



HAL
open science

Class Uncertainty: A Measure to Mitigate Class Imbalance

Zeynep Sonat Baltaci, Kemal Oksuz, Selim Kuzucu, Kivanc Tezoren, Berkin Kerim Konar, Alpay Ozkan, Sinan Kalkan, Emre Akbas

► **To cite this version:**

Zeynep Sonat Baltaci, Kemal Oksuz, Selim Kuzucu, Kivanc Tezoren, Berkin Kerim Konar, et al..
Class Uncertainty: A Measure to Mitigate Class Imbalance. 2023. hal-04317284

HAL Id: hal-04317284

<https://hal.science/hal-04317284>

Preprint submitted on 1 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Class Uncertainty: A Measure to Mitigate Class Imbalance

Zeynep Sonat Baltacı^{1,2}, Kemal Oksuz³, Selim Kuzucu¹, Kivanc Tezoren¹, Berkin Kerim Konar¹,
Alpay Ozkan¹, Emre Akbas^{1,4,†}, Sinan Kalkan^{1,4,†}

¹Dept. of Computer Engineering, METU, Ankara, Turkey

²LIGM, Ecole des Ponts, Univ Gustave Eiffel, CNRS, Marne-la-Vallée, France

³Five AI Ltd., United Kingdom

⁴Center for Robotics and Artificial Intelligence (ROMER), METU, Ankara, Turkey

sonat.baltaci@enpc.fr, kemal.oksuz@five.ai,

{selim.kuzucu, kivanc.tezoren, berkin.konar, alpay.ozkan, eakbas, skalkan}@metu.edu.tr

Abstract—Class-wise characteristics of training examples affect the performance of deep classifiers. A well-studied example is when the number of training examples of classes follows a long-tailed distribution, a situation that is likely to yield sub-optimal performance for under-represented classes. This class imbalance problem is conventionally addressed by approaches relying on the class-wise cardinality of training examples, such as data resampling. In this paper, we demonstrate that considering solely the cardinality of classes does not cover all issues causing class imbalance. To measure class imbalance, we propose CLASS UNCERTAINTY as the average predictive uncertainty of the training examples, and we show that this novel measure captures the differences across classes better than cardinality. We also curate SVCI-20 as a novel dataset in which the classes have equal number of training examples but they differ in terms of their hardness; thereby causing a type of class imbalance which cannot be addressed by the approaches relying on cardinality. We incorporate our CLASS UNCERTAINTY measure into a diverse set of ten class imbalance mitigation methods to demonstrate its effectiveness on long-tailed datasets as well as on our SVCI-20. Code and datasets will be made available.

Index Terms—long-tailed visual recognition, class imbalance, class uncertainty, predictive uncertainty

I. INTRODUCTION

The performance of data-driven, learning-based methods relies heavily on dataset characteristics. One characteristic that has received significant attention is the unequal cardinality of examples per class¹ [1], [2], [3], [4], [5], [6], [7], [8], [9], [10], [11], [12], [13], [14], [15]. This attention originates from the fact that the frequencies of object classes follow a long-tailed distribution in nature, which also manifests itself in many visual recognition datasets [16], [17], [18], [19], [20], [21]. Accordingly, long-tailed distribution of cardinalities is generally considered as the major reason for many performance issues encountered on such datasets. This common phenomenon may be more precisely described as “cardinality imbalance”.

While cardinality imbalance is the most evident type of imbalance to observe, it does not completely capture or explain

the spectrum of factors that cause performance gaps among classes. Other contributing factors include hardness of examples or classes [22], [23] and discrepancy between the training objective and the evaluation measure [24], [25]. If such factors, which we will refer to generally as “class imbalance”, are not addressed properly, they lead to sub-optimal performance for certain classes.

In this paper, we challenge the aforementioned convention to rely on the cardinality for alleviating the class imbalance and make the following main contributions:

- For setting an important milestone into the class imbalance literature, we define what a class imbalance measure is and outline important features that it should have.
- We introduce a new measure, CLASS UNCERTAINTY, as the average predictive uncertainty of training examples in a class. We extensively show its effectiveness compared to using cardinality. Specifically, incorporating our measure into a diverse set of ten methods from resampling, loss reweighting, margin adjustment and multi-stage training; we observe better or on-par performance compared to class cardinality on long-tailed datasets as highlighted in Fig. 1(a). As an example, CLASS UNCERTAINTY drastically reduces the top-1 error of the recent LDAM method by 3.5 – 4.0% on the CIFAR-100-LT dataset.
- To facilitate studies with a broader perspective than cardinality, we curate a new dataset, SVCI-20, where the classes have equal cardinality but different hardness.

We present these contributions in the following sections: Section II includes the background and discusses why minimizing the negative log-likelihood is not robust to class imbalance. It also provides a simple taxonomy of the existing imbalance mitigation methods with a brief discussion of such methods. Section III defines what a class imbalance measure is, analyzes using class cardinality as an imbalance measure and proposes our CLASS UNCERTAINTY. Section IV demonstrates how to incorporate our measure into the existing mitigation methods. Section V presents our experimental results for class uncertainty on long-tailed datasets and our novel

Manuscript received March 31, 2023.

[†]Equal contribution for senior authorship.

¹Throughout the paper, we use *cardinality* to refer to the number of training examples for a class.

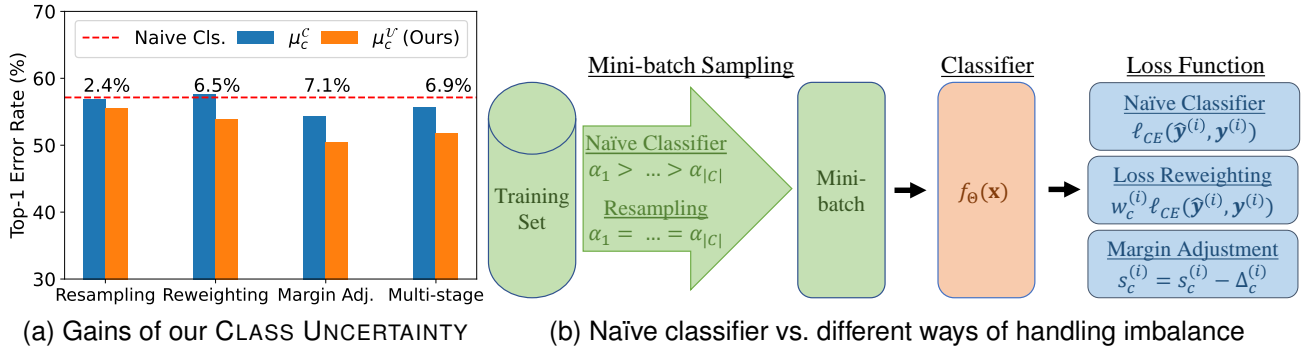


Fig. 1. (a) Using our CLASS UNCERTAINTY improves top-1 error rate of all aforementioned methods (CIFAR-100-LT with an IR of 50 using ResNet-32). In particular, we obtain (i) the sampling probability of each class in resampling methods; (ii) the weights of the loss for each class; (iii) the margins to be enforced around each class; or (iv) again sampling probability in the second-stage of multi-stage training strategy using CLASS UNCERTAINTY. The numbers on top of the histograms show relative gain over cardinality-based methods. Baselines: Progressively-balanced sampling [3], class-balanced Focal Loss [5], LDAM with reweighting [11] and deferred resampling. (b) Naïve classifier is not robust to class imbalance, giving rise to a plethora of mitigation methods: (1) “resampling” samples a balanced set of examples; (2) “reweighting” assigns a weight for each class ($w_c^{(i)}$); (3) “margin-adjustment” methods assign different margins ($\Delta_c^{(i)}$) to the logits ($s_c^{(i)}$) of different classes; and (4) “multi-stage training” first trains a naïve classifiers followed by a resampling or reweighting method to re-train or fine-tune the classifier. These methods generally rely on the cardinality of the classes.

semantically-imbalanced dataset, SVCI-20. Finally, Section VI concludes the paper.

II. BACKGROUND

A. Is Naïve Classifier Robust to Class Imbalance?

Consider a classifier that minimizes the negative log-likelihood (NLL) of training examples, which is the prevalent approach in the deep learning literature [26], [27], [28], [29], [30], [31]. We term this classifier as the *naïve classifier*. Such a classifier, $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$, maps an image $\mathbf{x}^{(i)} \in \mathcal{X}$ to a categorical distribution $\hat{\mathbf{y}}^{(i)} \in \mathcal{Y}$ such that $\hat{y}_c^{(i)}$, c -th element of $\hat{\mathbf{y}}^{(i)}$, represents the predicted probability for class $c \in \mathcal{C}$. Typically, $\hat{y}_c^{(i)}$ is obtained through the softmax function:

$$\hat{y}_c^{(i)} = \frac{\exp(s_c^{(i)})}{\sum_{c' \in \mathcal{C}} \exp(s_{c'}^{(i)})}, \quad (1)$$

with $s_c^{(i)}$ being the predicted logit for class c . Given $\mathbf{y}^{(i)} \in \{0, 1\}^{|\mathcal{C}|} \subseteq \mathcal{Y}$, the label of the i th example, $f_\theta(\cdot)$ is conventionally trained by minimizing Cross Entropy Loss ($\ell_{CE}(\cdot, \cdot)$) using mini-batch Stochastic Gradient Descent (SGD) [32]. Formally, the loss of a mini-batch is:

$$\mathcal{L} = \frac{1}{m} \left(\sum_{i=1}^m \ell_{CE}(\mathbf{y}^{(i)}, \hat{\mathbf{y}}^{(i)}) \right), \quad (2)$$

where m is the number of examples in the mini-batch.

By definition, minimizing the NLL over training using mini-batch SGD data does not prioritize any example; thereby making the naïve classifier sensitive to the class characteristics. As a result, the larger cardinality a class has, the more it is represented in the computed NLL. Specifically, as illustrated in Fig. 1(b), the random sampling of mini batches would naturally cause the sampling probability of an example from class c (α_c) to be proportional to the cardinality of the class. To state more formally, $\alpha_1 > \alpha_2 > \dots > \alpha_{|C|}$ holds true, assuming the class indices are ordered with respect to (wrt.) the cardinality in decreasing order. As a result, the

loss function does not focus on the examples that are under-represented either. Consequently, in the case of an imbalance between different classes, training is prone to overfit to the over-represented classes.

A prominent failure example of the naïve classifier is observed in training object detectors [18]. While training typical object detectors, there are prohibitively more negative examples from background compared to very few positives from foreground (i.e., objects). Consequently, unless an imbalance mitigation method is employed while training an object detector, the naïve classifier completely ignores the positive examples, the under-represented classes in this context [22], [25]. Therefore, the naïve classifier is not robust against class imbalance, a problem that remains still open despite a plethora of mitigation techniques.

B. Mitigating Class Imbalance

We categorize the existing class imbalance mitigation techniques into four groups. An overview is given in Fig. 1(b).

1) *Resampling*: This set of methods addresses class imbalance during mini-batch sampling. One standard approach is to equalize the sampling probabilities of classes:

$$\alpha_1 = \alpha_2 = \dots = \alpha_{|C|}, \quad (3)$$

known as the classical oversampling or class-balanced (CB) resampling [1]. As this approach is prone to overfitting in favor of tail classes, several mitigation strategies have been proposed. Peng et al. [2] combined CB and random sampling with a weight parameter to obtain α_c whereas Kang et al. [3] introduced progressive-balanced (PB) resampling, which initializes $\alpha_c s$ with random sampling before gradually getting replaced by CB resampling.

2) *Loss Reweighting*: This common group of approaches [4], [22], [5], [6], [7], [8], [9], [10] reweights the loss values of the examples to promote certain classes or examples:

$$\mathcal{L} = \frac{1}{m} \sum_{i=1}^m w_c^{(i)} \ell_{CE}(\hat{\mathbf{y}}^{(i)}, \mathbf{y}^{(i)}), \quad (4)$$

where $w_c^{(i)}$ is the weight of the example i from class c . Examples include the Cost-sensitive Cross Entropy (CSCE) Loss [4] relying on the inverse of the class cardinalities as $w_c^{(i)}$; and the Class-balanced Loss [5] introducing an “effective” number of examples for each class, which practically smooths the weights of CSCE. An alternative approach, Focal Loss [22], assigns example-specific weight based on prediction confidence of the classifier as an indicator of the hardness.

3) *Margin Adjustment*: These relatively recent methods [11], [12], [13], [14], [15] introduce margin terms to the logit of each class, which effectively demote the over-represented classes. This is achieved by subtracting a class-specific margin $\Delta_c^{(i)}$ from the predicted logit $s_c^{(i)}$ while computing Eq. (1). Specifically, we generalize the existing margin adjustment methods as:

$$\hat{y}_c^{(i)} = \frac{\exp(s_c^{(i)} - \Delta_c^{(i)})}{\exp(s_c^{(i)} - \Delta_c^{(i)}) + \sum_{k \in \mathcal{C} \setminus c} \exp(s_k^{(i)} - \Delta_k^{(i)} \beta)}, \quad (5)$$

where $\beta \in \{0, 1\}$ determines whether or not to take into account the margins assigned to other classes for class c . The examples include: (i) LDAM [11] with $\beta = 0$, $\Delta_c^{(i)} = \tau/N_c^{1/4}$ and τ is a hyper-parameter; (ii) Logit-adjusted Loss [13] with $\beta = 1$ and $\Delta_c^{(i)} = -\kappa \log(N_i / \sum_{k=1}^{\mathcal{C}} N_j)$ and κ being its hyper-parameter; and (iii) Distributional Robustness Loss (DRO) [15], which unlike other methods, enforces margins at the feature-level.

4) *Multi-stage Training*: Multi-stage training methods are based on the idea of decoupling the learning between representation and classification [3]. Specifically, such methods first train the classifier with no mitigation method, and subsequently continue by training either only the last layer [3] or the whole network [11], [33] whilst leveraging an imbalance mitigation method such as CB resampling or CSCE. Several methods are built upon this training strategy, owing to its simplicity and effectiveness [34], [15], [13], [35], [9].

Discussion As evidenced by the overview of the related work above, *existing methods mostly rely on “class cardinality” to measure class imbalance in a dataset*. Specifically, among the aforementioned methods, (i) resampling methods; (ii) loss reweighting methods except Focal Loss; (iii) margin adjustment methods except a specific setting of DRO-LT Loss; and (iv) multi-stage training methods are all based upon class cardinality. *In this paper, we challenge this ubiquitously adopted approach of using “class cardinality” as the sole basis for imbalance mitigation*.

III. CLASS UNCERTAINTY: A NOVEL MEASURE OF CLASS IMBALANCE

In this section, following our viewpoint of a class imbalance measure, we present an analysis on class cardinality, remark its limitations and propose CLASS UNCERTAINTY to measure and mitigate class imbalance.

A. Measuring Class Imbalance

As summarized in the previous section, current literature on class imbalance relies heavily on class cardinality, making

it the de-facto measure of class imbalance. A better measure promises gains for different types of mitigation techniques.

Definition 3.1 (CLASS IMBALANCE MEASURE) *A class imbalance measure is a function μ which takes in a dataset $\mathcal{D} \subseteq \mathcal{X} \times \mathcal{Y}$ with $|\mathcal{C}|$ classes as input and returns a $|\mathcal{C}|$ -dimensional vector. For convenience, we reuse μ to denote the output of this function, which is in $\mathbb{R}^{|\mathcal{C}|}$.*

Moreover, we introduce the following two important features (IF) that a class imbalance measure should have:

- **(IF1) Ability to capture different amounts of imbalance.** An imbalance measure should be able to quantify how well a class is represented by its examples. As such, we expect a measure to:
 - **(IF1a)** be higher for an under-represented class than an over-represented class, and
 - **(IF1b)** decrease for the over-represented classes (or increase for the under-represented classes) as the dataset becomes more imbalanced.
- **(IF2) Robustness to the amount of ineffective examples.** An imbalance measure should be robust to the presence of ‘ineffective’ examples² that might have populated the dataset. In essence, we desire our measure to capture imbalance not in terms of quantity (cardinality) but quality.

These important features allow us to get more insights on class cardinality and our proposed measure. In these analyses, for **(IF1a)** it is not trivial to identify over- and under-represented classes as we argue that class imbalance is not limited to cardinality imbalance. Therefore, we use the class-wise test errors of the naïve classifier as a valid proxy to the underlying imbalance. For **(IF1b)**, we change the imbalance ratio (IR) of the dataset and plot how different measures vary. As for **(IF2)**, we duplicate the training examples in favor of under-represented classes and exploit these duplicates as ineffective examples. Please refer to Section V-A for details.

B. CLASS CARDINALITY IMBALANCE MEASURE and Its Analysis

Considering that cardinality-based methods [1], [4], [5] promote classes with fewer examples during training, we formulate the underlying measure as follows:

Definition 3.2 (CLASS CARDINALITY IMBALANCE MEASURE) *CLASS CARDINALITY IMBALANCE MEASURE, μ_c^C , is an imbalance measure which is inversely proportional to N_c , the number of training examples in class c . Thus, the imbalance measured for the class c is:*

$$\mu_c^C = \frac{1}{|\mathcal{C}|} \frac{N}{N_c}, \quad (6)$$

where N is the cardinality of the entire training set and $|\mathcal{C}|$ (the number of classes in the dataset) normalizes the measure such that $\sum_{c \in \mathcal{C}} \mu_c^C = 1$.

²Following Cui et al. [5], an example can be named effective if it covers sufficiently different space than other examples.

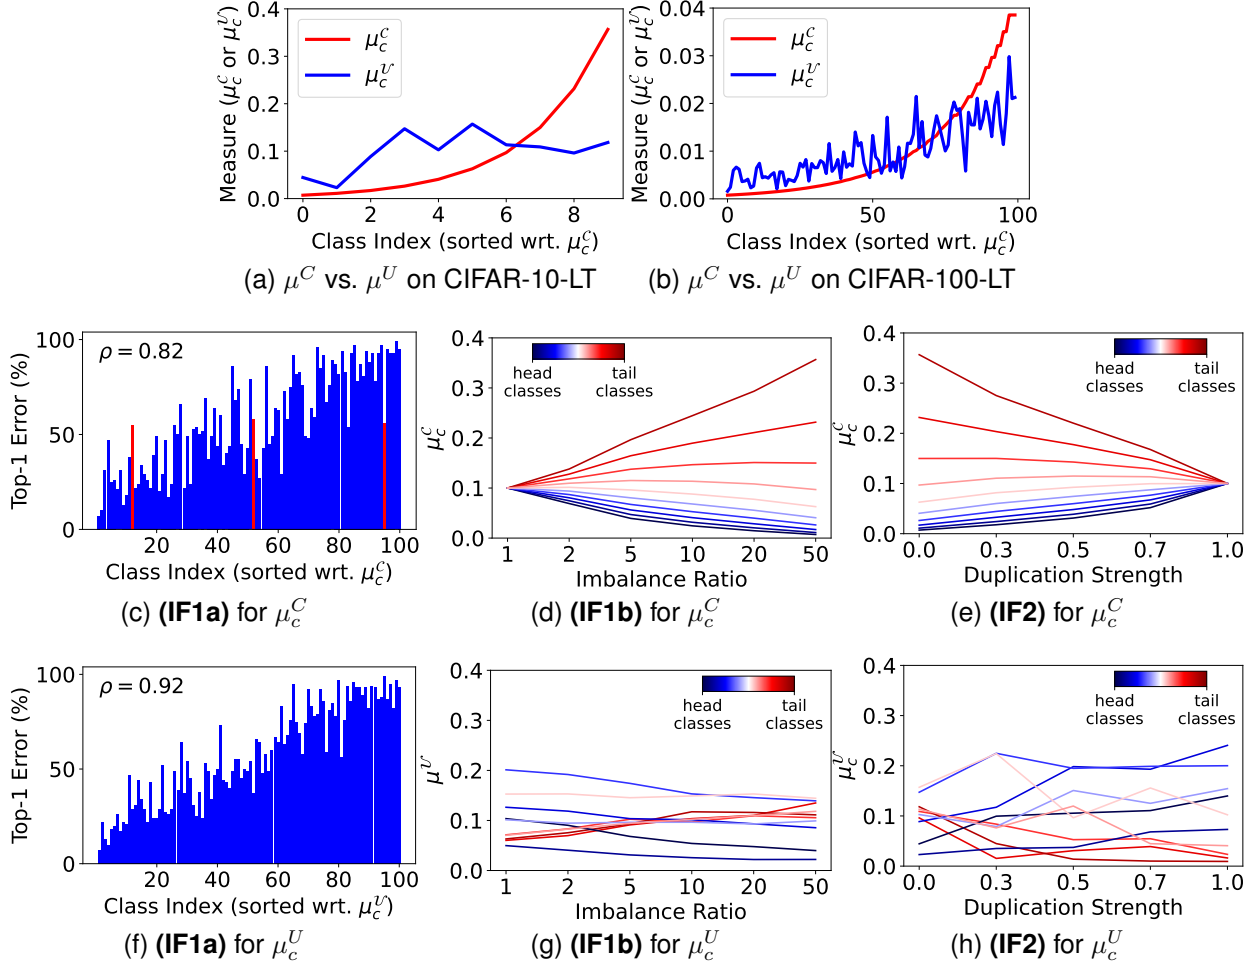


Fig. 2. (a,b) CLASS CARDINALITY IMBALANCE MEASURE (μ_c^C) vs. CLASS UNCERTAINTY (μ_c^U) on CIFAR-10-LT and CIFAR-100-LT with an IR of 50. While μ_c^C is entirely determined by cardinality (from left to right, class cardinality decreases), μ_c^U presents more diverse values. (c-h) An analysis of μ_c^C and μ_c^U in terms of important features. For capturing the amount of imbalance among classes ((c) and (f)), Spearman correlation coefficient (ρ) between the imbalance measure and the test accuracy of a naïve classifier on CIFAR-100-LT is notably higher for μ_c^U compared to μ_c^C , which implies that μ_c^U captures imbalance better. As IR increases, both can capture the change in IR in (d) and (g) by promoting the tail classes (in red) or demoting the head ones (in blue), while the relative change across classes in μ_c^C is more drastic. In (e) and (h), μ_c^C is highly sensitive to the duplicated training examples, which do not contain new information. μ_c^U is relatively less sensitive especially after the duplication strength of 0.3. Therefore, μ_c^U is a better alternative to μ_c^C when considering these important features.

The normalization enables us to provide an analysis independent from the range of the measure. Hence, if the class c has fewer examples, then the corresponding μ_c^C increases relative to the others. As an example, red curves in Fig. 2(a,b) demonstrate μ_c^C for CIFAR-10-LT and CIFAR-100-LT [5].

Analysis Here we analyse CLASS CARDINALITY IMBALANCE MEASURE based on the important features from Section III-A. In order to examine (IF1a), the extent μ_c^C captures the relations among different classes, we plot class indices wrt. μ_c^C vs. the test set error of class c in Fig. 2(c). The figure demonstrates that the error of a class *generally increases* as μ_c^C increases. However, there are some classes with very different number of examples but similar test errors. For example, three classes represented by the red histograms in Fig. 2(c) have very similar errors, but they have 323, 66 and 12 examples in the training set respectively; implying that μ_c^C does not necessarily quantify the imbalance in a class perfectly. In terms of (IF1b), Fig. 2(d) shows how μ_c^C changes as IR increases; suggesting

that the relative change between under- (red tones) and over-represented classes (blue tones) is drastic as the IR increases. Finally for (IF2), Fig. 2(e) suggests that μ_c^C is very sensitive to duplicates, the ineffective samples. Overall, while μ_c^C has benefits and has been used extensively in the literature, our brief analysis shows that it has limitations and offers room for improvement.

C. Our Proposal: CLASS UNCERTAINTY to Measure & Mitigate Imbalance

Considering the limitations of CLASS CARDINALITY IMBALANCE MEASURE (Section III-B), we seek to design a measure of class imbalance that not only considers the cardinality of a class but also reflects the semantic information such as hardness etc. This motivation leads us to using predictive uncertainty as a promising measure considering that it captures both epistemic uncertainty reflecting the lack of data at a point in the input space and aleatoric uncertainty,

representing ambiguity and noise. More specifically, we argue that predictive uncertainty is suitable as a measure of class imbalance due to the following two main reasons:

- Defined as a measure of lack of data, the epistemic component of the predictive uncertainty by definition is a suitable measure for capturing imbalance in data across classes. An important benefit of epistemic uncertainty is that the notion of sufficiency of data is not determined by the number of examples but rather with the behavior of the machine learning model.
- Determining the noise or the ambiguity, the aleatoric component of the predictive uncertainty provides information of the example itself. To illustrate, Kendall and Gal [36] use depth regression task, a closely related computer vision task, to elaborate on aleatoric uncertainty. They demonstrate that this type of uncertainty captures inherently difficult aspects such as large depths and reflective surfaces. Hence, this and alike observations on the aleatoric uncertainty align with our expectation to include the semantic hardness of example in the imbalance measure.

Since the deep learning-based approaches are over-confident in their predictions [37], we train T different models on the same training set, where typically $T = 5$, and quantify the uncertainty given the T predictions from these models, known as Deep Ensembles (DE) [38]. Formally, denoting the prediction vector of the t th model by $\hat{\mathbf{p}}^{(i),t}$ and the probability for class c by $\hat{p}_c^{(i),t} \in \hat{\mathbf{p}}^{(i),t}$, the predictive uncertainty of the deep ensemble on example i , $u^{(i)}$, can be quantified as,

$$u^{(i)} = - \sum_{c \in \mathcal{C}} \hat{p}_c^{(i)} \log \hat{p}_c^{(i)}, \quad (7)$$

where $\hat{p}_c^{(i)} = \frac{1}{T} \hat{p}_c^{(i),t}$.

To adopt predictive uncertainty as a class imbalance measure and obtain a *class-level* imbalance measure given *example-level* predictive uncertainties, we aggregate example-level predictive uncertainties from the training set simply by the class-wise average:

Definition 3.3 (CLASS UNCERTAINTY IMBALANCE MEASURE) CLASS UNCERTAINTY IMBALANCE MEASURE, μ^U , is the average predictive uncertainties of the training examples for each class. Note that unlike μ^C , μ^U is a function of not only the dataset \mathcal{D} but also a model. Specifically, given $u^{(i)}$ as the predictive uncertainty (defined in Eq. 7) of the i^{th} training example of class c and N_c as the cardinality of class c , the unnormalized imbalance in class c is:

$$\tilde{\mu}_c^U = \frac{1}{N_c} \sum_{i=1}^{N_c} u^{(i)}, \quad (8)$$

which is then normalized similar to Eq. (6):

$$\mu_c^U = \frac{\tilde{\mu}_c^U}{\sum_{c' \in \mathcal{C}} \tilde{\mu}_{c'}^U}. \quad (9)$$

Blue curves in Fig. 2(a,b) illustrate μ_c^U for CIFAR-10-LT and CIFAR-100-LT in which μ_c^U presents more diverse values than μ_c^C that are not entirely determined by the cardinality of the classes.

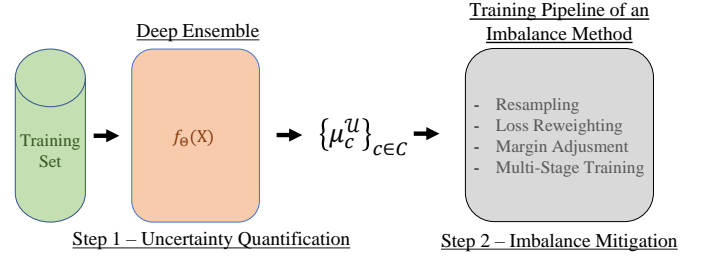


Fig. 3. High-level overview of our strategy. We first obtain CLASS UNCERTAINTY (μ_c^U) as the measure of imbalance using the predictive uncertainties of training examples based on DE. Then, we incorporate the resulting class imbalance measure μ_c^U into various methods during training the classifier.

Analysis We now investigate how CLASS UNCERTAINTY behaves wrt. our important features in comparison to CLASS CARDINALITY IMBALANCE MEASURE. While quantifying uncertainty, we leverage a Deep Ensemble (DE) [38] whose effectiveness in yielding reliable and calibrated uncertainty estimates is proven over a wide range of datasets and models [38], [39], [40]. Firstly for (IF1a), Fig. 2(f) depicts how class-wise test error changes wrt. CLASS UNCERTAINTY. As a quantitative summary, the Spearman correlation coefficient ρ between class-wise test errors and the CLASS UNCERTAINTY measure is 0.92, while it is 0.82 for the CLASS CARDINALITY IMBALANCE MEASURE in Fig. 2(c). This suggests that CLASS UNCERTAINTY captures the imbalance among classes better than CLASS CARDINALITY IMBALANCE MEASURE. Secondly, Fig. 2(g) demonstrates for (IF1b) that as the IR increases, the imbalance assigned to under-represented classes increases while that of over-represented ones decreases as expected from a good imbalance measure. Finally, to compare the robustness of the measures for (IF2), comparing Fig. 2(e) and Fig. 2(h) suggests that CLASS UNCERTAINTY is affected from duplicate examples less than CLASS CARDINALITY IMBALANCE MEASURE and it is a more robust measure. This analysis suggests that CLASS UNCERTAINTY is a better alternative to CLASS CARDINALITY IMBALANCE MEASURE for quantifying class imbalance.

IV. MITIGATING CLASS IMBALANCE USING CLASS UNCERTAINTY

In this section we incorporate our measure into a diverse set of imbalance mitigation methods following the taxonomy we presented in Section II-B. Fig. 3 visualizes the high-level scheme of our architecture to utilize predictive uncertainty to mitigate class imbalance. Specifically, having obtained class-wise uncertainty measures μ_c^U (Eq. 9) based on the predictive uncertainty estimates on the imbalanced training set, here we define α_c in Eq. (3) for resampling methods; $w_c^{(i)}$ in Eq. (4) for reweighting methods; $\Delta_c^{(i)}$ in Eq. (5) for margin-based methods; and finally the training strategy for multi-stage methods.

Uncertainty-based resampling (UBRs): For resampling, we define α_c in Eq. (3) for each class to either promote or demote it based on its uncertainty. The more uncertainty the examples from a class have, the more probability mass is. Since μ_c^U

defined in (9) is already a valid probability distribution, we simply set $\alpha_c = \mu_c^U$.

Uncertainty-based reweighting (UBRw): We define the weight of an example in class c , i.e. $w_c^{(i)}$ in Eq. (4), by multiplying μ_c^U by the total number of classes in order to ensure that the weights of the classes sum up to $|\mathcal{C}|$, which is the case once no reweighting is used and $w_c^{(i)} = 1$ for all c . The resulting $w_c^{(i)}$ of the example i from the class c is, then, defined by

$$w_c^{(i)} = \mu_c^U \times |\mathcal{C}|. \quad (10)$$

Uncertainty-based margins (UBM): We incorporate margin-based uncertainties into LDAM [11] and DRO-LT Loss [15]. Different from UBRs and UBRm, for LDAM we use the unnormalized class uncertainties from (8) considering that these methods offer normalizing the margins in their formulations and set the class margin $\Delta_c^{(i)}$ as:

$$\Delta_c^{(i)} = \frac{\tau \times \tilde{\mu}_c^U}{\max_{c \in \mathcal{C}}(\tilde{\mu}_c^U)}, \quad (11)$$

with $\tau = 0.5$ as in the original work. As for DRO-LT Loss [15], we again simply use $\Delta_c^{(i)} = \mu_c^U$.

Uncertainty-based multi-stage training: Once we have the resampling and reweighting methods, incorporating them into two-stage training is straightforward: We train models without any imbalance mitigation strategy up to a certain epoch, and then fine-tune them using UBRw or UBRs.

V. EXPERIMENTAL RESULTS

A. Implementation Details

1) *Obtaining Long-tailed Datasets*: Similar to the literature [33], while obtaining CIFAR-10-LT and CIFAR-100-LT, we first estimate the number of examples for each class based on their indices as follows:

$$N_c = \bar{N} \times \frac{1}{\text{IR}^{|\mathcal{C}|-1}}, \quad (12)$$

such that \bar{N} is the number of examples in each class in a balanced training set such as CIFAR-10 or CIFAR-100, IR is the imbalance ratio $\in \{50, 100\}$, and finally $|\mathcal{C}|$ is the total number of classes. Then, we randomly select N_c examples from each class to obtain the long-tailed versions of the datasets.

2) *Analyzing **IF1(a)***: As suggested in Section I, we use the class-wise accuracy of the naïve classifier on the test set as a proxy to the class imbalance. We assert that these class-wise accuracies provides us an indicator of the *effects of the imbalance*. We use the naïve classifier as the mitigation methods are conventionally built to mitigate the effects of the class imbalance on this naïve classifier by modifying different steps on its pipeline (see Eq. (4), (3) and (5)). Note that these modifications require at least a proper ranking of the classes by the underlying imbalance measure such that they are promoted or demoted as intended, which is why we adopted the Spearman ranking correlation coefficient for quantification purposes. Furthermore, the accuracy/error on the *test set* can not simply be employed as an imbalance measure since it requires access to the test set during training, which

is not possible. Note that this is also similar to the uncertainty calibration evaluation in which the methods are either tuned on the training set [41], [42] or validation set, and then evaluated against the test set accuracy [43], [44]. We use CIFAR-100-LT with an IR of 50 for this analysis; enabling us to estimate ρ for a larger population compared to CIFAR-10-LT.

3) *Analyzing **IF1(b)***: Here, we simply increase the IR of the dataset, which corresponds to removing some examples from the classes in different numbers depending on IR. Particularly, as IR increases, the number of examples removed from the under-represented classes increase; making the resulting dataset more imbalanced. Therefore, considering that cardinality is one of the contributors of the class imbalance (but not the only one as previously discussed), this provides us an easily-controllable and applicable test bed to analyze this requirement. Specifically, we adopt CIFAR-10-LT with IR of $\{1, 2, 10, 20, 50\}$. Different from **IF1(a)**, here we prefer CIFAR-10-LT; allowing us to demonstrate the individual class behaviours in a way to be followed more easily.

4) *Analyzing **IF2***: Here, we aim to investigate the effect of duplicating examples, which basically corresponds to using oversampling to alleviate the impact of an imbalanced dataset. In order not to be limited for a single setting of the oversampling and cover the effect more thoroughly, we interpolate the sampling probability of the class c denoted by α_c in Eq. (3). Following Peng et al. [45], we achieve this by setting the sampling probability α_c as follows:

$$\alpha_c = (\alpha_c^{\mathcal{R}})^{\lambda} (\alpha_c^{\mathcal{N}})^{1-\lambda}, \quad (13)$$

such that $\alpha_c^{\mathcal{R}}$ is the probability that an example will be sampled from the class c once oversampling is used (e.g., $\forall c, \alpha_c^{\mathcal{R}} = 0.1$ for CIFAR-10 with 10 classes), $\alpha_c^{\mathcal{N}}$ is the same probability but when no imbalance technique is used and finally $\lambda \in [0, 1]$ is the interpolation factor, which we term as the duplication strength. Note that when $\lambda = 0$, the dataset is in its original setting and $\lambda = 1$ implies the dataset is oversampled such that all classes having equal number of examples. We set the duplication strength to $\{0, 0.3, 0.5, 0.7, 1.0\}$ and similar to **IF1(b)**, we use CIFAR-10-LT to present the behaviour of the classes clearly.

5) *Mitigating Class Imbalance*: In our implementation, unless otherwise specified, we exploit the common settings in the literature and implement our models upon the common benchmark as the official implementation of ‘Bag-of-Tricks’ [33] by keeping the settings of the baseline models, train a ResNet-32 [29] and adopt random horizontal flipping and random cropping for data augmentation. For optimization, we employ SGD with momentum using a batch size of 128 on a single GPU. We tune the learning rate of our method by using grid-search generally only in 0.1 increments between $[0, 0.7]$. We also note that as SVCI-20 is a new dataset, we also tuned the learning rates of the baseline methods such as naïve classifier or LDAM to provide a fair comparison. All of the hyper-parameters including the learning rates that we tuned for our method will be made public with the release of our code upon acceptance.

TABLE I

COMPARISON OF USING CLASS UNCERTAINTY INSTEAD OF CARDINALITY WITHIN RESAMPLING, REWEIGHTING, MARGIN ADJUSTMENT, AND MULTI-STAGE TRAINING METHODS. UBRs: UNCERTAINTY-BASED RESAMPLING, UBRw: UNCERTAINTY-BASED REWEIGHTING, UBM: UNCERTAINTY-BASED MARGINS. FOR EACH GROUP, WE DIRECTLY TAKE THE RESULT OF COMPETITIVE METHODS FROM THE CORRESPONDING PAPERS (INDICATED WITH *) TO SHOW THAT OUR IMPLEMENTATION IS INLINE WITH THE PREVIOUSLY REPORTED RESULTS. AMONG 20 DIFFERENT SETTINGS (4 DATASETS AND 5 DIFFERENT SET OF METHODS), USING OUR CLASS UNCERTAINTY ACHIEVES THE BEST RESULTS (IN BOLD) IN 16 DIFFERENT CASES. BESIDES, THE TOP SCORES IN THE TABLE (UNDERLINED AND BOLD) ARE OBTAINED WITH CLASS UNCERTAINTY AS WELL.

Dataset		CIFAR-10-LT		CIFAR-100-LT	
Imbalance Ratio (IR)		50	100	50	100
Naïve Classifier		24.26±0.35	30.18±0.52	57.14±0.08	61.48±0.40
Resampling	CB Resampling [3]	23.33±0.14	30.33±0.30	61.13±0.52	66.34±0.30
	PB Resampling* [33]	25.03	32.91	57.09	61.41
	PB Resampling[3]	23.86±0.36	29.62±0.49	56.87±0.13	60.40±0.57
	UBRs (Ours)	23.32±0.11	29.09±0.37	66.12±0.73	72.53±0.74
	PB UBRs (Ours)	22.61±0.70	28.67±0.91	55.48±0.08	59.28±0.47
Reweighting	Focal Loss [22]	23.91±0.36	30.49±0.68	57.50±0.07	61.32±0.58
	Class-balanced Loss* [5]	20.73	25.43	54.68	60.40
	Class-balanced Focal Loss [5]	21.29±0.39	26.64±0.59	57.60±0.48	61.41±0.35
	UBRw (Ours)	22.51±0.61	26.72±0.59	55.23±0.53	59.23±0.42
	UBRw Focal Loss (Ours)	21.63±0.42	26.54±0.28	53.85±0.71	58.33±0.49
Margin Adj.	LDAM [11]	23.31±0.11	27.95±0.20	57.72±0.27	61.22±0.10
	LDAM+Reweighting* [11]	N/A	22.97	N/A	57.96
	LDAM+Reweighting [11]	20.32±0.02	23.90±0.28	54.27±0.43	57.98±0.20
	DRO-LT [15]	14.63±0.10	17.99±0.11	46.75±0.08	51.85±0.13
	UBM LDAM (Ours)	22.74±0.23	27.31±0.39	55.50±0.49	59.19±0.07
	UBM LDAM+Reweighting (Ours)	20.20±0.29	21.80±0.33	50.40±0.10	54.55±0.14
	UBM DRO (Ours)	14.54±0.25	17.68±0.31	46.54±0.11	51.65±0.08
Multi-stage Methods	<i>Two-stage Resampling:</i>				
	CB Resampling [3]	21.47±0.36	27.00±0.21	54.39±0.15	58.83±0.14
	PB Resampling* [33]	24.58	33.48	56.93	61.35
	PB Resampling [3]	22.56±0.10	28.99±0.72	54.96±0.26	59.27±0.30
	CAM-based Resampling* [33]	18.66	22.62	53.56	57.70
	UBRs (Ours)	19.06±0.14	22.50±0.14	51.82±0.44	56.24±0.46
	<i>Two-stage Reweighting:</i>				
	Cost-sensitive Cross Entropy [4]	21.03±0.46	26.11±0.22	54.41±0.25	58.78±0.13
	Focal Loss [22]	22.03±0.27	28.51±0.68	57.23±0.05	61.70±0.30
	Class-balanced Focal Loss* [33]	20.81	25.31	54.57	58.92
	Class-balanced Focal Loss [5]	20.45±0.27	26.08±0.34	57.09±0.55	59.90±0.19
	UBRw (Ours)	22.21±0.35	26.41±0.14	52.90±0.06	57.25±0.17
	UBRw Focal Loss (Ours)	22.22±0.22	25.71±0.21	53.27±0.08	57.89±0.13

B. Results on Long-tailed Datasets

1) *Datasets and Performance Measure:* Following the literature [33], [15], [11], [5], we use the CIFAR-10-LT and CIFAR-100-LT datasets. These are obtained by systematically reducing the number of examples for certain classes in CIFAR-10 and CIFAR-100 [46] such that the datasets follow a long-tailed distribution (see Section V-A for details). Specifically, we use IRs of 50 and 100 following Zhang et al. [33]. For testing, we use the original (balanced) test sets of CIFAR-10 and CIFAR-100 with 10K examples and report top-1 error.

2) *Resampling:* We first compare our Uncertainty-based Resampling (UBRs) with CB resampling and PB resampling [3], both of which rely on class cardinality as discussed in Section II-B. Similarly, we train the classifier by (i) simply using the class-wise sampling probabilities obtained through CLASS UNCERTAINTY (UBRs); and (ii) inspired from PB, gradually changing the weights towards our uncertainty-based weights (PB UBRs). Table I suggests that UBRs performs better than both methods on CIFAR-10-LT but not on CIFAR-100-LT, which is a more challenging dataset. For example, while the class with minimum cardinality has 50 training examples for CIFAR-10-LT, it is only 5 for CIFAR-100-LT. Note that using our PB UBRs outperforms all resampling

methods consistently. For example, *PB UBRs decreases the top-1 error rate of PB by 1.5% on CIFAR-100-LT on both IRs.* These results suggest that our PB UBRs is more effective than existing resampling methods relying on class cardinality.

3) *Loss Reweighting:* Here, we consider two common baselines: (i) Focal Loss [22], as a hardness-based method [18], and (ii) Class-balanced Loss, a cardinality-based method using the “effective” number of examples in each class [5]. Note that Class-balanced Loss is shown to be a stronger approach compared to directly relying on the number of examples [5], [33]. Table I displays that:

- Weighting the naïve classifier using our CLASS UNCERTAINTY (UBRw) outperforms Focal Loss and consistently improves performance (up to 3.77%).
- Incorporating CLASS UNCERTAINTY into Focal Loss (similar to Class-balanced Loss) outperforms Class-Balanced Loss by up to more than 2% on CIFAR-100-LT (60.40 vs. 58.33 with IR of 100).

These results demonstrate the effectiveness of our approach in loss reweighting methods.

4) *Margin Adjustment:* We replace the cardinality-based margins by our CLASS UNCERTAINTY in three different methods: (i) LDAM [11]; (ii) LDAM combined by reweighting

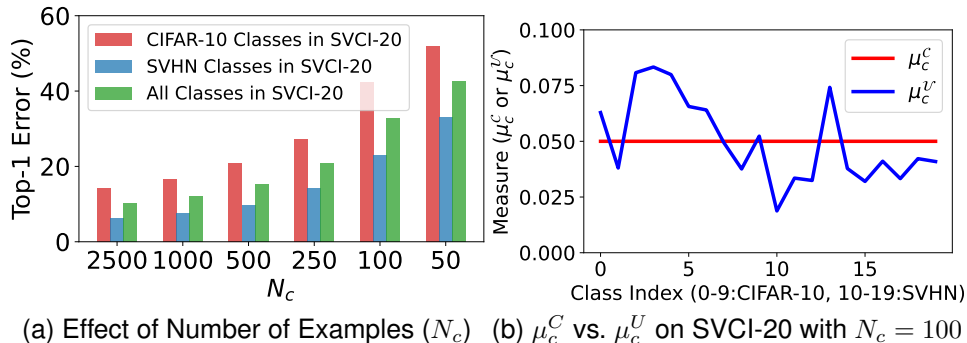


Fig. 4. (a) As the number of examples from each class (N_c) decreases, the error gap between CIFAR-10 and SVHN classes in SVCI-20 dataset increases for Naïve Classifier; implying each setting corresponds to a different semantic imbalance level. (b) While CLASS CARDINALITY IMBALANCE MEASURE (μ_c^C) does not prioritize any class on SVCI-20, our CLASS UNCERTAINTY (μ_c^U) differentiates between them. Generally, μ_c^U assigns more priority to CIFAR-10 classes (class index less than or equal to 9) as the harder classes.

[11] as a more competitive approach; and (iii) DRO-LT Loss [15] with cardinality-based margins. We observe in Table I that:

- Using margins based on CLASS UNCERTAINTY in LDAM (UBM LDAM) consistently outperforms LDAM for its both settings. Specifically, our gains are significant on CIFAR-100-LT, a similar situation to what we have observed for resampling and loss weighting methods. For example, UBM LDAM with reweighting improves its counterpart between 3.5–4% top-1 error on CIFAR-100-LT.
- As for DRO-LT Loss, UBM improves the baseline consistently across all four settings; reaching the top performance in Table I.

5) *Multi-stage Training*: We incorporate our imbalance measure into three multi-stage training methods following Zhang et al. [33]. That is, we train a baseline model without any imbalance mitigation for the first 160 epochs, then use UBRs, UBRw and UBRw Focal Loss with our CLASS UNCERTAINTY in the second stage. In addition to the baselines from the previous sections, we also compare our method with CAM-based sampling using class activation maps for resampling [33] and CSCE [4] relying on class cardinalities. Table I shows that:

- Over eight different dataset and IR combinations, our methods perform the best on six of them with other two being relatively on-par with common baselines.
- Similar to what we observed in resampling and loss reweighting, using CLASS UNCERTAINTY in the second stage also improves the performance especially for CIFAR-100-LT. For example, UBRs performs around 1.5–2.0% top-1 error better than CAM-based resampling, its closest counterpart.
- As for loss weighting, the naïve UBRw performs the best in CIFAR-100-LT with arguably similar results on CIFAR-10-LT.

We conducted ablation studies to explore the flexibility of our approach by combining CLASS UNCERTAINTY with effective number of examples [5] or using DUQ [47] for more reliable uncertainties. These findings suggest potential for future improvements (see supplementary material).

TABLE II
PERFORMANCE GAIN OF USING CLASS UNCERTAINTY. WHILE USING CLASS CARDINALITY-BASED METHODS GENERALLY CORRESPONDS TO THE NAÏVE CLASSIFIER, OUR CLASS UNCERTAINTY APPROACH IS APPLICABLE TO ANY EXISTING CLASS IMBALANCE MITIGATION METHODS; PROVIDING NOTABLE GAINS ON OUR SEMANTICALLY-IMBALANCED SETTINGS. NOTE THAT THE PERFORMANCES ON THE FIRST FOUR ROWS ARE THE SAME AS CLASS CARDINALITIES ARE THE SAME.

Dataset	SVCI-20			
	N_c	250	100	50
Naïve Classifier		22.04 ± 0.65	33.56 ± 1.40	46.17 ± 3.10
<i>Resamp. & Reweight.:</i>				
PB Rs [3]		22.04±0.65	33.56±1.40	46.17±3.10
Class-balanced Loss [5]		22.04±0.65	33.56±1.40	46.17±3.10
Two-stage CSCE Rw [4]		22.04±0.65	33.56±1.40	46.17±3.10
Focal Loss [22]		21.65±0.04	33.61±0.97	50.37±0.48
PB UBRs (Ours)		21.38±0.34	32.78±0.68	45.31±1.30
<i>Margin Adjustment:</i>				
LDAM [11]		20.72±0.80	30.90±0.39	43.36±0.57
UBM LDAM (Ours)		19.87±0.47	31.20±0.09	42.17±0.39

C. Results on Semantically-Imbalanced Datasets

1) *Dataset Curation*: Here, we introduce SVCI-20, a balanced dataset in terms of cardinality but semantically-imbalanced dataset, to test our models more thoroughly and foster research in this direction. Specifically, SVCI-20 combines CIFAR-10 [46] including 10 classes with SVHN [48] comprising of 10 digits taken from house numbers; totalling to 20 classes. Arguably, SVHN is an ‘easier’ dataset than CIFAR-10 considering that: (i) Existing work reports a higher accuracy for SVHN, e.g., DenseNet [31] with a depth of 40 obtains 7.00 top-1 error on CIFAR-10, while it has 1.79 top-1 error on SVHN; and (ii) from a model-free perspective, the intrinsic dimension of SVHN is $1.5\times$ smaller than that of CIFAR-10 [49]. Therefore, combining these datasets into SVCI-20 increases the semantic gap among classes and using uniform class cardinalities removes the impact of cardinality imbalance, both of which are fit for our purpose. We measure the performance using the average top-1 error of classes.

2) *Effect of Class Cardinality on Dataset Curation*: Here, we seek to construct SVCI-20 by setting N_c (the number of examples taken from classes) properly. Note from Fig. 4(a) that as N_c decreases: (i) Top-1 error increases since training

data decreases for all classes; and more importantly (ii) the error gap among the classes of CIFAR-10 and those of SVHN in SVCI-20 also increases. As an example, when $N_c = 1000$, their gap is $\sim 9\%$ while it is around 20% for $N_c = 50$. Based on these observations, we randomly sample 250, 100 and 50 training examples (i.e., N_c s for all $c \in \mathcal{C}$) to further analyse different levels of semantic imbalance, as we did with different IRs in long-tailed datasets.

3) *Class Cardinality on SVCI-20*: As all classes have equal number of training examples, cardinality is indistinctive across classes (the red curve in Fig. 4(b)). Consequently, except for Focal Loss and LDAM, the methods reduce to the Naïve Classifier; effectively ignoring this type of class imbalance (Table II).

4) *CLASS UNCERTAINTY on SVCI-20*: On the SVCI-20 dataset, CLASS UNCERTAINTY is expected to be more distinctive across classes compared to Class Cardinality since it considers other aspects. Fig. 4(b) shows that CLASS UNCERTAINTY for CIFAR-10 classes is generally higher than those for SVHN classes; making our approach different from class cardinality. Consequently, (i) our PB UBRs achieves the best results among resampling and reweighting methods; and (ii) UBM LDAM outperforms LDAM in two out of three cases $\sim 1\%$ while being on par for $N_c = 100$ (Table II).

VI. CONCLUSION

In this work, we demonstrate that CLASS CARDINALITY IMBALANCE MEASURE, as the de-facto measure of class imbalance, has some limitations, and our extensive analyses and experiments show that CLASS UNCERTAINTY is a better alternative. Our approach offers two important benefits. For the former, our measure CLASS UNCERTAINTY can be considered as easy-to-integrate for imbalance mitigation techniques since we only tune the learning rate while incorporating our method, and in very rare cases, such as LDAM, search for the method-specific hyperparameters. As for the latter, our perspective to devise new measures is promising as we obtain better performance on our new test bed comprising of a semantically-imbalanced dataset with equal number of examples such that cardinality-based approaches would fail to capture the difference across classes.

It is also worth noting that one related work [12] employs epistemic uncertainty to mitigate imbalance, but in a specific way that fits into their proposed loss function. Instead, we consider using predictive uncertainty to mitigate class imbalance from a broader perspective by (i) proposing a new measure based on predictive uncertainty, (ii) extensively analyzing and (iii) using the measure with ten imbalance mitigation approaches with different characteristics in imbalance setups.

ACKNOWLEDGMENTS

This work was partially supported by the Scientific and Technological Research Council of Turkey (TÜBİTAK) through the project titled “Addressing Class Imbalance in Visual Recognition Problems by Measuring Class Imbalance and Using Epistemic Uncertainty (DENGE)” (project no. 120E494). We also gratefully acknowledge the computational

resources kindly provided by TÜBİTAK ULAKBİM High Performance and Grid Computing Center (TRUBA) and METU Robotics and Artificial Intelligence Center (ROMER). Dr. Akbas is supported by the “Young Scientist Awards Program (BAGEP)” of Science Academy, Turkey.

REFERENCES

- [1] L. Shen, Z. Lin, and Q. Huang, “Relay backpropagation for effective learning of deep convolutional neural networks,” in *ECCV*, 2016.
- [2] J. Peng, X. Bu, M. Sun, Z. Zhang, T. Tan, and J. Yan, “Large-scale object detection in the wild from imbalanced multi-labels,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [3] B. Kang, S. Xie, M. Rohrbach, Z. Yan, A. Gordo, J. Feng, and Y. Kalantidis, “Decoupling representation and classifier for long-tailed recognition,” in *International Conference on Learning Representations (ICLR)*, 2020.
- [4] N. Japkowicz and S. Stephen, “The class imbalance problem: A systematic study,” *Intell. Data Anal.*, vol. 6, no. 5, p. 429–449, 2002.
- [5] Y. Cui, M. Jia, T.-Y. Lin, Y. Song, and S. Belongie, “Class-balanced loss based on effective number of samples,” in *CVPR*, 2019.
- [6] M. Ren, W. Zeng, B. Yang, and R. Urtasun, “Learning to reweight examples for robust deep learning,” in *Proceedings of the 35th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, J. Dy and A. Krause, Eds., vol. 80. PMLR, 2018, pp. 4334–4343. [Online]. Available: <https://proceedings.mlr.press/v80/ren18a.html>
- [7] J. Shu, Q. Xie, L. Yi, Q. Zhao, S. Zhou, Z. Xu, and D. Meng, “Meta-weight-net: Learning an explicit mapping for sample weighting,” in *NeurIPS*, 2019.
- [8] M. A. Jamal, M. Brown, M. Yang, L. Wang, and B. Gong, “Rethinking class-balanced methods for long-tailed visual recognition from a domain adaptation perspective,” *CoRR*, vol. abs/2003.10780, 2020. [Online]. Available: <https://arxiv.org/abs/2003.10780>
- [9] S. Park, J. Lim, Y. Jeon, and J. Y. Choi, “Influence-balanced loss for imbalanced visual classification,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021, pp. 735–744.
- [10] K. R. M. Fernando and C. P. Tsokos, “Dynamically weighted balanced loss: Class imbalanced learning and confidence calibration of deep neural networks,” *IEEE Trans. Neural Netw. Learn. Syst.*, vol. PP, pp. 1–12, Jan. 2021.
- [11] K. Cao, C. Wei, A. Gaidon, N. Arechiga, and T. Ma, “Learning imbalanced datasets with label-distribution-aware margin loss,” in *Advances in Neural Information Processing Systems*, 2019.
- [12] S. Khan, M. Hayat, S. W. Zamir, J. Shen, and L. Shao, “Striking the right balance with uncertainty,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 103–112.
- [13] A. K. Menon, S. Jayasumana, A. S. Rawat, H. Jain, A. Veit, and S. Kumar, “Long-tail learning via logit adjustment,” in *International Conference on Learning Representations*, 2021. [Online]. Available: <https://openreview.net/forum?id=37nvvqkCo5>
- [14] G. R. Kini, O. Paraskevas, S. Oymak, and C. Thrampoulidis, “Label-imbalanced and group-sensitive classification under overparameterization,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 18 970–18 983, 2021.
- [15] D. Samuel and G. Chechik, “Distributional robustness loss for long-tail learning,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [16] Z. Liu, Z. Miao, X. Zhan, J. Wang, B. Gong, and S. X. Yu, “Large-scale long-tailed recognition in an open world,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [17] L. Yang, H. Jiang, Q. Song, and J. Guo, “A survey on long-tailed visual recognition,” *International Journal of Computer Vision*, pp. 1–36, 2022.
- [18] K. Oksuz, B. C. Cam, S. Kalkan, and E. Akbas, “Imbalance problems in object detection: A review,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 43, no. 10, pp. 3388–3415, 2020.
- [19] B. Krawczyk, “Learning from imbalanced data: open challenges and future directions,” *Progress in Artificial Intelligence*, vol. 5, pp. 221–232, 2016.
- [20] S. Das, S. S. Mullick, and I. Zelinka, “On supervised class-imbalanced learning: An updated perspective and some key challenges,” *IEEE Transactions on Artificial Intelligence*, pp. 1–1, 2022.

- [21] J. Cheong, S. Kalkan, and H. Gunes, “The hitchhiker’s guide to bias and fairness in facial affective signal processing: Overview and techniques,” *IEEE Signal Processing Magazine*, vol. 38, no. 6, pp. 39–49, 2021.
- [22] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, “Focal loss for dense object detection,” in *2017 IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 2999–3007.
- [23] A. Shrivastava, A. Gupta, and R. Girshick, “Training region-based object detectors with online hard example mining,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [24] K. Chen, W. Lin, J. Li, J. See, J. Wang, and J. Zou, “Ap-loss for accurate one-stage object detection,” *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, pp. 1–1, 2020.
- [25] K. Oksuz, B. C. Cam, E. Akbas, and S. Kalkan, “A ranking-based, balanced loss function unifying classification and localisation in object detection,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [26] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” in *The International Conference on Learning Representations (ICLR)*, 2015.
- [27] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, “Going deeper with convolutions,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [28] S. Ioffe and C. Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” in *The International Conference on Machine Learning (ICML)*, 2015.
- [29] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [30] S. Xie, R. B. Girshick, P. Dollár, Z. Tu, and K. He, “Aggregated residual transformations for deep neural networks,” *arXiv*, vol. 1611.05431, 2016.
- [31] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, “Densely connected convolutional networks,” in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 2261–2269.
- [32] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Berlin, Heidelberg: Springer-Verlag, 2006.
- [33] Y. Zhang, X. Wei, B. Zhou, and J. Wu, “Bag of tricks for long-tailed visual recognition with deep convolutional neural networks,” in *AAAI*, 2021, pp. 3447–3455.
- [34] S. Zhang, Z. Li, S. Yan, X. He, and J. Sun, “Distribution alignment: A unified framework for long-tail visual recognition,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 2361–2370.
- [35] Z. Zhong, J. Cui, S. Liu, and J. Jia, “Improving calibration for long-tailed recognition,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 16489–16498.
- [36] A. Kendall and Y. Gal, “What uncertainties do we need in bayesian deep learning for computer vision?” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [37] J. Mukhoti, V. Kulharia, A. Sanyal, S. Golodetz, P. Torr, and P. Dokania, “Calibrating deep neural networks using focal loss,” in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 15288–15299. [Online]. Available: <https://proceedings.neurips.cc/paper/2020/file/aeb7b30ef1d024a76f21a1d40e30c302-Paper.pdf>
- [38] B. Lakshminarayanan, A. Pritzel, and C. Blundell, “Simple and scalable predictive uncertainty estimation using deep ensembles,” in *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30. Curran Associates, Inc., 2017. [Online]. Available: <https://proceedings.neurips.cc/paper/2017/file/9ef2ed4b7fd2c810847ffa5fa85bce38-Paper.pdf>
- [39] F. Pinto, H. Yang, S.-N. Lim, P. H. S. Torr, and P. K. Dokania, “Regmixup: Mixup as a regularizer can surprisingly improve accuracy and out distribution robustness,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [40] V. Kuleshov, N. Fenner, and S. Ermon, “Accurate uncertainties for deep learning using calibrated regression,” in *International Conference on Machine Learning (ICML)*, 2018.
- [41] A. Dave, P. Dollár, D. Ramanan, A. Kirillov, and R. B. Girshick, “Evaluating large-vocabulary object detectors: The devil is in the details,” *arXiv e-prints:2102.01066*, 2021.
- [42] T.-Y. Pan, C. Zhang, Y. Li, H. Hu, D. Xuan, S. Changpinyo, B. Gong, and W.-L. Chao, “On model calibration for long-tailed object detection and instance segmentation,” in *Advances in Neural Information Processing Systems (NeurIPS)*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds., vol. 34. Curran Associates, Inc., 2021, pp. 2529–2542.
- [43] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, “On calibration of modern neural networks,” in *Proceedings of the 34th International Conference on Machine Learning (ICML)*, ser. Proceedings of Machine Learning Research, D. Precup and Y. W. Teh, Eds., vol. 70. PMLR, 2017, pp. 1321–1330.
- [44] A. Kumar, P. S. Liang, and T. Ma, “Verified uncertainty calibration,” in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 32, 2019.
- [45] J. Peng, X. Bu, M. Sun, Z. Zhang, T. Tan, and J. Yan, “Large-scale object detection in the wild from imbalanced multi-labels,” *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9706–9715, 2020.
- [46] A. Krizhevsky, “Learning multiple layers of features from tiny images,” *Tech. Rep.*, 2009.
- [47] J. Van Amersfoort, L. Smith, Y. W. Teh, and Y. Gal, “Uncertainty estimation using a single deep deterministic neural network,” in *International conference on machine learning*. PMLR, 2020, pp. 9690–9700.
- [48] Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, and A. Y. Ng, “Reading digits in natural images with unsupervised feature learning,” in *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.
- [49] P. Pope, C. Zhu, A. Abdelkader, M. Goldblum, and T. Goldstein, “The intrinsic dimension of images and its impact on learning,” in *International Conference on Learning Representations (ICLR)*, 2021.

Supplementary Material for Class Uncertainty: A Measure to Mitigate Class Imbalance

Zeynep Sonat Baltacı^{1,2}, Kemal Oksuz³, Selim Kuzucu¹, Kivanc Tezoren¹, Berkin Kerim Konar¹,
Alpay Ozkan¹, Emre Akbas^{1,4,†}, Sinan Kalkan^{1,4,†}

¹Dept. of Computer Engineering, METU, Ankara, Turkey

²LIGM, Ecole des Ponts, Univ Gustave Eiffel, CNRS, Marne-la-Vallée, France

³Five AI Ltd., United Kingdom

⁴Center for Robotics and Artificial Intelligence (ROMER), METU, Ankara, Turkey

sonat.baltaci@enpc.fr, kemal.oksuz@five.ai

{selim.kuzucu, kivanc.tezoren, berkin.konar, alpay.ozkan, eakbas, skalkan}@metu.edu.tr

arXiv:2311.14090v1 [cs.LG] 23 Nov 2023

I. FURTHER ANALYSES

A. Using Different Uncertainty Estimation Methods

The uncertainty estimation method plays an important role for our CLASS UNCERTAINTY. Although we employ DE [3] to quantify predictive uncertainty due to its robustness over varying datasets, here we show that more reliable uncertainty quantification methods may bring in an improvement. As an initial study, we employ DUQ [5], which arguably yields more reliable uncertainties compared to DE on CIFAR-10 [4]. Fig. 1 compares the CLASS UNCERTAINTIES obtained from DE and DUQ, indicating a similar trend between the two with DUQ promoting under-represented classes more. Table I depicts that using DUQ with FL performs better on CIFAR-10-LT [1]. Note that we only provide results on CIFAR-10 since DUQ has not been applied to CIFAR-100 yet and scaling DUQ to larger datasets is beyond the scope of our work. This result shows that our approach promises further room for improvement with more reliable uncertainty estimation techniques.



Fig. 1. μ_c^U for DE vs. DUQ. DE and DUQ results in a similar trend in CLASS UNCERTAINTY (PU_c), while the uncertainties from DUQ promoting under-represented classes slightly more.

B. Combining CLASS UNCERTAINTY and Cardinality

We finally investigate the benefit of combining CLASS UNCERTAINTY with class cardinality. Specifically, we obtain

Manuscript received March 31, 2023.

[†]Equal contribution for senior authorship.

TABLE I

COMPARING DIFFERENT UNCERTAINTY QUANTIFICATION METHODS. COMBINING THE CLASS UNCERTAINTIES OF DUQ WITH FOCAL LOSS YIELDS THE BEST PERFORMANCE. COMPARED TO THE NAÏVE CLASSIFIER, THE IMPROVEMENT IS BETWEEN 3 TO 4 TOP-1 ERROR.

Method	Unc. Qua.	Imbalance Ratio (IR)	
		50	100
Naïve Cls.	N/A	24.26±0.35	30.18±0.52
UBRw	DE	22.51±0.61	26.72±0.59
UBRw	DUQ	22.84±0.36	27.60±0.89
UBRw FL	DE	21.63±0.42	26.54±0.28
UBRw FL	DUQ	20.90±0.17	25.22±0.87

the weights in Eq. (4) (see the main article) by combining our UBRw Focal Loss [2] and Class-balanced Focal Loss (CB FL) [1] simply by,

$$w_c^{(i)} = \gamma w_c^{(U)} + (1 - \gamma) w_c^{(C)}, \quad (1)$$

where $w_c^{(U)}$ and $w_c^{(C)}$ are the weights estimated by our UBRw FL and CB FL with $\gamma \in [0, 1]$ being the weighting coefficient. Table II shows on both CIFAR-10-LT and CIFAR-100-LT that the training benefits from this combination. These results suggest the potential benefit of combining multiple imbalance measures.

TABLE II

PERFORMANCE GAIN OF USING CLASS UNCERTAINTY COMPARED TO NAÏVE CLASSIFIER AND FOCAL LOSS. COMBINING CLASS UNCERTAINTIES WITH THE EFFECTIVE NUMBER OF SAMPLES PROPOSED BY CUI ET AL. [1] PROMISES FURTHER IMPROVEMENT; SUGGESTING BETTER MEASURES TO CAPTURE CLASS IMBALANCE CAN BE DEvised WITH A MORE THOROUGH INVESTIGATION ON COMBINING DIFFERENT MEASURES. $\gamma = 0.0$ AND $\gamma = 1.0$ CORRESPONDS TO CB FL AND UBRw FL, RESPECTIVELY.

Dataset	CIFAR-10-LT		CIFAR-100-LT	
	50	100	50	100
$\gamma = 0.0$	21.29±0.39	26.64±0.59	57.60±0.48	61.41±0.35
$\gamma = 0.3$	20.22±0.20	25.39±0.09	54.24±0.57	59.61±0.49
$\gamma = 0.5$	20.51±0.37	25.27±0.47	53.73±0.54	58.30±0.30
$\gamma = 0.7$	20.74±0.51	25.82±0.20	53.33±0.23	58.26±0.29
$\gamma = 1.0$	21.63±0.42	26.54±0.28	53.85±0.71	58.33±0.49

REFERENCES

- [1] Y. Cui, M. Jia, T.-Y. Lin, Y. Song, and S. Belongie, “Class-balanced loss based on effective number of samples,” in *CVPR*, 2019.
- [2] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, “Focal loss for dense object detection,” in *2017 IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 2999–3007.
- [3] B. Lakshminarayanan, A. Pritzel, and C. Blundell, “Simple and scalable predictive uncertainty estimation using deep ensembles,” in *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30. Curran Associates, Inc., 2017. [Online]. Available: <https://proceedings.neurips.cc/paper/2017/file/9ef2ed4b7fd2c810847ffa5fa85bce38-Paper.pdf>
- [4] A. Krizhevsky, “Learning multiple layers of features from tiny images,” Tech. Rep., 2009.
- [5] J. Van Amersfoort, L. Smith, Y. W. Teh, and Y. Gal, “Uncertainty estimation using a single deep deterministic neural network,” in *International conference on machine learning*. PMLR, 2020, pp. 9690–9700.