bertinetto, L.; Henriques, J. F.; Valmadre, J.; Torr, P.; and Vedaldi, A. 2016. Learning feed-forward one-shot learners. In *NeurIPS*.
- Bolei, Z.; Aditya, K.; Agata, L.; Aude, O.; and Antonio, T. 2016. Learning Deep Features for Discriminative Localization. In *CVPR*.
- Doersch, C.; Gupta, A.; and Zisserman, A. 2020. CrossTransformers: spatially-aware few-shot transfer. In *NeurIPS*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2021. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*.
- Gidaris, S.; and Komodakis, N. 2019. Generating classification weights with gnn denoising autoencoders for few-shot learning. In *CVPR*.
- Hongguang, Z.; Piotr, K.; Songlei, J.; Hongdong, L.; and Philip, T., H. S. 2021. Rethinking Class Relations: Absolute-relative Supervised and Unsupervised Few-shot Learning. In *CVPR*.
- Hou, R.; Chang, H.; Bingpeng, M.; Shan, S.; and Chen, X. 2019. Cross Attention Network for Few-shot Classification. In *NeurIPS*.
- Jiangtao, X.; Fei, L.; Jiaming, L.; Qilong, W.; and Peihua, L. 2022. Joint Distribution Matters: Deep Brownian Distance Covariance for Few-Shot Classification. In *CVPR*.
- Jinxiang, L.; Siqian, Y.; Guannan, J.; Xi, W.; Yuxi, L.; Zihui, J.; Xiaochen, C.; Jun, L.; Bin-Bin, G.; Wei, Z.; Yuan, X.; and Chengjie, W. 2022a. Rethinking the Metric in Few-shot Learning: From an Adaptive Multi-Distance Perspective. In *ACMMM*.
- Jinxiang, L.; Siqian, Y.; Wenlong, L.; Yi, Z.; Zhongyi, H.; Wenlong, W.; Jun, L.; Bin-Bin, G.; and Chengjie, W. 2022b. tSF: Transformer-Based Semantic Filter for Few-Shot Learning. In *ECCV*.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *NeurIPS*.
- Liu, C.; Fu, Y.; Xu, C.; Yang, S.; Li, J.; Wang, C.; and Zhang, L. 2021. Learning a Few-shot Embedding Model with Contrastive Learning. In *AAAI*.
- Ravi, S.; and Larochelle, H. 2017. Optimization as a model for few-shot learning. In *ICLR*.
- Rizve, M. N.; Khan, S.; Khan, F. S.; and Shah, M. 2021. Exploring Complementary Strengths of Invariant and Equivariant Representations for Few-Shot Learning. In *CVPR*.
- Rusu, A. A.; Rao, D.; Sygnowski, J.; Vinyals, O.; Pascanu, R.; Osindero, S.; and Hadsell, R. 2019. Meta-learning with latent embedding optimization. In *ICLR*.
- Sarlin, P.-E.; DeTone, D.; Malisiewicz, T.; and Rabinovich, A. 2020. Superglue: Learning feature matching with graph neural networks. In *CVPR*.
- Snell, J.; Swersky, K.; and Zemel, R. 2017. Prototypical networks for few-shot learning. In *NeurIPS*.
- Sung, F.; Yang, Y.; Zhang, L.; Xiang, T.; Torr, P. H.; and Hospedales, T. M. 2018. Learning to compare: Relation network for few-shot learning. In *CVPR*.
- Tian, Y.; Wang, Y.; Krishnan, D.; Tenenbaum, J. B.; and Isola, P. 2020. Rethinking few-shot image classification: a good embedding is all you need? In *ECCV*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *NeurIPS*.
- Vinyals, O.; Blundell, C.; Lillicrap, T.; Wierstra, D.; et al. 2016. Matching networks for one shot learning. In *NeurIPS*.
- Wang, W.; Xie, E.; Li, X.; Fan, D.-P.; Song, K.; Liang, D.; Lu, T.; Luo, P.; and Shao, L. 2021. Pyramid Vision Transformer: A Versatile Backbone for Dense Prediction without Convolutions. In *ICCV*.
- Xu, C.; Fu, Y.; Liu, C.; Wang, C.; Li, J.; Huang, F.; Zhang, L.; and Xue, X. 2021. Learning Dynamic Alignment via Meta-filter for Few-shot Learning. In *CVPR*.
- Yang, L.; Weifeng, Z.; Chao, X.; Tu, Z.; Deng, C.; and Xiaofei, H. 2022. Learning To Affiliate: Mutual Centralized Learning for Few-Shot Classification. In *CVPR*.
- Ye, H.-J.; Hu, H.; Zhan, D.-C.; and Sha, F. 2020a. Few-Shot Learning via Embedding Adaptation with Set-to-Set Functions. In *CVPR*.
- Ye, H.-J.; Hu, H.; Zhan, D.-C.; and Sha, F. 2020b. Few-shot learning via embedding adaptation with set-to-set functions. In *CVPR*.
- Zhang, C.; Cai, Y.; Lin, G.; and Shen, C. 2020. DeepEMD: Few-Shot Image Classification with Differentiable Earth Mover's Distance and Structured Classifiers. In *CVPR*.
- Zhengyu, C.; Jixie, G.; Heshen, Z.; Siteng, H.; and Donglin, W. 2021. Pareto Self-Supervised Training for Few-Shot Learning. In *CVPR*.
- Zhiqiang, S.; Zechun, L.; Jie, Q.; Marios, S.; and Kwang-Ting, C. 2021. Partial Is Better Than All: Revisiting Fine-tuning Strategy for Few-shot Learning. In *AAAI*.