

399 **A Algorithm for Episode Loss**

Algorithm 1: Training episode loss computation for adaptive cross-modality few-shot learning. M is the total number of classes in the training set, N is the number of classes in every episode, K is the number of supports for each class, K_Q is the number of queries for each class, \mathcal{W} is the pretrained label embedding dictionary.

Input: Training set $\mathcal{D}_{\text{train}} = \{(\mathbf{x}_i, y_i)\}_i, y_i \in \{1, \dots, M\}$. $\mathcal{D}_{\text{train}}^c = \{(\mathbf{x}_i, y_i) \in \mathcal{D}_{\text{train}} \mid y_i = c\}$.
Output: Episodic loss $\mathcal{L}(\theta)$ for sampled episode e .
 {Select N classes for episode e }
 $C \leftarrow \text{RandomSample}(\{1, \dots, M\}, N)$
 {Compute cross-modal prototypes}
for c in C **do**
 $\mathcal{S}_e^c \leftarrow \text{RandomSample}(\mathcal{D}_{\text{train}}^c, K)$
 $\mathcal{Q}_e^c \leftarrow \text{RandomSample}(\mathcal{D}_{\text{train}}^c \setminus \mathcal{S}_e^c, K_Q)$
 $\mathbf{p}_c \leftarrow \frac{1}{|\mathcal{S}_e^c|} \sum_{(s_i, y_i) \in \mathcal{S}_e^c} f(s_i)$
 $\mathbf{e}_c \leftarrow \text{LookUp}(c, \mathcal{W})$
 $\mathbf{w}_c \leftarrow g(\mathbf{e}_c)$
 $\lambda_c \leftarrow \frac{1}{1 + \exp(-h(\mathbf{w}_c))}$
 $\mathbf{p}'_c \leftarrow \lambda_c \cdot \mathbf{p}_c + (1 - \lambda_c) \cdot \mathbf{w}_c$
end for
 {Compute loss}
 $\mathcal{L}(\theta) \leftarrow 0$
for c in C **do**
 for (q_t, y_t) in \mathcal{Q}_e^c **do**
 $\mathcal{L}(\theta) \leftarrow \mathcal{L}(\theta) + \frac{1}{N \cdot K} [d(f(q_t), \mathbf{p}'_c)] + \log \sum_k \exp(-d(f(q_t), \mathbf{p}'_k))$
 end for
end for

400 **B Descriptions of data sets**

401 **miniImageNet.** This dataset is a subset of ImageNet ILSVRC12 dataset [40]. It contains 100
 402 randomly sampled categories, each with 600 images of size 84×84 . For fair comparison with other
 403 methods, we use the same split proposed by Ravi et al. [38], which contains 64 categories for training,
 404 16 for validation and 20 for test.

405 **tieredImageNet.** This dataset is a larger subset of ImageNet than *miniImageNet*. It contains 34
 406 high-level category nodes (779,165 images in total) that are split in 20 for training, 6 for validation
 407 and 8 for test. This leads to 351 actual categories for training, 97 for validation and 160 for the
 408 test. There are more than 1,000 images for each class. The train/val/test split is done according to
 409 their higher-level label hierarchy. According to Ren et al. [39], splitting near the root of ImageNet
 410 hierarchy results in a more realistic (and challenging) scenario with training and test categories that
 411 are less similar.

412 **CUB-200.** Caltech-UCSD-Birds 200-2011 (CUB-200) [55] is a fine-grained and medium scale
 413 dataset with respect to both number of images and number of classes, *i.e.* 11,788 images from 200
 414 different types of birds annotated with 312 attributes [58]. We chose the split proposed by Xian *et*
 415 *al.* [58]. We used the 312-dimensional hand-crafted attribution as the semantic modality for fair
 416 comparison with other published modality alignment methods.

417 **Word embeddings.** We use GloVe [37] to extract the semantic embeddings for the category labels.
 418 GloVe is an unsupervised approach based on word-word co-occurrence statistics from large text
 419 corpora. We use the Common Crawl version trained on 840B tokens. The embeddings are of
 420 dimension 300. When a category has multiple (synonym) annotations, we consider the first one. If
 421 the first one is not present in GloVe’s vocabulary we use the second. If there is no annotation in
 422 GloVe’s vocabulary for a category (4 cases in *tieredImageNet*), we randomly sample each dimension
 423 of the embedding from a uniform distribution with the range $(-1, 1)$. If an annotation contains more

424 than one word, the embedding is generated by averaging them. We also experimented with fastText
425 embeddings [17] and observed similar performances.

426 C Baselines

427 For modality alignment baselines, we follow CADA-VAE [44]’s few-shot experimental setting.
428 During training, we randomly sample N -shot images for the test classes, and add them in the training
429 data to train the alignment model. During test, we compare the image query and the class embedding
430 candidates in the aligned space to make decisions as in ZSL and GZSL.

431 For the meta-learning extensions of modality alignment methods, instead of including the N -shot
432 images into training data, we follow the standard episode training (explained in Section 3) of metric-
433 based meta-learning approach and train models only with samples from training classes. Moreover,
434 during training, we add an additional loss illustrated in Equation 1 and 3, to ensure the metric space
435 learned on the visual side matching the few-shot test scenario. At test, we employ the standard
436 few-shot testing approach (described in Appendix D) and calculate the prototype representations of
437 test classes as follows:

$$\mathbf{p}_c = \frac{\sum_i \mathbf{r}_i^c + \mathbf{w}_c}{N + 1}, \quad (7)$$

438 where \mathbf{r}_i is the representation of the i -th support image. For both training and test, we need a visual
439 representation space to calculate prototype representations. For DeVISE, they are calculated in its
440 visual space before the transformer [9]. For both ReViSE and CADA-VAE, prototype representations
441 are calculated in the latent space. For f-CLSWGAN, they are calculated in the discriminator’s input
442 space.

443 D Implementation Details of AM3 Experiments

444 We model the visual feature extractor f with a ResNet-12 [12], which has shown to be very effective
445 for few-shot classification [35]. This network produces embeddings of dimension 512. We use this
446 backbone in all the modality-alignment baselines mentioned above and in AM3 implementations
447 (with both backbones). We call *ProtoNets++* the prototypical network [47] implementation with this
448 more powerful backbone.

449 The semantic transformation g is a neural network with one hidden layer with 300 units which
450 also outputs a 512-dimensional representation. The transformation h of the mixture mechanism
451 also contains one hidden layer with 300 units and outputs a single scalar for λ_c . On both g and h
452 networks, we use ReLU non-linearity [10] and dropout [49] (we set the dropout coefficient to be 0.7
453 on *miniImageNet* and 0.9 on *tieredImageNet*).

454 The model is trained with stochastic gradient descent with momentum [51]. We use an initial learning
455 rate of 0.1 and a fixed momentum coefficient of 0.9. On *miniImageNet*, we train every model for
456 30,000 iterations and anneal the learning rate by a factor of ten at iterations 15,000, 17,500 and
457 19,000. On *tieredImageNet*, models are trained for 80,000 iterations and the learning rate is reduced
458 by a factor of ten at iteration 40,000, 50,000, 60,000.

459 The training procedure composes a few-shot training batch from several tasks, where a task is a fixed
460 selection of 5 classes. We found empirically that the best number of tasks per batch are 5, 2 and 1
461 for 1-shot, 5-shot and 10-shot, respectively. The number of query per batch is 24 for 1-shot, 32 for
462 5-shot and 64 for 10-shot. All our experiments are evaluated following the standard approach of
463 few-shot classification: we randomly sample 1,000 tasks from the test set each having 100 random
464 query samples, and average the performance of the model on them.

465 All hyperparameters were chosen based on accuracy on validation set. All our results are reported
466 with an average over five independent run (with a fixed architecture and different random seeds) and
467 with 95% confidence intervals.

468 E Results on CUB-200

469 We also conduct experiments on CUB-200 to better compare with modality-alignment baselines
470 from ZSL. Table 3 shows the results. For 0-shot scenario, AM3 degrades to the simplest modality

471 alignment method that maps the text semantic space to the visual space. Therefore, without the
 472 adaptive mechanism, AM3 performs roughly the same with DeVISE, which indicates that the adaptive
 473 mechanism play the main role on the performance boost we observed in FSL. The results on other
 474 few-shot cases on CUB-200 are consistent with the other two few-shot learning data sets.

475 We also conduct generalized few-shot learning experiments as reported for CADA-VAE in [44] to
 476 compare AM3 with the published FSL results for CADA-VAE. Figure 4 shows that AM3-ProtoNets
 477 outperforms CADA-VAE in every case tested. We consider as a metric the harmonic mean (H-acc)
 478 between the accuracy of seen and unseen classes, as defined in [56, 44].

Model	Test Accuracy		
	0-shot	1-shot	5-shot
DeViSE [9]	52.0%	54.7%	60.4%
ReViSE [14]	55.2%	56.3%	63.7%
VZSL []	57.4%	60.8%	70.0%
CBPL [29]	61.9%	-	-
f-CLSWGAN [57]	62.1%	64.7%	73.7%
CADA-VAE [44]	61.7%	64.9%	71.9%
ProtoNets	-	68.8%	76.4%
AM3-ProtoNets	51.3%	73.6%	79.9%
TADAM [35]	-	69.2%	78.6%
AM3-TADAM	50.7%	74.1%	79.7%

Table 3: Few-shot classification accuracy on *unseen-test* split of CUB-200.

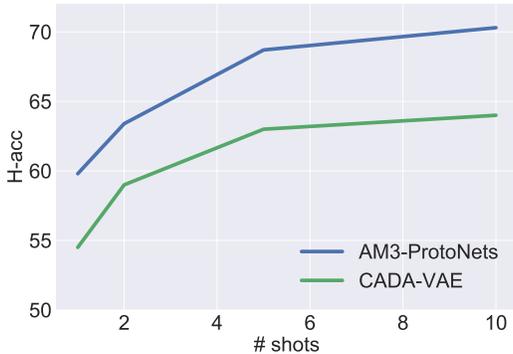


Figure 4: H-acc of generalized few-shot learning on CUB-200.

479 F Ablation study on the input of the adaptive mechanism

480 We also perform an ablation study to see how the adaptive mechanism performs with respect to
 481 different features. Table 4 shows results, on both datasets, of our method with three different inputs
 482 for the adaptive mixing network h : (i) the raw GloVe embedding ($h(\mathbf{e})$), (ii) the visual representation
 483 ($h(\mathbf{p})$) and (iii) a concatenation of both the query and the language embedding ($h(\mathbf{q}, \mathbf{w})$).

484 We observe that conditioning on transformed GloVe features performs better than on the raw features.
 485 Also, conditioning on semantic features performs better than when conditioning on visual ones,
 486 suggesting that the former space has a more appropriate structure to the adaptive mechanism than
 487 the latter. Finally, we note that conditioning on the query and semantic embeddings helps with the
 488 ProtoNets++ backbone but not with TADAM.

Method	ProtoNets++		TADAM	
	1-shot	5-shot	1-shot	5-shot
$h(\mathbf{e})$	61.23	74.77	57.47	72.27
$h(\mathbf{p})$	64.48	74.80	64.93	77.60
$h(\mathbf{w}, \mathbf{q})$	66.12	75.83	53.23	56.70
$h(\mathbf{w})$ (AM3)	65.21	75.20	65.30	78.10

Table 4: Performance of our method when the adaptive mixing network is conditioned on different features. Last row is the original model.