
Delving into Deep Imbalanced Regression

Yuzhe Yang¹ Kaiwen Zha¹ Ying-Cong Chen¹ Hao Wang² Dina Katabi¹

Abstract

Real-world data often exhibit imbalanced distributions, where certain target values have significantly fewer observations. Existing techniques for dealing with imbalanced data focus on targets with categorical indices, i.e., different classes. However, many tasks involve continuous targets, where hard boundaries between classes do not exist. We define Deep Imbalanced Regression (DIR) as learning from such imbalanced data with continuous targets, dealing with potential missing data for certain target values, and generalizing to the entire target range. Motivated by the intrinsic difference between categorical and continuous label space, we propose distribution smoothing for both labels and features, which explicitly acknowledges the effects of nearby targets, and calibrates both label and learned feature distributions. We curate and benchmark large-scale DIR datasets from common real-world tasks in computer vision, natural language processing, and healthcare domains. Extensive experiments verify the superior performance of our strategies. Our work fills the gap in benchmarks and techniques for practical imbalanced regression problems. Code and data are available at: <https://github.com/YyzHarry/imbalanced-regression>.

1. Introduction

Data imbalance is ubiquitous and inherent in the real world. Rather than preserving an ideal uniform distribution over each category, the data often exhibit skewed distributions with a long tail (Buda et al., 2018; Liu et al., 2019), where certain target values have significantly fewer observations. This phenomenon poses great challenges for deep recognition models, and has motivated many prior techniques for addressing data imbalance (Cao et al., 2019; Cui et al., 2019; Huang et al., 2019; Liu et al., 2019; Tang et al., 2020).

¹MIT Computer Science & Artificial Intelligence Laboratory
²Department of Computer Science, Rutgers University. Correspondence to: Yuzhe Yang <yuzhe@mit.edu>.

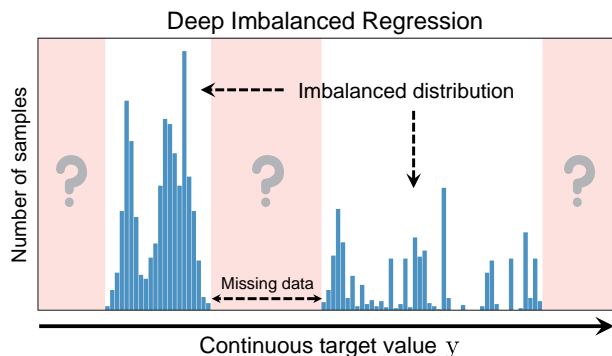


Figure 1. Deep Imbalanced Regression (DIR) aims to learn from imbalanced data with continuous targets, tackle potential missing data for certain regions, and generalize to the entire target range.

Existing solutions for learning from imbalanced data, however, focus on targets with categorical indices, i.e., the targets are different classes. However, many real-world tasks involve continuous and even infinite target values. For example, in vision applications, one needs to infer the age of different people based on their visual appearances, where age is a continuous target and can be highly imbalanced. Treating different ages as distinct classes is unlikely to yield the best results because it does not take advantage of the similarity between people with nearby ages. Similar issues happen in medical applications since many health metrics including heart rate, blood pressure, and oxygen saturation, are continuous and often have skewed distributions across patient populations.

In this work, we systematically investigate *Deep Imbalanced Regression* (DIR) arising in real-world settings (see Fig. 1). We define DIR as learning continuous targets from natural imbalanced data, dealing with potentially missing data for certain target values, and generalizing to a test set that is balanced over the entire range of continuous target values. This definition is analogous to the class imbalance problem (Liu et al., 2019), but focuses on the continuous setting.

DIR brings new challenges distinct from its classification counterpart. First, given continuous (potentially infinite) target values, the hard boundaries between classes no longer exist, causing ambiguity when directly applying traditional imbalanced classification methods such as re-sampling and re-weighting. Moreover, continuous labels inherently possess a meaningful distance between targets, which has im-

plication for how we should interpret data imbalance. For example, say two target labels t_1 and t_2 have a small number of observations in training data. However, t_1 is in a highly represented neighborhood (i.e., there are many samples in the range $[t_1 - \Delta, t_1 + \Delta]$), while t_2 is in a weakly represented neighborhood. In this case, t_1 does not suffer from the same level of imbalance as t_2 . Finally, unlike classification, certain target values may have no data at all, which motivates the need for target extrapolation & interpolation.

In this paper, we propose two simple yet effective methods for addressing DIR: label distribution smoothing (LDS) and feature distribution smoothing (FDS). A key idea underlying both approaches is to leverage the similarity between nearby targets by employing a kernel distribution to perform explicit distribution smoothing in the label and feature spaces. Both techniques can be easily embedded into existing deep networks and allow optimization in an end-to-end fashion. We verify that our techniques not only successfully calibrate for the intrinsic underlying imbalance, but also provide large and consistent gains when combined with other methods.

To support practical evaluation of imbalanced regression, we curate and benchmark large-scale DIR datasets for common real-world tasks in computer vision, natural language processing, and healthcare. They range from single-value prediction such as age, text similarity score, health condition score, to dense-value prediction such as depth. We further set up benchmarks for proper DIR performance evaluation.

Our contributions are as follows:

- We formally define the DIR task as learning from imbalanced data with continuous targets, and generalizing to the entire target range. DIR provides thorough and unbiased evaluation of learning algorithms in practical settings.
- We develop two simple, effective, and interpretable algorithms for DIR, LDS and FDS, which exploit the similarity between nearby targets in both label and feature space.
- We curate benchmark DIR datasets in different domains: computer vision, natural language processing, and healthcare. We set up strong baselines as well as benchmarks for proper DIR performance evaluation.
- Extensive experiments on large-scale DIR datasets verify the consistent and superior performance of our strategies.

2. Related Work

Imbalanced Classification. Much prior work has focused on the imbalanced classification problem (also referred to as long-tailed recognition (Liu et al., 2019)). Past solutions can be divided into data-based and model-based solutions: Data-based solutions either over-sample the minority class or under-sample the majority (Chawla et al., 2002; García & Herrera, 2009; He et al., 2008). For example, SMOTE generates synthetic samples for minority classes by linearly

interpolating samples in the same class (Chawla et al., 2002). Model-based solutions include re-weighting or adjusting the loss function to compensate for class imbalance (Cao et al., 2019; Cui et al., 2019; Dong et al., 2019; Huang et al., 2016; 2019), and leveraging relevant learning paradigms, including transfer learning (Yin et al., 2019), metric learning (Zhang et al., 2017), meta-learning (Shu et al., 2019), and two-stage training (Kang et al., 2020). Recent studies have also discovered that semi-supervised learning and self-supervised learning lead to better imbalanced classification results (Yang & Xu, 2020). In contrast to these past work, we identify the limitations of applying class imbalance methods to regression problems, and introduce new techniques particularly suitable for learning continuous target values.

Imbalanced Regression. Regression over imbalanced data is not as well explored. Most of the work on this topic is a direct adaptation of the SMOTE algorithm to regression scenarios (Branco et al., 2017; 2018; Torgo et al., 2013). Synthetic samples are created for pre-defined rare target regions by either directly interpolating both inputs and targets (Torgo et al., 2013), or using Gaussian noise augmentation (Branco et al., 2017). A bagging-based ensemble method that incorporates multiple data pre-processing steps has also been introduced (Branco et al., 2018). However, there exist several intrinsic drawbacks for these methods. First, they fail to take the distance between targets into account, and rather heuristically divide the dataset into rare and frequent sets, then plug in classification-based methods. Moreover, modern data is of extremely high dimension (e.g., images and physiological signals); linear interpolation of two samples of such data does not lead to meaningful new synthetic samples. Our methods are intrinsically different from past work in their approach. They can be combined with existing methods to improve their performance, as we show in Sec. 4. Further, our approaches are tested on large-scale real-world datasets in computer vision, NLP, and healthcare.

3. Methods

Problem Setting. Let $\{(\mathbf{x}_i, y_i)\}_{i=1}^N$ be a training set, where $\mathbf{x}_i \in \mathbb{R}^d$ denotes the input and $y_i \in \mathbb{R}$ is the label, which is a continuous target. We introduce an additional structure for the label space \mathcal{Y} , where we divide \mathcal{Y} into B groups (bins) with equal intervals, i.e., $[y_0, y_1), [y_1, y_2), \dots, [y_{B-1}, y_B)$. Throughout the paper, we use $b \in \mathcal{B}$ to denote the group index of the target value, where $\mathcal{B} = \{1, \dots, B\} \subset \mathbb{Z}^+$ is the index space. In practice, the defined bins reflect a minimum resolution we care for grouping data in a regression task. For instance, in age estimation, we could define $\delta y \triangleq y_{b+1} - y_b = 1$, showing a minimum age difference of 1 is of interest. Finally, we denote $\mathbf{z} = f(\mathbf{x}; \theta)$ as the feature for \mathbf{x} , where $f(\mathbf{x}; \theta)$ is parameterized by a deep neural network model with parameter θ . The final prediction \hat{y} is given by a regression function $g(\cdot)$ that operates over \mathbf{z} .

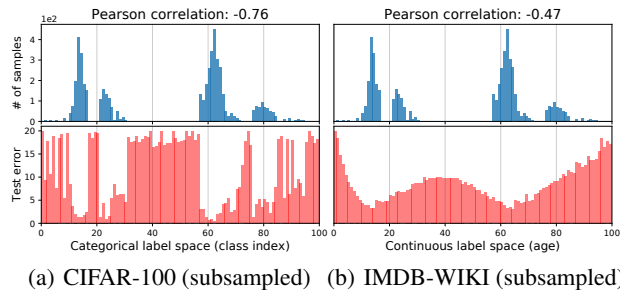


Figure 2. Comparison on the test error distribution (bottom) using same training label distribution (top) on two different datasets: (a) CIFAR-100, a classification task with categorical label space. (b) IMDB-WIKI, a regression task with continuous label space.

3.1. Label Distribution Smoothing

We start by showing an example to demonstrate the difference between classification and regression when imbalance comes into the picture.

Motivating Example. We employ two datasets: (1) CIFAR-100 (Krizhevsky et al., 2009), which is a 100-class classification dataset, and (2) the IMDB-WIKI dataset (Rothe et al., 2018), which is a large-scale image dataset for age estimation from visual appearance. The two datasets have intrinsically different label space: CIFAR-100 exhibits *categorical label space* where the target is class index, while IMDB-WIKI has a *continuous label space* where the target is age. We limit the age range to $0 \sim 99$ so that the two datasets have the same label range, and subsample them to simulate data imbalance, while ensuring they have exactly the same label density distribution (Fig. 2). We make both test sets balanced. We then train a plain ResNet-50 model on the two datasets, and plot their test error distributions.

We observe from Fig. 2(a) that the error distribution *correlates* with label density distribution. Specifically, the test error as a function of class index has a high negative Pearson correlation with the label density distribution (i.e., -0.76) in the categorical label space. The phenomenon is expected, as majority classes with more samples are better learned than minority classes. Interestingly however, as Fig. 2(b) shows, the error distribution is very different for IMDB-WIKI with continuous label space, even when the label density distribution is the same as CIFAR-100. In particular, the error distribution is much smoother and no longer correlates well with the label density distribution (-0.47).

The reason why this example is interesting is that all imbalanced learning methods, directly or indirectly, operate by compensating for the imbalance in the *empirical* label density distribution. This works well for class imbalance, but for continuous labels the empirical density does not accurately reflect the imbalance as seen by the neural network. Hence, compensating for data imbalance based on empirical label density is inaccurate for the continuous label space.

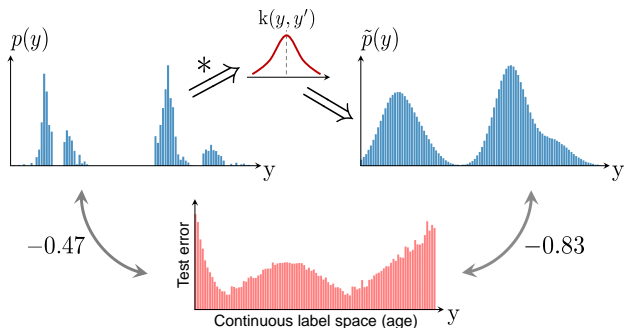


Figure 3. Label distribution smoothing (LDS) convolves a symmetric kernel with the empirical label density to estimate the effective label density distribution that accounts for the continuity of labels.

LDS for Imbalanced Data Density Estimation. The above example shows that, in the continuous case, the empirical label distribution does not reflect the real label density distribution. This is because of the dependence between data samples at nearby labels (e.g., images of close ages). In fact, there is a significant literature in statistics on how to estimate the expected density in such cases (Parzen, 1962). Thus, Label Distribution Smoothing (LDS) advocates the use of kernel density estimation to learn the effective imbalance in datasets that corresponds to continuous targets.

LDS convolves a symmetric kernel with the empirical density distribution to extract a kernel-smoothed version that accounts for the overlap in information of data samples of nearby labels. A symmetric kernel is any kernel that satisfies: $k(y, y') = k(y', y)$ and $\nabla_y k(y, y') + \nabla_{y'} k(y', y) = 0$, $\forall y, y' \in \mathcal{Y}$. Note that a Gaussian or a Laplacian kernel is a symmetric kernel, while $k(y, y') = yy'$ is not. The symmetric kernel characterizes the similarity between target values y' and any y w.r.t. their distance in the target space. Thus, LDS computes the *effective label density distribution* as:

$$\tilde{p}(y') \triangleq \int_{\mathcal{Y}} k(y, y') p(y) dy, \quad (1)$$

where $p(y)$ is the number of appearances of label of y in the training data, and $\tilde{p}(y')$ is the effective density of label y' .

Fig. 3 illustrates LDS and how it smooths the label density distribution. Further, it shows that the resulting label density computed by LDS correlates well with the error distribution (-0.83). This demonstrates that LDS captures the real imbalance that affects regression problems.

Now that the effective label density is available, techniques for addressing class imbalance problems can be directly adapted to the DIR context. For example, a straightforward adaptation can be the cost-sensitive re-weighting method, where we re-weight the loss function by multiplying it by the inverse of the LDS estimated label density for each target. We show in Sec. 4 that LDS can be seamlessly incorporated with a wide range of techniques to boost DIR performance.

3.2. Feature Distribution Smoothing

We are motivated by the intuition that continuity in the target space should create a corresponding continuity in the feature space. That is, if the model works properly and the data is balanced, one expects the feature statistics corresponding to nearby targets to be close to each other.

Motivating Example. We use an illustrative example to highlight the impact of data imbalance on feature statistics in DIR. Again, we use a plain model trained on the images in the IMDB-WIKI dataset to infer a person’s age from visual appearance. We focus on the learned feature space, i.e., \mathbf{z} . We use a minimum bin size of 1, i.e., $y_{b+1} - y_b = 1$, and group features with the same target value in the same bin. We then compute the feature statistics (i.e., mean and variance) with respect to the data in each bin, which we denote as $\{\boldsymbol{\mu}_b, \boldsymbol{\Sigma}_b\}_{b=1}^B$. To visualize the similarity between feature statistics, we select an anchor bin b_0 , and calculate the cosine similarity of the feature statistics between b_0 and all other bins. The results are summarized in Fig. 4 for $b_0 = 30$. The figure also shows the regions with different data densities using the colors purple, yellow, and pink.

Fig. 4 shows that the feature statistics around $b_0 = 30$ are highly similar to their values at $b_0 = 30$. Specifically, the cosine similarity of the feature mean and feature variance for all bins between age 25 and 35 are within a few percent from their values at age 30 (the anchor age). Further, the similarity gets higher for tighter ranges around the anchor. Note that bin 30 falls in the high shot region. In fact, it is among the few bins that have the most samples. So, the figure confirms the intuition that when there is enough data, and for continuous targets, the feature statistics are similar to nearby bins. Interestingly, the figure also shows the problem with regions that have very few data samples, like the age range 0 to 6 years (shown in pink). Note that the mean and variance in this range show unexpectedly high similarity to age 30. In fact, it is shocking that the feature statistics at age 30 are more similar to age 1 than age 17. This unjustified similarity is due to data imbalance. Specifically, since there are not enough images for ages 0 to 6, this range thus inherits its priors from the range with the maximum amount of data, which is the range around age 30.

FDS Algorithm. Inspired by these observations, we propose feature distribution smoothing (FDS), which performs distribution smoothing on the feature space, i.e., transfers the feature statistics between nearby target bins. This procedure aims to calibrate the potentially biased estimates of feature distribution, especially for underrepresented target values (e.g., medium- and few-shot groups) in training data.

FDS is performed by first estimating the statistics of each bin. Without loss of generality, we substitute variance with covariance to reflect also the relationship between the vari-

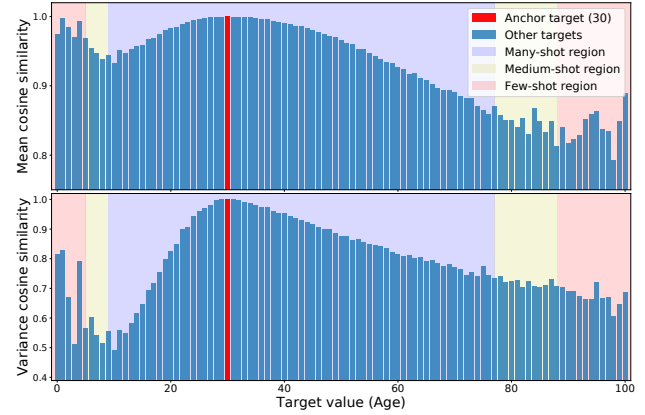


Figure 4. Feature statistics similarity for age 30. **Top:** Cosine similarity of the feature mean at a particular age w.r.t. its value at the anchor age. **Bottom:** Cosine similarity of the feature variance at a particular age w.r.t. its value at the anchor age. The color of the background refers to the data density in a particular target range. The figure shows that nearby ages have close similarities; However, it also shows that there is unjustified similarity between images at ages 0 to 6 and age 30, due to data imbalance.

ous feature elements within \mathbf{z} :

$$\boldsymbol{\mu}_b = \frac{1}{N_b} \sum_{i=1}^{N_b} \mathbf{z}_i, \quad (2)$$

$$\boldsymbol{\Sigma}_b = \frac{1}{N_b - 1} \sum_{i=1}^{N_b} (\mathbf{z}_i - \boldsymbol{\mu}_b)(\mathbf{z}_i - \boldsymbol{\mu}_b)^\top, \quad (3)$$

where N_b is the total number of samples in b -th bin. Given the feature statistics, we employ again a symmetric kernel $k(y_b, y_{b'})$ to smooth the distribution of the feature mean and covariance over the target bins \mathcal{B} . This results in a smoothed version of the statistics:

$$\tilde{\boldsymbol{\mu}}_b = \sum_{b' \in \mathcal{B}} k(y_b, y_{b'}) \boldsymbol{\mu}_{b'}, \quad (4)$$

$$\tilde{\boldsymbol{\Sigma}}_b = \sum_{b' \in \mathcal{B}} k(y_b, y_{b'}) \boldsymbol{\Sigma}_{b'}. \quad (5)$$

With both $\{\boldsymbol{\mu}_b, \boldsymbol{\Sigma}_b\}$ and $\{\tilde{\boldsymbol{\mu}}_b, \tilde{\boldsymbol{\Sigma}}_b\}$, we then follow the standard whitening and re-coloring procedure (Sun et al., 2016) to calibrate the feature representation for each input sample:

$$\tilde{\mathbf{z}} = \tilde{\boldsymbol{\Sigma}}_b^{-\frac{1}{2}} \boldsymbol{\Sigma}_b^{-\frac{1}{2}} (\mathbf{z} - \boldsymbol{\mu}_b) + \tilde{\boldsymbol{\mu}}_b. \quad (6)$$

We integrate FDS into deep networks by inserting a feature calibration layer after the final feature map. To train the model, we employ a *momentum update* of the running statistics $\{\boldsymbol{\mu}_b, \boldsymbol{\Sigma}_b\}$ across each epoch. Correspondingly, the smoothed statistics $\{\tilde{\boldsymbol{\mu}}_b, \tilde{\boldsymbol{\Sigma}}_b\}$ are updated across different epochs but fixed within each training epoch. The momentum update, which performs an exponential moving average (EMA) of running statistics, results in more stable and accurate estimations of the feature statistics during training. The

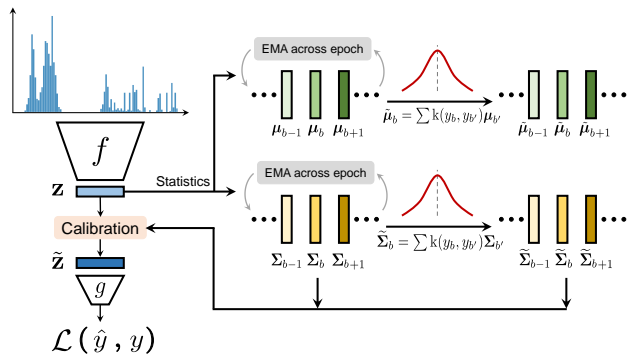


Figure 5. Feature distribution smoothing (FDS) introduces a feature calibration layer that uses kernel smoothing to smooth the distributions of feature mean and covariance over the target space.

calibrated features \tilde{z} are then passed to the final regression function and used to compute the loss.

We note that FDS can be integrated with any neural network model, as well as any past work on improving label imbalance. In Sec. 4, we integrate FDS with a variety of prior techniques for addressing data imbalance, and demonstrate that it consistently improves performance.

4. Benchmarking DIR

Datasets. We curate five DIR benchmarks that span computer vision, natural language processing, and healthcare. Fig. 6 shows the label density distribution of these datasets, and their level of imbalance.

- *IMDB-WIKI-DIR (age)*: We construct IMDB-WIKI-DIR using the IMDB-WIKI dataset (Rothe et al., 2018), which contains 523.0K face images and the corresponding ages. We filter out unqualified images, and manually construct balanced validation and test set over the supported ages. The length of each bin is 1 year, with a minimum age of 0 and a maximum age of 186. The number of images per bin varies between 1 and 7149, exhibiting significant data imbalance. Overall, the curated dataset has 191.5K images for training, 11.0K images for validation and testing.
- *AgeDB-DIR (age)*: AgeDB-DIR is constructed in a similar manner from the AgeDB dataset (Moschoglou et al., 2017). It contains 12.2K images for training, with a minimum age of 0 and a maximum age of 101, and maximum bin density of 353 images and minimum bin density of 1. The validation and test set are balanced with 2.1K images.
- *STS-B-DIR (text similarity score)*: We construct STS-B-DIR from the Semantic Textual Similarity Benchmark (Cer et al., 2017; Wang et al., 2018), which is a collection of sentence pairs drawn from news headlines, video and image captions, and natural language inference data. Each pair is annotated by multiple annotators with an averaged

continuous similarity score from 0 to 5. From the original training set of 7.2K pairs, we create a training set with 5.2K pairs, and balanced validation set and test set of 1K pairs each. The length of each bin is 0.1.

- *NYUD2-DIR (depth)*: We create NYUD2-DIR based on the NYU Depth Dataset V2 (Nathan Silberman & Fergus, 2012), which provides images and depth maps for different indoor scenes. The depth maps have an upper bound of 10 meters and we set the bin length as 0.1 meter. Following standard practices (Bhat et al., 2020; Hu et al., 2019), we use 50K images for training and 654 images for testing. We randomly select 9357 test pixels for each bin to make the test set balanced.
- *SHHS-DIR (health condition score)*: We create SHHS-DIR based on the SHHS dataset (Quan et al., 1997), which contains full-night Polysomnography (PSG) from 2651 subjects. Available PSG signals include Electroencephalography (EEG), Electrocardiography (ECG), and breathing signals (airflow, abdomen, and thorax), which are used as inputs. The dataset also includes the 36-Item Short Form Health Survey (SF-36) (Ware Jr & Sherbourne, 1992) for each subject, where a General Health score is extracted. The score is used as the target value with a minimum score of 0 and maximum of 100.

Network Architectures. We employ ResNet-50 (He et al., 2016) as our backbone network for IMDB-WIKI-DIR and AgeDB-DIR. Following (Wang et al., 2018), we adopt the same BiLSTM + GloVe word embeddings baseline for STS-B-DIR. For NYUD2-DIR, we use ResNet-50-based encoder-decoder architecture introduced in (Hu et al., 2019). Finally, for SHHS-DIR, we use the same CNN-RNN architecture with ResNet block for PSG signals as in (Wang et al., 2019).

Baselines. Since the literature has only a few proposals for DIR, in addition to past work on imbalanced regression (Branco et al., 2017; Torgo et al., 2013), we adapt a few imbalanced classification methods for regression, and propose a strong set of baselines. Below, we describe the baselines, and how we can combine LDS with each method. For FDS, it can be directly integrated with any baseline as a calibration layer, as described in Sec. 3.2.

- *Vanilla model*: We use term **VANILLA** to denote a model that does not include any technique for dealing with imbalanced data. To combine the vanilla model with LDS, we re-weight the loss function by multiplying it by the inverse of the LDS estimated density for each target bin.
- *Synthetic samples*: We choose existing methods for imbalanced regression, including **SMOTER** (Torgo et al., 2013) and **SMOBN** (Branco et al., 2017). SMOTER first defines frequent and rare regions using the original label density, and creates synthetic samples for pre-defined rare

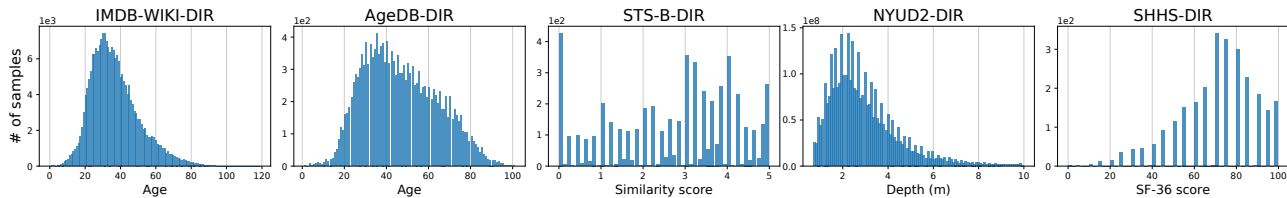


Figure 6. Overview of training set label distribution for five DIR datasets. They range from single-value prediction such as age, textual similarity score, and health condition score, to dense-value prediction such as depth estimation. More details are provided in Appendix B.

regions by linearly interpolating both inputs and targets. SMOGN further adds Gaussian noise to SMOTER. We note that LDS can be directly used for a better estimation of label density when dividing the target space.

- *Error-aware loss*: Inspired by the Focal loss (Lin et al., 2017) for classification, we propose a regression version called **Focal-R**, where the scaling factor is replaced by a continuous function that maps the absolute error into $[0, 1]$. Precisely, Focal-R loss based on L_1 distance can be written as $\frac{1}{n} \sum_{i=1}^n \sigma(|\beta e_i|)^\gamma e_i$, where e_i is the L_1 error for i -th sample, $\sigma(\cdot)$ is the Sigmoid function, and β, γ are hyper-parameters. To combine Focal-R with LDS, we multiply the loss with the inverse frequency of the estimated label density.
- *Two-stage training*: Following (Kang et al., 2020) where feature and classifier are decoupled and trained in two stages, we propose a regression version called regressor re-training (**RRT**), where in the first stage we train the encoder normally, and in the second stage freeze the encoder and re-train the regressor $g(\cdot)$ with inverse re-weighting. When adding LDS, the re-weighting in the second stage is based on the label density estimated through LDS.
- *Cost-sensitive re-weighting*: Since we divide the target space into finite bins, classic re-weighting methods can be directly plugged in. We adopt two re-weighting schemes based on the label distribution: inverse-frequency weighting (**INV**) and its square-root weighting variant (**SQINV**). When combining with LDS, instead of using the original label density, we use the LDS estimated target density.

Evaluation Process and Metrics. Following (Liu et al., 2019), we divide the target space into three disjoint subsets: *many-shot region* (bins with over 100 training samples), *medium-shot region* (bins with 20~100 training samples), and *few-shot region* (bins with under 20 training samples), and report results on these subsets, as well as overall performance. We also refer to regions with no training samples as *zero-shot*, and investigate the ability of our techniques to generalize to zero-shot regions in Sec. 4.2. For metrics, we use common metrics for regression, such as the mean-average-error (MAE), mean-squared-error (MSE), and Pearson correlation. We further propose another metric, called error Geometric Mean (**GM**), and is defined as $(\prod_{i=1}^n e_i)^{\frac{1}{n}}$ for better prediction fairness.

Table 1. Benchmarking results on IMDB-WIKI-DIR.

Metrics	MAE ↓				GM ↓			
	All	Many	Med.	Few	All	Many	Med.	Few
VANILLA	8.06	7.23	15.12	26.33	4.57	4.17	10.59	20.46
SMOTER (Torgo et al., 2013)	8.14	7.42	14.15	25.28	4.64	4.30	9.05	19.46
SMOGN (Branco et al., 2017)	8.03	7.30	14.02	25.93	4.63	4.30	8.74	20.12
SMOGN + LDS	8.02	7.39	13.71	23.22	4.63	4.39	8.71	15.80
SMOGN + FDS	8.03	7.35	14.06	23.44	4.65	4.33	8.87	16.00
SMOGN + LDS + FDS	7.97	7.38	13.22	22.95	4.59	4.39	7.84	14.94
FOCAL-R	7.97	7.12	15.14	26.96	4.49	4.10	10.37	21.20
FOCAL-R + LDS	7.90	7.10	14.72	25.84	4.47	4.09	10.11	19.14
FOCAL-R + FDS	7.96	7.14	14.71	26.06	4.51	4.12	10.16	19.56
FOCAL-R + LDS + FDS	7.88	7.10	14.08	25.75	4.47	4.11	9.32	18.67
RRT	7.81	7.07	14.06	25.13	4.35	4.03	8.91	16.96
RRT + LDS	7.79	7.08	13.76	24.64	4.34	4.02	8.72	16.92
RRT + FDS	7.65	7.02	12.68	23.85	4.31	4.03	7.58	16.28
RRT + LDS + FDS	7.65	7.06	12.41	23.51	4.31	4.07	7.17	15.44
SQINV	7.87	7.24	12.44	22.76	4.47	4.22	7.25	15.10
SQINV + LDS	7.83	7.31	12.43	22.51	4.42	4.19	7.00	13.94
SQINV + FDS	7.83	7.23	12.60	22.37	4.42	4.20	6.93	13.48
SQINV + LDS + FDS	7.78	7.20	12.61	22.19	4.37	4.12	7.39	12.61
OURS (BEST) VS. VANILLA	+0.41	+0.21	+2.71	+4.14	+0.26	+0.15	+3.66	+7.85

4.1. Main Results

We report the main results in this section for all DIR datasets. All training details, hyper-parameter settings, and additional results are provided in Appendix C and D.

Inferring Age from Images: IMDB-WIKI-DIR & AgeDB-DIR. We report the performance of different methods in Table 1 and 2, respectively. For each dataset, we group the baselines into four sections to reflect their different strategies. First, as both tables indicate, when applied to modern high-dimensional data like images, SMOTER and SMOGN can actually degrade the performance in comparison to the vanilla model. Moreover, within each group, adding either LDS, FDS, or both leads to performance gains, while LDS + FDS often achieves the best results. Finally, when compared to the vanilla model, using our LDS and FDS maintains or slightly improves the performance overall and on the many-shot regions, while substantially boosting the performance for the medium-shot and few-shot regions.

Inferring Text Similarity Score: STS-B-DIR. Table 3 shows the results, where similar observations can be made on STS-B-DIR. Again, both SMOTER and SMOGN perform worse than the vanilla model. In contrast, both LDS and FDS consistently and substantially improve the results for various methods, especially in medium- and few-shot re-

Table 2. Benchmarking results on AgeDB-DIR.

Metrics	MAE ↓				GM ↓			
	All	Many	Med.	Few	All	Many	Med.	Few
VANILLA	7.77	6.62	9.55	13.67	5.05	4.23	7.01	10.75
SMOTER (Torgo et al., 2013)	8.16	7.39	8.65	12.28	5.21	4.65	5.69	8.49
SMOBN (Branco et al., 2017)	8.26	7.64	9.01	12.09	5.36	4.90	6.19	8.44
SMOBN + LDS	7.96	7.44	8.64	11.77	5.03	4.68	5.69	7.98
SMOBN + FDS	8.06	7.52	8.75	11.89	5.02	4.66	5.63	8.02
SMOBN + LDS + FDS	7.90	7.32	8.51	11.19	4.98	4.64	5.41	7.35
FOCAL-R	7.64	6.68	9.22	13.00	4.90	4.26	6.39	9.52
FOCAL-R + LDS	7.56	6.67	8.82	12.40	4.82	4.27	5.87	8.83
FOCAL-R + FDS	7.65	6.89	8.70	11.92	4.83	4.32	5.89	8.04
FOCAL-R + LDS + FDS	7.47	6.69	8.30	12.55	4.71	4.25	5.36	8.59
RRT	7.74	6.98	8.79	11.99	5.00	4.50	5.88	8.63
RRT + LDS	7.72	7.00	8.75	11.62	4.98	4.54	5.71	8.27
RRT + FDS	7.70	6.95	8.76	11.86	4.82	4.32	5.83	8.08
RRT + LDS + FDS	7.66	6.99	8.60	11.32	4.80	4.42	5.53	6.99
SQINV	7.81	7.16	8.80	11.20	4.99	4.57	5.73	7.77
SQINV + LDS	7.67	6.98	8.86	10.89	4.85	4.39	5.80	7.45
SQINV + FDS	7.69	7.10	8.86	9.98	4.83	4.41	5.97	6.29
SQINV + LDS + FDS	7.55	7.01	8.24	10.79	4.72	4.36	5.45	6.79
OURS (BEST) VS. VANILLA	+0.30	-0.05	+1.31	+3.69	+0.34	-0.02	+1.65	+4.46

Table 3. Benchmarking results on STS-B-DIR.

Metrics	MSE ↓				Pearson correlation (%) ↑			
	All	Many	Med.	Few	All	Many	Med.	Few
VANILLA	0.974	0.851	1.520	0.984	74.2	72.0	62.7	75.2
SMOTER (Torgo et al., 2013)	1.046	0.924	1.542	1.154	72.6	69.3	65.3	70.6
SMOBN (Branco et al., 2017)	0.990	0.896	1.327	1.175	73.2	70.4	65.5	69.2
SMOBN + LDS	0.962	0.880	1.242	1.155	74.0	71.5	65.2	69.8
SMOBN + FDS	0.987	0.945	1.101	1.153	73.0	69.6	68.5	69.9
SMOBN + LDS + FDS	0.950	0.851	1.327	1.095	74.6	72.1	65.9	71.7
FOCAL-R	0.951	0.843	1.425	0.957	74.6	72.3	61.8	76.4
FOCAL-R + LDS	0.930	0.807	1.449	0.993	75.7	73.9	62.4	75.4
FOCAL-R + FDS	0.920	0.855	1.169	1.008	75.1	72.6	66.4	74.7
FOCAL-R + LDS + FDS	0.940	0.849	1.358	0.916	74.9	72.2	66.3	77.3
RRT	0.964	0.842	1.503	0.978	74.5	72.4	62.3	75.4
RRT + LDS	0.916	0.817	1.344	0.945	75.7	73.5	64.1	76.6
RRT + FDS	0.929	0.857	1.209	1.025	74.9	72.1	67.2	74.0
RRT + LDS + FDS	0.903	0.806	1.323	0.936	76.0	73.8	65.2	76.7
INV	1.005	0.894	1.482	1.046	72.8	70.3	62.5	73.2
INV + LDS	0.914	0.819	1.319	0.955	75.6	73.4	63.8	76.2
INV + FDS	0.927	0.851	1.225	1.012	75.0	72.4	66.6	74.2
INV + LDS + FDS	0.907	0.802	1.363	0.942	76.0	74.0	65.2	76.6
OURS (BEST) VS. VANILLA	+0.71	+0.49	+4.19	+0.68	+1.8	+2.0	+5.8	+2.1

gions. The advantage is even more profound under *Pearson correlation*, which is commonly used for this NLP task.

Inferring Depth: NYUD2-DIR. For NYUD2-DIR, which is a dense regression task, we verify from Table 4 that adding LDS and FDS significantly improves the results. We note that the vanilla model can inevitably overfit to the many-shot regions during training. FDS and LDS help alleviate this effect, and generalize better to all regions, with minor degradation in the many-shot region but significant boosts for other regions.

Inferring Health Score: SHHS-DIR. Table 5 reports the results on SHHS-DIR. Since SMOTER and SMOGN are not directly applicable to this medical data, we skip them for

Table 4. Benchmarking results on NYUD2-DIR.

Metrics	RMSE ↓				δ_1 ↑			
	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
VANILLA + FDS	1.442	0.615	0.940	2.059	0.681	0.760	0.695	0.596
VANILLA + LDS + FDS	1.338	0.670	0.851	1.880	0.705	0.730	0.764	0.655
OURS (BEST) VS. VANILLA	+1.39	-0.24	+1.01	+2.43	+0.28	-0.17	+0.71	+0.85

Table 5. Benchmarking results on SHHS-DIR.

Metrics	MAE ↓				GM ↓			
	All	Many	Med.	Few	All	Many	Med.	Few
VANILLA	15.36	12.47	13.98	16.94	10.63	8.04	9.59	12.20
FOCAL-R	14.67	11.70	13.69	17.06	9.98	7.93	8.85	11.95
FOCAL-R + LDS	14.49	12.01	12.43	16.57	9.98	7.89	8.59	11.40
FOCAL-R + FDS	14.18	11.06	13.56	15.99	9.45	6.95	8.81	11.13
FOCAL-R + LDS + FDS	14.02	11.08	12.24	15.49	9.32	7.18	8.10	10.39
RRT	14.78	12.43	14.01	16.48	10.12	8.05	9.71	11.96
RRT + LDS	14.56	12.08	13.44	16.45	9.89	7.85	9.18	11.82
RRT + FDS	14.36	11.97	13.33	16.08	9.74	7.54	9.20	11.31
RRT + LDS + FDS	14.33	11.96	12.47	15.92	9.63	7.35	8.74	11.17
INV	14.39	11.84	13.12	16.02	9.34	7.73	8.49	11.20
INV + LDS	14.14	11.66	12.77	16.05	9.26	7.64	8.18	11.32
INV + FDS	13.91	11.12	12.29	15.53	8.94	6.91	7.79	10.65
INV + LDS + FDS	13.76	11.12	12.18	15.07	8.70	6.94	7.60	10.18
OURS (BEST) VS. VANILLA	+1.60	+1.41	+1.80	+1.87	+1.93	+1.13	+1.99	+2.02

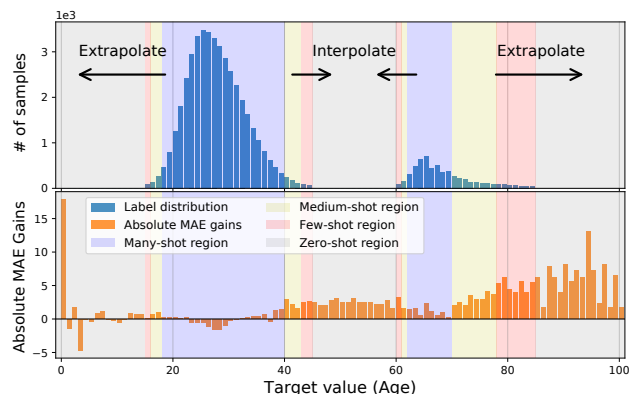
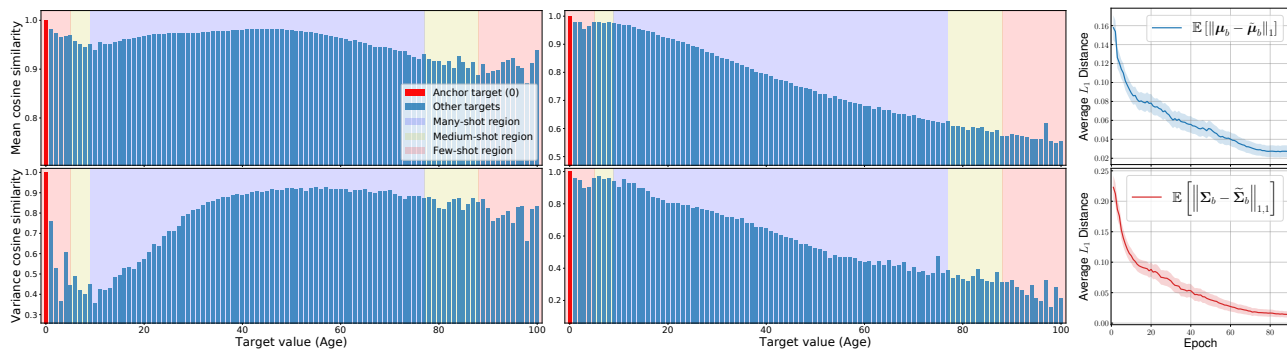


Figure 7. The absolute MAE gains of LDS + FDS over the vanilla model, on a curated subset of IMDB-WIKI-DIR with certain target values having no training data. We establish notable performance gains w.r.t. all regions, especially for extrapolation & interpolation.

this dataset. The results again confirm the effectiveness of both FDS and LDS when applied for real-world imbalanced regression tasks, where by combining FDS and LDS we often get the highest gains over all tested regions.

4.2. Further Analysis

Extrapolation & Interpolation. In real-world DIR tasks, certain target values can have no data at all (e.g., see SHHS-DIR and STS-B-DIR in Fig. 6). This motivates the need for target extrapolation and interpolation. We curate a subset from the training set of IMDB-WIKI-DIR, which has no



(a) Feature statistics similarity for age 0, without FDS (b) Feature statistics similarity for age 0, with FDS (c) Statistics change
 Figure 8. Analysis on how FDS works. (a) & (b) Feature statistics similarity for anchor age 0, using model trained without and with FDS. (c) L_1 distance between the running statistics $\{\mu_b, \Sigma_b\}$ and the smoothed statistics $\{\tilde{\mu}_b, \tilde{\Sigma}_b\}$ during training.

Table 6. Interpolation & extrapolation results on the curated subset of IMDB-WIKI-DIR. Using LDS and FDS, the generalization results on zero-shot regions can be consistently improved.

Metrics	MAE ↓				GM ↓			
	All	w/ data	Interp.	Extrap.	All	w/ data	Interp.	Extrap.
VANILLA	11.72	9.32	16.13	18.19	7.44	5.33	14.41	16.74
VANILLA + LDS	10.54	8.31	14.14	17.38	6.50	4.67	12.13	15.36
VANILLA + FDS	11.40	8.97	15.83	18.01	7.18	5.12	14.02	16.48
VANILLA + LDS + FDS	10.27	8.11	13.71	17.02	6.33	4.55	11.71	15.13
OURS (BEST) vs. VANILLA	+1.45	+1.21	+2.42	+1.17	+1.11	+0.78	+2.70	+1.61

training data in certain regions (Fig. 7), but evaluate on the original testset for zero-shot generalization analysis.

As Table 6 shows, compared to the vanilla model, LDS and FDS can both improve the results not only on regions that have data, but also achieve larger gains on those without data. Specifically, substantial improvements are established for both target interpolation and extrapolation, where interpolation enjoys larger boosts.

We further visualize the absolute MAE gains of our method over vanilla model in Fig. 7. Our method provides a comprehensive treatment to the many, medium, few, as well as zero-shot regions, achieving remarkable performance gains.

Understanding FDS. We investigate how FDS influences the feature statistics. In Fig. 8(a) and 8(b) we plot the similarity of the feature statistics for anchor age 0, using model trained without and with FDS. As the figure indicates, since age 0 lies in the few-shot region, the feature statistics can have a large bias, i.e., age 0 shares large similarity with region 40 ~ 80 as in Fig. 8(a). In contrast, when FDS is added, the statistics are better calibrated, resulting in a high similarity only in its neighborhood, and a gradually decreasing similarity score as target value becomes larger. We further visualize the L_1 distance between the running statistics $\{\mu_b, \Sigma_b\}$ and the smoothed statistics $\{\tilde{\mu}_b, \tilde{\Sigma}_b\}$ during training in Fig. 8(c). Interestingly, the average L_1 distance becomes smaller and gradually diminishes as the

training evolves, indicating that the model learns to generate features that are more accurate even without smoothing, and finally the smoothing module can be removed during inference. We provide more results for different anchor ages in Appendix E.7, where similar effects can be observed.

Ablation: Kernel type for LDS & FDS (Appendix E.1).

We study the effects of different kernel types for LDS and FDS when applying distribution smoothing. We select three different kernel types, i.e., Gaussian, Laplacian, and Triangular kernel, and evaluate their influences on both LDS and FDS. In general, all kernel types lead to notable gains (e.g., 3.7% ~ 6.2% relative MSE gains on STS-B-DIR), with the Gaussian kernel often delivering the best results.

Ablation: Different regression loss functions (Appendix E.2).

We investigate the influence of different training loss functions on LDS and FDS. We select three common losses used for regression tasks, i.e., L_1 loss, MSE loss, and the Huber loss (also referred to as smoothed L_1 loss). We find that similar results are obtained for all losses, indicating that both LDS and FDS are robust to different loss functions.

Ablation: Hyper-parameter for LDS & FDS (Appendix E.3).

We investigate the effects of hyper-parameters on both LDS and FDS. As we mainly employ the Gaussian kernel for distribution smoothing, we extensively study different choices of the kernel size l and standard deviation σ . Interestingly, we find LDS and FDS are surprisingly robust to different hyper-parameters in a given range, and obtain similar gains. For example, on STS-B-DIR with $l \in \{5, 9, 15\}$ and $\sigma \in \{1, 2, 3\}$, overall MSE gains range from 3.3% to 6.2%, with $l = 5$ and $\sigma = 2$ exhibiting the best results.

Ablation: Robustness to diverse skewed label densities (Appendix E.4).

We curate different imbalanced distributions for IMDB-WIKI-DIR by combining different number of disjoint skewed Gaussian distributions over the target space, with potential missing data in certain target regions, and evaluate the robustness of FDS and LDS to the distribution change. We verify that even under different imbalanced

label distributions, LDS and FDS consistently boost the performance across all regions compared to the vanilla model, with relative MAE gains ranging from 8.8% to 12.4%.

Comparisons to imbalanced classification methods (Appendix E.6). Finally, to gain more insights on the intrinsic difference between imbalanced classification & imbalanced regression problems, we directly apply existing imbalanced classification schemes on several appropriate DIR datasets, and show empirical comparisons with imbalanced regression approaches. We demonstrate in Appendix E.6 that LDS and FDS outperform imbalanced classification schemes by a large margin, where the errors for few-shot regions can be reduced by up to 50% to 60%. Interestingly, the results also show that imbalanced classification schemes often perform *worse* than even the vanilla regression model, which confirms that regression requires different approaches for data imbalance than simply applying classification methods. We note that imbalanced classification methods could fail on regression problems for several reasons. First, they ignore the similarity between data samples that are close w.r.t. the continuous target. Moreover, classification cannot extrapolate or interpolate in the continuous label space, therefore unable to deal with missing data in certain target regions.

5. Conclusion

We introduce the DIR task that learns from natural imbalanced data with continuous targets, and generalizes to the entire target range. We propose two simple and effective algorithms for DIR that exploit the similarity between nearby targets in both label and feature spaces. Extensive results on five curated large-scale real-world DIR benchmarks confirm the superior performance of our methods. Our work fills the gap in benchmarks and techniques for practical DIR tasks.

References

- Bhat, S. F., Alhashim, I., and Wonka, P. Adabins: Depth estimation using adaptive bins. *arXiv preprint arXiv:2011.14141*, 2020.
- Branco, P., Torgo, L., and Ribeiro, R. P. Smogn: a pre-processing approach for imbalanced regression. In *First international workshop on learning with imbalanced domains: Theory and applications*, pp. 36–50. PMLR, 2017.
- Branco, P., Torgo, L., and Ribeiro, R. P. Rebagg: Resampled bagging for imbalanced regression. In *Second International Workshop on Learning with Imbalanced Domains: Theory and Applications*, pp. 67–81. PMLR, 2018.
- Buda, M., Maki, A., and Mazurowski, M. A. A systematic study of the class imbalance problem in convolutional neural networks. *Neural Networks*, 106:249–259, 2018.
- Cao, K., Wei, C., Gaidon, A., Arechiga, N., and Ma, T. Learning imbalanced datasets with label-distribution-aware margin loss. In *NeurIPS*, 2019.
- Cer, D., Diab, M., Agirre, E., Lopez-Gazpio, I., and Specia, L. Semeval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation. In *Proceedings of the 11th International Workshop on Semantic Evaluation*, pp. 1–14, 2017.
- Chawla, N. V., Bowyer, K. W., Hall, L. O., and Kegelmeyer, W. P. Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–357, 2002.
- Cui, Y., Jia, M., Lin, T.-Y., Song, Y., and Belongie, S. Class-balanced loss based on effective number of samples. In *CVPR*, 2019.
- Dong, Q., Gong, S., and Zhu, X. Imbalanced deep learning by minority class incremental rectification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(6):1367–1381, Jun 2019.
- Eigen, D., Puhrsch, C., and Fergus, R. Depth map prediction from a single image using a multi-scale deep network. *NeurIPS*, 2014.
- García, S. and Herrera, F. Evolutionary undersampling for classification with imbalanced datasets: Proposals and taxonomy. *Evolutionary computation*, 17(3):275–306, 2009.
- Gardner, M., Grus, J., Neumann, M., Tafjord, O., Dasigi, P., Liu, N. F., Peters, M., Schmitz, M., and Zettlemoyer, L. S. Allennlp: A deep semantic natural language processing platform. 2017.
- He, H., Bai, Y., Garcia, E. A., and Li, S. Adasyn: Adaptive synthetic sampling approach for imbalanced learning. In *IEEE international joint conference on neural networks*, pp. 1322–1328, 2008.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *CVPR*, 2016.
- Hu, J., Ozay, M., Zhang, Y., and Okatani, T. Revisiting single image depth estimation: Toward higher resolution maps with accurate object boundaries. In *WACV*, 2019.
- Huang, C., Li, Y., Change Loy, C., and Tang, X. Learning deep representation for imbalanced classification. In *CVPR*, 2016.
- Huang, C., Li, Y., Chen, C. L., and Tang, X. Deep imbalanced learning for face recognition and attribute prediction. *IEEE transactions on pattern analysis and machine intelligence*, 2019.

- Kang, B., Xie, S., Rohrbach, M., Yan, Z., Gordo, A., Feng, J., and Kalantidis, Y. Decoupling representation and classifier for long-tailed recognition. *ICLR*, 2020.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009.
- Lei, T., Zhang, Y., Wang, S. I., Dai, H., and Artzi, Y. Simple recurrent units for highly parallelizable recurrence. In *EMNLP*, pp. 4470–4481, 2018.
- Lin, T.-Y., Goyal, P., Girshick, R., He, K., and Dollár, P. Focal loss for dense object detection. In *ICCV*, pp. 2980–2988, 2017.
- Liu, Z., Miao, Z., Zhan, X., Wang, J., Gong, B., and Yu, S. X. Large-scale long-tailed recognition in an open world. In *CVPR*, 2019.
- Loper, E. and Bird, S. Nltk: The natural language toolkit. *arXiv preprint cs/0205028*, 2002.
- Moschoglou, S., Papaioannou, A., Sagonas, C., Deng, J., Kotsia, I., and Zafeiriou, S. Agedb: The first manually collected, in-the-wild age database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshop*, volume 2, pp. 5, 2017.
- Nathan Silberman, Derek Hoiem, P. K. and Fergus, R. Indoor segmentation and support inference from rgb-d images. In *ECCV*, 2012.
- Parzen, E. On estimation of a probability density function and mode. *The annals of mathematical statistics*, 33(3): 1065–1076, 1962.
- Pennington, J., Socher, R., and Manning, C. D. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pp. 1532–1543, 2014.
- Quan, S. F., Howard, B. V., Iber, C., Kiley, J. P., Nieto, F. J., O’Connor, G. T., Rapoport, D. M., Redline, S., Robbins, J., Samet, J. M., et al. The sleep heart health study: design, rationale, and methods. *Sleep*, 20(12):1077–1085, 1997.
- Rothe, R., Timofte, R., and Gool, L. V. Deep expectation of real and apparent age from a single image without facial landmarks. *International Journal of Computer Vision*, 126(2-4):144–157, 2018.
- Shu, J., Xie, Q., Yi, L., Zhao, Q., Zhou, S., Xu, Z., and Meng, D. Meta-weight-net: Learning an explicit mapping for sample weighting. *arXiv preprint arXiv:1902.07379*, 2019.
- Sun, B., Feng, J., and Saenko, K. Return of frustratingly easy domain adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, 2016.
- Tang, K., Huang, J., and Zhang, H. Long-tailed classification by keeping the good and removing the bad momentum causal effect. In *NeurIPS*, 2020.
- Torgo, L., Ribeiro, R. P., Pfahringer, B., and Branco, P. Smote for regression. In *Portuguese conference on artificial intelligence*, pp. 378–389. Springer, 2013.
- Verma, V., Lamb, A., Beckham, C., Najafi, A., Mitliagkas, I., Lopez-Paz, D., and Bengio, Y. Manifold mixup: Better representations by interpolating hidden states. In *International Conference on Machine Learning*, 2019.
- Wang, A., Singh, A., Michael, J., Hill, F., Levy, O., and Bowman, S. R. Glue: A multi-task benchmark and analysis platform for natural language understanding. *EMNLP 2018*, pp. 353, 2018.
- Wang, H., Mao, C., He, H., Zhao, M., Jaakkola, T. S., and Katabi, D. Bidirectional inference networks: A class of deep bayesian networks for health profiling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 766–773, 2019.
- Ware Jr, J. E. and Sherbourne, C. D. The mos 36-item short-form health survey (sf-36): I. conceptual framework and item selection. *Medical care*, pp. 473–483, 1992.
- Yang, Y. and Xu, Z. Rethinking the value of labels for improving class-imbalanced learning. In *NeurIPS*, 2020.
- Yin, X., Yu, X., Sohn, K., Liu, X., and Chandraker, M. Feature transfer learning for face recognition with under-represented data. In *In Proceeding of IEEE Computer Vision and Pattern Recognition*, Long Beach, CA, June 2019.
- Zhang, H., Cisse, M., Dauphin, Y. N., and Lopez-Paz, D. mixup: Beyond empirical risk minimization. In *ICLR*, 2018.
- Zhang, X., Fang, Z., Wen, Y., Li, Z., and Qiao, Y. Range loss for deep face recognition with long-tailed training data. In *ICCV*, 2017.