

Learnable Adaptive Cosine Estimator (LACE) for Image Classification

Joshua Peeples Connor H. McCurley Sarah Walker Dylan Stewart Alina Zare
Department of Electrical & Computer Engineering, University of Florida
{jpeeples, cmccurley, sarah.walker, d.stewart, azare}@ufl.edu

Abstract

In this work, we propose a new loss to improve feature discriminability and classification performance. Motivated by the adaptive cosine/coherence estimator [43] (ACE), our proposed method incorporates angular information that is inherently learned by artificial neural networks. Our learnable ACE (LACE) transforms the data into a new “whitened” space that improves the inter-class separability and intra-class compactness. We compare our LACE to alternative state-of-the-art softmax-based and feature regularization approaches. Our results show that the proposed method can serve as a viable alternative to cross entropy and angular softmax approaches. Our code is publicly available.¹

1. Introduction

Feature embedding models for image classification should map data from the same class onto a single point that is metrically far from embeddings of alternative classes [1, 36, 2, 55, 34]. State-of-the-art (SOA) classification networks typically use softmax or embedding losses to guide this feature learning [45, 14, 27, 46, 53]. Yet, *softmax loss* [54, 51] has no inherent property to encourage intra-class compactness or inter-class variation and *embedding loss* [45, 15, 25, 24] methods have proven difficult and untimely to train due to the necessity of hard mining schemes [61, 30, 29, 15, 45]. Consequentially, many studies have investigated variants of both softmax and embedding losses to encourage better discriminating performance [59].

The works of [61, 6, 13, 16, 60, 30, 40] adopt multi-loss learning to increase feature discriminability. While these approaches increase classification performance over traditional softmax loss, they fail to promote both intra-class compactness and inter-class separability. Several works attempt to utilize joint supervision by combining softmax with an auxiliary, Euclidean-based margin loss [50, 52, 61]. Feature representations guided by softmax loss often show

angular separability [29, 28, 59, 9]. To maintain the intrinsic consistency of the features learned by softmax loss neural networks, it is instead favorable to use angular metrics [5] to promote class compactness and separability as opposed to Euclidean-based approaches [45, 25, 4, 47]. Angular softmax losses [29, 28, 59, 9] were recently proposed to take advantage of this knowledge. As with the previous approaches, these losses fail to promote both inter-class separability and intra-class compactness or rely heavily on hyperparameters which can be difficult to select *a priori*.

In this paper, we introduce an alternative loss which adopts the discriminative adaptive cosine/coherence estimator (ACE) [66, 12, 49] to encourage angular separation among classes. As illustrated in Figure 1, the *Learnable ACE (LACE) loss* takes features from a model and whitens them according to a global background distribution. The ACE detection statistic is computed between the samples and each of the C class representatives in the whitened space. Multi-class classification is performed by taking the softmax over ACE outputs. The proposed method demonstrates three key features: 1) it promotes both inter-class separability and intra-class compactness without the need of a separation hyperparameter, 2) unlike comparable methods, LACE enforces strong regularization constraints on the feature embeddings through a whitening transformation and 3) target representatives and data whitening statistics can be updated easily using backpropagation. The proposed LACE loss complements the angular features inherently learned with softmax loss and can be used in joint supervision tasks.

Our contributions can be summarized as follows:

- Apply our whitening transformation based on “background” information to improve feature discriminability - unlike previous approaches
- Ability to learn target signatures and non-target statistics (*i.e.*, background mean and covariance matrix) through backpropagation
- Outperform other angular regularization approaches without hyperparameters (*i.e.*, scale and margin)

¹<https://github.com/GatorSense/LACE>

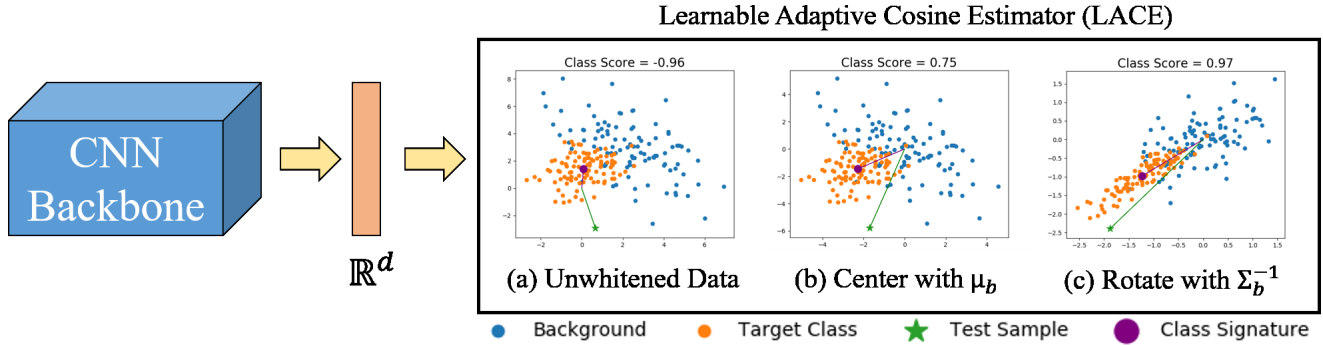


Figure 1: Network with Learnable ACE (LACE) loss. The backbone architecture (e.g., ResNet18 [14]) extracts convolutional feature maps and the features are aggregated (i.e., global average pooling) to be a \mathbb{R}^d feature vector. Here we show a toy example of the transformation performed in LACE. By subtracting the background mean (μ_b) and multiplying by the inverse background covariance matrix (Σ_b), we can improve the estimated score for a new test sample. In LACE, the features are projected into a whitened data space where the C target classes (orange) are compared with the non-target classes (blue). The LACE scores for each sample are then passed into the softmax function and the loss is computed as shown in Equation 4. Network parameters (i.e., target signatures, background mean, and background covariance) are updated via backpropagation. In Figure 1a, many background samples (i.e., samples not in target class) would obtain a high cosine similarity with the target class signature (i.e., class average in this example), while at least half of the true target samples would not. In Figure 1b, which shows background mean subtraction, the cosine similarity between the target signature and true targets drastically increases (-0.96 to 0.75) while the similarity with background samples decreases. The cosine similarity between target samples and the target class signature is further improved post whitening rotation (0.75 to 0.97) as shown in Figure 1c. Once whitening is completed, the background data should be centered and spherical (i.e., identity covariance and zero mean vector). In our figure, the background data may not appear spherical due to scale of axes.

2. Background and Related Work

Softmax Given a set of N training input features $\mathbf{X} \in \mathbb{R}^{N \times D}$ for dimensionality D , a deep CNN will condense the data into a lower dimension d , (typically 512 or 1024), through an embedding function, f . This gives embedding vectors $\mathbf{Z} = f(\mathbf{X}) \in \mathbb{R}^{N \times d}$, which are often passed to a fully-connected layer(s) for classification with cross-entropy + softmax loss or compared to alternative sample embeddings directly with an embedding loss [23, 22]. If $\hat{\mathbf{Y}} \in \mathbb{R}^{N \times C}$ are the sample predictions for C classes, and $\mathbf{Y} \in \mathbb{R}^{N \times C}$ are the true class labels, then the *cross-entropy loss* for class c is given by

$$L_{CE} = \frac{1}{N} \sum_{n=1}^N -\log \left(\frac{\exp \hat{y}_{nc}}{\sum_{j=1}^C \exp \hat{y}_{nj}} \right), \quad (1)$$

where \hat{y}_{nc} denotes the c -th activation of the vector of predicted class scores $\hat{\mathbf{y}}_n$ for sample embedding \mathbf{z}_n . In the softmax loss, $\hat{\mathbf{y}}_n$ contains the activations of a fully connected layer $\mathbf{W} \in \mathbb{R}^{d \times C}$. Thus, $\hat{y}_{nc} = \mathbf{W}_c^T \mathbf{z}_n + b_c$, where $\mathbf{W}_c \in \mathbb{R}^d$ is the c -th column of weight matrix \mathbf{W} and b_c is a bias term. Because this activation is the result of a dot product, [30] showed that the softmax can be reformulated in terms of the angle θ_c ($0 \leq \theta_c \leq \pi$) between vectors \mathbf{W}_c and \mathbf{z}_n . As shown in Equation 2, the embedding features

inferred by a deep CNN with softmax loss are encouraged to show angular distributions (similar to [30], the bias term is omitted to simplify analysis):

$$L_{CE} = \frac{1}{N} \sum_{n=1}^N -\log \left(\frac{\exp \|\mathbf{W}_c\| \|\mathbf{z}_n\| \cos(\theta_c)}{\sum_{j=1}^C \exp \|\mathbf{W}_j\| \|\mathbf{z}_n\| \cos(\theta_j)} \right). \quad (2)$$

Methods using cross entropy + softmax loss are widely used and obtain top scores for many image classification tasks, such as common object and land cover classification, object localization and face recognition [21, 9, 31, 48, 37, 44, 56, 57].

Angular Losses Motivated by the angular feature distributions shown by softmax loss [5, 39], angular margin losses have recently been utilized with much success [30, 9, 59, 8]. Large-margin softmax (L-Softmax) [30] was among the first to characterize softmax loss in terms of angles. L-Softmax promotes inter-class separability by enforcing an angular margin of separation on the classes. Angular softmax (A-softmax) [29] improves upon L-Softmax by projecting the Euclidean features into an angular feature space and introducing a parameter for inter-class separation. Large-margin cosine loss (LMCL) [59] shares the same fundamental ideas of A-Softmax, but ensures that each class is

separated by a uniform margin. Additive angular margin loss [9] directly considers both angle and arc to promote inter-class separability on the normalized hypersphere. Finally, angular margin contrastive loss [5] utilizes training concepts commonly used with ranking losses to jointly optimize an angular loss along with softmax loss. Similarly to Euclidean-based embedding loss methods, angular losses have been utilized primarily in facial recognition tasks [8, 68].

Data Whitening Whitening is a powerful transformation that is used to normalize data. Whitening consist of scaling, rotating, and shifting the data to achieve zero means, unit variances, and decorrelated features [20]. An example of data whitening in deep learning models is batch normalization [20]. Batch normalization has been shown to not only improve model performance, but also training stability. Extensions to the baseline batch normalization approach have been proposed to further improve computational costs and performance of data whitening in deep neural networks [19, 32, 11, 67, 18]. In addition to applying data whitening within neural networks, loss functions that incorporate the whitening transformation have been introduced to achieve better performance [12]. However, all previous approaches apply the data transformation based on representatives from all the data (e.g., via mini-batch). To improve inter-class separation and intra-class compactness of the data samples, the whitening transform can be computed based on “background” data [66, 35], or samples not in the target class, to improve data discriminability. Our work investigates use of data whitening with background statistics on classification performance in deep neural networks.

ACE Many approaches in binary classification are statistical methods in which the target, or class of interest, and background information (all other classes) are modeled as random variables distributed according to an underlying probability distribution [66, 36, 55]. Thus, the classification problem can be posed as a binary hypothesis test: target absent (H_0) or target present (H_1), and a classifier can be designed using the generalized likelihood ratio test (GLRT) [58, 33]. The adaptive cosine/coherence estimator (ACE) [43] is a popular and effective method for performing statistical binary classification which can be formulated as the square-root of the GLRT,

$$D_{ACE}(\mathbf{x}, \mathbf{s}) = \frac{\mathbf{s}^T \Sigma_b^{-1} (\mathbf{x} - \boldsymbol{\mu}_b)}{\sqrt{\mathbf{s}^T \Sigma_b^{-1} \mathbf{s} \sqrt{(\mathbf{x} - \boldsymbol{\mu}_b)^T \Sigma_b^{-1} (\mathbf{x} - \boldsymbol{\mu}_b)}}, \quad (3)$$

where \mathbf{x} is a sample feature vector and \mathbf{s} is an *a priori* target class representative. The background class distribution is parameterized by the mean vector ($\boldsymbol{\mu}_b$) and inverse covariance (Σ_b^{-1}). The response of the ACE classifier, D , is essentially the dot product between a sample and known

class representative in a whitened coordinate space. Similarly to the mentioned angular margin approaches, ACE measures the angular or cosine similarity of a sample with a known class exemplar. This can be considered as comparing the spectral shape of the feature vectors. ACE, however, measures this discrepancy in a coordinate space where data is whitened by the statistics of the background class distribution. While inherently binary, ACE can be utilized in a multi-class setting by forming C one-versus-all classifiers. By applying softmax over all classifier responses and taking the maximum output, our proposed method provides a hard class label for new test samples.

3. Proposed Method

A brief overview of the image classification method using LACE is shown in Figure 1. A CNN feature extraction backbone (e.g., ResNet18 [14]) is used to extract features from the input images. The features are then aggregated (i.e., global average pooling) to provide a deep feature representation in the form of d length vectors. The backbone architecture was jointly trained from scratch using LACE to encourage the deep features to adhere to our metric and not be biased towards the initial pre-training of ImageNet. The LACE score for each sample is then computed and passed into a softmax activation function to produce a predicted probability for each class. In this way, training our approach encourages that feature embedding positions for an image are angularly similar to their corresponding class representative.

3.1. LACE Loss

Inspired by ACE [43, 66] and cross entropy, our LACE objective function is the following,

$$L_{LACE} = \frac{1}{B} \sum_{n=1}^B -\log \left(\frac{\exp \hat{\mathbf{s}}_c^T \hat{\mathbf{x}}_n}{\sum_{j=1}^C \exp \hat{\mathbf{s}}_j^T \hat{\mathbf{x}}_n} \right), \quad (4)$$

where B is the mini-batch size and C is the number of classes. The inner product shown here is simply the ACE similarity which stems from an expansion of Equation 3. This is derived as $\hat{\mathbf{s}}_c = \frac{\hat{\mathbf{s}}}{\|\hat{\mathbf{s}}\|}$, $\hat{\mathbf{x}} = \frac{\hat{\mathbf{x}}}{\|\hat{\mathbf{x}}\|}$, where $\hat{\mathbf{s}} = D^{-\frac{1}{2}} U^T \mathbf{s}$ and $\hat{\mathbf{x}} = D^{-\frac{1}{2}} U^T (\mathbf{x} - \boldsymbol{\mu}_b)$. Here U and D are the eigenvectors and eigenvalues of the inverse background covariance matrix, Σ_b^{-1} , respectively. Thus, the ACE value is the cosine similarity between a test point, \mathbf{x} , and a class representative, \mathbf{s}_c , in a whitened coordinate space. The objective shown in Equation 4 maximizes the ACE similarity between a sample and its corresponding class representative. By computing the loss on the ACE similarity, we encourage the network to learn discriminative features that promote samples belonging to the same class to be “near” one another in the whitened coordinate space. Since ACE is

generally used for binary classification, we extend LACE to include C target signatures. Similar to the weight matrix in a standard fully connected layer, LACE learns a signature matrix, $\mathbf{S} \in \mathbb{R}^{d \times C}$, that is used for multi-class classification.

As can be observed, our LACE loss requires three components which can be learned during training: 1) background mean, 2) background covariance and 3) target class representatives. Unlike other angular softmax alternatives, we do not need to set hyperparameters such as scale and margin beforehand. Therefore, we randomly initialize the target class signature matrix, background mean, and background covariance. The parameters for our proposed LACE method are updated via backpropagation. The derivatives for 1) background mean 2) background covariance and 3) target class representatives are provided in the supplementary material. Pseudo-code for our approach is shown in Algorithm 1.

Algorithm 1 Mini-batch training of LACE loss

Require: $\mathbf{X} \in \mathbb{R}^{N \times D}$ training inputs, $\mathbf{Y} \in \mathbb{R}^{N \times C}$ corresponding class labels, $\mathbf{Z} \in \mathbb{R}^{N \times d}$ deep feature embeddings of \mathbf{X} , $f_{\theta}(\mathbf{x})$ CNN feature embedding model with parameters θ , epochs T , mini-batch size B

- 1: **for** $t = 1$ to T **do**
- 2: **for** each minibatch B **do**
- 3: $\mathbf{Z}_B \leftarrow f_{\theta}(\mathbf{X}_B)$
 Whiten data, \mathbf{Z}_B , using μ_b, Σ_b
 Compute L_{LACE} according to Eq. 4
 Update $\theta, \mu_b, \Sigma_b, \mathbf{S}$ using backpropagation
- 4: **end for**
- 5: **end for**
- 6: **return** f_{θ}

3.2. Implementation and Training Details

Our code was implemented in Python 3.8 with PyTorch 1.9. All models were trained for a max of 300 epochs using the Adam optimizer with mini-batch size of $n = 256$ ($n = 128$ for backbone architecture experiments in Section 4.3) and maximum learning rate of .001. Early stopping was used to terminate training if the validation loss did not improve within 10 epochs. We use the Adam momentum parameters of $\beta_1 = .9$ and $\beta_2 = .999$. Experiments were ran on two 2080 TI GPUs. For LACE, we enforced the inverse covariance matrix to be a positive semi-definite (PSD) matrix. To satisfy this constraint, we performed matrix multiplication between the background covariance matrix and its transpose. Singular value decomposition (SVD) was then performed on the resulting PSD outer product to compute the eigenvectors and eigenvalues (\mathbf{U} and \mathbf{D}) of the background covariance matrix.

A notable advantage of our approach is that LACE can be implemented easily and efficiently if you consider a single, global background class. In this scenario, the target class signature matrix \mathbf{S} and whitening rotation matrix act as weight matrices while the background mean vector acts as a bias. Thus, LACE scores can be easily computed over all classes in the feed-forward portion of the neural network.

4. Experimental Results

We investigated four image datasets: FashionMNIST [62], CIFAR10 [26], SVHN [38], and CIFAR100 [26]. For CIFAR100, we used the coarse (*i.e.*, 20 classes) and sparse (*i.e.*, 100 classes) labels. We evaluated each comparison approach quantitatively through test classification accuracy. We held 10% of the training data for validation and applied each trained model to the holdout test set. We performed three runs of the same random initialization for each method. The CNN backbone was trained from scratch for each loss function to promote the features learned by the model to adhere to the chosen metric. The average test classification accuracy is reported across the different runs. We followed the same training procedure and data augmentation as Xue et. al. [64]. The proposed method was compared to several SOA angular softmax and feature regularization approaches: arcface [10, 9], cosface [59], spherface [29], center loss [61], and angular margin contrastive (AMC) loss [5]. The hyperparameters for the angular margin (m) and scale (γ) for comparison approaches were set to the following values: arcface ($m = 0.05$, $\gamma = 30$), cosface ($m = 0.4$, $\gamma = 30.0$), and spherface ($m = 64.0$, $\gamma = 1.35$). For AMC and center loss, the weighting terms on the contrastive loss was set to .1 and .9 on the softmax loss (the weights on both terms were constrained to sum to 1). The backbone architecture for the SOA comparisons was ResNet18 [14].

4.1. Comparison with SOA Approaches

The performances for our proposed approaches and SOA angular softmax/feature regularization methods are shown in Table 1. The results of LACE indicate the effectiveness of the data whitening when comparing against the baseline softmax approach. We can infer that that the background statistics improve discriminability between target and non-target samples. For CIFAR10, FashionMNIST, and SVHN, the LACE approach overall achieved statistically significant performance improvements in comparison to the baseline softmax models. LACE also outperforms the other angular softmax alternatives. Additionally, LACE outperforms feature regularization approaches based on Euclidean (*i.e.*, center loss) and angular (*i.e.*, AMC) information. As a results, we can optimize for only a single objective term that will jointly maximize classification performance and feature regularization.

Table 1: Test classification accuracy on each dataset for softmax, angular softmax alternatives, feature regularization methods, and our proposed approach (LACE) with ResNet18 as the backbone architecture. The best average result is bolded and the second best result is underlined.

Method	CIFAR10	FashionMNIST	SVHN	CIFAR100 (Coarse)	CIFAR100 (Sparse)
Softmax	88.69±0.21	93.96±0.08	94.68±0.15	76.38±0.76	66.86±0.76
Arcface [9]	88.75±0.27	93.71±0.55	94.93±0.32	75.78±1.10	66.94±0.80
Cosface [59]	88.84±0.23	<u>94.28±0.04</u>	<u>95.28±0.13</u>	77.04±0.40	<u>68.84±0.08</u>
Sphereface [29]	88.80±0.51	94.17±0.16	94.64±0.47	75.74±0.37	65.12±0.78
AMC [5]	88.40±0.51	94.09±0.30	94.98±0.23	<u>77.20±0.28</u>	67.68±0.22
Center Loss [61]	<u>89.02±0.39</u>	94.29±0.09	94.75±0.50	77.77±0.10	70.03±0.90
LACE (ours)	90.26±0.19	94.27±0.13	95.69±0.20	63.32±1.00	33.12±0.56

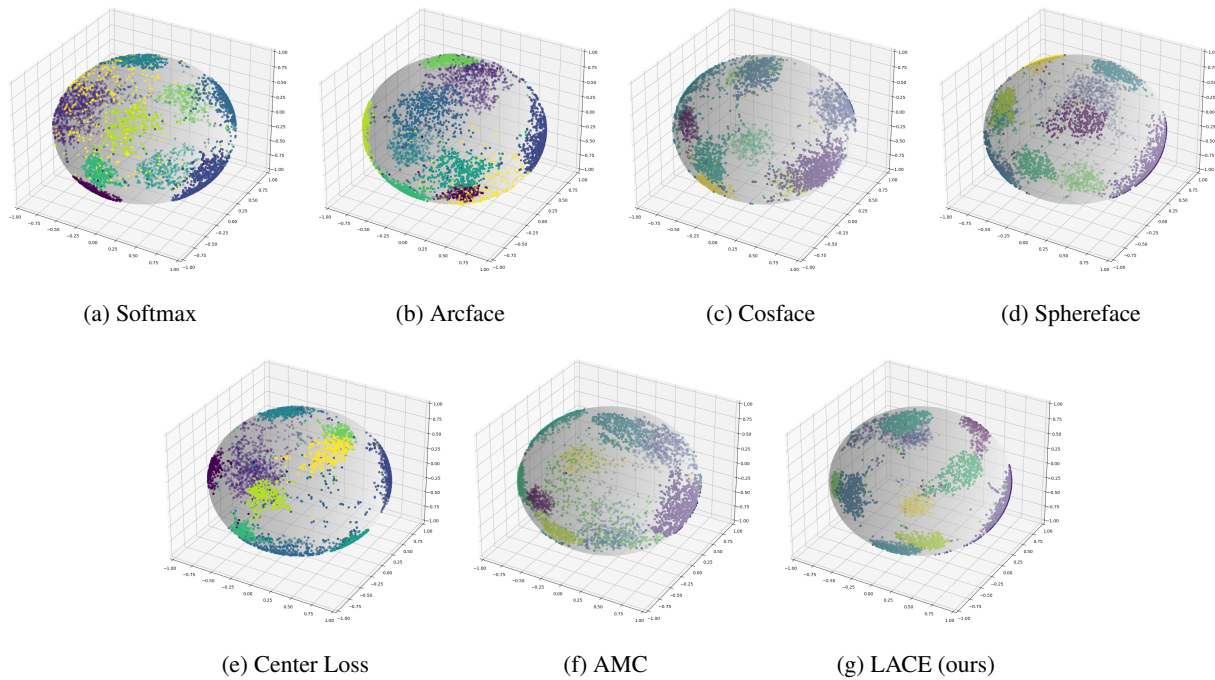


Figure 2: 3D t-SNE projection of features extracted from the penultimate layer for each method on the same 5,000 test images from SVHN. The projections are constrained to lie on the unit sphere. Our method improves intra-class compactness and inter-class separation in comparison to all other comparison methods.

For CIFAR100 (coarse and sparse), our proposed method did not perform as well in comparison to other SOA approaches. A possible reason for this is that there are a total of 20 or 100 classes in this dataset. As a result, it will be difficult to learn background statistics that improve performance for all classes. We can also verify this observation in that the performance of LACE significantly decreases once the number of classes increases from 20 to 100 for CIFAR100. For the remaining three datasets, each was comprised of only 10 classes. LACE could be further improved by modeling additional background statistics to account for the various classes.

The qualitative results for each method on SVHN test

data are shown in Figure 2. We projected the features from the original \mathbb{R}^{512} to \mathbb{R}^3 using t-SNE with Cosine as the distance metric. LACE improves the angular class compactness and separation in comparison to the other approaches. For softmax in Figure 2a, there is no direct component in the objective to consider intra-class and inter-class compactness. We observe the same here in the visualization of the features from t-SNE. Center loss, as shown in Figure 2e, improves intra-class compactness, but there is still considerable overlap with data samples from different classes in the projected space (the classes maybe more separable in higher dimensions). For LACE and other approaches that consider angular information, the features appear more

Table 2: Cluster validity scores for the Resnet18 backbone on FashionMNIST. The best average result is bolded and the second best result is underlined. As shown in each metric, LACE improves the inter-class separation and intra-class compactness in comparison to other methods.

Method	Silhouette [42] (1 = best)	Davies-Bouldin [7] (lowest = best)	Calinski-Harabasz [3] (highest = best)
Softmax	0.58±0.01	1.07±0.03	3055±251
Arcface [9]	0.64±0.01	0.98±0.00	3768±259
Cosface [59]	0.71±0.01	0.98±0.05	3085±30
Sphereface [29]	0.65±0.01	0.92±0.01	5070±384
AMC [5]	0.71±0.01	0.89±0.03	4703±122
Center Loss [61]	0.71±0.01	0.98±0.05	3085±30
LACE (ours)	<u>0.77±0.00</u>	<u>0.69±0.00</u>	5737±195
LACE (pre-whitened) (ours)	0.83±0.00	0.68±0.00	<u>5640±174</u>

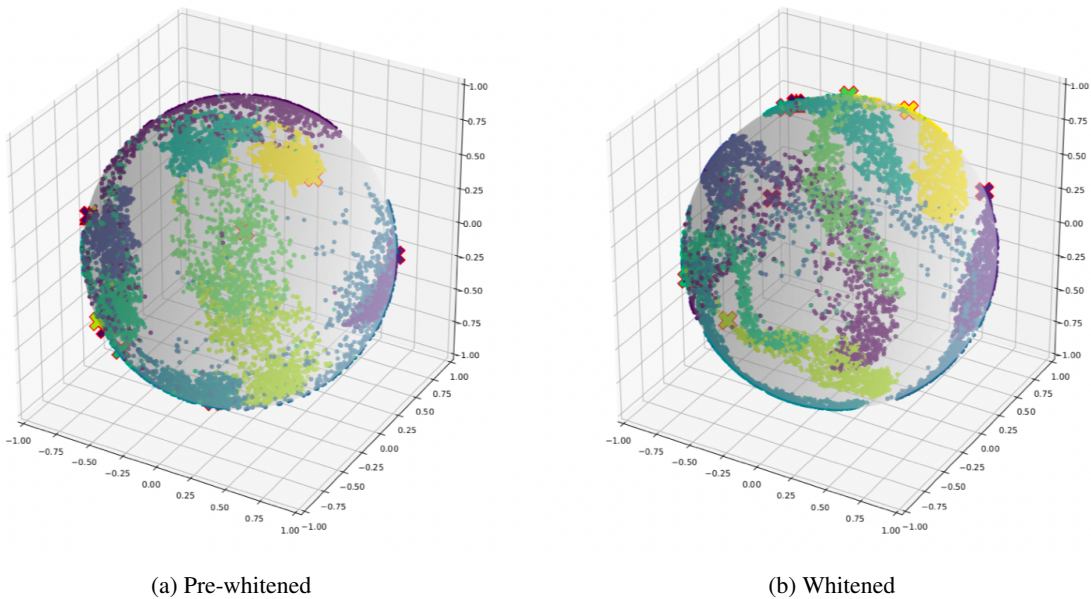


Figure 3: t-SNE visualization of FashionMNIST test data (a) before and (b) after LACE whitening. Learned target signatures are depicted by ‘X’s outlined in red. As seen here, data whitening with the learned background statistics reduces feature variance across multiple dimensions. This makes features more separable according to an angular metric, as opposed to Euclidean distance.

separable and compact with LACE achieving the minimal intra-class and maximum inter-class variation. We further evaluate the discriminative power of the features learned by LACE through cluster validity metrics.

4.2. Clustering Performance

SOA approaches in the literature aim to improve softmax loss by either reducing intra-class scatter or increasing between-class separation. In this work, five cluster validity metrics were used to compare the induced distributions of classes. Namely, Silhouette [42], Davies-Bouldin (DBI)

[7] and Calinski-Harabasz (also called Variance Ratio Criterion) [3] scores were computed to quantify both inter-cluster separation and within-cluster compactness, while homogeneity and completeness [41] measured the accuracy of the induced clustering. Cluster validity scores for each algorithm on FashionMNIST are shown in Table 2. As can be seen, LACE features (both whitened and pre-whitened) consistently outperformed the alternatives for the three metrics measuring intra-class compactness and inter-class separation. Homogeneity and completeness essentially measure the label accuracy of samples assigned to clusters. Per-

Table 3: Test classification accuracy on each dataset for different backbone architectures with CIFAR10. The best average result is bolded and the second best result is underlined.

Method	ResNet50 [14]	ResNeXt-50 32x4d [63]	WideResNet50-2 [65]	DenseNet121 [17]
Softmax	89.86±0.45	91.29±0.83	89.96±0.17	90.82±0.27
Arcface [9]	<u>89.88±0.40</u>	90.51±0.29	89.66±0.36	90.86±0.31
Cosface [59]	90.16±0.23	<u>90.86±0.36</u>	89.76±0.25	91.05±0.40
Sphereface [29]	89.58±0.28	90.57±0.53	90.55±0.35	90.67±0.55
AMC [5]	89.21±0.44	90.71±0.50	89.37±0.36	90.95±0.38
Center Loss [61]	89.77±0.37	89.75±1.03	<u>90.18±0.47</u>	<u>91.00±0.29</u>
LACE (ours)	87.76±0.14	87.06±1.46	85.51±0.43	84.49±0.37

formance for these two metrics was very close across all approaches (0.98), so the results are negligible. Thus, the networks trained with LACE loss not only obtained the best overall classification accuracy, but also discovered feature representations whose class distributions were the most compact and separable.

An interesting observation from Table 2 is that the cluster validity scores on LACE features were better before whitening. This result might be explained by analysis of Figure 3, which shows t-SNE visualizations of FashionMNIST test points before and after whitening. The figure shows that whitening with the learned background statistics makes the induced clusters less circular. This results in a trade-off where within-class compactness is lost, according to the chosen metrics, but between class separation is improved along the angular dimension. Additionally, in the whitened space, cluster variance is increased along an angle but decreased in two other dimensions. Thus, the LACE whitening might sacrifice the traditional ideal of projecting all data of the same class onto a single point to improve discriminability in the angular sense. This result matches intuition of classification with ACE.

4.3. Ablation Study

Backbone Architectures To further evaluate our proposed method, we investigate applying different backbone architectures to learn and extract features of varying dimensionality. We used four additional backbone architectures in addition to ResNet18: ResNet50 [14], ResNeXt-50 32x4d [63], WideResNet50-2 [65], and DenseNet121 [17]. The feature dimensionality of ResNet50, ResNeXt-50 32x4d, and WideResNet50-2 was \mathbb{R}^{2048} while DenseNet121 was \mathbb{R}^{1024} . We present results of these experiments for CIFAR10 in Table 3. From the results, we observe that performance is stable (*i.e.*, small standard deviations) across all methods. The results demonstrate the robustness of our method across different backbone architectures. LACE can thus be integrated seamlessly into various deep learning models to learn discriminative features to improve intra-class compactness and inter-class separation.

In comparison to Table 1, the results of the other methods improved while LACE performed slightly below the previous results with ResNet18. For each of the new backbone architectures, the feature dimensionality increased in comparison to ResNet18 (\mathbb{R}^{512}). As stated previously, it will become more difficult to learn global background statistics that optimize classification performance for all classes as dimensionality increases. Another observation is that LACE performs best with a smaller (*i.e.*, fewer parameters) model. As a result, models in future work could potentially use LACE as a method to maintain and/or improve performance while reducing the number of parameters needed by deeper artificial neural networks.

Impact of LACE components In order to provide more insight into our proposed method, we performed an ablation study of the different components of our approach. LACE is comprised of two major components: the background mean (μ_b) and background covariance (Σ_b). We evaluated each component of LACE for the CIFAR10 dataset with ResNet18 as the backbone architecture and the test accuracy is shown in Table 4. We observe that with only L_2 normalization, the performance is not statistically significant in comparison to the baseline softmax results in Table 1. The typical cross entropy loss does not directly account for inter-class relationships in the data. By including the

Table 4: We evaluated the components of LACE that made the most impact on performance by using CIFAR10: background mean (μ_b) and background covariance (Σ_b). The features and target signature matrix were L_2 normalized for each combination. A \checkmark indicates that component was active during the experimental runs.

μ_b	Σ_b	CIFAR10
		85.49±3.21
	\checkmark	87.93±0.14
\checkmark		87.26±1.54
\checkmark	\checkmark	90.26±0.19

background statistics in our loss, we can increase the separation between data samples in different classes and map samples in the same class closer to one another.

Including the background covariance improved the performance of the baseline model in conjunction with the $L2$ normalization. The model is able to learn meaningful rotations of the data that can account for both inter-class and intra-class relationships to maximize classification performance. The background mean also improved performance of baseline results as well. The background mean serves as a way to shift the data as we saw in the toy example in Figure 1. The background mean provides additional flexibility to the model. Using the background mean and covariance separately lead to similar performance with the background covariance achieving a slightly higher average test accuracy. Since artificial neural networks are inherently learning these “angular” feature distributions, the background covariance matrix will be the most powerful component of the whitened data transformation in LACE. When both background statistics are used, we achieve optimal performance for CIFAR10.

4.4. Computational Cost of LACE

LACE has shown to improve performance, but we must also consider the additional cost of the proposed method. In comparison to the traditional softmax output layer, LACE introduces two additional components: background mean and background covariance. By adding the background statistics, we introduce $d(d + 1)$ parameters to be learned (d and d^2 for the background mean and covariance respectively). As the feature dimensionality increases, the computational costs of the whitening operation will increase. However, as shown in other whitening works [49, 19, 67], various optimization approaches can be used to improve the cost of computing the background statistics of the proposed method.

5. Conclusion

We presented a new loss function, LACE, that can be used to improve image classification performance. Our proposed method uses the angular features inherently learned by artificial neural networks, tuning them to become more discriminative data representations. Qualitative and quantitative analysis show that LACE outperforms and/or performs comparatively with other SOA angular softmax and feature regularization approaches. Additionally, LACE does not need an angular margin and/or scale set *a priori* since all parameters are learned during training.

Our proposed approach outperforms or achieves comparable performance to alternative SOA angular margin approaches in several benchmark image classification tasks. A notable observation is that LACE typically achieves better performance when implemented with smaller models and

fewer classes. To improve performance with larger models, several approaches could be investigated. Instead of assuming a global background distribution, the authors believe performance can be further improved through one-vs-all or one-vs-one classification where each class has its own corresponding background mean and covariance. Additionally, alternative update approaches could be implemented to update the LACE parameters (background means, background covariance, target signatures). Future work will explore parameter estimation from the data, directly, as well as employing physical constraints on the possible solutions to adhere more closely to the data. This could be beneficial for learning both discriminative and physically meaningful representations. Finally, as with batch normalization, LACE may be applied at different portions of the deep learning model. Future investigation will explore how whitening with background statistics within the model affect classification performance and training stability.

Acknowledgments

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-1842473. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

References

- [1] Y. Bengio, A. C. Courville, and P. Vincent. Unsupervised feature learning and deep learning: A review and new perspectives. *CoRR*, abs/1206.5538, 2012.
- [2] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, New York, 2006.
- [3] T. Caliński and J. Harabasz. A dendrite method for cluster analysis. *Communications in Statistics*, 3(1):1–27, 1974.
- [4] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations, 2020.
- [5] Hongjun Choi, Anirudh Som, and Pavan Turaga. Amc-loss: Angular margin contrastive loss for improved explainability in image classification, 2020.
- [6] S. Chopra, R. Hadsell, and Y. LeCun. Learning a similarity metric discriminatively, with application to face verification. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*, volume 1, pages 539–546 vol. 1, 2005.
- [7] David L. Davies and Donald W. Bouldin. A cluster separation measure. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-1(2):224–227, 1979.

- [8] Jiankang Deng, J. Guo, Tongliang Liu, Mingming Gong, and S. Zafeiriou. Sub-center arcface: Boosting face recognition by large-scale noisy web faces. In *ECCV*, 2020.
- [9] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition, 2019.
- [10] Jiankang Deng, Jia Guo, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. *CoRR*, abs/1801.07698, 2019.
- [11] Guillaume Desjardins, Karen Simonyan, Razvan Pascanu, and Koray Kavukcuoglu. Natural neural networks. *arXiv preprint arXiv:1507.00210*, 2015.
- [12] Aleksandr Ermolov, Aliaksandr Siarohin, Enver Sangineto, and Nicu Sebe. Whitening for self-supervised representation learning. *CoRR*, abs/2007.06346, 2020.
- [13] R. Hadsell, S. Chopra, and Y. LeCun. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 1735–1742, 2006.
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition, 2015.
- [15] Alexander Hermans, Lucas Beyer, and Bastian Leibe. In defense of the triplet loss for person re-identification. *CoRR*, abs/1703.07737, 2017.
- [16] Elad Hoffer and Nir Ailon. Deep metric learning using triplet network. In Aasa Feragen, Marcello Pelillo, and Marco Loog, editors, *Similarity-Based Pattern Recognition*, pages 84–92, Cham, 2015. Springer International Publishing.
- [17] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- [18] Lei Huang, Dawei Yang, Bo Lang, and Jia Deng. Decorrelated batch normalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 791–800, 2018.
- [19] Lei Huang, Yi Zhou, Fan Zhu, Li Liu, and Ling Shao. Iterative normalization: Beyond standardization towards efficient whitening. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4874–4883, 2019.
- [20] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. PMLR, 2015.
- [21] Sen Jia, Shuguo Jiang, Zhijie Lin, Nanying Li, Meng Xu, and Shiqi Yu. A survey: Deep learning for hyperspectral image classification with few labeled samples. *Neurocomputing*, 448:179–204, 2021.
- [22] Mahmut KAYA and Hasan Şakir BİLGE. Deep metric learning: A survey. *Symmetry*, 11(9), 2019.
- [23] Asifullah Khan, Anabia Sohail, Umme Zahoora, and Aqsa Saeed Qureshi. A survey of the recent architectures of deep convolutional neural networks. *Artificial Intelligence Review*, 53(8):5455–5516, Apr 2020.
- [24] Gregory Koch. *Siamese Neural Networks for One-Shot Image Recognition*. PhD thesis, University of Toronto, Toronto, Canada, 2015.
- [25] Gregory R. Koch. Siamese neural networks for one-shot image recognition. 2015.
- [26] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [27] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc., 2012.
- [28] Yutian Li, Feng Gao, Zhijian Ou, and Jiasong Sun. Angular softmax loss for end-to-end speaker verification, 2018.
- [29] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Sphereface: Deep hypersphere embedding for face recognition. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6738–6746, 2017.
- [30] Weiyang Liu, Yandong Wen, Zhiding Yu, and Meng Yang. Large-margin softmax loss for convolutional neural networks, 2017.
- [31] D. Lu and Q. Weng. A survey of image classification methods and techniques for improving classification performance. *International Journal of Remote Sensing*, 28(5):823–870, 2007.
- [32] Ping Luo. Learning deep architectures via generalized whitened neural networks. In *International Conference on Machine Learning*, pages 2238–2246. PMLR, 2017.
- [33] D. Manolakis, M. Pieper, E. Truslow, T. Cooley, M. Brueggeman, and S. Lipson. The remarkable success of adaptive cosine estimator in hyperspectral target detection. In Sylvia S. Shen and Paul E. Lewis, editors, *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XIX*, volume 8743, pages 1 – 13. International Society for Optics and Photonics, SPIE, 2013.

- [34] Connor H. McCurley. University of florida research proposal: Discriminative manifold embedding with imprecise, uncertain and ambiguous data, 2021.
- [35] S.K. Meerdink, J. Bocinsky, A. Zare, N. Kroeger, C. H. McCurley, D. Shats, and P.D. Gader. Multi-target multiple instance learning for hyperspectral target detection. *IEEE Transaction on Geoscience and Remote Sensing (TGRS)*, 2021.
- [36] K. P. Murphy. *Machine Learning: A Probabilistic Perspective*. The MIT Press, 2012.
- [37] Siddhartha Sankar Nath, Girish Mishra, Jajnayaseni Kar, Sayan Chakraborty, and Nilanjan Dey. A survey of image classification methods and techniques. In *2014 International Conference on Control, Instrumentation, Communication and Computational Technologies (ICCICCT)*, pages 554–557, 2014.
- [38] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. 2011.
- [39] Joshua Peeples, Sarah Walker, Connor McCurley, Alina Zare, and James Keller. Divergence regulated encoder network for joint dimensionality reduction and classification. *arXiv preprint arXiv:2012.15764*, 2020.
- [40] Ce Qi and Fei Su. Contrastive-center loss for deep neural networks, 2017.
- [41] Andrew Rosenberg and Julia Hirschberg. V-measure: A conditional entropy-based external cluster evaluation measure. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 410–420, Prague, Czech Republic, June 2007. Association for Computational Linguistics.
- [42] Peter J. Rousseeuw. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20:53–65, 1987.
- [43] L.L. Scharf and L.T. McWhorter. Adaptive matched subspace detectors and adaptive coherence estimators. In *Conference Record of The Thirtieth Asilomar Conference on Signals, Systems and Computers*, pages 1114–1117 vol.2, 1996.
- [44] Lars Schmarje, Monty Santarossa, Simon-Martin Schröder, and Reinhard Koch. A survey on semi-, self- and unsupervised learning for image classification, 2021.
- [45] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. *CoRR*, abs/1503.03832, 2015.
- [46] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition, 2015.
- [47] Kihyuk Sohn. Improved deep metric learning with multi-class n-pair loss objective. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 1857–1865. Curran Associates, Inc., 2016.
- [48] M. Sornam, Kavitha Muthusubash, and V. Vanitha. A survey on image classification and activity recognition using deep convolutional neural network architecture. In *2017 Ninth International Conference on Advanced Computing (ICoAC)*, pages 121–126, 2017.
- [49] Jianlin Su, Jiarun Cao, Weijie Liu, and Yangyiwen Ou. Whitening sentence representations for better semantics and faster retrieval. *CoRR*, abs/2103.15316, 2021.
- [50] Yi Sun, Yuheng Chen, Xiaogang Wang, and Xiaoou Tang. Deep learning face representation by joint identification-verification. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2, NIPS’14*, page 1988–1996, Cambridge, MA, USA, 2014. MIT Press.
- [51] Yi Sun, Xiaogang Wang, and Xiaoou Tang. Deep learning face representation from predicting 10,000 classes. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1891–1898, 2014.
- [52] Yi Sun, Xiaogang Wang, and Xiaoou Tang. Sparsifying neural network connections for face recognition, 2015.
- [53] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–9, 2015.
- [54] Yaniv Taigman, Ming Yang, Marc’Aurelio Ranzato, and Lior Wolf. Deepface: Closing the gap to human-level performance in face verification. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1701–1708, 2014.
- [55] Sergios Theodoridis and Konstantinos Koutroumbas. *Pattern Recognition, Fourth Edition*. Academic Press, Inc., USA, 4th edition, 2008.
- [56] Hugo Touvron, Andrea Vedaldi, Matthijs Douze, and Hervé Jégou. Fixing the train-test resolution discrepancy, 2020.
- [57] Hugo Touvron, Andrea Vedaldi, Matthijs Douze, and Hervé Jégou. Fixing the train-test resolution discrepancy: Fixefficientnet, 2020.

- [58] François Vincent and Olivier Besson. Non zero mean adaptive cosine estimator and application to hyperspectral imaging. *IEEE Signal Processing Letters*, 27:1989–1993, 2020.
- [59] Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei Liu. Cosface: Large margin cosine loss for deep face recognition. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5265–5274, 2018.
- [60] Jiang Wang, Yang Song, Thomas Leung, Chuck Rosenberg, Jingbin Wang, James Philbin, Bo Chen, and Ying Wu. Learning fine-grained image similarity with deep ranking. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1386–1393, 2014.
- [61] Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. A discriminative feature learning approach for deep face recognition. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision – ECCV 2016*, pages 499–515, Cham, 2016. Springer International Publishing.
- [62] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- [63] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017.
- [64] Jia Xue, Hang Zhang, and Kristin Dana. Deep texture manifold for ground terrain recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 558–567, 2018.
- [65] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.
- [66] Alina Zare, Changzhe Jiao, and Taylor Glenn. Discriminative multiple instance hyperspectral target characterization. *IEEE Trans. Pattern Anal. Mach. Inteli.*, 40(10):2342–2354, Oct. 2018.
- [67] Shengdong Zhang, Ehsan Nezhadarya, Homa Fashandi, Jiayi Liu, Darin Graham, and Mohak Shah. Stochastic whitening batch normalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10978–10987, 2021.
- [68] Xiao Zhang, Rui Zhao, Yu Qiao, Xiaogang Wang, and Hongsheng Li. Adacos: Adaptively scaling cosine logits for effectively learning deep face representations, 2019.