

Balanced MSE for Imbalanced Visual Regression

Jiawei Ren¹, Mingyuan Zhang¹, Cunjun Yu², Ziwei Liu¹ 

¹S-Lab, Nanyang Technological University

²School of Computing, National University of Singapore

{jiawei011, mingyuan001}@e.ntu.edu.sg, cunjun.yu@comp.nus.edu.sg, ziwei.liu@ntu.edu.sg

Abstract

Data imbalance exists ubiquitously in real-world visual regressions, e.g., age estimation and pose estimation, hurting the model’s generalizability and fairness. Thus, imbalanced regression gains increasing research attention recently. Compared to imbalanced classification, imbalanced regression focuses on continuous labels, which can be boundless and high-dimensional and hence more challenging. In this work, we identify that the widely used Mean Square Error (MSE) loss function can be ineffective in imbalanced regression. We revisit MSE from a statistical view and propose a novel loss function, *Balanced MSE*, to accommodate the imbalanced training label distribution. We further design multiple implementations of *Balanced MSE* to tackle different real-world scenarios, particularly including the one that requires no prior knowledge about the training label distribution. Moreover, to the best of our knowledge, *Balanced MSE* is the first general solution to high-dimensional imbalanced regression in modern context. Extensive experiments on both synthetic and three real-world benchmarks demonstrate the effectiveness of *Balanced MSE*. Code and models are available at github.com/jiawei-ren/BalancedMSE.

1. Introduction

Visual regression, where models learn to predict continuous labels, is one of the most fundamental tasks in machine learning. However, in real-world applications, data imbalance is widely encountered, hurting the model’s generalizability and fairness. For example, age estimation predicts people’s age from their visual appearance, where age is a continuous label. In practice, most of the training images are from adults, while very few images are from children and senior adults. As a result, models trained from such an imbalanced dataset can have inferior performance on under-represented groups [38]. Therefore, imbalanced regression

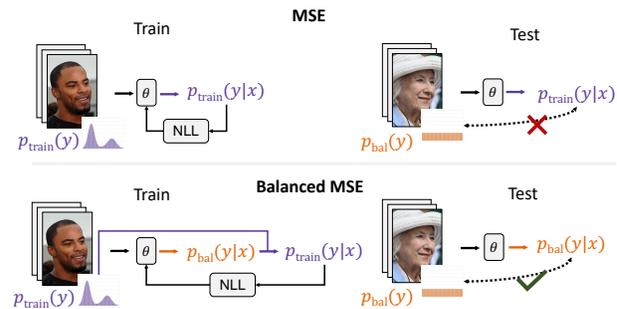


Figure 1. Comparison between MSE and Balanced MSE. MSE is equivalent to NLL on a prediction distribution, where the regressor θ ’s prediction is the distribution mean. MSE lets the regressor model $p_{\text{train}}(\mathbf{y}|\mathbf{x})$, which is not suitable to infer on the test set due to a shift between the training label distribution $p_{\text{train}}(\mathbf{y})$ and the balanced test label distribution $p_{\text{bal}}(\mathbf{y})$. In comparison, Balanced MSE leverages $p_{\text{train}}(\mathbf{y})$ to make a statistical conversion from $p_{\text{bal}}(\mathbf{y}|\mathbf{x})$ to $p_{\text{train}}(\mathbf{y}|\mathbf{x})$, thus allowing the regressor to model the desired $p_{\text{bal}}(\mathbf{y}|\mathbf{x})$ by still minimizing NLL of $p_{\text{train}}(\mathbf{y}|\mathbf{x})$.

gains increasing research attention. A fresh imbalanced regression benchmark [38] in the modern deep learning context has been curated recently as well. Compared with imbalanced and long-tailed classification [4, 12, 24] that studies categorical labels, imbalanced regression focuses on continuous labels, which can be boundless, high-dimensional, and hence more challenging.

Unlike imbalanced classification that has been widely discussed [15, 20, 35, 39], imbalanced regression is under-explored. Previous works [1, 33] focus on synthesizing samples for rare labels, which have limited feasibility in modern deep learning where inputs are always high-dimensional. Recent research [31, 38] focuses on loss reweighting. Reweighting assigns larger loss weights to rare samples and smaller loss weights to frequent samples. [31, 38] estimates the training label distributions using kernel density estimation (KDE) and reweight losses accordingly. However, prior works [5, 36] show that reweighting has limited effectiveness on imbalanced classification. In a following case study, we validate this finding in imbalanced

 Corresponding author.

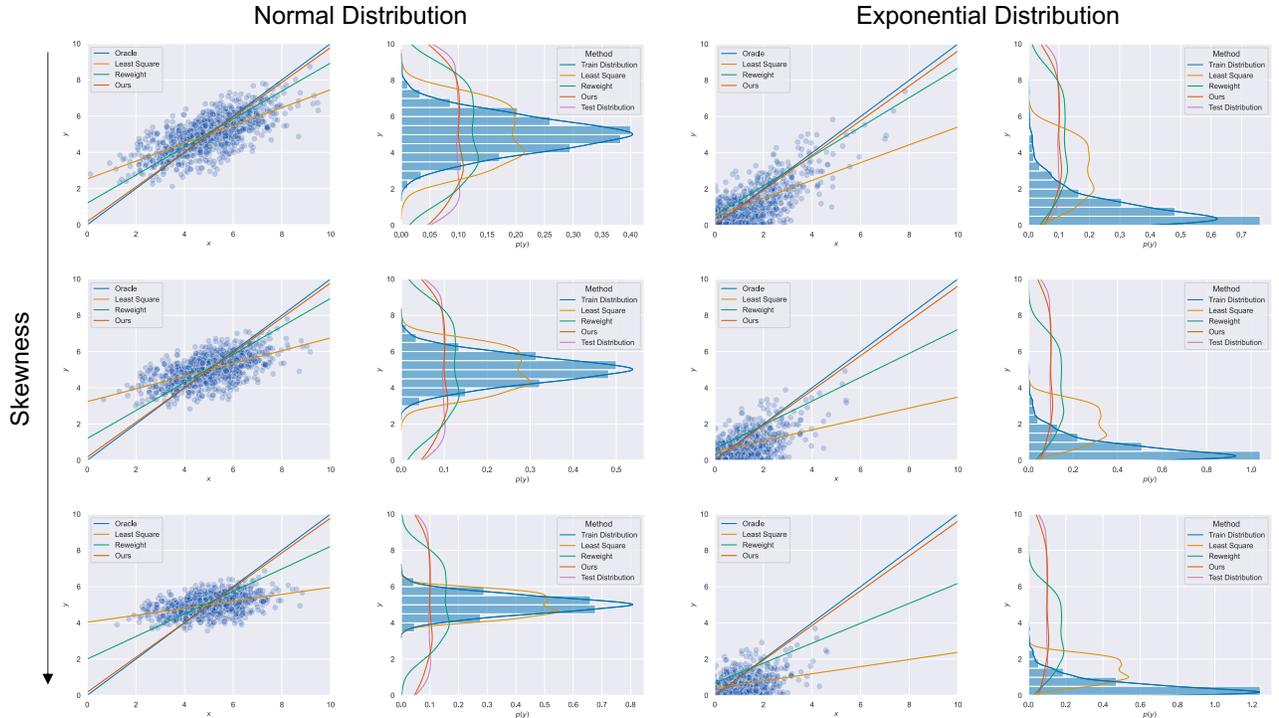


Figure 2. Comparison of Balanced MSE and existing methods on a 1-D imbalanced linear regression synthetic benchmark. Column 1 and 3 are visualizations of regression results: points represent training data, x is input and y is label; (x, y) is noisily generated by an oracle linear relation (in blue) and is artificially label-imbalanced; other lines represent different regressors, the closer to the oracle the better. Column 2 and 4 are visualizations of label distributions: blue shaded histogram represents the training label distribution $p_{\text{train}}(y)$, it gets more skewed from top to bottom; purple histogram represents the test label distribution, which is balanced; other histograms are the marginal label distributions predicted by different regressors on the test set, the closer to the test distribution the better. Although reweighting (in green) is closer to the oracle (in blue) compared with least square (in yellow), it suffers a larger error when $p_{\text{train}}(y)$ gets more skewed. Our method (in red), Balanced MSE, makes the estimation closest to the oracle and has a uniform marginal label distribution on the test set.

regression as well. To sum up, *imbalanced regression is still in an early stage and lacks an effective approach.*

To mitigate the gap, we present a statistically principled loss function, **Balanced MSE**, for imbalanced regression. We revisit Mean Square Error (MSE), the standard loss function in regression, from a statistical view. We identify that MSE carries the label imbalance into predictions, which leads to inferior performance on rare labels. We propose Balanced MSE to restore a balanced prediction by leveraging the training label distribution prior to make a statistical conversion. Moreover, we provide various implementation options for Balanced MSE, including the one that estimates the training label distribution online and requires no additional prior knowledge, making Balanced MSE applicable to different real-world scenarios.

Balanced MSE shows clear advantages over existing methods both theoretically and practically. As a motivating example, we compare Balanced MSE with reweighting using an 1-D linear regression synthetic benchmark shown in Fig. 2. Regressors trained with Balanced MSE show a consistent performance that is invariant to the skewness of the training label distribution. On the contrary, reweighting suffers from a significantly larger prediction error when the

training label distribution gets more skewed.

We further demonstrate Balanced MSE’s empirical success on existing real-world benchmarks [38], including age estimation and depth estimation. Note that existing imbalanced regression benchmarks only consider uni-dimensional label space, *e.g.*, age and depth. However, labels sometimes have more than one dimension in real-world applications. To close the gap, we propose a new multi-dimensional imbalanced regression benchmark on Human Mesh Recovery (HMR) [18], which is an important task that estimates 3D human meshes from monocular images. We extend the standard metrics of HMR (*e.g.*, mean per joint position error (MPJPE)) to balanced metrics so that we evaluate the regression performance on human meshes with different rarity fairly. We call the new imbalanced regression benchmark Imbalanced HMR (IHMR). We show that Balanced MSE delivers strong empirical results on both uni- and multi-dimensional benchmarks. *To the best of our knowledge, Balanced MSE is the first general solution to high-dimensional imbalanced regression in modern context.*

In summary, our contributions are three-fold: **1)** We identify the ineffectiveness of MSE in imbalanced regression and propose a statistically principled loss function,

Balanced MSE, that leverages the training label distribution prior to restore a balanced prediction. **2)** We devise various implementation options of Balanced MSE to tackle different real-world scenarios, including the one that estimates the training label distribution online and requires no prior knowledge beforehand. **3)** We propose a new multi-dimensional benchmark IHMR, and show that Balanced MSE achieves state-of-the-art performance on both uni- and multi-dimensional real-world benchmarks.

2. Related Works

Imbalanced & Long-Tailed Classification. Many techniques have been explored for imbalanced & long-tailed classification, for example, resampling [7, 9, 13, 21] and reweighting [6, 10, 16, 17]. Here, we focus on the logit adjustment techniques, which are the most relevant to this work. Recent works [27, 29, 32] show that modifying the logits in the mapping function, e.g., Softmax or Sigmoid, by an offset proportional to $\log p_{\text{train}}(y)$ gives the Bayes-optimal estimation of the $p_{\text{bal}}(y|x)$. The logit adjustment techniques can work as either a train-time loss function or a test-time adjustment. [34] further develops an online version that accumulates the statistics of label distribution during training instead of requiring statistics of all training labels ahead of time.

Imbalanced Regression. Imbalanced regression is relatively under-explored. Earlier works [1, 33] focus on resampling and synthesizing new samples for rare labels. Further work [2] ensembles regressors trained under different resampling policies. Extending their method towards high-dimensional observations like images is non-trivial. Recent research [31, 38] proposes to estimate the empirical training distribution with KDE and then apply the standard reweighting technique. [38] proposes a feature level smoothing as well, which is complementary to this work.

3. Methodology

3.1. Problem Setting

We study a regression task. We consider input $\mathbf{x} \in X$ and label $\mathbf{y} \in Y = \mathbb{R}^d$. Different from [38] which only discusses uni-dimensional ($d = 1$) regressions, we discuss multi-dimensional ($d > 1$) regressions in this paper as well.

Normally, both the training set and test set are drawn from the same joint distribution. However, when the label distribution is highly skewed, a model may learn a trivial solution by always predicting the frequent labels. The trivial model will still have a low error rate on the test set [27]. To address the issue, either a balanced evaluation metric or a balanced test set is employed to fairly evaluate a model’s performance on samples with different rarities. Moreover, it can be shown that using a balanced metric on an arbitrary

test set is equivalent to using an overall metric on a balanced test set that hypothetically exists [3].

Therefore, imbalanced regression assumes that the training set and test set are drawn from different joint distributions, $p_{\text{train}}(\mathbf{x}, \mathbf{y})$ and $p_{\text{bal}}(\mathbf{x}, \mathbf{y})$ respectively, where the training set’s label distribution $p_{\text{train}}(\mathbf{y})$ is skewed and the balanced test set’s label distribution $p_{\text{bal}}(\mathbf{y})$ is uniform [38]. The label-conditional probability $p(\mathbf{x}|\mathbf{y})$ is assumed to be the same in both training and testing. Instead of learning $p_{\text{train}}(\mathbf{y}|\mathbf{x})$, imbalanced regression’s goal is to estimate $p_{\text{bal}}(\mathbf{y}|\mathbf{x})$ to better perform on the balanced test set. A similar setting is generally adopted in imbalanced classification literature [6, 15, 27, 29] as well.

3.2. Revisiting Mean Square Error

In this section, we revisit Mean Square Error (MSE) from a statistical view. MSE loss is the most commonly used loss function in regression. For a predicted label \mathbf{y}_{pred} and a target label \mathbf{y} , the MSE loss can be written as

$$\text{MSE}(\mathbf{y}, \mathbf{y}_{\text{pred}}) = \|\mathbf{y} - \mathbf{y}_{\text{pred}}\|_2^2, \quad (3.1)$$

where $\|\cdot\|_2$ denotes the L2 norm. It is well known that minimizing MSE can be equivalent to maximum likelihood estimation in regression [28]. The prediction of a regressor \mathbf{y}_{pred} can be considered as the mean of a noisy prediction distribution, which is modeled as a Gaussian distribution in the classic probabilistic interpretation [26]:

$$p(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta}) = \mathcal{N}(\mathbf{y}; \mathbf{y}_{\text{pred}}, \sigma_{\text{noise}}^2 \mathbf{I}), \quad (3.2)$$

where $\boldsymbol{\theta}$ is the regressor’s parameter, \mathbf{y}_{pred} is the regressor’s prediction and σ_{noise} is the scale of an i.i.d. error term $\epsilon \sim \mathcal{N}(0, \sigma_{\text{noise}}^2 \mathbf{I})$. It is easy to show that MSE equals to the Negative Log Likelihood (NLL) loss of the prediction distribution $p(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta})$ [28]. Therefore, a regressor trained using MSE in fact learns to model $p_{\text{train}}(\mathbf{y}|\mathbf{x})$.

However, as mentioned in the problem setting, we are interested in estimating $p_{\text{bal}}(\mathbf{y}|\mathbf{x})$ instead of $p_{\text{train}}(\mathbf{y}|\mathbf{x})$. Due to a shift from the long-tailed training distribution $p_{\text{train}}(\mathbf{y})$ to the balanced test distribution $p_{\text{bal}}(\mathbf{y})$, there is a mismatch between $p_{\text{train}}(\mathbf{y}|\mathbf{x})$ and $p_{\text{bal}}(\mathbf{y}|\mathbf{x})$. By Bayes’ Rule, we have $p_{\text{train}}(\mathbf{y}|\mathbf{x}) \propto p(\mathbf{x}|\mathbf{y}) \cdot p_{\text{train}}(\mathbf{y})$ and $p_{\text{bal}}(\mathbf{y}|\mathbf{x}) \propto p(\mathbf{x}|\mathbf{y}) \cdot p_{\text{bal}}(\mathbf{y})$. By change of variables, we have:

$$\frac{p_{\text{train}}(\mathbf{y}|\mathbf{x})}{p_{\text{bal}}(\mathbf{y}|\mathbf{x})} \propto \frac{p_{\text{train}}(\mathbf{y})}{p_{\text{bal}}(\mathbf{y})} \quad (3.3)$$

Eq. 3.3 quantifies that the ratio between $p_{\text{train}}(\mathbf{y}|\mathbf{x})$ and $p_{\text{bal}}(\mathbf{y}|\mathbf{x})$ is proportional to $p_{\text{train}}(\mathbf{y})$, which is lower when a label rarely appears in the training set. Therefore, *a regressor trained with MSE will underestimate on rare labels.*

Although this mismatch is a well-known observation in imbalanced classification [15, 27, 29, 32], we are the first to address the mismatch in imbalanced regression. Different

from classification tasks that model prediction distribution explicitly by Softmax scores, regression tasks model the prediction distribution implicitly: only the mean of the prediction distribution plays a role in both training and testing. Therefore, the probabilistic meaning of imbalanced regression has been constantly overlooked by existing research. Our work is an initial attempt to rethink imbalanced regression in a statistical framework. We will show how the statistical insight sheds light on imbalanced regression in the following sections.

3.3. Balanced MSE

We propose Balanced MSE to restore $p_{\text{bal}}(\mathbf{y}|\mathbf{x})$. First, we discuss a statistical conversion from $p_{\text{bal}}(\mathbf{y}|\mathbf{x})$ to $p_{\text{train}}(\mathbf{y}|\mathbf{x})$ using the training label distribution $p_{\text{train}}(\mathbf{y})$.

Theorem 1 (Statistical Conversion) *Let $p_{\text{train}}(\mathbf{x}, \mathbf{y})$ be the training distribution where $p_{\text{train}}(\mathbf{y})$ is imbalanced and $p_{\text{bal}}(\mathbf{x}, \mathbf{y})$ be the balanced test distribution where $p_{\text{bal}}(\mathbf{y})$ is uniform. $p_{\text{train}}(\mathbf{x}, \mathbf{y})$ and $p_{\text{bal}}(\mathbf{x}, \mathbf{y})$ have the same label-conditional distribution $p(\mathbf{x}|\mathbf{y})$. $p_{\text{train}}(\mathbf{y}|\mathbf{x})$ can always be expressed by $p_{\text{bal}}(\mathbf{y}|\mathbf{x})$ and $p_{\text{train}}(\mathbf{y})$ as:*

$$p_{\text{train}}(\mathbf{y}|\mathbf{x}) = \frac{p_{\text{bal}}(\mathbf{y}|\mathbf{x}) \cdot p_{\text{train}}(\mathbf{y})}{\int_Y p_{\text{bal}}(\mathbf{y}'|\mathbf{x}) \cdot p_{\text{train}}(\mathbf{y}') d\mathbf{y}'}. \quad (3.4)$$

The proof can be found in the supplementary materials. [Theorem 1](#) allows us to estimate $p_{\text{bal}}(\mathbf{y}|\mathbf{x})$ by minimizing the NLL loss of $p_{\text{train}}(\mathbf{y}|\mathbf{x})$. Specifically, we let the regressor directly estimate the desired $p_{\text{bal}}(\mathbf{y}|\mathbf{x})$, *i.e.*,

$$p_{\text{bal}}(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta}) = \mathcal{N}(\mathbf{y}; \mathbf{y}_{\text{pred}}, \sigma_{\text{noise}}^2 \mathbf{I}) \quad (3.5)$$

As illustrated in [Fig. 1](#), in training, we first predict $p_{\text{bal}}(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta})$, convert it into $p_{\text{train}}(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta})$ using [Eq. 3.4](#) and then compute the NLL loss to update $\boldsymbol{\theta}$; in testing, we skip the conversion and directly output the regressor's prediction $p_{\text{bal}}(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta})$. We name the NLL loss of the converted conditional probability as Balanced MSE.

Definition 3.1 (Balanced MSE) *For a regressor's prediction \mathbf{y}_{pred} , and a training label distribution prior $p_{\text{train}}(\mathbf{y})$, the Balanced MSE loss is defined as:*

$$\begin{aligned} L &= -\log p_{\text{train}}(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta}) \\ &= -\log \frac{p_{\text{bal}}(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta}) \cdot p_{\text{train}}(\mathbf{y})}{\int_Y p_{\text{bal}}(\mathbf{y}'|\mathbf{x}; \boldsymbol{\theta}) \cdot p_{\text{train}}(\mathbf{y}') d\mathbf{y}'} \\ &\cong -\log \mathcal{N}(\mathbf{y}; \mathbf{y}_{\text{pred}}, \sigma_{\text{noise}}^2 \mathbf{I}) \\ &\quad + \log \int_Y \mathcal{N}(\mathbf{y}'; \mathbf{y}_{\text{pred}}, \sigma_{\text{noise}}^2 \mathbf{I}) \cdot p_{\text{train}}(\mathbf{y}') d\mathbf{y}', \end{aligned} \quad (3.6)$$

where \cong hides a constant term $-\log p_{\text{train}}(\mathbf{y})$.

Balanced MSE has two parts: the first part is equivalent to the standard MSE loss, and the second part is a new balancing term, where an integral needs to be computed. We

show in the supplementary materials that the new balancing term equals a constant when the training label distribution $p_{\text{train}}(\mathbf{y})$ is uniform. Therefore, the standard MSE loss can be viewed as a special case of Balanced MSE.

Balanced MSE closes the distribution mismatch between training and testing, thus being a statistically principled loss function for imbalanced regression. In the following sections, we discuss 1) Balanced MSE's connection with imbalanced classification and 2) how to implement Balanced MSE in practice.

3.4. Connection with Imbalanced Classification

We show that Balanced MSE has an underlying connection with existing solutions in imbalanced classification. [Theorem 1](#) is true not only in imbalanced regression but also in imbalanced classification. In imbalanced classification, the label space Y is one-dimensional and discrete, the integral on Y can be written into summation, [Eq. 3.4](#) becomes:

$$p_{\text{train}}(y|\mathbf{x}) = \frac{p_{\text{bal}}(y|\mathbf{x}) \cdot p_{\text{train}}(y)}{\sum_{y' \in Y} p_{\text{bal}}(y'|\mathbf{x}) \cdot p_{\text{train}}(y')}. \quad (3.7)$$

Usually, Softmax is employed to convert model outputs into a prediction distribution in classification. When using Softmax to express the desired $p_{\text{bal}}(y|\mathbf{x})$, we have:

$$p_{\text{bal}}(y|\mathbf{x}; \boldsymbol{\theta}) = \frac{\exp(\eta[y])}{\sum_{y' \in Y} \exp(\eta[y'])}, \quad (3.8)$$

where $\eta[y]$ is the model's output on class y . Plugging the Softmax expression in [Eq. 3.8](#) into [Eq. 3.7](#), we have

$$p_{\text{train}}(y|\mathbf{x}; \boldsymbol{\theta}) = \frac{\exp(\eta[y]) \cdot p_{\text{train}}(y)}{\sum_{y' \in Y} \exp(\eta[y']) \cdot p_{\text{train}}(y')}, \quad (3.9)$$

which achieves the same form as the logit adjustment techniques in imbalanced classification literature [[15](#), [27](#), [29](#)]. Therefore, Balanced MSE and the logit adjustment techniques can be viewed as two different instantiations of [Theorem 1](#) on imbalanced regression and imbalanced classification respectively. *Our work provides a unified statistical view of both imbalanced classification and regression for the first time.* We hope that the unified perspective can help future research to further bridge the two tasks.

3.5. Implementation Options

We discuss how to implement Balanced MSE in practice. The integral in Balanced MSE ([Eq. 3.6](#)) can be difficult to compute. In the following sections, we provide both closed-form options and numerical options to calculate the integral. Particularly, the option Batch-based Monte-Carlo (BMC) requires no prior knowledge of the training label distribution and hence can be more generally deployed in real-world applications.

3.5.1 Closed-form Options

In this section, we aim to find a closed-form expression for the integral $\int_Y \mathcal{N}(\mathbf{y}; \mathbf{y}_{\text{pred}}, \sigma_{\text{noise}}^2 \mathbf{I}) \cdot p_{\text{train}}(\mathbf{y}) d\mathbf{y}$. The main challenge is how to express $p_{\text{train}}(\mathbf{y})$ to make the integral tractable. Here, we discuss a viable option, which is to express $p_{\text{train}}(\mathbf{y})$ as a Gaussian Mixture Model (GMM).

GMM-based Analytical Integration (GAI). The advantage of employing GMM is the fact that the product of two Gaussians is an unnormalized Gaussian. Concretely, let us have $p_{\text{train}}(\mathbf{y})$ expressed by a Gaussian Mixture:

$$p_{\text{train}}(\mathbf{y}) = \sum_{i=1}^K \phi_i \mathcal{N}(\mathbf{y}; \boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i), \quad (3.10)$$

where K is the number of Gaussian components, ϕ , $\boldsymbol{\mu}$, $\boldsymbol{\Sigma}$ are the weights, means and covariances of the GMM. Since the product of two Gaussians is an unnormalized Gaussian, we have:

$$\begin{aligned} & \int_Y \mathcal{N}(\mathbf{y}; \mathbf{y}_{\text{pred}}, \sigma_{\text{noise}}^2 \mathbf{I}) \cdot \sum_{i=1}^K \phi_i \mathcal{N}(\mathbf{y}; \boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i) d\mathbf{y} \\ &= \sum_{i=1}^K \phi_i S_i \int_Y \mathcal{N}(\mathbf{y}; \tilde{\boldsymbol{\mu}}_i, \tilde{\boldsymbol{\Sigma}}_i) d\mathbf{y}. \end{aligned} \quad (3.11)$$

where S , $\tilde{\boldsymbol{\mu}}$, $\tilde{\boldsymbol{\Sigma}}$ are the norms, means, and covariances of the new unnormalized Gaussian. Now, the integral is on Gaussian distribution and can be trivially solved. We leave the detailed derivation in the supplementary material. The final loss form is:

$$\begin{aligned} L &= -\log \mathcal{N}(\mathbf{y}; \mathbf{y}_{\text{pred}}, \sigma_{\text{noise}}^2 \mathbf{I}) \\ &+ \log \sum_{i=1}^K \phi_i \cdot \mathcal{N}(\mathbf{y}_{\text{pred}}; \boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i + \sigma_{\text{noise}}^2 \mathbf{I}). \end{aligned} \quad (3.12)$$

3.5.2 Numerical Options

The closed-form solution above imposes a constraint on the modeling of $p_{\text{train}}(\mathbf{y})$. However, in modern deep learning tasks, $p_{\text{train}}(\mathbf{y})$ could be very high-dimensional and has a complex underlying distribution. With the constraint on the distribution modeling, analytically expressing $p_{\text{train}}(\mathbf{y})$ could be challenging. Therefore, we discuss a few numerical approaches, which could be more generally applicable to all types of label data but could bear a larger variance in optimization. In essence, we use Monte Carlo Method (MCM) to approximate $p_{\text{train}}(\mathbf{y})$:

$$\begin{aligned} & \int_Y \mathcal{N}(\mathbf{y}; \mathbf{y}_{\text{pred}}, \sigma_{\text{noise}}^2 \mathbf{I}) \cdot p_{\text{train}}(\mathbf{y}) d\mathbf{y} \\ &= \mathbb{E}_{\mathbf{y} \sim p_{\text{train}}(\mathbf{y})} [\mathcal{N}(\mathbf{y}; \mathbf{y}_{\text{pred}}, \sigma_{\text{noise}}^2 \mathbf{I})] \\ &\approx \frac{1}{N} \sum_{i=1}^N \mathcal{N}(\mathbf{y}_{(i)}; \mathbf{y}_{\text{pred}}, \sigma_{\text{noise}}^2 \mathbf{I}). \end{aligned} \quad (3.13)$$

Batch-based Monte-Carlo (BMC). BMC requires no prior knowledge on $p_{\text{train}}(\mathbf{y})$. It treats all labels in a training batch as random samples from $p_{\text{train}}(\mathbf{y})$. For labels in a training batch $B_{\mathbf{y}} = \{\mathbf{y}_{(1)}, \mathbf{y}_{(2)}, \dots, \mathbf{y}_{(N)}\}$, the loss will be:

$$\begin{aligned} L &= -\log \mathcal{N}(\mathbf{y}; \mathbf{y}_{\text{pred}}, \sigma_{\text{noise}}^2 \mathbf{I}) \\ &+ \log \sum_{i=1}^N \mathcal{N}(\mathbf{y}_{(i)}; \mathbf{y}_{\text{pred}}, \sigma_{\text{noise}}^2 \mathbf{I}). \end{aligned} \quad (3.14)$$

Furthermore, BMC in Eq. 3.14 can be rewritten like Softmax with temperature:

$$L = -\log \frac{\exp(-\|\mathbf{y}_{\text{pred}} - \mathbf{y}\|_2^2 / \tau)}{\sum_{\mathbf{y}' \in B_{\mathbf{y}}} \exp(-\|\mathbf{y}_{\text{pred}} - \mathbf{y}'\|_2^2 / \tau)}, \quad (3.15)$$

where $\tau = 2\sigma_{\text{noise}}^2$ is a temperature coefficient.

BMC is easy to implement. Interestingly, its form in Eq. 3.15 is equivalent to *classifying within a batch*, and shows similarity to contrastive loss functions used in self-training [8, 14]. The similarity could potentially be connected with self-training’s effectiveness in imbalanced learning [19, 37], which we leave for future discussion.

Bin-based Numerical Integration (BNI). Although the “bin” based idea mainly applies to uni-dimensional label space, it allows us to leverage recent progress on estimating label densities using KDE [31, 38]. These prior works first divide the label space into evenly distributed bins, then use KDE to estimate the $p_{\text{train}}(\mathbf{y})$ at the bin centers. We may directly use their results to make a numerical integration. For N bin centers $\{\mathbf{y}_{(1)}, \mathbf{y}_{(2)}, \dots, \mathbf{y}_{(N)}\}$, the loss is:

$$\begin{aligned} L &= -\log \mathcal{N}(\mathbf{y}; \mathbf{y}_{\text{pred}}, \sigma_{\text{noise}}^2 \mathbf{I}) \\ &+ \log \sum_{i=1}^N p_{\text{train}}(\mathbf{y}_{(i)}) \cdot \mathcal{N}(\mathbf{y}_{(i)}; \mathbf{y}_{\text{pred}}, \sigma_{\text{noise}}^2 \mathbf{I}). \end{aligned} \quad (3.16)$$

3.5.3 Finding Optimal Noise Scale

Unlike the standard MSE loss, the noise scale σ_{noise} makes a difference in the proposed method. Locating an optimal noise scale is thus important. A hyper-parameter search on σ_{noise} will be affordable given that σ_{noise} is defined in \mathbb{R}_+ and bounded by the square root of train-time and test-time MSEs. However, in this paper, instead of using hyper-parameter search, we jointly optimize σ_{noise} with \mathbf{y}_{pred} during model training.

We observe that we can obtain near-optimal σ_{noise} by simply setting σ_{noise} as a learnable parameter. A comparison between using the ground truth noise scale and using the jointly learned σ_{noise} is shown in the supplementary material. Therefore, no additional hyper-parameter tuning is required by Balanced MSE, making Balanced MSE more friendly to practitioners. We adopt the joint optimization paradigm in all empirical analyses unless specified.

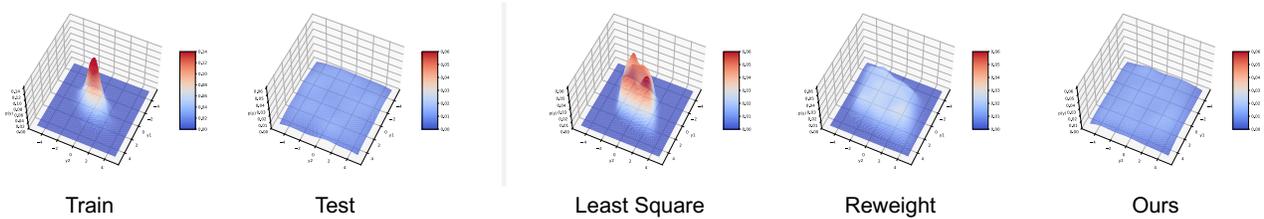


Figure 3. Comparison of marginal label distributions on 2D linear regression. Least square and reweighting show visible bias towards the high-frequency area around the center. In comparison, Balanced MSE achieves the closest marginal label distribution to the uniform test distribution.

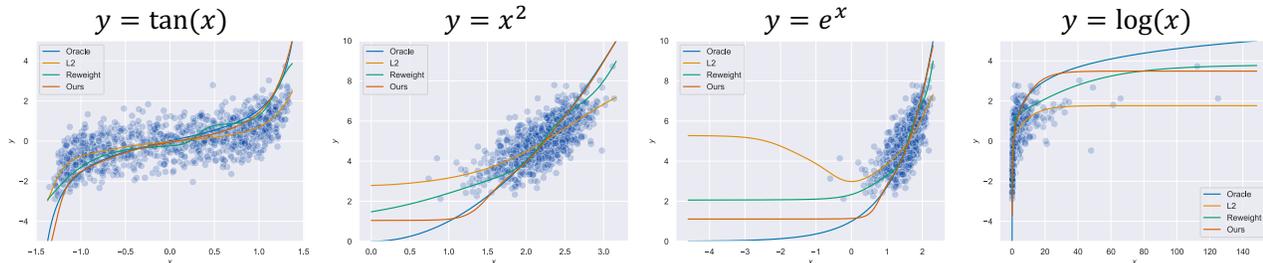


Figure 4. Qualitative comparison for nonlinear regression. Four nonlinear functions are studied. Balanced MSE (in red) gives the closest estimation to the oracle (in blue).

4. Experiments

4.1. Synthetic Benchmarks

We construct a simple one-dimensional linear imbalanced regression dataset, with the training label distribution being normal or exponential and skewed to various extents. We train a one-layer linear regressor on the imbalanced training set and test on a uniform test set with no additive noise. We compare three types of regressors: a least-square estimator, a linear regressor inversely reweighted by the true $p_{\text{train}}(\mathbf{y})$ as described in [38], and Balanced MSE’s closed-form option GAI with true noise scale. We show the visualized results in Fig. 2. We observe that the reweighted regressor shows increasingly larger error when the training distribution becomes more skewed. In comparison, Balanced MSE gives an accurate estimation that is robust to different levels of skewness.

We further compare the three methods on a two-dimensional regression. The training label distribution is set as a Multivariate Normal (MVN) distribution. We visualize the marginal label distributions in Fig. 3, where Balanced MSE achieves a marginal label distribution closest to uniform. For nonlinear regressions, Balanced MSE achieves a consistent effectiveness as well, as shown in Fig. 4. We provide another experiment on random seeds in the supplementary material to demonstrate Balanced MSE’s robustness to noise.

Despite recent works [31, 38] focusing on estimating the training label distribution, our synthetic benchmark shows that *the bottleneck of existing techniques is reweighting*. Even given the true label distribution, reweighting fails to

find the optimal estimator in all settings. Our conclusion aligns with recent research that shows reweighting’s incapability on imbalanced classification [5, 36]. In comparison, our proposed Balanced MSE is robust to different skewness of the training distribution and noise, meanwhile applicable to nonlinear and multi-dimensional regressions.

We provide quantitative results for the above described synthetic benchmark in the supplementary material, where we compare different implementation options and choices of noise scale as well. The results show that the numerical option achieves comparable results with closed-form options. Moreover, the jointly-optimized noise scale achieves near-optimal results in most cases.

4.2. Real-World Benchmarks

4.2.1 Datasets and Settings

Age and Depth Estimation. We select two representative tasks from [38]’s DIR benchmark. We estimate ages from face images on the IMDB-WIKI-DIR dataset and estimate depth maps from images of indoor scenes on the NYUD2-DIR dataset.

Imbalanced Human Mesh Recovery (IHMR). IHMR is a new, multi-dimensional imbalanced regression benchmark. We estimate human meshes from images, where the mesh is represented by a parametric human model known as SMPL [25]. Typically, SMPL model has two parameters: $\theta \in \mathbb{R}^{24 \times 3}$ represents the rotation of 23 body joints and 1 global orientation and $\beta \in \mathbb{R}^{10}$ represents the 10 PCA components for body shape. Therefore, the label space of IHMR is multi-dimensional. Aligned with recent works [30], we observe that the distribution of human meshes is long-tailed.

Table 1. Comparison experiment on IMDB-WIKI-DIR. †: MAE metric reported in [38]. Best results are bolded.

Method	bMAE↓				MAE↓			
	All	Many	Med.	Few	All	Many	Med.	Few
Vanilla†	13.92	7.32	15.93	32.78	8.06	7.23	15.12	26.33
RRT†	13.12	7.27	14.03	30.48	7.81	7.07	14.06	25.13
RRT+LDS†	13.09	7.30	14.05	30.26	7.79	7.08	13.76	24.64
Ours (BMC)	12.69	7.59	12.90	28.28	8.08	7.52	12.47	23.29
Ours (GAI)	12.66	7.65	12.68	28.14	8.12	7.58	12.27	23.05

Table 2. Comparison experiment on NYUD2-DIR. †: reported in [38]. Best results are bolded.

Method	RMSE↓				$\delta_1 \uparrow$			
	All	Many	Med.	Few	All	Many	Med.	Few
Vanilla†	1.477	0.591	0.952	2.123	0.677	0.777	0.693	0.570
Vanilla + LDS†	1.387	0.671	0.913	1.954	0.672	0.701	0.706	0.630
Ours (BNI)	1.283	0.787	0.870	1.736	0.694	0.622	0.806	0.723
Ours (GAI)	1.251	0.692	0.959	1.703	0.702	0.676	0.734	0.715

We show a visualization of training distribution in the supplementary material. Following [23], we train on a combination of 3D and 2D human datasets and test on an in-the-wild 3D dataset. Detailed settings can be found in the supplementary material.

4.2.2 Evaluation Metrics

The DIR benchmark [38] uses primarily overall metric, *e.g.*, Mean Absolute Error (MAE) to report the performance on the benchmark. This is on the assumption that the test dataset is perfectly balanced. However, we observe visible tails in IMDB-WIKI-DIR’s test set as shown in the supplementary material. To fairly measure the model’s performance on tail labels, we follow the idea of balanced metrics [3], and divide the label space into a finite number of even sub-regions, compute the average inside the sub-regions, and take the mean overall sub-regions. We name it “balanced-” (“b-”) metric, *e.g.*, bMAE.

Age & Depth Estimation. We primarily report bMAE on IMDB-WIKI-DIR. NYUD2-DIR’s test set is balanced, we follow [38] and report RMSE.

Imbalanced Human Mesh Recovery (IHMR). We extend HMR’s evaluation metrics to balanced metrics in IHMR. We evenly divide the label space into 100 sub-regions according to their vertex-based distances to the mean parameter and compute balanced metrics as described above. Following [30], we primarily report balanced mean per-vertex position error (bMPVPE). We also report balanced mean per-joint position error (bMPJPE) and balanced Procrustes-aligned mean per joint position error (bPA-MPJPE). We include the “tail 5%” metric and the “tail 10%” metric to show performance on extreme poses as well.

4.2.3 Comparison Results

Tab. 1 shows a comparison with state-of-the-art (SOTA) methods on age estimation. Regressor Re-Training (RRT) [38] first trains the feature extractor normally and retrain the last linear layer using inverse re-weighting. RRT+LDS is an improved version of RRT, where training label distribution is estimated using Label Distribution Smoothing [38]. RRT and RRT+LDS are the best performing regressor learning methods on IMDB-WIKI-DIR in [38]’s benchmark. Balanced MSE substantially outperforms the previous methods. Notably, the BMC option outperforms SOTAs with a large margin without relying on the pre-processed training label distribution. We further analyze the bMAE gain in Fig. 5 and observe an effective trade-off between frequent and rare labels towards a balanced estimation. Note that we do not include Feature Distribution Smoothing (FDS) δ_1 in the comparison since it works on feature learning and should be complementary to our method.

Tab. 2 shows comparison with SOTA on depth estimation. Note that depth map has an inter-pixel dependency, the pixel-wise error σ_{noise} can be under-estimated and the BMC can give an inaccurate estimation to $p_{\text{train}}(\mathbf{y})$. We set a fixed σ_{noise} to 1, and use BNI for numerical option evaluation. Compared with the SOTA, both closed-form and numerical implementations achieve clear improvements.

Tab. 3 shows a comparison between Balanced MSE and existing HMR methods. Balanced MSE outperforms the baseline by a large margin on the main metric bMPVPE (-3.4). We show the qualitative comparison in Fig. 6. PM-Net [30] achieves better results on tail-5% bMPVPE, by designing prototypes and adaptively selecting them as the initialization for SMPL regression. It is noteworthy that

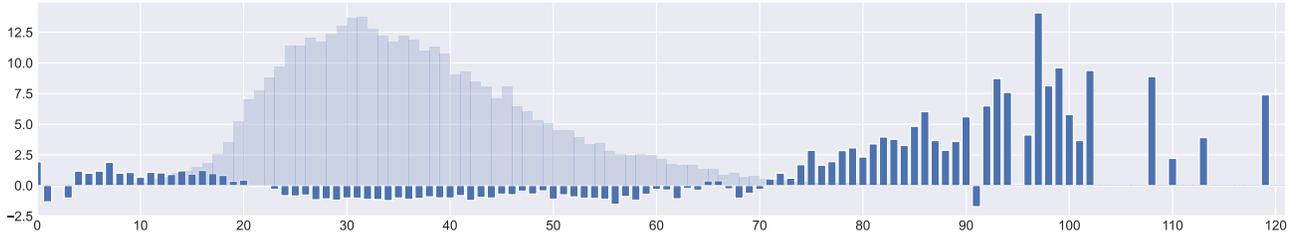


Figure 5. Balanced MSE’s bMAE gain over the baseline. The light blue area in the background shows the training label histogram of IMDB-WIKI-DIR. Balanced MSE improves the performance on tail labels (age < 20 and > 70) substantially.

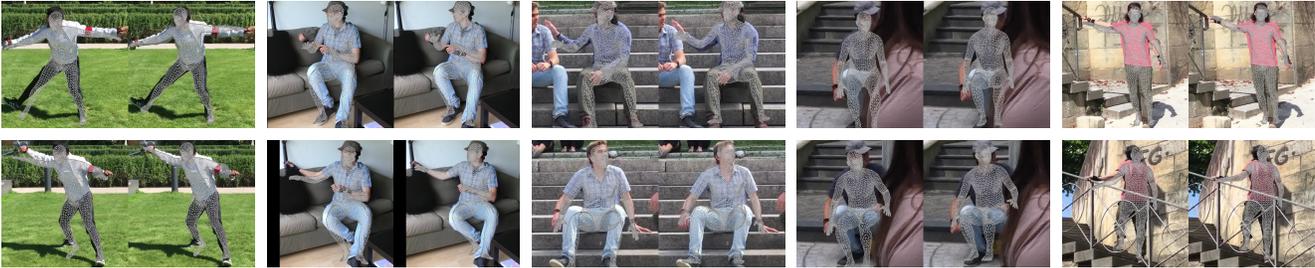


Figure 6. Qualitative comparison of Balanced MSE and the baseline, SPIN-RT. Left: SPIN-RT. Right: Balanced MSE. We observe that the baseline’s predictions are less stretched out. They bias towards the mean pose, particularly for poses like raising arms and bending legs. In comparison, our method effectively eliminates the bias and recovers rare poses.

Table 3. Comparison experiment on Imbalanced Human Mesh Recovery. †: reported in [30]. SPIN-RT: keep the SPIN’s feature extractor fixed and retrain the last linear regression layers. Best results are bolded.

Method	bMPVPE↓			bMPJPE↓			bPA-MPJPE↓		
	All	10%	5%	All	10%	5%	All	10%	5%
SPIN†	-	130.0	130.6	-	-	-	-	-	-
PM-Net†	-	124.9	126.4	-	-	-	-	-	-
SPIN-RT	116.1	127.0	130.5	99.58	113.5	114.5	66.53	77.71	76.66
Ours (BMC)	113.9	128.6	129.6	97.87	113.7	113.0	65.90	77.73	76.35
Ours (GAI)	112.7	122.9	128.1	96.70	108.8	111.9	64.69	74.04	74.35

PM-Net improves the regression initialization and should be complementary to our method.

5. Discussion and Conclusion

In conclusion, we revisit MSE’s probabilistic interpretation and identify its ineffectiveness in imbalanced regression. We therefore propose a statistically principled loss function, Balanced MSE, for imbalanced regression. We further discuss various implementation options of Balanced MSE, including both closed-form options and numerical options. Balanced MSE outperforms existing methods on various uni- and multi-dimensional imbalanced regression benchmarks.

Future Works. Future works may use Balanced MSE as a bridge to introduce more approaches that are developed on the imbalanced classification to the imbalanced regression. For example, Eq. 3.15 can be viewed as Softmax with temperature. Margin-based methods might be introduced to

adjust the pair-wise distances as well. One may also leverage deep generative models, *e.g.*, VAE [22] and GAN [11], to better model $p_{\text{train}}(\mathbf{y})$.

Broader Impacts. Our method only addresses the bias brought by imbalanced label distribution. However, there are still other types of biases in a training dataset besides the mentioned label imbalance. A regressor may still learn those biases and make predictions that produce negative social impacts even with the proposed method applied.

6. Acknowledgment

This work is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG2-PhD-2021-08-018), NTU NAP, MOE AcRF Tier 2 (T2EP20221-0033), and under the RIE2020 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s).

References

- [1] Paula Branco, Luís Torgo, and Rita P Ribeiro. Smogn: a pre-processing approach for imbalanced regression. In *First international workshop on learning with imbalanced domains: Theory and applications*, pages 36–50. PMLR, 2017. 1, 3
- [2] Paula Branco, Luís Torgo, and Rita P Ribeiro. Rebagg: Resampled bagging for imbalanced regression. In *Second International Workshop on Learning with Imbalanced Domains: Theory and Applications*, pages 67–81. PMLR, 2018. 3
- [3] Kay Henning Brodersen, Cheng Soon Ong, Klaas Enno Stephan, and Joachim M Buhmann. The balanced accuracy and its posterior distribution. In *2010 20th international conference on pattern recognition*, pages 3121–3124. IEEE, 2010. 3, 7
- [4] Mateusz Buda, Atsuto Maki, and Maciej A Mazurowski. A systematic study of the class imbalance problem in convolutional neural networks. *Neural Networks*, 106:249–259, 2018. 1
- [5] Jonathon Byrd and Zachary Lipton. What is the effect of importance weighting in deep learning? In *International Conference on Machine Learning*, pages 872–881. PMLR, 2019. 1, 6
- [6] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Aréchéiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. In *NeurIPS*, pages 1565–1576, 2019. 3
- [7] Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–357, 2002. 3
- [8] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. 5
- [9] Peng Chu, Xiao Bian, Shaopeng Liu, and Haibin Ling. Feature space augmentation for long-tailed data. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIX 16*, pages 694–710. Springer, 2020. 3
- [10] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9268–9277, 2019. 3
- [11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. 8
- [12] Agrim Gupta, Piotr Dollár, and Ross B. Girshick. Lvis: A dataset for large vocabulary instance segmentation. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5351–5359, 2019. 1
- [13] Haibo He and Eduardo A Garcia. Learning from imbalanced data. *IEEE Transactions on knowledge and data engineering*, 21(9):1263–1284, 2009. 3
- [14] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9726–9735. Computer Vision Foundation / IEEE, 2020. 5
- [15] Youngkyu Hong, Seungju Han, Kwanghee Choi, Seokjun Seo, Beomsu Kim, and Buru Chang. Disentangling label distribution for long-tailed visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6626–6636, 2021. 1, 3, 4
- [16] Chen Huang, Yining Li, Chen Change Loy, and Xiaoou Tang. Learning deep representation for imbalanced classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5375–5384, 2016. 3
- [17] Muhammad Abdullah Jamal, Matthew Brown, Ming-Hsuan Yang, Liqiang Wang, and Boqing Gong. Rethinking class-balanced methods for long-tailed visual recognition from a domain adaptation perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7610–7619, 2020. 3
- [18] Angjoo Kanazawa, Michael J. Black, David W. Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [19] Bingyi Kang, Yu Li, Sa Xie, Zehuan Yuan, and Jiashi Feng. Exploring balanced feature spaces for representation learning. In *ICLR*. OpenReview.net, 2021. 5
- [20] Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. Decoupling representation and classifier for long-tailed recognition. In *ICLR*. OpenReview.net, 2020. 1
- [21] Jaehyung Kim, Jongheon Jeong, and Jinwoo Shin. M2m: Imbalanced classification via major-to-minor translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13896–13905, 2020. 3
- [22] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *ICLR*, 2014. 8
- [23] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2252–2261, 2019. 7
- [24] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X. Yu. Large-scale long-tailed recognition in an open world. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2532–2541, 2019. 1
- [25] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM transactions on graphics (TOG)*, 34(6):1–16, 2015. 6
- [26] Peter McCullagh and John A. Nelder. *Generalized Linear Models*. Springer, 1989. 3
- [27] Aditya Krishna Menon, Sadeep Jayasumana, Ankit Singh Rawat, Himanshu Jain, Andreas Veit, and Sanjiv Kumar. Long-tail learning via logit adjustment. In *ICLR*. OpenReview.net, 2021. 3, 4
- [28] David A Nix and Andreas S Weigend. Estimating the mean and variance of the target probability distribution. In *Proceedings of 1994 IEEE international conference on neural*

networks (ICNN'94), volume 1, pages 55–60. IEEE, 1994. [3](#)

- [29] Jiawei Ren, Cunjun Yu, Shunan Sheng, Xiao Ma, Haiyu Zhao, Shuai Yi, and Hongsheng Li. Balanced meta-softmax for long-tailed visual recognition. In *NeurIPS*, 2020. [3](#), [4](#)
- [30] Yu Rong, Ziwei Liu, and Chen Change Loy. Chasing the tail in monocular 3d human reconstruction with prototype memory. *arXiv preprint arXiv:2012.14739*, 2020. [6](#), [7](#), [8](#)
- [31] Michael Steininger, Konstantin Kobs, Padraig Davidson, Anna Krause, and Andreas Hotho. Density-based weighting for imbalanced regression. *Machine Learning*, pages 1–25, 2021. [1](#), [3](#), [5](#), [6](#)
- [32] Junjiao Tian, Yen-Cheng Liu, Nathaniel Glaser, Yen-Chang Hsu, and Zsolt Kira. Posterior re-calibration for imbalanced datasets. In *NeurIPS*, 2020. [3](#)
- [33] Luís Torgo, Rita P Ribeiro, Bernhard Pfahringer, and Paula Branco. Smote for regression. In *Portuguese conference on artificial intelligence*, pages 378–389. Springer, 2013. [1](#), [3](#)
- [34] Jiaqi Wang, Wenwei Zhang, Yuhang Zang, Yuhang Cao, Jiangmiao Pang, Tao Gong, Kai Chen, Ziwei Liu, Chen Change Loy, and Dahua Lin. Seesaw loss for long-tailed instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9695–9704, 2021. [3](#)
- [35] Xudong Wang, Long Lian, Zhongqi Miao, Ziwei Liu, and Stella Yu. Long-tailed recognition by routing diverse distribution-aware experts. In *ICLR*. OpenReview.net, 2021. [1](#)
- [36] Da Xu, Yuting Ye, and Chuanwei Ruan. Understanding the role of importance weighting for deep learning. In *ICLR*. OpenReview.net, 2021. [1](#), [6](#)
- [37] Yuzhe Yang and Zhi Xu. Rethinking the value of labels for improving class-imbalanced learning. In *NeurIPS*, 2020. [5](#)
- [38] Yuzhe Yang, Kaiwen Zha, Ying-Cong Chen, Hao Wang, and Dina Katabi. Delving into deep imbalanced regression. In *International Conference on Machine Learning (ICML)*, 2021. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [39] Boyan Zhou, Quan Cui, Xiu-Shen Wei, and Zhao-Min Chen. BBN: bilateral-branch network with cumulative learning for long-tailed visual recognition. In *CVPR*, pages 9716–9725. Computer Vision Foundation / IEEE, 2020. [1](#)