

LA-Net: Landmark-Aware Learning for Reliable Facial Expression Recognition under Label Noise

Zhiyu Wu, Jinshi Cui*

School of Intelligence Science and Technology, Peking University

wuzhiyu@pku.edu.cn, cjs@cis.pku.edu.cn

Abstract

Facial expression recognition (FER) remains a challenging task due to the ambiguity of expressions. The derived noisy labels significantly harm the performance in real-world scenarios. To address this issue, we present a new FER model named Landmark-Aware Net (LA-Net), which leverages facial landmarks to mitigate the impact of label noise from two perspectives. Firstly, LA-Net uses landmark information to suppress the uncertainty in expression space and constructs the label distribution of each sample by neighborhood aggregation, which in turn improves the quality of training supervision. Secondly, the model incorporates landmark information into expression representations using the devised expression-landmark contrastive loss. The enhanced expression feature extractor can be less susceptible to label noise. Our method can be integrated with any deep neural network for better training supervision without introducing extra inference costs. We conduct extensive experiments on both in-the-wild datasets and synthetic noisy datasets and demonstrate that LA-Net achieves state-of-the-art performance.

1. Introduction

Facial expression recognition (FER) is a fundamental task in the computer vision community, as it holds significance in various applications, such as human-computer interaction and fatigue driving detection. Given its wide applicability, researchers have developed multiple FER models in recent years [32, 20, 24]. Despite the progress made so far, FER remains challenging due to noisy labels. Specifically, individuals may interpret one expression differently, resulting in inconsistent annotations. Besides, facial expressions have inherent inter-class similarity, which further exacerbates the problem of label noise. Accordingly, mitigating label noise has become one of the primary tasks in FER.

Mainstream noise-tolerant FER models can be divided

into two groups: selecting clean samples and using label distributions as auxiliary training targets. For example, SCN [39] estimates the uncertainty of each sample and corrects the mislabeled samples during training. Moreover, DMUE [33] utilizes multiple learnable branches to estimate the label distribution of each sample. While these methods promote FER performance under label noise, they still encounter some problems. Firstly, sample selection methods assume that the neural network will fit clean samples before overfitting noisy labels [1, 44], which can lead to confusion between hard samples and noisy samples. Secondly, label distributions calculated based solely on expression information can still be noisy, since corrupt labels will disrupt both class and feature spaces [23].

To address the above problems, this paper proposes a new noise-tolerant FER model, called Landmark-Aware Net (LA-Net). The overview of LA-Net is plotted in Fig. 1. LA-Net leverages facial landmark information to combat label noise based on the underlying assumption that expressions with analogous landmark patterns probably belong to the same emotion category. The model comprises two key modules: label distribution estimation (LDE) and expression-landmark contrastive loss (EL Loss).

LDE calculates the label distribution of each sample and uses it as an auxiliary supervision signal. Based on the assumption that expressions should have similar emotions to their neighbors in feature space, LDE identifies neighbors in both expression and landmark spaces for each sample. The landmark information is utilized to correct the errors in the expression space. The module then learns pairwise contribution scores and performs neighborhood aggregation to obtain target label distributions. Furthermore, to mitigate the impact of batch division on online aggregation, the target label distributions are summed over previous epochs using exponential moving average (EMA).

EL Loss incorporates landmark information into expression representations to develop a robust backbone that is less susceptible to label noise. The algorithm treats landmarks and expressions as two views of facial images and establishes interactions between them via supervised con-

*Corresponding author

trastive learning (SCL) [18]. However, traditional SCL uses one-hot labels to select positive and negative pairs, thus performing poorly in the presence of label noise. Accordingly, our EL Loss designs a new pair selection strategy based on the label distributions to enable noise-tolerant SCL. Specifically, it first assigns pseudo-labels to confident images and takes the rest as unsupervised samples. The module then uses the expression and landmark features of the same images, or images with the identical pseudo-label, as positive pairs and all other combinations as negative pairs.

The proposed modules are solely used during training and thus incur no extra costs in deployment. Overall, our contributions can be summarized as follows:

(1) We present a landmark-aware FER model, named LA-Net, which leverages facial landmarks to alleviate the label noise issue.

(2) The LDE module uses landmark information to correct the errors in expression space and finds a set of neighbors to construct the label distribution of each sample.

(3) EL Loss devises noise-tolerant supervised contrastive learning and strengthens the expression feature extractor via expression-landmark interactions.

(4) LA-Net achieves state-of-the-art performance on both in-the-wild datasets and synthetic noisy datasets.

2. Related Work

2.1. Facial Expression Recognition

Researchers have proposed many algorithms to improve FER performance [32, 20, 24]. At the outset, handcrafted features including HOG [8] and SIFT [28] are applied to analyze expressions, whereas they perform poorly in the presence of strong illumination changes, large pose variations, and occlusions. Subsequently, learning-based methods advance the research mainly in two aspects: combating label noise and locating key areas. To address label ambiguity, SCN [39] evaluates the uncertainty of each sample and corrects mislabeled training samples on-the-fly. DMUE [33], on the other hand, applies multiple branches to calculate the label distribution of each sample. Regarding region-based models, RAN [40] uses self-attention to capture the importance of each facial area. TransFER [42] further introduces a dropout-like strategy to extract diverse key areas.

Concurrent with our work, LDLVA [21] also utilizes auxiliary facial information (valance-arousal) to alleviate label noise. Nonetheless, they use the extra information in a plug-and-play manner, while we take it as an auxiliary task and leverage supervised contrastive learning adapted to the label noise scenario to develop a robust feature extractor. Moreover, we combine the landmark and expression information to construct label distributions for training images and mitigate the impact of batch division by exponential moving average (EMA).

2.2. Learning with Noisy Labels

Label noise is a common issue in datasets, making learning with noisy labels a crucial topic in the community. Current research can be categorized into two groups: designing robust loss functions and selecting clean samples for training. Regarding the noise-tolerant loss functions, Zhang *et al.* [47] devise generalized cross-entropy loss to suppress label noise. Moreover, Patrini *et al.* [29] propose a loss correction approach for robust training. The sample selection strategy is based on the small-loss assumption that neural networks fit clean samples before overfitting noisy labels. Specifically, Han *et al.* [14] develop a co-teaching method that selects low-loss samples in two models to filter errors caused by noisy labels. Malach *et al.* [26] propose to train two models and perform updates only in case of disagreement between them, thereby keeping the effective number of noisy labels seen throughout the training process at a constant rate.

This paper aims to address the label noise in FER. The proposed LA-Net leverages facial landmarks to construct label distributions for training samples as well as strengthen the expression feature extractor via noise-tolerant supervised contrastive learning.

2.3. Contrastive Learning

Recently, Contrastive Learning (CL) [41, 15, 7, 6, 37, 18], which is based on the Siamese network [4], achieved great progress in unsupervised learning. Two prominent CL frameworks are SimCLR [6] and MoCo [15, 7]. SimCLR takes two random augmented views of the same image as positive pairs and different images as negative pairs, forming an instance discrimination task. Moreover, MoCo maintains a memory bank to increase negative samples and turns one branch of the Siamese network into a momentum encoder to improve the consistency of the memory bank. In addition to the self-supervised scenario, Khosla *et al.* [18] utilize the contrastive loss in the supervised setting (SCL) and achieve leading performance on multiple benchmarks.

Our EL Loss draws inspiration from supervised contrastive learning and implements interactions between landmarks and expressions to develop a more robust feature extractor. Moreover, we utilize label distribution as an alternative criterion for pair selection, which makes our method more resistant to label noise.

3. Methodology

We first introduce the notations that will be used in this paper. Let \mathbf{x} be the instance variable and \mathbf{x}_i be the i -th sample. We denote the class set as $\mathcal{Y} = \{y_1, \dots, y_C\}$, where C is the number of classes. Let $\mathbf{l}_i = (l_i^{y_1}, \dots, l_i^{y_C})$ with $l_i^{y_j} \in \{0, 1\}$ and $\|\mathbf{l}_i\|_1 = 1$ indicate the one-hot label of sample \mathbf{x}_i . The label distribution of sample \mathbf{x}_i is denoted

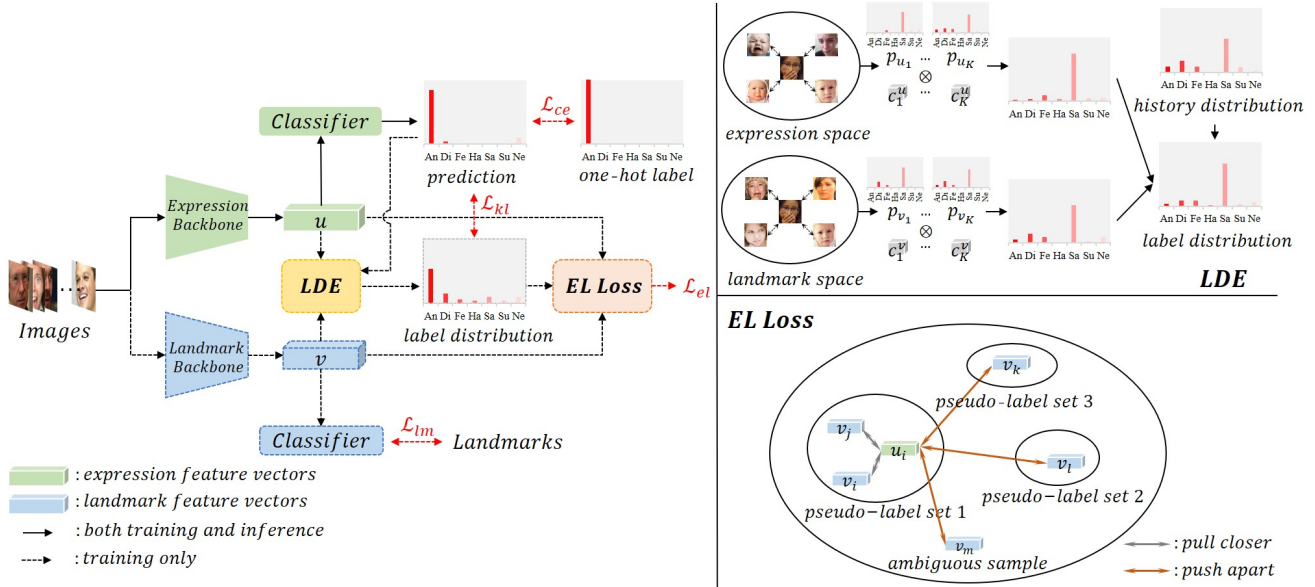


Figure 1. We present the pipeline of LA-Net on the left, with dashed lines indicating components used only during training and solid lines indicating those used in both training and inference. We zoom in on the structure of label distribution estimation (LDE) and expression-landmark contrastive loss (EL Loss) at the upper right and lower right of the figure, respectively. Regarding EL Loss, we provide an example of selecting positive and negative pairs for the expression feature of sample x_i .

as $\mathbf{d}_i = (d_i^{y_1}, \dots, d_i^{y_C})$, where $d_i^{y_j} \in [0, 1]$ and $\|\mathbf{d}_i\|_1 = 1$. Let \mathbf{u}_i and \mathbf{v}_i be the expression and landmark features of sample x_i , respectively. The prediction of sample x_i for the FER task is denoted as $\mathbf{p}_i = (p_i^{y_1}, \dots, p_i^{y_C})$, where $p_i^{y_j} \in [0, 1]$ and $\|\mathbf{p}_i\|_1 = 1$.

The overview of LA-Net is shown in Fig. 1. The model consists of three main parts: backbone, landmark distribution estimation (LDE), and expression-landmark contrastive loss (EL Loss). Specifically, LA-Net first uses two backbones to extract the expression and landmark features respectively. The landmark localization part employs a fully connected layer as the classifier and minimizes the mean square error, denoted as \mathcal{L}_{lm} , during training. LDE identifies $2K$ neighbors (K in expression space and K in landmark space) for each sample and performs neighborhood aggregation to generate target label distributions, which in turn improve the quality of training supervision. Besides, EL Loss considers the similarity between landmarks and expressions and incorporates landmark information into expression representations using noise-tolerant supervised contrastive learning. In the following sections, we will describe the LDE module and EL Loss in detail.

3.1. Label Distribution Estimation

Given the presence of noisy labels, label distribution is a better descriptor for facial expressions than the one-hot label. To construct the target label distribution of each sample, previous works hold that expressions should have similar emotions to their neighbors in the feature space or a

supporting space [21, 5]. Nonetheless, extreme label noise will disrupt both class and feature spaces. Accordingly, the LDE module uses landmark information to correct the errors in the expression space based on the assumption that images with similar landmark patterns should be assigned to the same emotion class.

Given the mini-batch $\mathcal{D}_{batch} = \{(x_i, \mathbf{l}_i) | i = 1, 2, \dots, n\}$, LDE calculates the label distributions in four steps. Firstly, for image x_i , LDE adopts the K -Nearest Neighbor algorithm to identify its neighbors in the expression space, denoted as $N^u(i)$, based on the cosine similarity.

$$s_{i,j}^u = \frac{\mathbf{u}_i \mathbf{u}_j^T}{\|\mathbf{u}_i\| \|\mathbf{u}_j\|} \quad (1)$$

$$N^u(i) = KNN(\mathcal{D}_{batch}; K; s_i^u) \quad (2)$$

where $s_{i,j}^u$ denotes the similarity between x_i and x_j in the expression space. The neighbor set in the landmark space, denoted as $N^v(i)$, can be generated in an identical way. Then, LDE evaluates the pairwise contribution scores in the expression space as follows:

$$c_{i,j}^u = \text{Sigmoid}(f([g_1(\mathbf{u}_i; \theta_1), g_2(\mathbf{u}_j; \theta_2)]; \theta)) \quad (3)$$

where $c_{i,j}^u$ denotes the contribution of x_j to x_i in the expression space, \square denotes concatenation, f , g_1 , and g_2 are three MLPs with learnable parameters θ , θ_1 , and θ_2 . The contribution scores in the landmark space, denoted as $c_{i,j}^v$, can be obtained in the same way. Based on the scores, the module performs neighborhood aggregation in each space

and produces the target label distributions by mean pooling.

$$\mathbf{d}_i = \frac{1}{2} \left(\frac{\sum_{k \in N^u(i)} C_{i,k}^u \mathbf{p}_k}{\sum_{k \in N^u(i)} C_{i,k}^u} + \frac{\sum_{k \in N^v(i)} C_{i,k}^v \mathbf{p}_k}{\sum_{k \in N^v(i)} C_{i,k}^v} \right) \quad (4)$$

As a result, the label distributions incorporate both expression and landmark information and thus are less susceptible to noisy labels. However, the above online aggregation approach can be affected by batch division, resulting in noisy and erratic label distributions in some cases. In particular, one batch may contain an excessive number of corrupt labels, causing the neighborhood of each sample to be extremely noisy. To address this issue, LDE sums up the targets over previous epochs using exponential moving average (EMA). This allows us to compute the target label distribution of image \mathbf{x}_i in the e -th epoch as follows:

$$d_i^{[e]} = \omega d_i^{[e-1]} + (1 - \omega) \mathbf{d}_i \quad (5)$$

where ω denotes the decay of previous targets. LA-Net leverages both one-hot labels and label distributions as supervision and minimizes the following loss functions:

$$\mathcal{L}_{ce} = -\frac{1}{n} \sum_{i=1}^n \sum_{j=1}^C l_i^{y_j} \log(p_i^{y_j}) \quad (6)$$

$$\mathcal{L}_{kl} = -\frac{1}{n} \sum_{i=1}^n \sum_{j=1}^C d_i^{[e], y_j} \log \frac{d_i^{[e], y_j}}{p_i^{y_j}} \quad (7)$$

3.2. Expression-Landmark Contrastive Loss

The LDE module estimates label distributions for training samples to enhance resistance to corrupt labels. Besides, a knowledgeable feature extractor can also help mitigate the label noise. As such, LA-Net leverages facial landmarks to strengthen the expression feature extractor. Specifically, we create expression-landmark pairs and utilize supervised contrastive learning (SCL) [18], denoted as expression-landmark contrastive loss (EL Loss), to implement interactions between these two facial modalities. However, traditional SCL constructs positive and negative pairs based on the one-hot labels in datasets, thus performing poorly in the presence of label noise. To address this problem, EL Loss uses label distributions as an alternative criterion for pair selection, enabling noise-tolerant supervised contrastive learning. We describe the pair selection process for an anchor expression feature vector in Fig. 1, and provide further details below.

Given sample \mathbf{x}_i , EL Loss first derives its pseudo-label \hat{l}_i from the target label distribution in the current epoch:

$$\hat{l}_i = \begin{cases} \operatorname{argmax}(d_i^{[e]}) & \text{if } \max(d_i^{[e]}) > \delta \\ -1 & \text{otherwise} \end{cases} \quad (8)$$

where δ denotes the confidence threshold for separating confident and ambiguous samples. We then treat the expression and landmark features of confident samples with the same pseudo-label as positive pairs. Moreover, to suppress the uncertainty of ambiguous samples, only expression and landmark features of the same ambiguous image are utilized as positive pairs. In other words, the indication set $\mathcal{I}(i)$ for the positive pairs of sample \mathbf{x}_i is computed by:

$$\mathcal{I}(i) = \begin{cases} \{j \mid \hat{l}_j = \hat{l}_i\} & \text{if } \hat{l}_i \neq -1 \\ \{i\} & \text{otherwise} \end{cases} \quad (9)$$

In addition, we take the remaining sample combinations as negative pairs and denote them as $\bar{\mathcal{I}}$. Following traditional contrastive learning [15, 7, 6], the module then performs non-linear projection to obtain query expression features, namely $q_i^u = g(\mathbf{u}_i; \theta)$. Here, g is a two-layer perceptron with learnable parameters θ . Query landmark features q_i^v , key expression features k_i^u , and key landmark features k_i^v are calculated similarly. Finally, EL Loss implements interactions between expressions and landmarks by pulling positive pairs closer and pushing negative pairs farther apart.

$$\mathcal{L}_1 = \frac{1}{n} \sum_{i=1}^n \frac{-1}{|\mathcal{I}(i)|} \sum_{j \in \mathcal{I}(i)} \log \frac{\exp(q_i^u \cdot k_j^v / \tau)}{\sum_{m \in \bar{\mathcal{I}}(i) \cup \{j\}} \exp(q_i^u \cdot k_m^v / \tau)} \quad (10)$$

$$\mathcal{L}_2 = \frac{1}{n} \sum_{i=1}^n \frac{-1}{|\mathcal{I}(i)|} \sum_{j \in \mathcal{I}(i)} \log \frac{\exp(q_i^v \cdot k_j^u / \tau)}{\sum_{m \in \bar{\mathcal{I}}(i) \cup \{j\}} \exp(q_i^v \cdot k_m^u / \tau)} \quad (11)$$

$$\mathcal{L}_{el} = \mathcal{L}_1 + \mathcal{L}_2 \quad (12)$$

where $a \cdot b$ denotes the cosine similarity between a and b , τ indicates a temperature parameter. Moreover, we follow the MoCo framework [15, 7] and use memory banks to increase negative pairs.

3.3. Overall Loss

LA-Net minimizes the following loss function during training:

$$\mathcal{L} = \mathcal{L}_{ce} + \mathcal{L}_{lm} + \alpha \mathcal{L}_{kl} + \beta \mathcal{L}_{el} \quad (13)$$

where α indicates the importance of label distributions and β denotes the contribution of EL Loss.

4. Experiments

4.1. Datasets

RAF-DB [25] consists of 30,000 images with basic or compound labels. Following previous works, we only use the images annotated with six basic expressions (*anger, disgust, fear, happiness, sadness, surprise*) and *neutral*, where 12,271 are used for training and 3,068 for testing.

Noise	Method	RAF-DB	FERPlus	AffectNet
10%	Baseline	81.01	83.29	57.24
	SCN [39]	82.18	84.28	58.58
	RUL [45]	86.17	86.93	60.54
	EAC [46]	88.02	87.03	61.11
	LA-Net (Ours)	88.75±0.11	88.02±0.08	62.85±0.13
20%	Baseline	77.98	82.34	55.89
	SCN [39]	80.10	83.17	57.25
	RUL [45]	84.32	85.05	59.01
	EAC [46]	86.05	86.07	60.29
	LA-Net (Ours)	87.12±0.16	86.85±0.14	61.72±0.21
30%	Baseline	75.50	79.77	52.16
	SCN [39]	77.46	82.47	55.05
	RUL [45]	82.06	83.90	56.93
	EAC [46]	84.42	85.44	58.91
	LA-Net (Ours)	85.33±0.18	86.01±0.19	60.82±0.21

Table 1. Evaluation (%) on synthetic noisy datasets.

FERPlus [3] is an improved version of FER2013 [12]. It contains 28,709/3,589/3,589 training/validation/testing images. Each image is labeled by ten experts and assigned to one of eight classes (six basic expressions, *neutral*, and *contempt*). We use the class with the most votes as the label.

AffectNet [27] is the largest and most challenging FER dataset. It contains more than 280K training images and 4,000 testing images, which are classified into the same eight classes as FERPlus. Previous works vary in the use of AffectNet (with or without *contempt*), and we utilize 8-class AffectNet by default in the experiments.

Following previous works, we report the overall accuracy of the testing set for all datasets.

4.2. Implementation Details

We perform experiments on 4 NVIDIA RTX 3080Ti GPUs and use ResNet-18 [16] pretrained on MS-Celeb-1M [13] as the backbone. Regarding data processing, we resize all images to 224×224 pixels and use HRNet [35, 38] to generate ground truth landmarks. To alleviate class imbalance, we adopt progressively balanced sampling [17]. Besides, on-the-fly data augmentation techniques, including random cropping, random horizontal flipping, random erasing, and random color jitter, are employed to enhance the generalization performance of LA-Net. During training, we set the batch size to 128 and use the Adam optimizer [19] with an initial learning rate of $1e-3$. For all datasets, we train the model for 80 epochs and linearly decrease the learning rate to 0. Moreover, we conduct grid searches to determine all parameters. Specifically, we set the number of neighbors K to 8, the decay of targets ω to 0.9, the temperature τ to 0.1, and the confidence threshold δ to 0.7. Besides, weights α and β in the overall objective are set to 1.0 and 0.1.

4.3. Experiments on Synthetic Noisy Datasets

Noisy labels, caused by ambiguous facial expressions, significantly harm FER performance in real-world scenarios. To address this issue, LA-Net leverages landmark infor-

LD	LM	EL	PL	RAF-DB (original)	RAF-DB (30% noise)
-	-	-	-	87.06	75.50
✓	-	-	-	88.43	79.17
✓	✓	-	-	89.04	82.37
✓	✓	✓	-	90.48	83.21
✓	✓	✓	✓	90.81	85.33

Table 2. Component analysis (%). (LD: label distribution generated using only expression information; LM: landmark information in LDE; EL: expression-landmark contrastive loss; PL: pseudo-labels in EL Loss.)

mation to estimate label distributions for training samples and strengthen the expression feature extractor. We compare LA-Net with previous noise-tolerant FER models on three synthetic noisy datasets with varying levels of noise (10%, 20%, and 30%). To inject synthetic noise, we follow prior studies [39, 45, 46] and randomly flip the one-hot label to other categories. Additionally, we repeat the experiments five times due to the randomness of label noise injection and report the mean and standard deviation of overall accuracy.

Tab. 1 shows that LA-Net consistently outperforms other noise-tolerant FER algorithms in all settings. Compared to the current state-of-the-art model EAC [46], LA-Net achieves an average improvement of 0.91%, 0.82%, and 1.70% on RAF-DB, FERPlus, and AffectNet, respectively. Additionally, LA-Net becomes more preferred as the noise ratio increases. Specifically, compared to the baseline, the model promotes the performance by 5.61% and 8.66% on AffectNet with 10% and 30% noise, respectively. In conclusion, the experimental results in Tab. 1 demonstrate that the proposed approach is effective in mitigating label noise.

In addition to the aforementioned symmetric noise, we further test the more challenging asymmetric noise, where the label is flipped to its most similar class based on the confusion matrix. As shown in Appendix A, LA-Net consistently outperforms previous models when dealing with asymmetric noise, indicating its effectiveness.

4.4. Ablation Study

Contribution of each component. To quantify the contribution of the proposed modules, we conduct an ablation study on both original and noisy RAF-DB. For simplicity, we refer to the performance on these two datasets as (a, b) in the following. As shown in Tab. 2, complete LDE (line 3) promotes performance by (1.98%, 6.87%) compared to the baseline in line 1. EL Loss (line 5), on the other hand, yields gains of (1.77%, 2.96%) in contrast to the model in line 3. These two modules leverage landmark information from different perspectives and mitigate real-world and synthetic noise effectively. Besides, comparing line 2 and line 3, landmark information in LDE brings the improvement of (0.61%, 3.20%), indicating its importance, espe-

Method	RAF-DB	AffectNet (7 classes)	AffectNet (8 classes)	FERPlus
RAN [40]	86.90	-	59.50	89.16
SCN [39]	87.03	-	60.23	89.39
DAFL [10]	87.78	65.20	-	-
KTN [22]	88.07	63.97	-	90.49
FDRL [30]	89.47	-	-	-
ARM [34]	90.42	65.20	61.33	-
DMUE [33]	88.76	-	62.84	88.64
EAC [46]	89.99	65.32	-	89.64
LDLVA [21]	90.51	66.23	-	-
TransFER [42]	90.91	66.23	-	90.83
LA-Net (ResNet-18)	90.81	67.09	64.24	91.39
LA-Net (ResNet-50)	91.56	67.60	64.54	91.78

Table 3. Comparison (%) with state-of-the-art methods on original datasets.




Figure 2. Performance on original and noisy RAF-DB with different number of neighbors in each space (K).

cially in the presence of severe noise. Additionally, the pseudo-labels (line 5) enable noise-tolerant supervised contrastive learning and achieve an improvement of 2.12% on corrupt RAF-DB compared to the EL Loss using one-hot labels (line 4). Overall, the results in Tab. 2 demonstrate the effectiveness of the proposed modules and the advantages of their combination in LA-Net.

Number of nearest neighbors in each space K . We explore the effect of K on the performance in Fig. 2. For original RAF-DB, LA-Net achieves optimal performance with $K=12$. Smaller or larger K can lead to slight degradation, as the former may result in inaccurate target label distributions, while the latter can cause excessive noisy neighbors. The choice of K has a similar effect on both original and noisy RAF-DB. The only exception is that LA-Net achieves the best performance on the noisy RAF-DB with $K=8$ since it contains more corrupt labels. Therefore, we set K to 8 throughout experiments to alleviate label noise.

4.5. Experiments on Original Datasets

While the original datasets are considered "clean", they inevitably contain noisy labels due to the ambiguity of expressions. Hence, we compare LA-Net with state-of-the-art methods on original datasets to verify its resistance to real-

Method	RAF-DB	AffectNet (7 classes)
AIR [2]	67.37	54.23
NAL [11]	84.22	55.97
IPA2LT [43]	83.80	57.85
LDL-ALSG [5]	85.53	59.35
LDLVA [21]	87.26	62.89
LA-Net (Ours)	88.10	65.43


Table 4. Experiments with inconsistent labels (%). Following previous works [43, 5, 21], we use the 7-class AffectNet to keep its label set consistent with RAF-DB.

world uncertainty. Moreover, we consider both 7-class and 8-class AffectNet for a fair comparison.

We compare the performance of LA-Net with several state-of-the-art FER models in Tab. 3. Among the existing models, TransFER [42] utilizes the powerful Vision Transformer [9] structure to improve performance, while the others enhance FER performance by locating key areas or alleviating label noise. Compared to existing models, LA-Net gains improvements of 0.65%, 1.37%, 1.70%, and 0.95% on RAF-DB, 7-class AffectNet, 8-class AffectNet, and FER-Plus, respectively. Moreover, the proposed model also exhibits substantial advantages in challenging cases, such as 7-class and 8-class AffectNet. Our findings indicate that LA-Net is a promising approach for handling ambiguous expressions in real-world scenarios.


4.6. Experiments with Inconsistent Labels

The performance of deep neural networks generally improves as more training samples are provided, but FER presents a unique challenge. Due to the subjective nature of interpreting facial expressions, individuals may assign different labels to the same expression, leading to inconsistent labels within a dataset and across different datasets. To evaluate the effectiveness of LA-Net in addressing this problem, we conduct a cross-dataset experiment following previous works [43, 5, 21]. Specifically, we train our model on a mixed dataset composed of the training samples from RAF-DB and AffectNet and evaluate it on the testing set



(a) Baseline model on RAF-DB with 30% noise (b) SCN [39] on RAF-DB with 30% noise (c) LA-Net on RAF-DB with 30% noise

Figure 3. Feature visualization on noisy RAF-DB. The baseline model memorizes most of the noisy labels, resulting in indistinguishable feature clusters. SCN [39] suppresses the noise to some extent, whereas it still overfits many corrupt labels. LA-Net pushes the noisy samples to the decision boundary and maintains clean clusters.



(a) Baseline model on RAF-DB with 30% noise (b) SCN [39] on RAF-DB with 30% noise (c) LA-Net on RAF-DB with 30% noise

Figure 4. Cross-entropy between predictions and one-hot labels after training.

of both datasets. Note that we use the 7-class AffectNet to ensure the identical class set between two datasets.

We will analyze the results in Tab. 4 from two perspectives. Firstly, compared to LDLVA [21], which is the best available model for addressing label inconsistency, LA-Net achieves the improvement of 0.84% and 2.54% on RAF-DB and AffectNet, respectively. Secondly, LA-Net helps alleviate the performance drop caused by cross-dataset inconsistency. Specifically, comparing Tab. 4 and Tab. 3, LA-Net encounters the degradation of 2.71% on RAF-DB and 1.66% on AffectNet, whereas LDLVA suffers from the drop of 3.25% and 3.34% on these two datasets, respectively. In summary, the improvement in two aspects demonstrates the effectiveness of LA-Net in dealing with inconsistent labels. Nonetheless, cross-dataset inconsistency remains a challenge for FER models and warrants further research.

4.7. Visualization Analysis

High-dimensional features. For an intuitive understanding of LA-Net, we use t-SNE [36] to plot the trained features of different models, including a baseline model, SCN [39], and our LA-Net, on RAF-DB with 30% noise. As shown in Fig. 3(a), the baseline model memorizes the noisy labels, resulting in adjacent feature clusters. SCN learns the uncertainty of each sample and proposes a relabeling mechanism

to address label noise. Nonetheless, as depicted in Fig. 3(b), it still overfits many noisy samples, such as the *disgust* expressions in the cluster of *anger*. In comparison, LA-Net, plotted in Fig. 3(c), pushes noisy samples to the decision boundary and generates clean clusters. Moreover, our model achieves improved separation across categories and better compactness within each class. Overall, LA-Net prevents overfitting noisy samples and forms easy-to-distinguish category clusters in the presence of severe noise, indicating its effectiveness in dealing with noisy labels.

Cross-entropy between predictions and one-hot labels. We evaluate the ability of LA-Net to handle noisy labels by calculating the cross-entropy between predictions and one-hot labels of the training samples after training. As shown in Fig. 4(a), the baseline model memorizes almost all noisy labels, resulting in poor generalization ability. SCN [39] alleviates the noise to some extent by relabeling, whereas it still overfits some corrupt labels since these noisy samples are not identified or mistakenly relabeled. In contrast, LA-Net, which leverages landmark information to construct label distributions and enhance expression feature extractor, could easily distinguish between clean and noisy samples after training. Overall, the results in Fig. 4 indicate that our model can suppress label noise to a large extent.

Target label distributions. We conduct a user study on




Figure 5. Comparison between labels, user study results, and generated label distributions. Please refer to the appendix for more results.

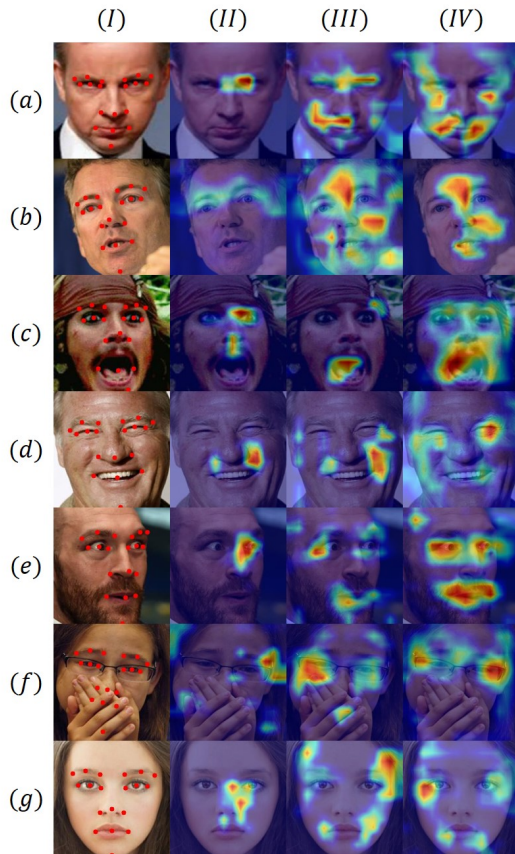


Figure 6. Attention maps of images from AffectNet. Row (a) - (g) denote anger, disgust, fear, happiness, sadness, surprise, and neutral respectively. Column (I) shows the images and the landmarks. Columns (II) - (IV) present the attention maps of three training strategies: (II) Baseline strategy; (III) Using LDE to generate target distribution; (IV) Complete LA-Net.

several images randomly selected from RAD.DB and AffectNet to diagnose the LDE module. We present part of the results in Fig. 5, and *more details are provided in Appendix C*. The left two expressions in Fig. 5 are found to be ambiguous in the user study, while the one-hot labels only provide one possible category. In contrast, LDE reveals all possible classes and generates label distributions that are consistent with the user study results. Additionally, the model correctly identifies the latent truth for the third image, which is mistakenly assigned to *happiness* in the dataset (highlighted

in red). In conclusion, the results indicate that LA-Net can achieve agreement with human perception to some extent and effectively mitigate label noise.

Attention visualization. To further analyze how landmark information benefits LA-Net, we randomly select some images from AffectNet and generate their attention maps using GradCAM [31]. Note that we perform visualization in the same setting as in inference mode. That is, no landmark information is provided when extracting attention maps.

Fig. 6 plots the attention maps of the selected images, where the categories are anger, disgust, fear, happiness, sadness, surprise, and neutral from top to bottom. Column (I) provides the images and their landmarks, and the rest presents the attention maps of three strategies. Comparing different columns, the baseline model typically focuses on a few areas, while the LDE module reveals more crucial areas by providing better supervision in the presence of label noise. Moreover, EL Loss contributes to locating more key regions, such as the eyes in (c, IV) and (d, IV), and boosts the attention of key areas, including the eyes and mouth in (e, IV). Overall, LA-Net generates regions of interest that align well with facial landmarks, indicating that it develops a knowledgeable feature extractor by incorporating landmark information into expression representations.

5. Conclusion

This paper introduces a new FER model named LA-Net, which aims to mitigate the impact of label noise by incorporating landmark information. LA-Net consists of two main modules, namely label distribution estimation and expression-landmark contrastive loss. The former uses landmark information to correct errors in expression space and estimates the label distribution of each sample, which in turn provides better supervision. The latter enables noise-tolerant supervised contrastive learning and develops a more robust feature extractor by performing interactions between expressions and landmarks. We conduct extensive experiments on multiple scenarios, whose results demonstrate the effectiveness of LA-Net in handling label noise and inconsistent labels. Additionally, we present various visualization studies to illustrate how landmark information enhances the performance of LA-Net.

References

- [1] Devansh Arpit, Stanisław Jastrzebski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, et al. A closer look at memorization in deep networks. In *International conference on machine learning*, pages 233–242. PMLR, 2017. 1
- [2] Samaneh Azadi, Jiashi Feng, Stefanie Jegelka, and Trevor Darrell. Auxiliary image regularization for deep cnns with noisy labels. *arXiv preprint arXiv:1511.07069*, 2015. 6
- [3] Emad Barsoum, Cha Zhang, Cristian Canton Ferrer, and Zhengyou Zhang. Training deep networks for facial expression recognition with crowd-sourced label distribution. In *Proceedings of the 18th ACM International Conference on Multimodal Interaction*, pages 279–283, 2016. 5
- [4] Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah. Signature verification using a” siamese” time delay neural network. *Advances in neural information processing systems*, 6, 1993. 2
- [5] Shikai Chen, Jianfeng Wang, Yuedong Chen, Zhongchao Shi, Xin Geng, and Yong Rui. Label distribution learning on auxiliary label space graphs for facial expression recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13984–13993, 2020. 3, 6
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. 2, 4
- [7] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. 2, 4
- [8] Navneet Dalal and Bill Triggs. Histograms of oriented gradients for human detection. In *2005 IEEE computer society conference on computer vision and pattern recognition (CVPR’05)*, volume 1, pages 886–893. Ieee, 2005. 2
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 6
- [10] Amir Hossein Farzaneh and Xiaojun Qi. Facial expression recognition in the wild via deep attentive center loss. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2402–2411, 2021. 6
- [11] Jacob Goldberger and Ehud Ben-Reuven. Training deep neural-networks using a noise adaptation layer. In *International conference on learning representations*, 2017. 6
- [12] Ian J Goodfellow, Dumitru Erhan, Pierre Luc Carrier, Aaron Courville, Mehdi Mirza, Ben Hamner, Will Cukierski, Yichuan Tang, David Thaler, Dong-Hyun Lee, et al. Challenges in representation learning: A report on three machine learning contests. In *International conference on neural information processing*, pages 117–124. Springer, 2013. 5
- [13] Yandong Guo, Lei Zhang, Yuxiao Hu, Xiaodong He, and Jianfeng Gao. Ms-celeb-1m: A dataset and benchmark for large-scale face recognition. In *European conference on computer vision*, pages 87–102. Springer, 2016. 5
- [14] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. *Advances in neural information processing systems*, 31, 2018. 2
- [15] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. 2, 4
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 5
- [17] Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. Decoupling representation and classifier for long-tailed recognition. *arXiv preprint arXiv:1910.09217*, 2019. 5
- [18] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in neural information processing systems*, 33:18661–18673, 2020. 2, 4
- [19] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [20] Jyoti Kumari, Reghunadhan Rajesh, and KM Pooja. Facial expression recognition: A survey. *Procedia computer science*, 58:486–491, 2015. 1, 2
- [21] Nhat Le, Khanh Nguyen, Quang Tran, Erman Tjiputra, Bac Le, and Anh Nguyen. Uncertainty-aware label distribution learning for facial expression recognition. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 6088–6097, 2023. 2, 3, 6, 7
- [22] Hangyu Li, Nannan Wang, Xinpeng Ding, Xi Yang, and Xinbo Gao. Adaptively learning facial expression representation via cf labels and distillation. *IEEE Transactions on Image Processing*, 30:2016–2028, 2021. 6
- [23] Jichang Li, Guanbin Li, Feng Liu, and Yizhou Yu. Neighborhood collective estimation for noisy label identification and correction. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXIV*, pages 128–145. Springer, 2022. 1
- [24] Shan Li and Weihong Deng. Deep facial expression recognition: A survey. *IEEE transactions on affective computing*, 2020. 1, 2
- [25] Shan Li, Weihong Deng, and JunPing Du. Reliable crowd-sourcing and deep locality-preserving learning for expression recognition in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2852–2861, 2017. 4
- [26] Eran Malach and Shai Shalev-Shwartz. Decoupling” when to update” from” how to update”. *Advances in neural information processing systems*, 30, 2017. 2

- [27] Ali Mollahosseini, Behzad Hasani, and Mohammad H Mahoor. Affectnet: A database for facial expression, valence, and arousal computing in the wild. IEEE Transactions on Affective Computing, 10(1):18–31, 2017. [5](#)
- [28] Pauline C Ng and Steven Henikoff. Sift: Predicting amino acid changes that affect protein function. Nucleic acids research, 31(13):3812–3814, 2003. [2](#)
- [29] Giorgio Patrini, Alessandro Rozza, Aditya Krishna Menon, Richard Nock, and Lizhen Qu. Making deep neural networks robust to label noise: A loss correction approach. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 1944–1952, 2017. [2](#)
- [30] Delian Ruan, Yan Yan, Shenqi Lai, Zhenhua Chai, Chunhua Shen, and Hanzi Wang. Feature decomposition and reconstruction learning for effective facial expression recognition. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 7660–7669, 2021. [6](#)
- [31] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In Proceedings of the IEEE international conference on computer vision, pages 618–626, 2017. [8](#)
- [32] Caifeng Shan, Shaogang Gong, and Peter W McOwan. Facial expression recognition based on local binary patterns: A comprehensive study. Image and vision Computing, 27(6):803–816, 2009. [1](#), [2](#)
- [33] Jiahui She, Yibo Hu, Hailin Shi, Jun Wang, Qiu Shen, and Tao Mei. Dive into ambiguity: latent distribution mining and pairwise uncertainty estimation for facial expression recognition. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 6248–6257, 2021. [1](#), [2](#), [6](#)
- [34] Jiawei Shi, Songhao Zhu, and Zhiwei Liang. Learning to amend facial expression representation via de-albino and affinity. arXiv preprint arXiv:2103.10189, 2021. [6](#)
- [35] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 5693–5703, 2019. [5](#)
- [36] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. Journal of machine learning research, 9(11), 2008. [7](#)
- [37] Haoqing Wang, Xun Guo, Zhi-Hong Deng, and Yan Lu. Rethinking minimal sufficient representation in contrastive learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 16041–16050, 2022. [2](#)
- [38] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, et al. Deep high-resolution representation learning for visual recognition. IEEE transactions on pattern analysis and machine intelligence, 43(10):3349–3364, 2020. [5](#)
- [39] Kai Wang, Xiaojiang Peng, Jianfei Yang, Shijian Lu, and Yu Qiao. Suppressing uncertainties for large-scale facial expression recognition. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 6897–6906, 2020. [1](#), [2](#), [5](#), [6](#), [7](#)
- [40] Kai Wang, Xiaojiang Peng, Jianfei Yang, Debin Meng, and Yu Qiao. Region attention networks for pose and occlusion robust facial expression recognition. IEEE Transactions on Image Processing, 29:4057–4069, 2020. [2](#), [6](#)
- [41] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 3733–3742, 2018. [2](#)
- [42] Fanglei Xue, Qiangchang Wang, and Guodong Guo. Transfer: Learning relation-aware facial expression representations with transformers. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 3601–3610, 2021. [2](#), [6](#)
- [43] Jiabei Zeng, Shiguang Shan, and Xilin Chen. Facial expression recognition with inconsistently annotated datasets. In Proceedings of the European conference on computer vision (ECCV), pages 222–237, 2018. [6](#)
- [44] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning (still) requires rethinking generalization. Communications of the ACM, 64(3):107–115, 2021. [1](#)
- [45] Yuhang Zhang, Chengrui Wang, and Weihong Deng. Relative uncertainty learning for facial expression recognition. Advances in Neural Information Processing Systems, 34:17616–17627, 2021. [5](#)
- [46] Yuhang Zhang, Chengrui Wang, Xu Ling, and Weihong Deng. Learn from all: Erasing attention consistency for noisy label facial expression recognition. In Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXVI, pages 418–434. Springer, 2022. [5](#), [6](#)
- [47] Zhilu Zhang and Mert Sabuncu. Generalized cross entropy loss for training deep neural networks with noisy labels. Advances in neural information processing systems, 31, 2018. [2](#)