

Striking the Right Balance with Uncertainty

Salman Khan* Munawar Hayat* Syed Waqas Zamir Jianbing Shen† Ling Shao
Inception Institute of Artificial Intelligence, UAE

firstname.lastname@inceptioniai.org

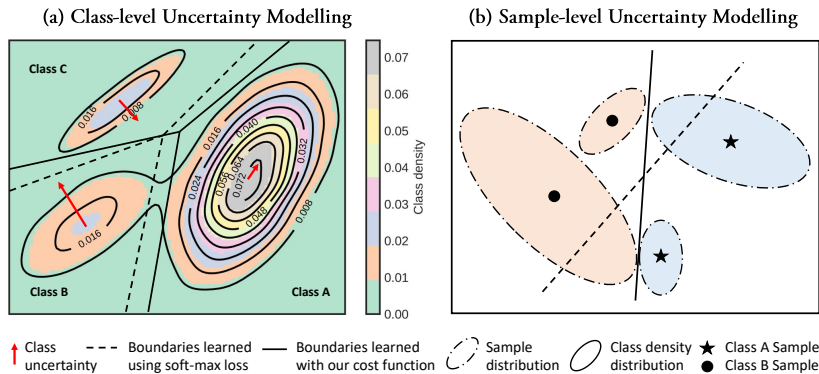


Figure 1: Imbalance learning with Bayesian uncertainty estimates. (a) We enforce class-level margin penalty based on class uncertainty. This pushes boundaries further away from rare classes B and C. (b) We also consider sample-level uncertainty that is modeled as a Gaussian distribution. The learned margins consider the confidence level of classifier to re-adjust boundaries (i.e., provide more room to uncertain samples). This improves the generalization ability of the proposed model for imbalanced classes.

Abstract

Learning unbiased models on imbalanced datasets is a significant challenge. Rare classes tend to get a concentrated representation in the classification space which hampers the generalization of learned boundaries to new test examples. In this paper, we demonstrate that the Bayesian uncertainty estimates directly correlate with the rarity of classes and the difficulty level of individual samples. Subsequently, we present a novel framework for uncertainty based class imbalance learning that follows two key insights: First, classification boundaries should be extended further away from a more uncertain (rare) class to avoid over-fitting and enhance its generalization. Second, each sample should be modeled as a multi-variate Gaussian distribution with a mean vector and a covariance matrix defined by the sample’s uncertainty. The learned boundaries should respect not only the individual samples but also their distribution in the feature space. Our proposed approach efficiently utilizes sample and class uncertainty information to learn robust features and more generalizable classifiers. We systematically study the class imbalance problem and derive a novel loss formulation for max-margin learning based on Bayesian uncertainty measure. The proposed method shows significant performance improvements on six benchmark datasets for face verification, attribute prediction, digit/object classification and skin lesion detection.

1. Introduction

Objects, events, actions and visual concepts appear with varying frequencies in real world imagery [38]. This often leads to highly skewed datasets where a few abundant classes outnumber several rare classes in a typical long-tail data distribution. The low amount of training data for infrequent classes makes it challenging to learn optimal classification boundaries in the feature space. Existing approaches to tackle class imbalance either modify data distribution [42, 7, 22] or introduce appropriate costs to re-weight class errors [25, 1, 40]. The popular data-level approaches are prone to over-fitting while the cost-sensitive learning requires careful choice of weights for successful training. Despite an overwhelming success of deep neural networks on computer vision problems, learning from highly imbalanced sets is still an open problem for deep learning [24].

This paper proposes a new direction towards learning balanced representations using deep neural networks (Fig. 1). We use a principled approach to integrate Bayesian uncertainty estimates for class imbalance learning at two distinct levels, i.e., category-level and individual sample-level. Our approach is based on the observation that rare classes have higher uncertainty in the prediction space and the associated classifier confidence levels are low. Therefore, the uncertainty estimates can be used to expand decision regions for less frequent classes so that classifier’s generalization to new test examples is improved. This concept is illustrated in Fig. 1. Since all samples within a class do

*Equal contribution, †Corresponding author

not have a uniform difficulty level, our approach also optimizes margins with respect to the uncertainty associated with individual samples. The basic intuition for both cases is the same: a classifier should assign larger regions to more uncertain (rare) samples/classes.

Related work: State-of-the-art deep imbalance learning methods mainly propose novel objective functions [17]. Khan *et al.* [25] presented a cost-sensitive loss for CNNs where class-specific weights were automatically learned. Huang *et al.* [21] suggested a combination of triplet and quintuplet losses to preserve local class structures. A modified softmax was proposed in [31] to maximize the angular margin, thus avoiding class imbalance. Quite recently, a meta-learning approach in [40] used selective instances for training on imbalanced sets. These methods have their respective limitations, e.g., [25] only considers class-level costs, [21] is not differentiable and requires heavy pre-processing for quintuplet creation, [31] can only maximize margin on the hypersphere surface and [40] used an additional validation set to assign sample weights. Concurrent to this work, [10] re-weights the loss by the inverse effective number of samples to learn balanced representations.

Contributions: Our approach is distinct in two ways: (a) this is the first work to link class imbalance with Bayesian uncertainty estimates [16], that have shown great promise on other tasks [23, 15, 52], and (b) we incorporate both class and sample-level confidence estimates to appropriately reshape learned boundaries. The paper therefore introduces the following major novelties. **(1)** A principled margin-enforcing formulation for softmax loss, underpinned by the Bayesian uncertainty estimates. **(2)** Sample modeling using multi-variate Gaussian distributions. The class boundaries are optimized to respect second order moments which improves generalization. **(3)** A fully differentiable loss formulation that can be easily plugged into existing architectures and used alongside other regularization techniques.

2. The Imbalance Problem

We begin with an in depth analysis of the imbalance problem and draw several insights which lead to our proposed framework. We base our analysis around softmax loss, which is the most popular objective function for classification. For brevity, we consider a simplistic case of binary classification with two classes A and B in the training set denoted by image-label pairs: $\mathcal{D} = \{\mathbf{x}^k, y^k\}_{k=1}^K$. The goal is to learn an optimal set of ‘representative vectors’ ($\mathbf{w}_A, \mathbf{w}_B$ for classes A and B , respectively) that lead to minimal empirical loss on set \mathcal{D} . The Class-Representative Vectors (CRV) define a loss-minimizing hyper-plane ‘ \mathbf{w} ’, which is the boundary between two classes, i.e., given a feature projection ‘ \mathbf{f} ’ corresponding to an input image \mathbf{x} , $\mathbf{f} \in \mathbf{w}$ iff $(\mathbf{w}_A - \mathbf{w}_B)\mathbf{f} = 0$ (ignoring unit biases). The class imbalance problem exists when the class frequencies

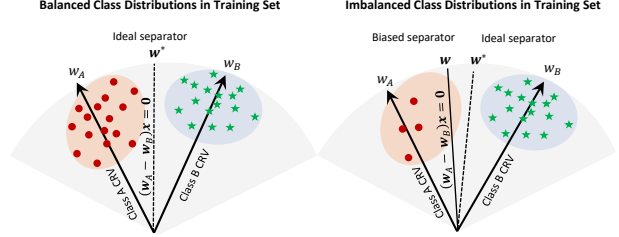


Figure 2: Illustration for the class imbalance problem. True class distributions are shown in green and red. Unbalanced distributions lead to a skewed classification boundary that is biased towards the minority class.

τ_A, τ_B are greatly mismatched in the set \mathcal{D} . As illustrated in Fig. 2, in such cases, the hypothesis (\mathbf{w}) learned on \mathcal{D} using a softmax loss can be biased towards the minority class and significantly different from the ideal separator (\mathbf{w}^*). Next, we breakdown the imbalance problem and explain underlying reasons.

2.1. Bias due to Empirical Loss Minimization

We consider \mathbf{w}^* to be an optimal boundary obtained by loss minimization with respect to the actual hidden distributions P_A and P_B of classes, i.e.:

$$\mathbf{w}^* = \arg \min_{\mathbf{w}} \mathcal{L}_P(\mathbf{w}), \text{ where,} \quad (1)$$

$$\mathcal{L}_P(\mathbf{w}) = \int_{\mathcal{R}_B^{\mathbf{w}}} P_A(\mathbf{f})d\mathbf{f} + \int_{\mathcal{R}_A^{\mathbf{w}}} P_B(\mathbf{f})d\mathbf{f},$$

and $\mathcal{R}_A^{\mathbf{w}}, \mathcal{R}_B^{\mathbf{w}}$ denote the classification regions for classes A and B , respectively. Given \mathcal{D} , the empirical loss calculated on the training set is:

$$\mathcal{L}_{\mathcal{D}}(\mathbf{w}) = \#\{\mathbf{x}^k | \mathbf{f}^k \in \mathcal{D}^A \wedge \mathbf{f}^k \in \mathcal{R}_B^{\mathbf{w}}\} + \#\{\mathbf{x}^k | \mathbf{f}^k \in \mathcal{D}^B \wedge \mathbf{f}^k \in \mathcal{R}_A^{\mathbf{w}}\} \quad (2)$$

Further, assume that the normalized class frequencies τ_A and τ_B are related as $\tau_A + \tau_B = 1$. Then, the expected empirical loss for any hypothesis \mathbf{w} is:

$$\mathbb{E}[\mathcal{L}_{\mathcal{D}}(\mathbf{w})] = \tau_A \int_{\mathcal{R}_B^{\mathbf{w}}} P_A(\mathbf{f})d\mathbf{f} + \tau_B \int_{\mathcal{R}_A^{\mathbf{w}}} P_B(\mathbf{f})d\mathbf{f}. \quad (3)$$

Note that $\tau_A \neq \tau_B \neq 0.5$ due to class imbalance and typically $|\tau_A - \tau_B| > 0.5$ in practical cases where a significant imbalance ratio exists. Next, we show that when large imbalance exists, the learned classification boundaries are biased towards minority classes.

Theorem 1. *For high imbalanced ratios, minimization of empirical loss results in a hypothesis $\hat{\mathbf{w}}$ that is highly likely to be biased towards the minority class ‘ z ’ such that $\mathcal{R}_z^{\hat{\mathbf{w}}} > \mathcal{R}_z^{\mathbf{w}^*}$. In other words, the classification region induced by the optimal separator is larger than the one induced by empirically learned boundary.*

Proof. According to Eq. 3, due to the imbalanced proportion among classes, \mathbf{w}^* is more likely to incur higher empirical error than an alternate hypothesis based on an empirical loss, i.e., for any $\hat{\mathbf{w}} : \mathcal{R}_{\hat{\mathbf{w}}}^z > \mathcal{R}_{\mathbf{w}^*}^z$, it is more likely that:

$$\begin{aligned} \mathcal{L}_{\mathcal{D}}(\mathbf{w}^*) &> \mathcal{L}_{\mathcal{D}}(\hat{\mathbf{w}}) \text{ because } \tau_A \int_{\mathcal{R}_{\mathbf{w}^*}^B} P_A(\mathbf{f})d\mathbf{f} + \\ \tau_B \int_{\mathcal{R}_{\mathbf{w}^*}^A} P_B(\mathbf{f})d\mathbf{f} &> \tau_A \int_{\mathcal{R}_{\hat{\mathbf{w}}}^B} P_A(\mathbf{f})d\mathbf{f} + \tau_B \int_{\mathcal{R}_{\hat{\mathbf{w}}}^A} P_B(\mathbf{f})d\mathbf{f}. \end{aligned} \quad (4)$$

Then, for a significant imbalance ratio such that $\tau_z \ll \tau_{\text{rest}}$, it directly follows that $\mathcal{R}_{z}^{\mathbf{w}^*} > \mathcal{R}_{z}^{\hat{\mathbf{w}}}$. Intuitively, this is a natural implication of imbalanced class distribution which forces the classifier to shift $\hat{\mathbf{w}}$ closer to minority classes because it reduces empirical error. The likelihood of classifier bias is directly proportional to the imbalance rate. \square

A common strategy to tackle data imbalance is through the introduction of cost-sensitive loss functions [25]. We briefly elaborate on the effect of these losses next and explain why this solution is sub-optimal.

2.2. Cost-sensitive Loss

From Eq. 4, one simple solution seems to be the introduction of costs to re-weight the minority class errors. Existing cost-sensitive losses (particularly those based on deep-networks [25]) adopt this idea and assign score-level penalty to the minority class predictions. This means that the classifier is forced to correctly classify training samples belonging to minority classes. This approach has certain limitations. **(1)** Appropriately tuning class specific costs is a challenging task as it requires domain-knowledge with costs usually fixed at the beginning and not dynamically changed during the course of training. **(2)** A more stringent caveat is that such costs do not affect the learned boundaries $\hat{\mathbf{w}}$ if the training samples are separable [53, 25]. Further, when the classes are non-separable, minority class representation in the dataset is directly proportional to its mis-classification probability, i.e.,: $\tau_z \propto \int_{\mathcal{R}_{rest}^{\hat{\mathbf{w}}}} P_z(\mathbf{f})d\mathbf{f}$. **(3)** Generally, these costs are not applied at test time and therefore class-boundaries are effectively unchanged. In summary, while this practice enforces the classifier to more accurately classify training samples belonging to minority classes, it does not enhance generalization capability of the learned model. This can be understood from the relation for generalization error and Fig. 2. For an empirical distribution Q for classes A and B , the generalization error (expected loss on the test set \mathcal{T}) is:

$$\mathbb{E}[\mathcal{L}_{\mathcal{T}}(\mathbf{w})] = \tau'_A \int_{\mathcal{R}_{\mathbf{w}}^B} Q_A(\mathbf{f})d\mathbf{f} + \tau'_B \int_{\mathcal{R}_{\mathbf{w}}^A} Q_B(\mathbf{f})d\mathbf{f}, \quad (5)$$

where τ'_A, τ'_B are the normalized frequencies on the test set. The paper aims to overcome these existing limitations by

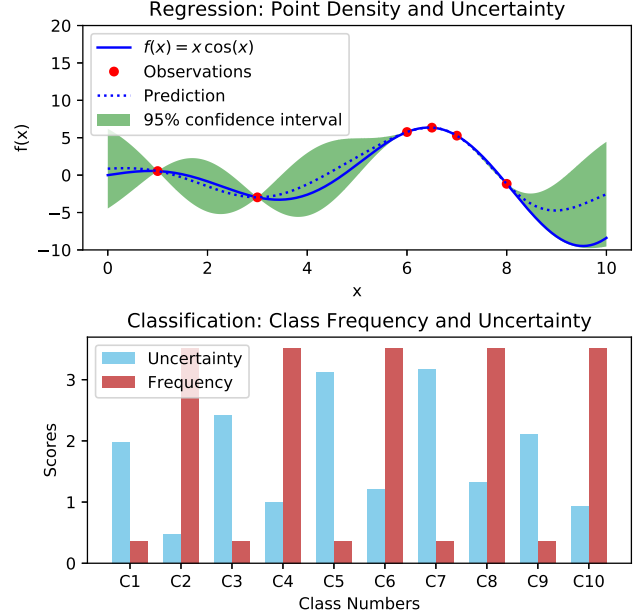


Figure 3: *Top:* One-dimensional Gaussian process regression using maximum-likelihood estimation. The lack of observations results in higher confidence levels. *Bottom:* The uncertainty estimates for imbalanced CIFAR-10 dataset. The uncertainty is higher for classes with less representation.

proposing a new loss formulation that seeks to simultaneously extend minority class boundaries and enforce margin constraints on less represented classes to achieve better generalization performance. We provide details of our technique in the next section.

3. Bayesian Uncertainty Estimates

Bayesian models can provide uncertainty estimates alongside output predictions. Given an input, the uncertainty estimates correspond to the confidence level for each outcome predicted by the model. We hypothesize that the confidence-level of predictions is directly related to the class representation in the training set. As illustrated in Fig. 3, under-represented classes in the training set lead to higher uncertainty and bigger confidence intervals. In contrast, well-represented classes are associated with less uncertainty and compact confidence intervals.

We use deep CNNs with dropout to obtain Bayesian uncertainty estimates. It has been proved that dropout-based deep networks provide an approximation to Gaussian process [16]. A Gaussian process is a Bayesian technique because it constructs a prior distribution over a family of functions \mathcal{F} [39]. This distribution is updated conditioned on observations, i.e., all the functions that are consistent with the labels are retained. At inference time, an output is ob-

tained from each of the functions and expectation is computed to generate the final prediction. The variance of these outputs gives an uncertainty estimate. In the following, we first provide an overview of dropout and then describe uncertainty computation using dropout.

Dropout: Dropout was originally proposed as a regularization measure for deep neural networks [46]. During training, a sub-network is sampled from the full network by randomly dropping a set of neurons. In this manner, each neuron is activated with a fixed probability ‘ p ’. At test time, full model is used for prediction and the output activations are multiplied with the probability p to obtain expectation. Suppose the network parameters are denoted by $\Theta = \{\theta_1, \dots, \theta_L\}$ for a total of L network layers. Then, by applying masks \mathbf{m} generated using i.i.d binary distributions, we can obtain N samples all corresponding to different network configurations $\hat{\Theta}$ that form an ensemble \mathcal{M} :

$$\mathcal{M} = \{\hat{\Theta}^i : i \in [1, N]\}, \text{ where, } \hat{\Theta}^i = \Theta \circ \mathbf{m}^i, \\ \mathbf{m}^i = \{m_l : l \in [1, L]\}, \text{ s.t., } m_l \sim \text{Bernoulli}(p) \quad (6)$$

Uncertainty: For each input \mathbf{x}_k , N model configurations are applied to obtain a set of outputs $\{\hat{\mathbf{y}}\}$. The expected output is calculated using Monte Carlo estimate for first moment ($\mathbb{E}_{q(\mathbf{y}|\mathbf{x})}[\mathbf{y}]$): $\mathbf{y} \approx \frac{1}{N} \sum_{i=1}^N \hat{\mathbf{y}}(\mathbf{x}; \hat{\Theta}^i)$, where q denotes an output distribution that approximates the intractable posterior distribution of deep Gaussian process. The uncertainty is calculated using the second moment ($\mathbb{V}_{q(\mathbf{y}|\mathbf{x})}[\mathbf{y}]$) through Monte Carlo estimation:

$$\mathbf{u} \approx \tau^{-1} \mathbf{I}_D + \frac{1}{N} \sum_{i=1}^N \hat{\mathbf{y}}^T \hat{\mathbf{y}} - \mathbb{E}_{q(\mathbf{y}|\mathbf{x})}[\mathbf{y}]^T \mathbb{E}_{q(\mathbf{y}|\mathbf{x})}[\mathbf{y}], \quad (7)$$

where τ is the model precision (a function of weight decay) and $\mathbf{I}_C \in \mathbb{R}^{C \times C}$ is an identity matrix where C denotes the number of classes.

4. Uncertainty based Max-margin Learning

The softmax loss can be computed for a given feature \mathbf{f} and its true class label y as follows:

$$\mathcal{L}_{sm} = -\log \left(\frac{\exp(\mathbf{w}_y^T \mathbf{f})}{\sum_j \exp(\mathbf{w}_j^T \mathbf{f})} \right), \quad (8)$$

where $j \in [1, C]$ (C is the number of classes). In the above loss formulation, we include the last fully connected layer within softmax loss which will be useful for our analysis later on. Further, for the sake of brevity, we do not mention bias in Eq. 8. Note that the dot-product $\mathbf{w}_y^T \mathbf{f}$ can also be expressed as $\mathbf{w}_y^T \mathbf{f} = \|\mathbf{w}_y\| \|\mathbf{f}\| \cos(\alpha_y)$. Therefore, if a class z is rare in the training set, for an input feature belonging to this class, the softmax loss enforces:

$$\|\mathbf{w}_z\| \|\mathbf{f}\| \cos(\alpha_z) > \|\mathbf{w}_{rest}\| \|\mathbf{f}\| \cos(\alpha_{rest}). \quad (9)$$

Intuitively, we would like to impose a large margin on more uncertain classes. Our experiments show that the class uncertainty is inversely proportional to its frequency in the training set, i.e., rare classes are more uncertain (Fig. 3). To improve generalization performance, we can impose a more strict constraint for uncertain classes:

$$\|\mathbf{w}_z\| \|\mathbf{f}\| \cos(m_z \alpha_z) > \|\mathbf{w}_{rest}\| \|\mathbf{f}\| \cos(m_{rest} \alpha_{rest}),$$

where, $m = \max(1, \lfloor 0.5 \mathbf{u}_y \rfloor)$, $\mathbf{u}_z > \mathbf{u}_{rest}$, $0 \leq \alpha_z \leq \frac{\pi}{\mathbf{u}_z}$ and $\mathbf{u}_z \in \mathbb{R}^+$. This implies that the classifier will try to separate rare classes by a more rigorous margin. The margin maximizing softmax loss [32] is defined as:

$$\mathcal{L}'_{sm} = -\log \left(\frac{\exp(\|\mathbf{w}_y\| \|\mathbf{f}\| \psi(\alpha_y))}{\sum_j \exp(\|\mathbf{w}_j\| \|\mathbf{f}\| \psi(\alpha_j))} \right), \quad (10)$$

where $\psi(\cdot)$ is a continuous and monotonically decreasing function in the range $[0, \pi]$:

$$\psi(\alpha_j) = \begin{cases} (-1)^r \cos(m\alpha_j) - 2r & \alpha_j = y \in \left[\frac{r\pi}{m}, \frac{(r+1)\pi}{m} \right] \\ \cos(\alpha_j) & j \neq y, \end{cases}$$

where $r \in [0, m-1]$ is an integer. The gradient back-propagation requires relations in terms of \mathbf{w} and \mathbf{f} , therefore we substitute $\cos(m\alpha_j)$ with its expansion in terms of Chebyshev polynomials of the first kind (T_m), i.e.,

$$\cos(m\alpha_j) = \sum_{t=0}^{\lfloor m/2 \rfloor} \binom{m}{2t} (\cos^2(\alpha_j) - 1)^t \cos(\alpha_j)^{m-2t}.$$

Here, $\cos(\alpha_j)$ is substituted with $\frac{\mathbf{w}_j^T \mathbf{f}}{\|\mathbf{w}_j\| \|\mathbf{f}\|}$. This gives us the max-margin formulation in terms of differentiable relations.

5. Sample-level Uncertainty Modeling

Although uncertainty driven class-level margin enforcement is important, not all samples in a class have equal difficulty level. The samples that can potentially be misclassified have larger uncertainties. We therefore propose a mechanism to incorporate sample-level uncertainty for imbalanced learning. Existing classification networks only use the mean representation from a distribution of samples to represent each training example. Inspired by [51], we propose to represent a single sample as a function of its first and second order moments. To this end, consider that the deep feature representations of input media is randomly sampled from a multi-variate Gaussian distribution: $\mathbf{f} \sim \mathcal{N}(\mu_f, \Sigma_f)$, where μ_f and Σ_f , respectively denote mean and covariance of the features. The softmax loss can be computed for a given feature \mathbf{f} and its true class label y using Eq. 8.

For an input feature to be correctly classified, its projection on the true class vector should give a maximum response (Eq. 9). In contrast, an example is classified into a

wrong category k if the following condition holds true:

$$\exists j \in [1, C] \text{ s.t., } \mathbf{w}_j^T \mathbf{f} < \mathbf{w}_y^T \mathbf{f}, \quad j \neq y. \quad (11)$$

We are interested in quantifying the probability of misclassification taking into account the distribution of each sample. It can provide a measure of confidence for loss estimates computed on the training samples. Direct computation of the misclassification probability using softmax loss in Eq. 8 is intractable and can only be approximated. Therefore, we introduce a simpler error function that models the essential loss behavior.

$$\mathcal{E}(\mathbf{f}) = \mathbf{w}_j^T \mathbf{f} - \mathbf{w}_y^T \mathbf{f}, \quad j \neq y. \quad (12)$$

The formulation can be used to exactly compute the misclassification probability as we will show next. Since the error function is a linear transformation of input feature $\mathbf{f} \sim \mathcal{N}(\mu_f, \Sigma_f)$, the resulting error distribution is also a uni-variate Gaussian variable. The first and second order statistics of the error distribution can be given in terms of μ_f, Σ_f as follows:

$$\begin{aligned} \mu_{\mathcal{E}} &= \mathbb{E}[\mathcal{E}(\mathbf{f})] = (\mathbf{w}_j - \mathbf{w}_y)^T \mu_f \\ \sigma_{\mathcal{E}}^2 &= \mathbb{E}[(\mathcal{E}(\mathbf{f}) - \mu_{\mathcal{E}})^2] = (\mathbf{w}_j - \mathbf{w}_y)^T \Sigma_f (\mathbf{w}_j - \mathbf{w}_y) \end{aligned} \quad (13)$$

Now, the misclassification probability for a feature \mathbf{f} can be linked with error distribution because $\mathcal{E}(\mathbf{f}) > 0$ denotes a misclassification. The complementary cumulative probability distribution function (CCDF) $\hat{\mathbb{F}}_{\mathcal{E}}$ is given as follows:

$$\begin{aligned} \hat{\mathbb{F}}_{\mathcal{E}}(0) &= \mathbb{P}(\mathcal{E}(\mathbf{f}) > 0) = 1 - \mathbb{P}(\mathcal{E}(\mathbf{f}) < 0) \\ &= 1 - \mathbb{P}\left(\frac{\mathcal{E}(\mathbf{f}) - \mu_{\mathcal{E}}}{\sigma_{\mathcal{E}}} < -\frac{\mu_{\mathcal{E}}}{\sigma_{\mathcal{E}}}\right) \\ \text{if } \mathbf{z} &= \frac{\mathcal{E}(\mathbf{f}) - \mu_{\mathcal{E}}}{\sigma_{\mathcal{E}}} \sim \mathcal{N}(0, 1) \\ \hat{\mathbb{F}}_{\mathcal{E}}(0) &= 1 - \mathbb{F}_{\mathbf{z}}\left(-\frac{\mu_{\mathcal{E}}}{\sigma_{\mathcal{E}}}\right) = \frac{1}{2} \left(1 + \operatorname{erf}\left[\frac{\mu_{\mathcal{E}}}{\sqrt{2}\sigma_{\mathcal{E}}}\right]\right), \end{aligned} \quad (14)$$

where $\mathbb{F}_{\mathbf{z}}$ denotes the cumulative probability distribution function (CDF). The probability estimates are then used to re-weight the loss values such that uncertainty is incorporated. The function $\psi(\cdot)$ is modified as follows to obtain an improved loss function in Eq. 10:

$$\psi(\alpha_j) = \begin{cases} \hat{\mathbb{F}}_{\mathcal{E}}(0)((-1)^r \cos(m\alpha_j) - 2r) & \alpha_{j=y} \in \left[\frac{r\pi}{m}, \frac{(r+1)\pi}{m}\right] \\ \cos(\alpha_j) & j \neq y \end{cases} \quad (15)$$

The loss defined above enforces a margin (m) between output predictions weighted by the cumulative probability ($\hat{\mathbb{F}}_{\mathcal{E}}(0)$). A higher uncertainty means a stricter margin based penalty for a class j . The modified loss function becomes equal to the original softmax loss when the uncertainty is zero and $m = 1$.

6. Experiments

6.1. Datasets

Face Recognition: Facial recognition datasets commonly exhibit large-imbalance which poses a significant challenge for classifier learning. Following [12], we use VGG2 dataset [5] with 3,141,890 images of 8,631 subjects to train our deep network. We evaluate the trained model on four large-scale datasets namely Labeled Faces in the Wild (LFW) [27] and YouTube Faces (YTF) [57], AgeDB [37] and Celebrities in Frontal Profile (CFP) [41]. LFW [27] contains 13,233 web-collected images belonging to 5,749 different identities, with large variations in pose, expression and illumination. We follow the standard protocol of ‘unrestricted with labeled outside data’. YTF [57] has 3,425 sequences of 195 subjects. We follow the standard evaluation protocol on 5,000 video pairs. AgeDB [37] dataset has 12,240 images of 440 subjects. The test set is divided into four groups with different year gaps (5, 10, 20 and 30 years). We only report the performance on the most challenging subset, AgeDB-30. CFP [41] has 500 subjects in total, each with 10 frontal and 4 profile images. In this paper, we only evaluate on the most challenging subset CFP-FP.

Skin Lesion Classification: Edinburgh Dermofit Image Library (DIL) consists of 1,300 high quality skin lesion images based on diagnosis from dermatologists and dermatopathologists. There are 10 types of lesions identified in this dataset including melanomas, seborrheic keratosis and basal cell carcinomas. The number of images in each category varies between 24 and 331 (mean 130, median 83). Similar to [3], we report results with 3-fold cross validation.

Digit/Object Classification: We evaluate on imbalanced MNIST and CIFAR-10 datasets for generic digit and object classification. Standard MNIST consists of 70,000 images of handwritten digits (0-9). Out of these, 60,000 images are used for training (600/class) and the remaining 10,000 for testing (100/class). CIFAR-10 contains 60,000 images belonging to 10 classes (6,000 images/class). The standard train/test split for each class is $\sim 83.3\%/16.7\%$ images. We evaluate our approach on the standard split as well as on an artificially created imbalanced split. To imbalance the training distribution, we randomly drop 90% of the samples for half of the classes.

Attribute Prediction: We use the large-scale CelebA dataset [33] for (multi-label) facial attribute prediction task. This dataset consists of 202,599 images belonging to 10,177 human identities. Each image is annotated with a diverse set of 40 binary attributes. There exists a significant imbalance in the training set with ratios up to 1:43. Following the standard protocol [33], we use 152,770 images for training, 10,000 for validation, and remaining 19,867 for testing. For evaluation, we report *Balanced Classification Accuracy* (BCA) defined as: $BCA = 0.5 \times \frac{t_p}{N_p} + 0.5 \times \frac{t_n}{N_n}$, where t_p

and t_n , respectively denote true positives and true negatives, and N_p and N_n are total number of positive and negative samples. This evaluation metric is more suitable for multi-label imbalanced learning tasks since it gives equal weight to both majority and minority classes. Other evaluation metrics used in the literature [33] which define accuracy as $\frac{t_p+t_n}{N_p+N_n}$ can be biased towards majority classes.

6.1.1 Implementation Details

The uncertainty estimates are applied progressively during training. We start with standard softmax ($m = 1$), followed by class-level uncertainty based max-margin learning and finally sample-level uncertainty modeling during the last 10 epochs. The compute intensive sample-level uncertainty estimates are therefore only done for few epochs. The proposed strategy can be related with curriculum learning, since it starts with a *simple* task by considering a balanced class distribution, and gradually introduces *harder* tasks by expanding or shrinking classification boundaries of different classes based upon their uncertainty estimates. For attribute prediction on CelebA dataset, the training times required for standard softmax and ours are ~ 3.4 and 4.6 hours, respectively, on a Dell Precision 7920 machine with TitanXp GPU. In our experiments, we fixed $m = 3$ since it gives relatively better results. Experiments on imbalanced CIFAR-10 for $m = 2, 3, 4$ achieve an accuracy of 80.2%, 80.6%, 80.5%, respectively. Values of $N \geq 5$ give stable uncertainty estimates. We fixed $N = 10$ for the optimal trade-off between reliable uncertainty and compute efficiency. An ablation study on different values of $N = 5, 10, 20, 40$ on imbalanced CIFAR-10 results in respective accuracies of 80.4%, 80.6%, 80.6%, 80.7%.

For Skin Lesion detection, we deploy ResNet-18 backbone with two fully connected layers (with intermediate rectified linear units non-linearities and dropout) inserted after the global pooling layer. For face verification tasks, we train Squeeze and Excitation (SE) networks [20] with ResNet-50 backbone. The face images are pre-processed to 112×112 using multi-task cascaded CNN [61]. After the network is trained on VGG-2 dataset, we use features extracted after global pooling layer for face verification evaluations. On imbalanced MNIST dataset, we use the same settings as in [25] to enable direct comparison with the recently proposed imbalanced learning technique [25]. For experiments on imbalanced CIFAR-10 dataset, we extract features from VGG16 [44] pre-trained on ImageNet [11]. A simple neural network with two hidden layers (512 neurons each) with dropouts is trained on the extracted features to get uncertainty estimates and perform classification. Training a network on VGG extracted features enables us to compare against traditional imbalanced learning techniques. Specifically, data level under-sampling & over-sampling methods

	Methods	LFW	YTF
Methods using non-public data	DeepFace [49]	97.35	91.4
	FaceNet [43]	99.63	95.4
	Web-scale [50]	98.37	-
	DeepID2+ [48]	99.47	93.2
	Baidu [30]	99.13	-
	Center Face [56]	99.28	94.9
	Marginal Loss [13]	99.48	96.0
	Noisy Softmax [8]	99.18	94.9
Novel Loss Functions	Softmax+Contrastive [47]	98.78	93.5
	Triplet Loss [43]	98.70	93.4
	Large-Margin Softmax [32]	99.10	94.0
	Center Loss [56]	99.05	94.4
	SphereFace [31]	99.42	95.0
	CosFace [55]	99.33	96.1
Imbalance Learning	Range Loss [63]	99.52	93.7
	Augmentation [36]	98.06	-
	Center Inv. Loss [58]	99.12	93.9
	Feature transfer [60]	99.37	-
	LMLE [21]	99.51	95.8
	This Paper	99.71	97.3

Table 1: Face Verification Performance on LFW and YTF datasets. We trained our model on VGG2 dataset. Most methods in the first cell use large-scale outside data that are not publicly available. The second cell includes novel loss functions. The state-of-the-art imbalanced learning methods are in the last group.

in Table 5 are used on VGG features, followed by the two layered NN for classification. For attribute prediction on CelebA dataset, we train a model with ResNet-50 backbone and two fully connected layers with dropout inserted after the global pooling and the first fully connected layers. The model is trained to minimize sum of binary cross entropy losses, using relatively smaller learning rates for the layers before the global pooling layer and larger rates for the layers inserted afterwards.

6.2. Results and Comparisons

Face Verification: We compare our approach with 20 recent and top-performing methods on LFW and YTF datasets (Table 1). We divide these methods into three categories: (a) methods that use large amounts of non-publicly available data sources to train their models, (b) methods that design novel loss functions for face verification and (c) methods that deal with data imbalance. We note that the performances on both LFW and YTF are currently saturated with many recent methods already surpassing human performance. Our method achieves competitive performance on both these datasets. Note that some of the compared methods used as much as 200M images [43] and an ensemble of 25 models [48] for training. Further evaluations on additional datasets show that the proposed method achieves verification accuracies of 97.0% and 94.4% on CFP-FP and

Methods	Imbalanced Split	Performances	
		Exp. 1 (5-classes)	Exp. 2 (10-classes)
Hierarchical-KNN [3]	✓	74.3 ± 2.5%	68.8 ± 2.0%
Hierarchical-Bayes [2]	✓	69.6 ± 0.4%	63.1 ± 0.6%
Flat-KNN [3]	✓	69.8 ± 1.6%	64.0 ± 1.3%
Baseline CNN [25]	✓	75.2 ± 2.7%	69.5 ± 2.3%
CoSen CNN [25]	✓	80.2 ± 2.5%	72.6 ± 1.6%
This paper	✓	95.7 ± 1.2%	86.9 ± 0.7%

Table 2: Experimental results for Skin Lesion Classification on the DIL dataset.

Methods	Imbalanced Split	Performances
Deeply Supervised Nets [29]	✗	99.6%
Generalized Pooling Func. [28]	✗	99.7%
Maxout NIN [6]	✗	99.8%
Baseline CNN [25]	✓	97.1%
CoSen CNN [25]	✓	98.4%
This paper	✓	98.7%

Table 3: Results for digit classification on MNIST.

AgeDB-30 datasets, respectively.

Skin Lesion Detection: The results for skin lesion detection are listed in Table 2. We perform a 3-fold cross validation to report the results. Compared to a recent cost-sensitive CNN approach [25], we obtain an impressive absolute performance gain of 15.5% and 14.3% on Experiments 1 (5 classes) and 2 (10 classes), respectively.

Digit Classification: For hand-written digit classification, we report our results on an imbalanced split of MNIST dataset in Table 3. For the sake of comparison, we also report some representative methods on the original balanced split of MNIST. However, the two set of techniques are not directly comparable since the bottom group uses $\sim 45\%$ less training data. Our technique outperforms other imbalance learning approaches.

Attribute Prediction: We report the multi-label attribute prediction results on CelebA dataset in Table 4. This dataset is particularly challenging as it exhibits a high imbalance with majority-to-minority ratio up to 1:43. We compare with nine recent state-of-the-art methods. These methods include techniques that specifically focus on class imbalance learning (Table 4, right block). Our approach performs significantly better compared to both normal and class imbalance learning methods. Specifically, we achieve top-most accuracy in 23/40 attributes and second-best accuracy in other 8/40 classes. In particular, since our method focuses on assigning a larger classification region to under-represented classes, we achieve more pronounced boost for the case of highly imbalanced classes. For example, out of the top 50% most imbalanced classes (with fewer samples), we achieve best performance in 16/20 (80%) cases. Qualitative examples for attribute prediction are shown in Fig. 4.

Object Classification: We compare against a number of popular imbalanced learning approaches including both

Attributes	Imbalance level	Triplet-kNN [43]	Panda [62]	ANet [34]	DeepID2 [47]	Over-sampling [14]		Under-sampling [14]		Threshold Adj. [9]	Cost-sensitive [19]	LMLE [21]	This paper
						77	78	69	78				
Attractive	1	83	85	87	78	77	78	69	78	88	77		
Mouth Open	2	92	93	96	89	89	87	89	89	96	93		
Smiling	2	92	98	97	89	90	90	88	90	99	90		
Wear Lipstick	3	91	97	95	92	92	91	89	91	99	94		
High Cheekbones	5	86	89	89	84	84	80	83	85	92	83		
Male	8	91	99	99	94	95	90	95	93	99	98		
Heavy Makeup	11	88	95	96	88	87	89	89	89	98	91		
Wavy Hair	18	77	78	81	73	70	70	77	75	83	84		
Oval Face	22	61	66	67	63	63	58	72	64	68	67		
Pointy Nose	22	61	67	69	66	67	63	72	65	72	72		
Arched Eyebrows	23	73	77	76	77	79	70	76	78	79	83		
Black Hair	26	82	84	90	83	84	80	86	85	92	90		
Big Lips	26	55	56	57	62	61	61	66	61	60	69		
Big Nose	27	68	72	78	73	73	76	76	74	80	79		
Young	28	75	78	84	76	75	80	24	75	87	72		
Straight Hair	29	63	66	69	65	66	61	73	67	73	81		
Brown Hair	30	76	85	83	79	82	76	81	84	87	86		
Bags Under Eyes	30	63	67	70	74	73	71	76	74	73	83		
Wear Earrings	31	69	77	83	75	76	70	76	76	83	89		
No Beard	33	82	87	93	88	88	88	15	88	96	85		
Bangs	35	81	92	90	91	90	88	93	90	98	95		
Blond Hair	35	81	91	90	90	90	85	92	89	99	95		
Bushy Eyebrows	36	68	74	82	78	80	75	84	79	82	87		
Wear Necklace	38	50	51	59	70	71	66	62	71	59	76		
Narrow Eyes	38	47	51	57	64	65	61	71	65	59	75		
5 o'clock Shadow	39	66	76	81	85	85	82	82	84	82	91		
Receding Hairline	42	60	67	70	81	82	79	83	81	96	86		
Wear Necktie	43	73	85	79	83	79	80	76	82	90	92		
Eyeglasses	44	82	88	95	92	91	85	95	91	98	99		
Rosy Cheeks	44	64	68	76	86	90	82	82	92	78	89		
Goatee	44	73	84	86	90	89	85	89	86	95	96		
Chubby	44	64	65	70	81	83	78	81	82	79	87		
Sideburns	44	71	81	79	89	90	80	89	90	88	95		
Blurry	45	43	50	56	74	76	68	78	76	59	87		
Wear Hat	45	84	90	90	90	89	90	95	90	99	97		
Double Chin	45	60	64	68	83	84	80	83	84	74	85		
Pale Skin	46	63	69	77	81	82	78	85	80	80	89		
Gray Hair	46	72	79	85	90	90	88	91	90	91	94		
Mustache	46	57	63	61	88	90	60	86	88	73	91		
Bald	48	75	74	73	93	92	79	93	93	90	95		
Overall	-	72	77	80	81	82	78	79	82	84	87		

Table 4: Multi-label attribute prediction results on CelebA dataset. The compared methods are divided into two categories (a) *left*: methods without class imbalance learning and (b) *right*: methods that focus on imbalance learning.

data-level (e.g., SMOTE NN [7], ADASYS NN [18]) and algorithm-level imbalance removal techniques (e.g., cost-sensitive SVM [4], NearMiss [35]), on imbalanced CIFAR-10 dataset (by retaining only 10% of the samples for 50% classes). Table 5 summarizes our results in terms of a di-

Methods	Performances					
	Accuracy	Precision	Recall	F1	G-Mean	IBA
Cost-sen SVM [4]	34.2	60.9±33.6	34.3±32.4	30.5±23.0	69.7±29.4	55.5±30.6
NearMiss [35]	63.5	66.1±14.9	63.5±19.7	62.8±14.2	79.2±9.5	61.8±14.8
SMOTE NN [7]	76.3	80.5±12.6	76.3±22.1	75.0±12.3	88.3±6.4	77.1±11.9
ADASYN NN [18]	76.4	80.4±12.2	76.4±21.1	75.2±10.6	88.2±6.0	76.9±11.4
Under-samp. Clustering [59]	75.7	78.3±12.2	75.7±10.3	75.9±6.9	87.0±6.8	74.8±12.2
Neighborhood Cleaning [26]	76.7	79.2±10.4	76.7±18.6	75.9±8.6	85.7±10.2	73.2±18.1
Instance Hardness [45]	67.8	76.3±19.8	67.8±18.7	67.6±10.8	79.9±11.1	63.4±17.4
This Paper	80.6	80.8±6.1	80.6±9.3	80.4±6.3	88.7±5.2	77.6±9.6

Table 5: Performance on imbalanced CIFAR-10. Standard deviation on class-specific performance is reported for each metric. Our method performs better compared to others on a diverse set of evaluation metrics.



Figure 4: Sample attribute predictions on the CelebA.

Method (↓)	Accuracy		
CNN + Softmax loss	97.2		
Dropout Rate (→)	0.3	0.5	0.7
CNN + UMM loss	98.0	98.3	98.2
CNN + UMM + SUM (Ours)	98.3	98.7	98.7

Table 6: Ablation study on imbalanced MNIST dataset.

verse range of metrics (e.g., accuracy, F1, G-mean, IBA, precision, recall). These metrics provide a comprehensive view on the performances of ours and other imbalanced techniques and are more suitable for imbalanced learning scenarios. Our results show that the proposed uncertainty based technique consistently performs better than the traditional imbalance removal methods. Particularly, for the case of evaluation metrics that give equal importance to rare classes (e.g., recall and F1 measure), our approach achieves significant performance boost of 3.9 and 4.5, respectively. Furthermore, the deviation of individual class scores is much lower compared to other techniques.

Ablation Study: We experiment with different variants of our approach in Table 6. Specifically, we report the performance of a simple baseline model with softmax loss and compare it with the uncertainty-based margin (UMM) enforcement and sample-level uncertainty modeling (SUM). We note a progression in performance from UMM to UMM+SUM. This is because the SUM penalizes hard examples which further helps in improving generalizability. We also study the effect of changing the dropout rate on our uncertainty-based approach. Increasing the dropout rate beyond 0.5 does not help, while a smaller rate of 0.2 or 0.3 results in a performance drop.

Other margin enforcing losses with uncertainty: We also experiment with recent variants of softmax loss that explicitly enforce margin constraints during classification along

Loss Type	Settings		Original	with Uncer.
	n	m		
Softmax (Baseline)	-	-	76.8	-
Additive Margin [54]	30	0.35	78.3	78.9
Arc Margin [12]	30	0.4	78.1	78.4
Large Margin [32]	-	4	79.4	79.9
Sphere Product [31]	-	4	76.2	77.5

Table 7: The behavior of other recent margin-based loss formulations with uncertainty. The accuracy values are reported for imbalanced CIFAR-10 dataset. ‘ n ’, ‘ m ’ stands for feature norm and margin, respectively.

with uncertainty estimates (Table 7). These methods include Additive Margin Softmax [54], Angular (Arc) Margin Softmax [12], Large Margin Softmax [32] and SphereFace [31]. As these loss functions have generally been proposed for face verification, we already compare with them in Table 1. However, here our main goal is to analyze if the uncertainty estimates help in learning better boundaries for the rare classes. To this end, we use exactly the same features for all techniques and use uncertainty estimates in place of their original parameter m settings. Since our application is different from face verification, we note a relatively lower performance from SphereFace and a higher performance from Large Margin Softmax. Overall, the uncertainty scores help in achieving discriminativeness between difficult classes and gives better performance in all cases.

7. Conclusion

We present a new approach to address class imbalance problem, underpinned by the Bayesian uncertainty estimates. We demonstrate that the classifier confidence levels are directly associated with: (a) the difficulty-level of individual samples and (b) the scarcity of the training data for under-represented classes. Our proposed approach utilizes uncertainty to enforce larger classification regions for rare classes and challenging training samples. This results in better generalization of learned classifier to new samples for less frequent classes. We achieve significant performance gains on several datasets for face verification, attribute prediction, object/digit classification and skin lesion detection.

References

- [1] R. Akbani, S. Kwek, and N. Japkowicz. Applying support vector machines to imbalanced datasets. In *European conference on machine learning*, pages 39–50. Springer, 2004.
- [2] L. Ballerini, R. B. Fisher, B. Aldridge, and J. Rees. Non-melanoma skin lesion classification using colour image data in a hierarchical k-nn classifier. In *Biomedical Imaging (ISBI), 2012 9th IEEE International Symposium on*, pages 358–361. IEEE, 2012.
- [3] L. Ballerini, R. B. Fisher, B. Aldridge, and J. Rees. A color and texture based hierarchical k-nn approach to the classification of non-melanoma skin lesions. In *Color Medical Image Analysis*, pages 63–86. Springer, 2013.
- [4] P. Cao, D. Zhao, and O. Zaiane. An optimized cost-sensitive svm for imbalanced data learning. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 280–292. Springer, 2013.
- [5] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman. Vggface2: A dataset for recognising faces across pose and age. In *International Conference on Automatic Face and Gesture Recognition*, 2018.
- [6] J.-R. Chang and Y.-S. Chen. Batch-normalized maxout network in network. *arXiv preprint arXiv:1511.02583*, 2015.
- [7] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer. Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–357, 2002.
- [8] B. Chen, W. Deng, and J. Du. Noisy softmax: Improving the generalization ability of dcnn via postponing the early softmax saturation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5372–5381, 2017.
- [9] J. Chen, C.-A. Tsai, H. Moon, H. Ahn, J. Young, and C.-H. Chen. Decision threshold adjustment in class prediction. *SAR and QSAR in Environmental Research*, 17(3):337–352, 2006.
- [10] Y. Cui, M. Jia, T.-Y. Lin, Y. Song, and S. Belongie. Class-balanced loss based on effective number of samples. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [11] J. Deng, W. Dong, R. Socher, L. Li, K. Li, and L. Fei-Fei. ImageNet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255, 2009.
- [12] J. Deng, J. Guo, and S. Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. *arXiv preprint arXiv:1801.07698*, 2018.
- [13] J. Deng, Y. Zhou, and S. Zafeiriou. Marginal loss for deep face recognition. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 2006–2014. IEEE, 2017.
- [14] C. Drummond and R. C. Holte. C4. 5, class imbalance, and cost sensitivity: Why under-sampling beats over-sampling. *ICML Workshops*, 2003.
- [15] R. Feinman, R. R. Curtin, S. Shintre, and A. B. Gardner. Detecting adversarial samples from artifacts. *arXiv preprint arXiv:1703.00410*, 2017.
- [16] Y. Gal and Z. Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059, 2016.
- [17] M. Hayat, S. Khan, W. Zamir, J. Shen, and L. Shao. Max-margin class imbalanced learning with gaussian affinity. *arXiv preprint arXiv:1901.07711*, 2019.
- [18] H. He, Y. Bai, E. A. Garcia, and S. Li. Adasyn: Adaptive synthetic sampling approach for imbalanced learning. In *Neural Networks, 2008. IJCNN 2008. (IEEE World Congress on Computational Intelligence)*. *IEEE International Joint Conference on*, pages 1322–1328. IEEE, 2008.
- [19] H. He and E. A. Garcia. Learning from imbalanced data. *IEEE Transactions on knowledge and data engineering*, 21(9):1263–1284, 2009.
- [20] J. Hu, L. Shen, and G. Sun. Squeeze-and-excitation networks. *arXiv preprint arXiv:1709.01507*, 7, 2017.
- [21] C. Huang, Y. Li, C. Change Loy, and X. Tang. Learning deep representation for imbalanced classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5375–5384, 2016.
- [22] P. Jeatrakul, K. W. Wong, and C. C. Fung. Classification of imbalanced data by combining the complementary neural network and smote algorithm. In *International Conference on Neural Information Processing*, pages 152–159. Springer, 2010.
- [23] A. Kendall, Y. Gal, and R. Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [24] S. Khan, H. Rahmani, S. A. A. Shah, and M. Bennamoun. A guide to convolutional neural networks for computer vision. *Synthesis Lectures on Computer Vision*, 8(1):1–207, 2018.
- [25] S. H. Khan, M. Hayat, M. Bennamoun, F. Sohel, and R. Togneri. Cost sensitive learning of deep feature representations from imbalanced data. *IEEE Transactions on Neural Networks and Learning Systems*, 2017.
- [26] J. Laurikkala. Improving identification of difficult small classes by balancing class distribution. In *Conference on Artificial Intelligence in Medicine in Europe*, pages 63–66. Springer, 2001.
- [27] G. B. H. E. Learned-Miller. Labeled faces in the wild: Updates and new reporting procedures. Technical Report UM-CS-2014-003, University of Massachusetts, Amherst, May 2014.
- [28] C.-Y. Lee, P. W. Gallagher, and Z. Tu. Generalizing pooling functions in convolutional neural networks: Mixed, gated, and tree. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, pages 464–472, 2016.
- [29] C.-Y. Lee, S. Xie, P. Gallagher, Z. Zhang, and Z. Tu. Deeply-supervised nets. 2015.
- [30] J. Liu, Y. Deng, T. Bai, Z. Wei, and C. Huang. Targeting ultimate accuracy: Face recognition via deep embedding. *arXiv preprint arXiv:1506.07310*, 2015.
- [31] W. Liu, Y. Wen, Z. Yu, M. Li, B. Raj, and L. Song. Sphreface: Deep hypersphere embedding for face recognition.

- [32] W. Liu, Y. Wen, Z. Yu, and M. Yang. Large-margin softmax loss for convolutional neural networks. In *International Conference on Machine Learning*, pages 507–516, 2016.
- [33] Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3730–3738, 2015.
- [34] Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, 2015.
- [35] I. Mani and I. Zhang. knn approach to unbalanced data distributions: a case study involving information extraction. In *Proceedings of workshop on learning from imbalanced datasets*, volume 126, 2003.
- [36] I. Masi, A. T. Trn, T. Hassner, J. T. Leksut, and G. Medioni. Do we really need to collect millions of faces for effective face recognition? In *European Conference on Computer Vision*, pages 579–596. Springer, 2016.
- [37] S. Moschoglou, A. Papaioannou, C. Sagonas, J. Deng, I. Kotisia, and S. Zafeiriou. Agedb: the first manually collected, in-the-wild age database. In *Proceedings of IEEE Intl Conf. on Computer Vision and Pattern Recognition (CVPR-W 2017)*, Honolulu, Hawaii, June 2017.
- [38] S. Rahman, S. Khan, and F. Porikli. Zero-shot object detection: Learning to simultaneously recognize and localize novel concepts. *arXiv preprint arXiv:1803.06049*, 2018.
- [39] C. E. Rasmussen. Gaussian processes in machine learning. In *Advanced lectures on machine learning*, pages 63–71. Springer, 2004.
- [40] M. Ren, W. Zeng, B. Yang, and R. Urtasun. Learning to reweight examples for robust deep learning. In *International Conference on Machine Learning*, 2018.
- [41] C. C. V. P. R. C. D. J. S. Sengupta, J.C. Cheng. Frontal to profile face verification in the wild. In *IEEE Conference on Applications of Computer Vision*, February 2016.
- [42] J. A. Sáez, J. Luengo, J. Stefanowski, and F. Herrera. Smote-tpf: Addressing the noisy and borderline examples problem in imbalanced classification by a re-sampling method with filtering. *Information Sciences*, 291:184–203, 2015.
- [43] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015.
- [44] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations (ICLR)*, 2015.
- [45] M. R. Smith, T. Martinez, and C. Giraud-Carrier. An instance level analysis of data complexity. *Machine learning*, 95(2):225–256, 2014.
- [46] N. Srivastava, G. E. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- [47] Y. Sun, Y. Chen, X. Wang, and X. Tang. Deep learning face representation by joint identification-verification. In *Advances in neural information processing systems*, pages 1988–1996, 2014.
- [48] Y. Sun, X. Wang, and X. Tang. Deeply learned face representations are sparse, selective, and robust. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2892–2900, 2015.
- [49] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf. Deepface: Closing the gap to human-level performance in face verification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1701–1708, 2014.
- [50] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf. Web-scale training for face identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2746–2754, 2015.
- [51] C. Tzelepis. *Maximum Margin Learning Under Uncertainty*. PhD thesis, Queen Mary University of London, 2018.
- [52] C. Tzelepis, V. Mezaris, and I. Patras. Linear maximum margin classifier for learning from uncertain data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- [53] B. C. Wallace, K. Small, C. E. Brodley, and T. A. Trikalinos. Class imbalance, redux. In *Data Mining (ICDM), 2011 IEEE 11th International Conference on*, pages 754–763. IEEE, 2011.
- [54] F. Wang, J. Cheng, W. Liu, and H. Liu. Additive margin softmax for face verification. *IEEE Signal Processing Letters*, 25(7):926–930, 2018.
- [55] H. Wang, Y. Wang, Z. Zhou, X. Ji, D. Gong, J. Zhou, Z. Li, and W. Liu. Cosface: Large margin cosine loss for deep face recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [56] Y. Wen, K. Zhang, Z. Li, and Y. Qiao. A discriminative feature learning approach for deep face recognition. In *European Conference on Computer Vision*, pages 499–515. Springer, 2016.
- [57] L. Wolf, T. Hassner, and I. Maoz. Face recognition in unconstrained videos with matched background similarity. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 529–534. IEEE, 2011.
- [58] Y. Wu, H. Liu, J. Li, and Y. Fu. Deep face recognition with center invariant loss. In *Proceedings of the on Thematic Workshops of ACM Multimedia 2017*, pages 408–414. ACM, 2017.
- [59] S.-J. Yen and Y.-S. Lee. Cluster-based under-sampling approaches for imbalanced data distributions. *Expert Systems with Applications*, 36(3):5718–5727, 2009.
- [60] X. Yin, X. Yu, K. Sohn, X. Liu, and M. Chandraker. Feature transfer learning for deep face recognition with long-tail data. *arXiv preprint arXiv:1803.09014*, 2018.
- [61] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Processing Letters*, 23(10):1499–1503, 2016.
- [62] N. Zhang, M. Paluri, M. Ranzato, T. Darrell, and L. Bourdev. Panda: Pose aligned networks for deep attribute modeling. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1637–1644, 2014.
- [63] X. Zhang, Z. Fang, Y. Wen, Z. Li, and Y. Qiao. Range loss for deep face recognition with long-tailed training data. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5409–5418, 2017.