Contents lists available at ScienceDirect







journal homepage: www.elsevier.com/locate/patcog

Discriminative feature generation for classification of imbalanced data

Sungho Suh^{a,b,1}, Paul Lukowicz^{b,c}, Yong Oh Lee^{a,*}

^a Smart Convergence Group, Korea Institute of Science and Technology Europe Forschungsgesellschaft mbH, Saarbrücken 66123, Germany
 ^b Department of Computer Science, TU Kaiserslautern, Kaiserslautern 67663, Germany
 ^c German Research Center for Artificial Intelligence (DFKI), Kaiserslautern 67663, Germany

ARTICLE INFO

Article history: Received 14 October 2020 Revised 28 January 2021 Accepted 2 September 2021 Available online 4 September 2021

Keywords: Imbalanced classification Generative adversarial networks Discriminative feature generation Transfer learning Feature map regularization

ABSTRACT

The data imbalance problem is a frequent bottleneck in the classification performance of neural networks. In this paper, we propose a novel supervised discriminative feature generation (DFG) method for a minority class dataset. DFG is based on the modified structure of a generative adversarial network consisting of four independent networks: generator, discriminator, feature extractor, and classifier. To augment the selected discriminative features of the minority class data by adopting an attention mechanism, the generator for the class-imbalanced target task is trained, and the feature extractor and classifier are regularized using the pre-trained features from a large source data. The experimental results show that the DFG generator enhances the augmentation of the label-preserved and diverse features, and the classification results are significantly improved on the target task. The feature generation model can contribute greatly to the development of data augmentation methods through discriminative feature generation and supervised attention methods.

© 2021 Elsevier Ltd. All rights reserved.

1. Introduction

Deep learning has achieved significant improvements in various tasks in computer vision applications, such as multimodal learning [1], image retrieval [2] and pose estimation [3], using open image datasets containing a large amount of data. However, the acquisition of large datasets is a challenge in real-world applications, especially if they pertain to new areas in deep learning. For training deep neural networks, the quantity of data is as important as the quality because training with only a small dataset often results in degradation of the neural network performance or introduces an over-fitting problem. Attempts have been made to establish a large training dataset in various fields; however, the annotation of a dataset still remains expensive, laborious, and time-consuming. Furthermore, the distribution of classes in the dataset is often imbalanced. Training a network under class-imbalanced conditions can produce a detrimental effect on neural network performance and a biased classification result [4-6].

Data augmentation enhances the size and quality of the training dataset. The methods of data augmentation at the input data level range from simple data augmentation, such as flip, shift, and ro-

* Corresponding author.

¹ https://github.com/opensuh/DFG/

tation, to generative models: the statistical generative models such as the synthetic minority oversampling technique (SMOTE) [7] and its variants [8–10], and the deep learning generative models based on generative adversarial networks (GAN) [11]. GANs, approximating the distribution of the real input data and generate realistic samples from a generative model has the powerful capacity to generate synthetic data so that many recent studies have shown that augmenting a small training dataset using a GAN can further improve the classification performance compared to statistical generative models in real class-imbalanced applications [12-17]. However, such data augmentation at input level using GAN showed the limited improvement in the performance, because balancing the data distribution have week relation to the enhancement of feature extraction in the minority dataset. Also, the state-of-art GAN has the limited generating capability on large-scale images [18,19]. When the size of input image is large, these augmentation method is not applicable.

In the data augmentation using GAN, the adversarial feature augmentation is considered. Unlike data augmentation at the input data level, adversarial feature augmentation generates domaininvariant features, increasing the size of the minority classes in the feature space without considering the modality of the input data [20]. This augmentation was adopted in supervised learning, especially for the class imbalance problem[21]. Although the feature augmentation method performs better than the unsuper-

E-mail address: yongoh.lee@kist-europe.de (Y.O. Lee).

vised domain adaptation method by generating domain-invariant and modality-free features, the generated feature maps from the outer fully connected layer do not have sufficient feature dimensions to represent the data distribution with a small amount of minority class data. In other words, generating features in convolutional layers has a greater capacity to improve classification performance in the case of the class-imbalanced conditions.

Taking these problems into consideration, we propose a novel discriminative feature generation (DFG) method using attention maps in the feature space. The proposed method is a combination of transfer learning and adversarial feature augmentation to complement their drawbacks. The baseline is transfer learning with regularizations, and the features are augmented using a GAN with the weight of the activation level of the feature maps in each class.

The framework is based on a modified GAN structure containing an independent classifier to improve the neural network performance. We extend the adversarial feature augmentation [22] in the following three manners: (1) Our GAN structure is composed of four independent networks: a feature generator, feature discriminator, feature extractor, and feature classifier. The more constituent networks enable the network to be trained for generating more effective features, especially for small sample data, in the feature space. (2) Similar to DELTA [20], we deploy regularizations to the feature extractor and the feature classifier for transfer learning to select the discriminative features from the outer layer outputs and transfer pre-trained knowledge from the largescale source dataset. (3) In the training phase, we employ a feature generation method in a small-sized target dataset, which augments the selected discriminative features by adopting the attention mechanisms for each class label through supervised learning.

Our main contribution is a novel feature generation and augmentation using the proposed GAN structure to unravel the data imbalanced problem and improve neural network performance. In the proposed GAN structure, a feature extractor and a feature classifier are included to train together with a feature generator and a feature discriminator. For the generation of meaningful features for classification of small-sized target data, transfer learning with regularization and class-wise attention is adopted. We evaluate the proposed method with various pairs of source and target dataset to show general applicability to classification in the class-imbalanced conditions.

The remainder of this paper is organized as follows. In Section 2, related works are summarized. The framework and training procedure are presented in Section 3. In Section 4, the benchmark datasets for verification and the experimental results are detailed. Section 5 concludes the paper and provides scope for future research.

2. Related work

GAN [11] is a generative model based on a min-max game theory scenario that pits two networks against one another. A generator network, *G*, competes against a discriminator network, *D*, that distinguishes between samples generated from *G* and samples from the training data. GAN can generate synthetic data close to the original; however, the training process of GAN has the instability of loss function convergence and the problem of mode collapse [23]. To avoid such problems, the Wasserstein GAN (WGAN) [24] and the WGAN with a gradient penalty (WGAN-GP) [25] using the Wasserstein-K distance as the loss function were proposed. Recently, GAN has been extended in several ways to control the generated properties [26–28] and its utilization, such as data augmentation [13,14,29].

Recently, GAN-based data augmentation methods were proposed to further improve the creation of augmented synthetic training data. Huang et al. [15] proposed stacked GAN (SGAN), which is trained to invert the hierarchical representations of a bottom-up discriminative network. Guo et al. [16] proposed a discriminative variational autoencoding adversarial network, which learns a latent two-component mixture distributor and alleviates the class imbalance for deep imbalanced learning. Cui et al. [30] proposed a class-balanced loss for long-tailed distributions. The class-balanced loss re-weights losses inversely with the effective number of samples per class. Although the generated synthetic data can balance the distribution between the classes, the classification performance is limited due to no guarantee enhancing the feature extraction ability of the classifier, when it is an independent network of the GAN. Even if GAN is trained with the classifier together like [31], GAN has the limitation to generate real-looking synthetic data on large-scale images. Not-qualified synthetic data is not able to improve the performance of the classification significantly.

Unlike GAN-based data augmentation at the input data level, Volpi et al. [22] proposed augmentation in the feature space. A feature extractor is trained with the source dataset under supervised learning, and then a feature generator for the unlabeled target dataset is trained in the convolutional GAN (CGAN) framework [26] against the feature extractor. Zhang et al. [21] developed a more general feature generation framework for imbalanced classification, inspired by the adversarial feature augmentation approach. These methods can generate domain-invariant features without considering the modality of the data. Consequently, they achieved better performance than unsupervised domain adaptation methods. However, the improvement of the classification under the class-imbalanced condition is still not significant. The dimension of the features in the feature extractor is considerably lower than the input data. Such a low dimension may not be sufficient to present the data distribution of a small amount of minority class data. The generated features should be domain-invariant and modality-free.

To generate key features from small or minority class data, transfer learning is able to be applied. Transfer learning is a machine learning method that transfers knowledge learned in a source task to a target task [32-34]. The weights of deep neural networks are first pre-trained on a large-scale dataset, which is called the source task, and then fine-tuned using the data from the target task with a small amount of data [33]. In the case of finetuning in convolutional neural networks, the weights in the first few convolution layers are fixed by pre-training. The last convolution layers are fine-tuned by the target task. During fine-tuning, the parameters of the target model can be driven far away from the pre-trained parameter values, leading to the incorporation of information relevant to the targeted problem and overfitting to the target task, so-called catastrophic forgetting [35]. This simple method cannot guarantee good performance because it may in many cases burden the network of the target task with irrelevant information. Yosinki et al. [36] quantified the transferability of features from each layer to learn transferable representations. For constraining catastrophic forgetting in inductive transfer learning, L²-norm regularization was proposed in [37]. The key concept of L^2 -SP is "starting point as reference" optimization, which tries to drive weight parameters to pre-trained values by regularizing the distance between the parameters of source and target tasks.

Instead of regularizing the weight parameters, deep learning transfer using a feature map with attention (DELTA), which is a regularized transfer learning framework, was proposed by Li et al. [20]. Inspired by knowledge distillation for model compression [38–40], DELTA employs the ideas of "inactivated channel re-usage" and feature map regularization with attention. These regularization approaches achieved significant improvement and alleviated the catastrophic forgetting problem by drawing weight parameters close to pre-trained values or aligning transferable channels in feature maps. It constrains the difference between the feature

maps generated by the convolution layers of the source and target networks with attention. DELTA selects the discriminative features from the outer layer outputs using a supervised attention mechanism. To further improve the performance of the classification with a small training dataset or class-imbalanced dataset, a DFG method using attention maps in the feature space is proposed in this paper.

3. Discriminative feature generation

Our goal is to generate discriminative features to improve the performance of the neural network on a small number of training data and the class-imbalanced dataset.

Unlike CGAN and ACGAN[26,27], the structure of the feature classifier to the independent network is employed in the proposed GAN structure. The proposed GAN structure is composed of four independent networks: a feature generator, feature discriminator, feature extractor, and feature classifier. The structure of the proposed GAN model is shown in Fig. 1. The feature extractor represents a first convolution layer or block in a neural network classifier, and the feature classifier represents the rest of the convolution layers and fully connected layers after the feature extractor in the neural network classifier. The feature extractor and classifier are regularized with transferred weights from the source model, similar to [20] for better feature extraction of small sampled classes in the target domain. First, we deploy the feature map regularization following the conceptual ideas of the knowledge distillation technique in the training procedure of GAN. Through this regularization, the feature extraction in the target domain is enhanced. Moreover, we train the feature generator with supervised attention models to generate discriminative features. Applying the supervised attention models increase the selectivity of key features for classification. The feature discriminator is trained to distinguish the real features from the feature extractor and the generated features from the feature generator, and the feature classifier is trained with the real features and the generated features. More detailed explanations about the proposed method are followed.

3.1. Training GAN model with four independent networks

For the stability of the training procedure and the quality of the generated data, we deploy the objective formulation of WGAN-GP for the feature discriminator and the feature generator. The loss functions of the feature discriminator and the feature generator are denoted by \mathbb{L}_D and \mathbb{L}_G , respectively.

$$\begin{split} \mathbb{L}_{D}(x,z,\hat{y};\theta_{D}) &= - \mathop{\mathbb{E}}_{(z,\hat{y})\sim(P_{z},\hat{Y})} \left[D(G(z,\hat{y})) \right] + \mathop{\mathbb{E}}_{x\sim X} \left[D(E(x)) \right] \\ &+ \lambda \mathop{\mathbb{E}}_{\hat{x}\sim P_{z}} \left[(\|\nabla_{\hat{x}} D(\hat{x})\|_{2} - 1)^{2} \right], \end{split}$$
(1)

$$\mathbb{L}_{G}(z,\hat{y};\theta_{G}) = - \mathop{\mathbb{E}}_{(z,\hat{y})\sim(P_{z},\hat{Y})} [D(G(z,\hat{y}))],$$
(2)

where z is the noise vector sampled from uniform distribution P_z , and feature generator G generates synthetic feature G(z). λ is the penalty coefficient, and $P_{\hat{x}}$ is the uniform sampling along straight lines between pairs of points from the real data distribution, P_r , and the generated data distribution. θ_D and θ_G are the parameters of the feature discriminator and the feature generator, respectively. Feature discriminator D is trained to minimize \mathbb{L}_D to distinguish between real and generated features. Generator G is trained to minimize \mathbb{L}_G .

To generate the distribution of features similar to the real features from the feature extractor, a total loss function of the feature classifier contains a loss function of the feature classifier on the features concatenating the real and generated features as well. The objective functions of the feature extractor and classifier are represented as follows:

$$\mathbb{L}_{E}(x, y; \theta_{E}) = \underset{(x, y) \sim (X, Y)}{\mathbb{E}} [-y \log C(E(x))],$$
(3)

$$\mathbb{L}_{C}(x, y, z, \hat{y}; \theta_{C}) = \alpha \mathbb{L}_{C}^{r}(x, y; \theta_{C}) + \beta \mathbb{L}_{C}^{r}(x, y; \theta_{C})$$

$$+ \gamma L_{C}^{c}(\alpha + \beta + \gamma = 1).$$

$$(4)$$

In (4), the loss functions of L_C^r , L_C^g , and L_C^c are defined as

$$\mathbb{L}_{C}^{r}(x, y; \theta_{C}) = \mathbb{E}_{(x, y) \sim (X, Y)}[-y \log C(E(x))],$$
(4a)



Fig. 1. Structure of the proposed model.

Algorithm 1 Training procedure of the proposed DFG method. We use the default values of $\lambda = 10$, $n_D = 5$, $n_{C1} = 2$, $n_{C2} = 10$, and $n_W = 10$ 2000. **Require:** Batch size m, learning rate η , hyperparameter for weight sum ρ , hyperparameters α , β , and γ , and a threshold value for filter weight δ . 1: **Initialize:** θ_E , θ_C from pre-trained source networks θ_{E_S} , θ_{C_S} . 2: $W_j(E_S; \theta_{E_S}) \leftarrow \operatorname{softmax}(-y \log C_S(E_{\theta_{E_S}^{\setminus j}}(x)) + y \log C_S(E_{\theta_{E_S}}(x)))$ 3: **for** $k = 1,..., n_l$ number of class labels **do** $W_i^*(E_S, y = k; \theta_{E_S}) \leftarrow W_j \text{ (if } W_j > \delta/n_l), W_i^*(E_S, y = k; \theta_{E_S}) \leftarrow 0 \text{ (else)}$ 4: 5: **for** step = 1, ..., number of training iteration **do for** $t = 1, ..., n_D$ **do** 6: Sample $\{x^{(i)}\}_{i=1}^m \sim P_r$ a batch from the real data. 7: Sample $\{z^{(i)}\}_{i=1}^{m} \sim P_z$ a generated batch and labels $\{\hat{y}^{(i)}\}_{i=1}^{m}$. Update feature discriminator *D* using Eq.(1): $\theta_D \leftarrow \theta_D - \eta_D \nabla_{\theta_D} \mathbb{L}_D(x, z, \hat{y}; \theta_D)$, 8: 9: Sample $\{z^{(i)}\}_{i=1}^{m} \sim P_{z}$ a generated batch and labels $\{\hat{y}^{(i)}\}_{i=1}^{m}$. Sample $\{x^{(i)}\}_{i=1}^{m} \sim P_{r}$ a batch from the real data and labels $\{y^{(i)}\}_{i=1}^{m}$. Update feature generator *G* using Eq. (2): $\theta_{G} \leftarrow \theta_{G} - \eta_{G} \nabla_{\theta_{G}} \mathbb{L}_{G}(z, \hat{y}, W_{j}^{*}; \theta_{G})$, 10: 11: 12: if $step \equiv 0 \pmod{n_{C1}}$ then 13: Concatenate real and generated features following Eq. (4c): \tilde{x} , \tilde{y} , 14: Update *G* using Eq. (4c): $\theta_G \leftarrow \theta_G - \eta_G \nabla_{\theta_G} \mathbb{L}_C^c(x, y, z, \hat{y}, W_i^*; \theta_G)$, 15: Update feature extractor *E* using Eq. (3): $\theta_E \leftarrow \theta_E - \eta_E \nabla_{\theta_E}^{\mathsf{J}} \mathbb{L}_E(x, y; \theta_E)$, Update feature classifier *C* using Eq. (4a): $\theta_C \leftarrow \theta_C - \eta_C \nabla_{\theta_C}^{\mathsf{J}} \mathbb{L}_C^{\mathsf{r}}(x, y; \theta_C)$, 16: 17: 18: if $step \equiv 0 \pmod{n_{C2}}$ then Concatenate real and generated features following Eq. (4c): \tilde{x} , \tilde{y} , 19: Update *C* using Eq. (4): $\theta_C \leftarrow \theta_C - \eta_C \nabla_{\theta_C} \mathbb{L}_C(x, y, z, \hat{y}, W_i^*; \theta_C)$, 20: **if** iter $\equiv 0 \pmod{n_W}$ **then** 21: Update filter weight using Eq.(5): $W_j(E; \theta_E) \leftarrow \operatorname{softmax}(-y \log C(E_{\theta_E^{\setminus j}}(x)) + y \log C(E_{\theta_E}(x))).$ 22: **for** $k = 1, ..., n_l$ number of class labels **do** 23: $W_j^*(E, y = k; \theta_E) \leftarrow \rho W_j(E_S, y = k; \theta_{E_S}) + (1 - \rho) W_j(E, y = k; \theta_E) \text{ Eq.(7) (if } W_j(E; \theta_E) > \delta/n_l),$ 24: $W_i^*(E, y = k; \theta_E) \leftarrow 0$ (else) 25:

$$\mathbb{L}^{g}_{\mathcal{C}}(z,\hat{y};\theta_{\mathcal{C}}) = \mathbb{E}_{(z,\hat{y})\sim(p_{z}(z),\hat{Y})}[-\hat{y}\log\mathcal{C}(G(z,\hat{y}))], \tag{4b}$$

$$\mathbb{L}_{C}^{c}(x, y, z, \hat{y}; \theta_{C}, \theta_{G}) = \mathbb{E}_{(\tilde{x}, \tilde{y}) \sim (\tilde{X}, \tilde{Y})} [-\tilde{y} \log C(\tilde{x})],$$

where $\tilde{x} = E(x) \oplus G(z, \hat{y}), \quad \tilde{y} = y \oplus \hat{y} \quad (4c)$

where \mathbb{L}_E and \mathbb{L}_C are the loss functions of the feature extractor and feature classifier, respectively. θ_E and θ_C are the parameters E and C, respectively. \oplus denotes a concatenation operation. \mathbb{L}_C^c is the loss function of the feature classifier on the feature maps concatenated with the real and generated features. α , β , and γ are hyperparameters that control the importance of the classification for the real and generated data. In \mathbb{L}_E and \mathbb{L}_C^r , the regularization term, which characterizes the differences between the source and target network for transfer learning, can be added.

Feature generator *G* (to minimize (2) and (4c)) and feature classifier *C* (to minimize (4) and (4a)) are trained simultaneously, whereas feature discriminator *D* and feature extractor *E* are trained to minimize (1) and (3), respectively. The reason why feature classifier *C* is optimized not only with the real data is to prevent the classifier from overfitting to the real data because the performance of the classification under the class-imbalanced condition.

In the training procedure, we define two training parameters. The generator learning parameter controls the ratio between \mathbb{L}_G and \mathbb{L}_G^c while optimizing the generator parameter, θ_G . The classifier learning parameter is used to control the balance of how much the classifier is trained from the feature extracted from the real data \mathbb{L}_G^r or from the real features combined with the generated features,

 \mathbb{L}_{C} . The training details of the proposed method using (1), (2), (3), and (4) with these parameters are summarized in Algorithm 1.

3.2. Discriminative feature generation by using the supervised attention model

To generate discriminative features, we adopt supervised attention mechanisms for each class label. To obtain the weights for feature maps, we propose a supervised attention method adopted from [20] for the generator network. Whereas the supervised attention method in [20] is calculated by averaging the filter weight for each filter, we utilize the fact that the importance of each filter varies from class to class. We transform the filter weights for a single data into class-wise filter weights and deactivate channels with low filter weights.

In the supervised attention method from [20], the weights of the features are characterized by the performance loss when removing the convolutional filter for each feature from the feature extractor network. For a conv2d layer in the feature extractor and generator, the parameter form is a four-dimensional tensor with the shape of (c_{i+1}, c_i, k_h, k_w) , where c_i denotes the number of channels of the *i*-th layer and (k_h, k_w) represents the size of the kernel. At the last convolution layer of the feature extractor, the dimension of the output feature maps is $(c_{i+1}, h_{i+1}, w_{i+1})$, which is matched with the dimension of the output of the feature generator. We can measure significance through the performance reduction of the feature extractor and classifier when the filter in the last convolution layer is disabled in the feature extractor network. In other words, because it usually causes higher performance loss to remove a filter with a greater capacity for discrimination, we can improve the performance of the classification by generating feature maps focusing on channels with high filter weights. The filter weight is calculated through the feature extractor and classifier and is input into the feature generator. The feature generator generates weighted feature maps with deactivated feature maps using the filter weight. The deactivated feature maps are denoted in black in Fig. 1.

The filter weight is expressed in (5), which is used to calculate the gap between the classification losses of the feature extractor and classifier on data (x, y) with and without the *j*-th filter in the last convolution layer of the feature extraction network.

$$W_{j}(E, x, y; \theta_{E}) = \operatorname{softmax}[-y \log C(E_{\theta_{E}^{\setminus j}}(x)) + y \log C(E_{\theta_{E}}(x))],$$
(5)

where $\theta_E^{\setminus j}$ denotes the modified parameter from θ_E with all elements of the *j*-th filter set to zero.

Because each class usually has different importance of the filter channel, the calculated filter weights for a single data point can be transformed into class-wise filter weights.

$$W_j(E,c;\theta_E) = n_f \times \frac{W_j(E,x_i,y_i|y_i=c;\theta_E)}{n_c}$$
(6)

where $W_j(E, c; \theta_E)$ ($W \in \mathcal{R}^{n_c \times n_f}$) is a class-wise filter weight, *c* is a class label, n_c is the number of images in the class, and n_f is the number of filters. By multiplying the number of filters n_f , we set the average value of the weights to 1. When transfer learning is used for the small training dataset and class imbalance problem, a filter weight on the source model can be obtained through feature extractor E_S with parameter θ_{E_S} and classifier C_S with parameter θ_{C_S} trained on the source dataset. Whereas the filter weight in [20] was calculated on the source model only, the proposed method updates the filter weight on the target model. In this case, the filter weight can be expressed in the form of a weighted sum between the filter weight on the source dataset and the target dataset.

$$W_i(E, c; \theta_E) = \rho W_i(E_S, c; \theta_{E_S}) + (1 - \rho) W_i(E, c; \theta_E), \tag{7}$$

By applying the weight to the filters of the last layer in the feature generator, the feature generator can generate discriminative features by following the supervised attention models. The operation of the last layer in the original feature generator is expressed as follows.

$$f^{l}(z,\hat{y}) = \operatorname{norm}[\operatorname{tanh}(\operatorname{TransConv}(f^{l-1}(z,\hat{y}),\theta^{l}_{G}) + b^{l})]$$
(8)

Note that f^l denotes the features of the *l*-th layer in the feature generator, norm() performs batch normalization, TransConv(f, θ) is a transpose convolution that takes f as the input with θ as a weight, and b^l is a bias vector of the *l*-th layer. To generate discriminative features by using the class-wise weights of a filter, the operation can be rewritten as follows.

$$\hat{f}_{j}^{l}(z,\hat{y}) = \operatorname{norm}[W_{j}^{*}(E,c=\hat{y};\theta_{E}) \\ \cdot \operatorname{tanh}(\operatorname{TransConv}_{i}(f^{l-1}(z,\hat{y}),\theta_{C}^{l}) + b_{i}^{l})]$$
(9)

 $\hat{f}^l(z, \hat{y})$ is a discriminative feature generated by the supervised attention model. To maximize the difference in the weights and inactivate unimportant filter channels, we set the weight values that are smaller than a threshold, δ , to zero and represent as W_j^* . δ is set to 0.95 of the average value of all filters for each class.

4. Experimental results

4.1. Datasets and architectures

To evaluate the proposed method, we used the following six benchmark datasets: Street View House Numbers (SVHN) [41], Fashion-MNIST (F-MNIST) [42], STL-10 [43], CINIC-10 [44], Caltech-256 [45], and Food-101 [46]. As a classifier model, LeNet5 [47] for

SVHN and F-MNIST, VGGNet-16 [48] for STL-10 and CINIC-10, and ResNet-50 [49] for Caltech-256 and Food-101 were employed. We used extended MNIST digits (EMNIST) [50], CIFAR-10 [51], and ImageNet [52] as the source domain for LeNet-5, VGG-16, and ResNet-50, respectively. The architecture of our feature discriminator and feature generator closely follow the deep convolutional GAN (DC-GAN) architecture model [53], which is an extended model of the GAN that uses three transpose convolution layers in the generator and three convolution layers in the feature discriminator. A summary of the benchmarked dataset is presented in Table 1. Appendix A contains additional information on benchmarks and hyperparameters. The source code to reproduce our experiments is available at https://github.com/opensuh/DFG.

To evaluate the proposed method under the class-imbalanced condition, we adopt the step imbalance type described in [6]. The step imbalance assumes that the classes are divided into minority and majority groups. Then, classes in the same group have the same number of data points, making a step in the data distribution plot. In this evaluation, we set the number of majority classes as two and eight to evaluate under imbalanced conditions. All combinations of class imbalance cases were tested in the evaluation. For small training datasets such as STL-10 and Caltech-256, we did not create an imbalanced condition. Under the class balanced condition, each experiment was repeated 10 times.

SVHN: To test the transferability in the same domain, we set EMNIST as the source task and SVHN as the target task, where both are single-digit images. To reduce the image dimension gap, the SVHN images are converted to grayscale. For the imbalance 10:1 ratio set-up in training, the number of majority class images is 5000, and one of the minority class images is 500 in SVHN.

F-MNIST: To test the transferability in the different domains, we set EMNIST as the source task and F-MNIST as the target task which has images of fashion and clothing items. For imbalanced 40:1 ratio set-up in training, the number of majority class images is 6000 and that of minority class images is 150 in F-MNIST.

STL-10: STL-10 images are resized to 32×32 to match the image dimension of CIFAR-10. In this case, class imbalance is not applied, but a small training dataset where 500 images per class are used in STL-10.

CINIC-10: CINIC-10 dataset contains 9000 training images in each class and we sampled 900 images from each minor class for imbalance 10:1 ratio.

Caltech-256: Caltech 256 contains 257 object categories and 30,607 images. In this study, we sampled 60 training samples for each category, referring to [20,37]. We resized the input images to 256 \times 256, followed by data augmentation operations of random mirror and random crop to 224 \times 224.

Food-101: Food-101 contains 101 food categories with 101,000 images. A total of 750 training images and 250 test images were provided for each class. We sampled 250 training samples from minority classes for an imbalance of 5:1 ratio set-up during training.

All the experiments were implemented using Python scripts in the PyTorch framework and tested on a Linux system. Training pro-

Summary of the benchmark datasets.

Table 1

Dataset	Classifier	# of training dataset
SVHN	LeNet5	73,237
F-MNIST	LeNet5	60,000
STL-10	VGGNet-16	5,000
CINIC-10	VGGNet-16	18,000
Caltech-256	ImageNet	30,607
Food-101	ImageNet	75,750

Table 2

Comparison results (in %) of our method with competing algorithms. The numbers are represented as *mean*±*std*. Four datasets are under the data-imbalanced condition, and two datasets are under the data-balanced conditions. The number of the training datasets is small.

Dataset (IR)	Original	Fine-tuning	DELTA	DIFA + cMWGAN	Ours(DFG)
SVHN (10:1)	76.57 ± 1.65	76.66 ± 0.65	78.47 ± 0.57	76.64 ± 1.32	$\textbf{80.81} \pm \textbf{0.25}$
F-MNIST (40:1)	77.13 ± 3.24	75.39 ± 3.54	78.91 ± 2.27	78.27 ± 2.32	$\textbf{82.65} \pm \textbf{0.28}$
STL-10 (1:1)	66.73 ± 0.67	72.89 ± 0.38	79.41 ± 0.27	80.08 ± 0.19	$\textbf{81.09} \pm \textbf{0.17}$
CINIC-10 (10:1)	58.14 ± 3.42	64.42 ± 0.81	68.89 ± 0.18	71.82 ± 0.35	$\textbf{72.09} \pm \textbf{0.20}$
Caltech-256 (1:1)	43.30 ± 0.34	81.99 ± 0.12	85.33 ± 0.15	82.23 ± 0.32	$\textbf{87.09} \pm \textbf{0.20}$
Food-101 (5:1)	30.17 ± 1.72	68.99 ± 0.42	72.03 ± 0.35	71.71 ± 0.08	$\textbf{76.00} \pm \textbf{0.36}$

cedures were performed on NVIDIA Tesla V100 GPUs. The training details can be found in Appendix A.

4.2. Results and comparisons

To establish a baseline for comparison, we denote "Original" as the performance of a classification model trained on the target dataset. "Original" is trained under the class-imbalanced condition in the case of SVHN, F-MNIST, CINIC-10, and Food-101. In the case of STL-10 and Caltech-256, "Original" is trained under the class-balanced condition that contains a small number of train data.

As mentioned, "Original" in Table 2 is trained without transfer learning and shows accuracies tested on the target test dataset achieved with classifiers trained on the target training dataset. For the class imbalance problem, in which classes are selected as majority classes, is an important factor in determining the performance of the neural network, and the deviation of result accuracies is large. Thus, accuracies under the class imbalanced condition show the mean and standard deviation.

The proposed method is compared with previous fine-tuning methods.

- Fine-tuning: Traditional transfer learning method by fixing weights before the last convolution block and fine-tuning the rest of them.
- DELTA [20]: The feature map regularization with attention. We used the publicly available source code and parameter settings at https://github.com/lixingjian/DELTA.
- DIFA+cMWGAN: The adversarial feature augmentation method based on [22]. We modified the structures of GAN with cMW-GAN [21].

For a fair comparison with the other methods, we conducted experiments on the same neural network architecture for classification, such as LeNet5, VGGNet-16, and ResNet-50. Table 2 shows the results of the proposed DFG method compared to the fine-tuning, DELTA, and DIFA+cMWGAN methods. The results are the top-1 classification accuracies. The results in Table 2 show that our method based on DFG leads to accuracies higher than other methods for three different classifiers and six different benchmark datasets.

4.3. Qualitative results and ablation study

Although the proposed DFG method shows higher accuracy than the other methods, we would like to better understand the benefits brought by DFG and visualize how closely the proposed DFG generated the distribution of features resembled that of the

Table 3					
Accuracy comparisons of	n	Caltech-256	dataset	with	different
network structures.					

Feature Extractor	Output Dimension	Accuracy (%)
1st block 2nd block 3rd block 4rd block	$56 \times 56 \times 256 28 \times 28 \times 512 14 \times 14 \times 1024 7 \times 7 \times 2048$	$\begin{array}{c} \textbf{87.09} \pm \textbf{0.20} \\ 86.39 \pm 0.16 \\ 86.67 \pm 0.18 \\ 84.48 \pm 0.37 \end{array}$

original in two-dimensional spaces. We apply PCA [54] and t-SNE [55] analyses of real and generated features. Fig. 2 shows comparisons between the real and generated features. Different colors indicate different classes in the left part, and two different colors indicate real and generated features in the right part. From a qualitative point of view, real and generated features are distributed similarly, and the inter-class structure is preserved.

Additionally, we analyzed the activation maps from the proposed DFG method using class activation maps [56]. From Figs. 3 and 4, we observe that the regions of the target object have higher activation than that of the original classifier and fine-tuning. Compared with the activation maps from the DELTA, the proposed method improved the concentration and activation degrees.

We carried out an ablation study to demonstrate the effectiveness of the proposed DFG method. In the proposed DFG method, the feature extractor represents a first convolution laver or block in a neural network classifier, and the feature classifier represents the rest of the convolution layers and fully connected layers after the feature extractor. The feature generator generates the same dimension of the output of the feature extractor. We evaluated the proposed network structure with variants of the feature extractor in the neural network classifier. Table 3 shows the comparisons on Caltech-256 dataset between the proposed structure with the first block in ResNet-50 and its three variants. We observe that the proposed structure outperforms other variants. By extracting and generating the features at the first block, the feature classifier can be trained deeply with the real and generated features and the feature generator can represent diverse features since the dimension of the feature map is large.

Moreover, we perform experiments to better understand how the filter weight changes during the training procedure. We investigate the sensitivity of the proposed loss function with respect to the parameters such as ρ value in (7) and a threshold value δ . Fig. 5 presents the distribution of the weights of two representative filters in the case of the SVHN dataset for two different hyperparameters: $\rho = 0$ and 0.75. The first row in Fig. 5 represents



Fig. 2. t-SNE visualization of feature distribution on (a) SVHN and (b) STL. In the left part, different colors indicate different classes. In the right part, red and cyan dots indicate real and generated features, respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

the distribution of the weights of two representative filters with $\rho = 0.75$, and the second row represents the distribution of the weights of the same filters with $\rho = 0$. The red straight line is the line of equality for all filters. For instance, if a weight value is placed above the line of equality, it means that the importance of the filter for the class is higher than the average of all filters. The case with $\rho = 0$ involves the training of the feature generator with the weights of the filters by updating from the feature extractor trained on the target dataset, and the case with $\rho = 0.75$ involves the training of the feature generator with the weights of the filters by two feature extractors: one is trained on the source dataset and the other is trained on the target dataset in the training procedure. We observed a significant difference between the two distributions of the weights. The distribution of the weights of the filters with $\rho = 0$ converges to the equality baseline, whereas the distribution of the weights of the filters with $\rho = 0.75$ maintains the characteristics of the weight for each class label. Our hypothesis of a possible cause of the difference is that the feature extractor trained on the target dataset tends to be trained to activate all filters and be overfitted to the target dataset. The feature extractor trained on the source dataset inactivates several convolution filters, and a small number of filters are activated. The average accuracy with $\rho = 0.75$ is 80.81% and that with $\rho = 0$ is 79.98%. The weights of the filters with $\rho = 0.75$, using the weight values from the feature extractor trained on the source dataset, prevent overfitting and help to generate discriminative features for each class label.

We have conducted experiments with hyperparameter ρ varying from 0 to 1.0 on the SVHN dataset using DFG. Fig. 6 shows the

effect of hyperparameter ρ on the classification accuracy. In this case, DFG with hyperparameter $\rho = 0.75$ outperformed that with other hyperparameters including the cases that the weights of the filters by the feature extractor only trained on the source or the target dataset. The test accuracy varies smoothly according to hyperparameter ρ , and combination of the two different weights of a filter yields a better accuracy.

Additionally, we have conducted experiments with the threshold parameter δ varying from 0 to 1.5 on the SVHN dataset using DFG. $\delta = 0$ means that the threshold is not applied to the weight values and $\delta = 1.0$ means the threshold value is set to the average value of all filters for each class. Fig. 7 shows the classification accuracy by changing the threshold parameter δ from 0 to 1.5. It shows that DFG with the threshold $\delta = 0.95$ provided that with other threshold values.

To measure the efficiency of the proposed model, the training time on Caltech-256 dataset is computed. Generally, the runtime of the deep learning model is more important than the training time, but the runtimes of the proposed method and other methods are the same because all the methods used the same classification model ResNet-50 for Caltech-256 dataset. Table 4 shows the training time comparison results, and the corresponding parameter amounts and memory. The proposed method took much more time than other methods because the parameter amount and GPU memory of the proposed method are larger than those of other methods. However, the test accuracy of the proposed DFG method provided higher than that of DELTA and DIFA+cMWGAN methods.



Fig. 3. Effect of the discriminative feature generation method on the Caltech-256 dataset: (a) original, (b) fine-tuning, (c) DELTA, and (d) DFG.

Table 4 Comparison of training time and parameter amount of the network structures.						
Method	Parameter amount	GPU training memory (MB)	Training time (sec)	Accuracy (%)		
DELTA	25,557,032	8,187	22,478	85.33 ± 0.15		
DIFA +cMWGAN	38,764,866	7,709	45,232	82.23 ± 0.32		
Ours(DFG)	73,590,722	11,597	62,473	87.09 ± 0.20		



(a)



(c)







(b)

(b)

(c)



(a)

(a)

(b)

(b)

(c)

(d)





(c)

(d)



Fig. 4. Effect of the discriminative feature generation method on the Food-101 dataset: (a) original, (b) fine-tuning, (c) DELTA, and (d) DFG.



Fig. 5. Distribution of the weights of two representative filters for the SVHN dataset. The red straight line is a baseline that indicates the equality of all filters. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)



Fig. 6. Classification accuracy on SVHN using DFG by changing hyperparameter ρ .



Fig. 7. Classification accuracy on SVHN using DFG by changing hyperparameter δ .

5. Conclusions

In this paper, we proposed a novel DFG method using attention maps in the feature space for enhancement in classification under data-imbalanced conditions and small-sampled dataset. The structure of the feature extractor and feature classifier network of the independent network was employed in the proposed GAN structure. To generate discriminative features, we adopted supervised attention mechanisms for each class label and applied the mechanisms to the feature generator network for the filter weights and feature activation. The experimental results showed that the proposed DFG method improved the accuracy of the classification models of LeNet-5, VGG-16, and ResNet-50 by 4.24%, 5.52%, 14.73%, 13.95%, 42.99%, and 45.83% on six different benchmark datasets, respectively. Comparing to the transfer learning method, DELTA, and the other adversarial feature augmentation methods, the proposed DFG method provided the largest improvement in classification performance. Furthermore, we visualized the distribution of generated features by applying PCA and t-SNE and analyzed the activation maps from the proposed method using class activation maps. The qualitative results showed that the features generated by the proposed method are distributed similarly to the real features and the proposed method improved the concentration and activation degrees in the class activation maps.

In the future, we will extend the proposed method to realworld problems where the data imbalance is serious such as document image analysis [17], prognostics and health management [12,14], and chemical property prediction [57]. Also, we will contemplate further accuracy improvement in the feature generation method. Although the proposed method improved the performance of the classification under the class-imbalanced conditions, we believe there is room for improvement as a gap still exists between our method and the balanced condition.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: None.

Declaration of Competing Interest

None.

Acknowledgments

This work was supported by the KIST Europe Institutional Program (Project No. 12120).

Appendix A. Network Architectures and Implementation Details

We implemented the proposed DFG method with LeNet-5 [47], VGGNet-16 [48], and ResNet-50 [49] as the feature extractor and feature classifier. We divided the original architecture of the classifiers into the feature extractor and feature classifier. The architecture of the feature extractor and feature classifier is presented in Tables A.5, A.7, and A.9. The output dimensions of the feature extractor and the feature generator should be the same. The architecture of our feature discriminator and feature generator closely follow the DCGAN architecture model [53], which is an extended model of the GAN that uses 3–4 transpose convolution layers in the generator and three convolution layers in the feature discriminator. The architecture of the generator and the discriminator for LeNet-5, VGGNet-16, and ResNet-50 are shown in Tables A.6, A.8, A.10, respectively.

We used extended MNIST digits (EMNIST) [50], CIFAR-10 [51], and ImageNet [52] as the source domain for LeNet-5, VGG-16, and ResNet-50, respectively. For transfer learning, we followed the same procedure as in [20] because of the close relationship between the behavior regularization in this study and that one. After adopting the pre-trained weight only for ResNet-50 and before fine-tuning the network with the target dataset, we replaced the last layer of the base network with random initialization in a suit

Table A1

Architecture	of	feature	extractor	and	feature	classifier	based	on
LeNet-5 used	l fo	r SVHN	and fashio	n-MN	VIST data	asets.		

Network	Layers	Act. Func.	Dimension
Feature	Input Image	-	$32 \times 32 \times 1$
Extractor	Conv 5×5	ReLU	$28\times28\times6$
	Max Pooling 2 \times 2	-	$14\times14\times6$
Feature	Input Feature	-	$14\times14\times6$
Classifier	Conv 5×5	ReLU	$10\times10\times16$
	Max Pooling 2 \times 2	-	$5 \times 5 \times 16$
	Fully Connected	ReLU	120
	Fully Connected	ReLU	84
	Fully Connected	-	10
	SoftMax	-	10
Feature Classifier	Max Pooling 2 × 2 Input Feature Conv 5 × 5 Max Pooling 2 × 2 Fully Connected Fully Connected Fully Connected SoftMax	- ReLU - ReLU ReLU - -	$ \begin{array}{c} 14 \times 14 \times 6\\ 14 \times 14 \times 6\\ 10 \times 10 \times 1\\ 5 \times 5 \times 16\\ 120\\ 84\\ 10\\ 10\\ \end{array} $

Table A2

Architecture of the generator and the discriminator for LeNet-5 classifier on SVHN and fashion-MNIST datasets.

Network	Layers	Act. func.	Dimension
Feature	Input	-	110
generator	FC	-	$4 \times 4 \times 128$
	BatchNorm	-	$4 \times 4 \times 128$
	DeConv 3 \times 3, st = 2	lReLU	$9 \times 9 \times 48$
	BatchNorm	-	$9 \times 9 \times 48$
	DeConv 3 \times 3, st = 1	lReLU	$11 \times 11 \times 12$
	BatchNorm	-	$11 \times 11 \times 12$
	DeConv 4 \times 4, st = 1	Tanh	$14 \times 14 \times 6$
	BatchNorm	-	$14 \times 14 \times 6$
Feature	Input	-	$14 \times 14 \times 6$
discriminator	Conv 5 \times 5, st = 2	lReLU	$7 \times 7 \times 32$
	Conv 5 \times 5, st = 2	lReLU	$4 \times 4 \times 64$
	Conv 5 \times 5, st = 2	lReLU	$2 \times 2 \times 128$
	FC	-	512

Table A3

Architecture of feature extractor and feature classifier based on VGGNet-16 used for STL-10 and CINIC-10 datasets.

Network	Layers	Act. Func.	Dimension
Feature	Input Image	-	$32\times32\times3$
Extractor	$2 \times Conv 3 \times 3$	ReLU	$32\times32\times64$
	Max Pooling 2 \times 2, st = 2	-	$16\times16\times64$
Feature	Input Feature	-	$16\times16\times64$
Classifier	$2 \times Conv 3 \times 3$	ReLU	$16 \times 16 \times 128$
	Max Pooling 2 \times 2, st = 2	-	$8 \times 8 \times 128$
	$3 \times Conv 3 \times 3$	ReLU	$8\ \times 8\ \times\ 256$
	Max Pooling 2 \times 2, st = 2	-	$4\times4\times256$
	$3 \times Conv 3 \times 3$	ReLU	$4\times4\times512$
	Max Pooling 2 \times 2, st = 2	-	$2\times2\times512$
	$3 \times Conv 3 \times 3$	ReLU	$2 \times 2 \times 512$
	Max Pooling 2 \times 2, st = 2	-	$1 \times 1 \times 512$
	Adaptive Avg Pooling	-	$7 \times 7 \times 512$
	Fully Connected	ReLU	4096
	Fully Connected	ReLU	512
	Fully Connected	-	10
	SoftMax	-	10

Table A4

Architecture of the generator and the discriminator for VGGNet-16 classifier on STL-10 and CINIC-10 datasets.

Network	Layers	Act. func.	Dimension
Feature	Input	-	110
generator	FC	-	$4 \times 4 \times 128$
	BatchNorm	-	$4 \times 4 \times 128$
	DeConv 4 \times 4, st = 2	lReLU	$8 \times 8 \times 128$
	BatchNorm	-	$8 \times 8 \times 128$
	DeConv 4 \times 4, st = 2	lReLU	$16\times16\times128$
	BatchNorm	-	$16\times16\times128$
	DeConv 4 \times 4, st = 1	lReLU	$17 \times 17 \times 64$
	BatchNorm	-	$17\times17\times64$
	DeConv 4 \times 4, st = 1	Tanh	$16\times16\times64$
	BatchNorm	-	$16\times16\times64$
Feature	Input	-	$16\times16\times64$
discriminator	Conv 5 \times 5, st = 2	lReLU	$8 \times 8 \times 32$
	Conv 5 \times 5, st = 2	lReLU	$4 \times 4 \times 64$
	Conv 5 \times 5, st = 2	lReLU	$2 \times 2 \times 128$
	FC	-	512

Table A5

Architecture of feature extractor and feature classifier based on ResNet-50 used for Caltech-256 and Food-101 datasets.

Network	Layers	Act. Func.	Dimension
Feature	Input Image	-	$224\times224\times3$
Extractor	Conv 7 \times 7 st = 2	ReLU	$112 \times 112 \times 64$
	BatchNorm	-	$112\times112\times64$
	Max Pooling 3×3 , st = 2	-	$56 \times 56 \times 64$
	Residual Block #1	ReLU	$56 \times 56 \times 256$
Feature	Input Feature	-	$56 \times 56 \times 256$
Classifier	Residual Block #2	ReLU	$28\times28\times512$
	Residual Block #3	ReLU	$14 \times 14 \times 1024$
	Residual Block #4	ReLU	$7 \times 7 \times 2048$
	Adaptive Avg Pooling	-	$1 \times 1 \times 2048$
	Fully Connected	-	No.Class
	SoftMax	-	No.Class

for the target dataset because the number of classes in the source dataset, ImageNet, is different from those in the target datasets, Caltech-256 [45] and Food-101 [46].

LeNet-5: LeNet-5 model was used for SVHN [41] and Fashion-MNIST (F-MNIST) [42]. To remove the image dimension difference, SVHN images were converted to grayscale and F-MNIST images were resized to 32×32 with zero-padding. The input images were normalized to zero-mean and the pixel values to range [-1, 1]. We did not use any data augmentation method for preprocessing. The weights of the feature extractor and feature classifier were pre-trained using EMNIST. The weights of the feature

Table A6

Architecture of the generator and the discriminator for ResNet-50 classifier on Caltech-256 and Food-101 datasets.

Network	Layers	Act. func.	Dimension
Feature	Input	-	4900 + No.Class
generator	FC	-	$4 \times 4 \times 512$
	BatchNorm	-	$4 \times 4 \times 512$
	DeConv 4 \times 4, st = 2	lReLU	$8 \times 8 \times 512$
	BatchNorm	-	$8 \times 8 \times 512$
	DeConv 4 \times 4, st = 2	lReLU	$16 \times 16 \times 256$
	BatchNorm	-	$16\times16\times256$
	DeConv 2 \times 2, st = 2	lReLU	$30\times30\times256$
	BatchNorm	-	$30\times30\times256$
	DeConv 2 \times 2, st = 2	Tanh	$56 \times 56 \times 256$
	BatchNorm	-	$56 \times 56 \times 256$
Feature	Input	-	$56 \times 56 \times 256$
discriminator	Conv 5 \times 5, st = 2	lReLU	$28\times28\times32$
	Conv 5 \times 5, st = 2	lReLU	$14\times14\times64$
	Conv 5 \times 5, st = 2	lReLU	$7 \times 7 \times 128$
	FC	-	6272
	FC	-	512

generator and feature discriminator were initialized using a Xavier procedure, and the biases were set to zero [58]. We used a batch size of 64 and the Adam optimizer for all networks with parameters $\beta_1 = 0.5$ and $\beta_2 = 0.9$ [59]. The learning rates of the feature extractor, feature classifier, feature generator, and feature discriminator for SVHN were 2.0×10^{-5} , 2.0×10^{-5} , 2.0×10^{-4} , and 1.0×10^{-4} , respectively (η_E , η_C , η_G , and η_D in Algorithm 1.). The learning rates for F-MNIST were $\eta_E = \eta_C = 1.0 \times 10^{-5}$ and $\eta_G = \eta_D = 1.0 \times 10^{-4}$. The training steps were 20,000 iterations. The parameters in Algorithm 1 were set to $n_{C1} = 2$, $n_{C2} = 10$, $n_W = 5000$, $\rho = 0.75$, $\delta = 0.95$, and $\alpha = \beta = \gamma = 0.333$.

VGGNet-16: VGGNet-16 model was used for STL-10 [43] and CINIC-10 [44]. CIFAR-10 (similar class labels to STL-10 and CINIC-10) was used for pre-training as the source domain. We downscaled the 96 \times 96 image dimension of STL-10 to match the 32 \times 32 dimension of CIFAR-10 and CINIC-10. The input images were normalized to zero-mean and the pixel values to range [-1, 1]. We applied data augmentation operations of random horizontal flip and random crop with 4-pixels padding. The weights of the feature extractor and feature classifier were pre-trained with CIFAR-10. The weights of the feature generator and the feature discriminator were initialized using a He procedure, and the biases were set to zero [60]. We used a batch size of 64 and the Adam optimizer for all networks with parameters $\beta_1 = 0.5$ and $\beta_2 = 0.9$. The learning rates for STL-10 were set to $\eta_E = \eta_C = 5.0 \times 10^{-5}$ and $\eta_G = \eta_D = 1.0 \times 10^{-4}$. The learning rates for CINIC-10 were set to $\eta_E = \eta_C = 2.0 \times 10^{-5}$ and $\eta_G = \eta_D = 1.0 \times 10^{-4}$. The training steps for STL-10 and CINIC-10 were 20,000 and 40,000 iterations, respectively. The parameters in Algorithm 1 were set to $n_{C1} = 2$, $n_{C2} = 10$, $n_W = 5000, \ \rho = 0.75, \ \delta = 0.95, \ \text{and} \ \alpha = \beta = \gamma = 0.333.$

ResNet-50: ResNet-50 model was used for Caltech-256 [45] and Food-101 [46]. ImageNet was used as the source domain for pretraining. We replaced the last fully connected layer of ResNet-50 because of the different number of classes in ImageNet, Caltech-256, and Food-101 (1000, 257, and 101, respectively). For ResNet-50, the input images were normalized to zero-mean and the pixel values to range [-1, 1] and resized to 256×256 . We applied the data augmentation operation of the random crop to 224×224 and random horizontal flip. For the Caltech-256 dataset, we sampled 60 training samples and 20 testing samples for each category, according to [20,37]. The weights of the feature extractor and the feature classifier were initialized by the model pre-trained on ImageNet, provided by Torchvision [61], and the feature generator and the feature discriminator were initialized using a He initialization, and the biases were set to zero. We used a batch size of 64 and the Adam optimizer with parameters $\beta_1 = 0.5$ and $\beta_2 = 0.9$ for the feature generator and the feature discriminator, and stochastic gradient descent with a momentum of 0.9 for optimizing the feature extractor and the feature classifier. The learning rate for the feature generator and the feature discriminator was set to $\eta_G = \eta_D = 1.0 \times 10^{-4}$, and the learning rate for the feature classifier started with 0.01 and was divided by 10 after 20,000 iterations. The training steps were completed at 40,000 iterations. The parameters in Algorithm 1 were set to $n_{C1} = 1$, $n_{C2} = 5$, $n_W = 2000$, $\rho = 0.75$, $\delta = 0.95$, and $\alpha = \beta = \gamma = 0.333$.

Additionally, we used the following publicly available source code to implement our benchmarks.

DELTA [20]: https://github.com/lixingjian/DELTA

DIFA+cMWGAN [21,22]:

https://github.com/ricvolpi/adversarial-feature-augmentation

References

- J. Yu, Y. Rui, D. Tao, Click prediction for web image reranking using multimodal sparse coding, IEEE Trans. Image Process. 23 (5) (2014) 2019–2032.
- [2] J. Yu, D. Tao, M. Wang, Y. Rui, Learning to rank using user clicks and visual features for image retrieval, IEEE Trans Cybern 45 (4) (2014) 767–779.
- [3] C. Hong, J. Yu, J. Zhang, X. Jin, K.-H. Lee, Multimodal face-pose estimation with multitask manifold deep learning, IEEE Trans. Ind. Inf. 15 (7) (2018) 3952–3961.
- [4] N. Japkowicz, S. Stephen, The class imbalance problem: a systematic study, Intell. Data Anal. 6 (5) (2002) 429–449.
- [5] J. Xie, Z. Qiu, The effect of imbalanced data sets on LDA: atheoretical and empirical analysis, Pattern Recognit 40 (2) (2007) 557–562.
- [6] M. Buda, A. Maki, M.A. Mazurowski, A systematic study of the class imbalance problem in convolutional neural networks, Neural Networks 106 (2018) 249–259.
- [7] N.V. Chawla, K.W. Bowyer, L.O. Hall, W.P. Kegelmeyer, Smote: synthetic minority over-sampling technique, Journal of artificial intelligence research 16 (2002) 321–357.
- [8] H. Han, W.-Y. Wang, B.-H. Mao, Borderline-SMOTE: a new over-sampling method in imbalanced data sets learning, in: International Conference on Intelligent Computing, Springer, 2005, pp. 878–887.
- [9] H. He, Y. Bai, E.A. Garcia, S. Li, Adasyn: Adaptive synthetic sampling approach for imbalanced learning, in: 2008 IEEE International Joint Conference on Neural Networks (IEEE World Congress on Computational Intelligence), IEEE, 2008, pp. 1322–1328.
- [10] S. Barua, M.M. Islam, X. Yao, K. Murase, Mwmote–majority weighted minority oversampling technique for imbalanced data set learning, IEEE Trans Knowl Data Eng 26 (2) (2012) 405–425.
- [11] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, Generative adversarial nets, in: Advances in neural information processing systems, 2014, pp. 2672–2680.
- [12] Y.O. Lee, J. Jo, J. Hwang, Application of deep neural network and generative adversarial network to industrial maintenance: A case study of induction motor fault detection, in: 2017 IEEE International Conference on Big Data (Big Data), IEEE, 2017, pp. 3248–3253.
- [13] G. Douzas, F. Bacao, Effective data generation for imbalanced learning using conditional generative adversarial networks, Expert Syst Appl 91 (2018) 464–471.
- [14] S. Suh, H. Lee, J. Jo, P. Lukowicz, Y.O. Lee, Generative oversampling method for imbalanced data on bearing fault detection and diagnosis, Applied Sciences 9 (4) (2019) 746.
- [15] X. Huang, Y. Li, O. Poursaeed, J. Hopcroft, S. Belongie, Stacked generative adversarial networks, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2017, pp. 5077–5086.
- [16] T. Guo, X. Zhu, Y. Wang, F. Chen, Discriminative sample generation for deep imbalanced learning, in: IJCAI, 2019, pp. 2406–2412.
- [17] S. Suh, J. Kim, P. Lukowicz, Y.O. Lee, Two-stage generative adversarial networks for document image binarization with color noise and background removal, arXiv preprint arXiv:2010.10103 (2020).
- [18] M. Lucic, K. Kurach, M. Michalski, S. Gelly, O. Bousquet, Are gans created equal? a large-scale study, in: Advances in Neural Information Processing Systems, 2018, pp. 700–709.
- [19] A. Brock, J. Donahue, K. Simonyan, Large scale GAN training for high fidelity natural image synthesis, in: International Conference on Learning Representations, 2019.
- [20] X. Li, H. Xiong, H. Wang, Y. Rao, L. Liu, J. Huan, DELTA: deep learning transfer using feature map with attention for convolutional networks, in: International Conference on Learning Representations, 2019.
- [21] Y. Zhang, B. Sun, Y. Xiao, R. Xiao, Y. Wei, Feature augmentation for imbalanced classification with conditional mixture wgans, Signal Process. Image Commun. 75 (2019) 89–99.

- [22] R. Volpi, P. Morerio, S. Savarese, V. Murino, Adversarial feature augmentation for unsupervised domain adaptation, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 5495–5504.
- [23] I. Goodfellow, Nips 2016 tutorial: generative adversarial networks, arXiv preprint arXiv:1701.00160 (2016).
- M. Arjovsky, S. Chintala, L. Bottou, Wasserstein generative adversarial networks, in: International Conference on Machine Learning, 2017, pp. 214–223.
 I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, A.C. Courville, Improved
- [25] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, A.C. Courville, Improved training of wasserstein gans, in: Advances in Neural Information Processing Systems, 2017, pp. 5767–5777.
- [26] M. Mirza, S. Osindero, Conditional generative adversarial nets, arXiv preprint arXiv:1411.1784 (2014).
- [27] A. Odena, C. Olah, J. Shlens, Conditional image synthesis with auxiliary classifier gans, in: Proceedings of the 34th International Conference on Machine Learning-Volume 70, JMLR. org, 2017, pp. 2642–2651.
- [28] M. Lee, J. Seok, Controllable generative adversarial network, IEEE Access 7 (2019) 28158–28169.
- [29] T.W. Cenggoro, et al., Deep learning for imbalance data classification using class expert generative adversarial network, Procedia Comput Sci 135 (2018) 60–67.
- [30] Y. Cui, M. Jia, T.-Y. Lin, Y. Song, S. Belongie, Class-balanced loss based on effective number of samples, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2019, pp. 9268–9277.
- [31] S. Suh, H. Lee, P. Lukowicz, Y.O. Lee, CEGAN: Classification enhancement generative adversarial networks for unraveling data imbalance problems, Neural Networks (2020).
- [32] R. Caruana, Multitask learning, Mach Learn 28 (1) (1997) 41-75.
- [33] S.J. Pan, Q. Yang, A survey on transfer learning, IEEE Trans Knowl Data Eng 22 (10) (2009) 1345–1359.
- [34] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, T. Darrell, Decaf: A deep convolutional activation feature for generic visual recognition, in: International Conference on Machine Learning, 2014, pp. 647–655.
- [35] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A.A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska, et al., Overcoming catastrophic forgetting in neural networks, Proceedings of the national academy of sciences 114 (13) (2017) 3521–3526.
- [36] J. Yosinski, J. Clune, Y. Bengio, H. Lipson, How transferable are features in deep neural networks? in: Advances in Neural Information Processing Systems, 2014, pp. 3320–3328.
- [37] X. Li, Y. Grandvalet, F. Davoine, Explicit inductive bias for transfer learning with convolutional networks, in: International Conference on Machine Learning, 2018, pp. 2825–2834.
- [38] G. Hinton, O. Vinyals, J. Dean, Distilling the knowledge in a neural network, arXiv preprint arXiv:1503.02531 (2015).
- [39] S. Zagoruyko, N. Komodakis, Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer, ICLR, 2017. https://arxiv.org/abs/1612.03928
- [40] J. Yim, D. Joo, J. Bae, J. Kim, A gift from knowledge distillation: Fast optimization, network minimization and transfer learning, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2017, pp. 4133–4141.
- [41] Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, A.Y. Ng, Reading digits in natural images with unsupervised feature learning (2011).
- [42] H. Xiao, K. Rasul, R. Vollgraf, Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, arXiv preprint arXiv:1708.07747 (2017).
- [43] A. Coates, A. Ng, H. Lee, An analysis of single-layer networks in unsupervised feature learning, in: Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics, 2011, pp. 215–223.
- [44] L.N. Darlow, E.J. Crowley, A. Antoniou, A.J. Storkey, Cinic-10 is not imagenet or cifar-10, arXiv preprint arXiv:1810.03505 (2018).
- [45] G. Griffin, A. Holub, P. Perona, Caltech-256 object category dataset(2007).
- [46] L. Bossard, M. Guillaumin, L. Van Gool, Food-101-mining discriminative components with random forests, in: European Conference on Computer Vision, Springer, 2014, pp. 446–461.
- [47] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, et al., Gradient-based learning applied to document recognition, Proc. IEEE 86 (11) (1998) 2278–2324.
- [48] K. Simonyan, A. Zisserman, Very deep convolutional networks for large-scale image recognition, arXiv preprint arXiv:1409.1556 (2014).
- [49] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 770–778.
- [50] G. Cohen, S. Afshar, J. Tapson, A. van Schaik, Emnist: an extension of mnist to handwritten letters, arXiv preprint arXiv:1702.05373 (2017).

- [51] A. Krizhevsky, G. Hinton, Learning multiple layers of features from tiny images, Technical Report, Citeseer, 2009.
- [52] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, L. Fei-Fei, Imagenet: A large-scale hierarchical image database, in: 2009 IEEE Conference on Computer Vision and Pattern Recognition, IEEE, 2009, pp. 248–255.
- [53] A. Radford, L. Metz, S. Chintala, Unsupervised representation learning with deep convolutional generative adversarial networks, arXiv preprint arXiv:1511.06434 (2015).
- [54] S. Wold, K. Esbensen, P. Geladi, Principal component analysis, Chemometrics and intelligent laboratory systems 2 (1-3) (1987) 37-52.
 [55] L.v.d. Maaten, G. Hinton, Visualizing data using t-SNE, Journal of machine
- [55] L.v.d. Maaten, G. Hinton, Visualizing data using t-SNE, Journal of machine learning research 9 (Nov) (2008) 2579–2605.
- [56] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, A. Torralba, Learning deep features for discriminative localization, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 2921–2929.
- [57] O. Prykhodko, S.V. Johansson, P.-C. Kotsias, J. Arús-Pous, E.J. Bjerrum, O. Engkvist, H. Chen, A de novo molecular generation method using latent vector based generative adversarial network, J Cheminform 11 (1) (2019) 1–13.
- [58] X. Glorot, Y. Bengio, Understanding the difficulty of training deep feedforward neural networks, in: Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, 2010, pp. 249–256.
- [59] D.P. Kingma, J. Ba, Adam: a method for stochastic optimization, arXiv preprint arXiv:1412.6980 (2014).
- [60] K. He, X. Zhang, S. Ren, J. Sun, Delving deep into rectifiers: Surpassing human-level performance on imagenet classification, in: The IEEE International Conference on Computer Vision (ICCV), 2015.
- [61] S. Marcel, Y. Rodriguez, Torchvision the machine-vision package of torch, in: Proceedings of the 18th ACM International Conference on Multimedia, in: MM '10, Association for Computing Machinery, New York, NY, USA, 2010, pp. 1485– 1488, doi:10.1145/1873951.1874254.



Sungho Suh is currently working toward the Ph.D. degree at the University of Kaiserslautern and is a Ph.D. student in the Smart Convergence Group, KIST Europe, Germany. He received the B.S. and M.S. degrees from the School of Electrical Engineering and Computer Science, Seoul National University, Seoul, South Korea, in 2009 and 2011, respectively. From 2011 to 2018, he was a software engineer at Samsung Electro-Mechanics, South Korea. His research interests include computer vision, machine vision, machine learning, and their applications.



Paul Lukowicz is both Scientific Director at the German Research Center for Artificial Intelligence (DFKI GmbH) and Professor of Computer Science at Kaiserslautern University (TUK) in Germany since 2012 where he heads the Embedded Intelligence group. His research focuses on context-aware ubiquitous and wearable systems including sensing, pattern recognition, system architectures, models of large-scale self-organized systems, and applications in areas ranging from healthcare through industry 4.0 to smart cities.



Yong Oh Lee is a group leader and principal researcher of the Smart Convergence Group at KIST Europe in Saarbrücken, Germany. He received the B.S. and M.S. degrees in electrical engineering from Yonsei University, Seoul, South Korea, in 2005 and 2007, respectively. He received Ph.D. degree in computer engineering from Texas A&M University, College Station, Texas, U.S.A., in 2012. Before joining KIST Europe, he worked at the IoT Vertical Solution Lab and and the next-generation communications team at the Samsung DMC R&D Center, Suwon-si, South Korea. His research interests are deep learning applications of computer vision and deep generative models for data imbalance problems.