

# LSSSED: A LARGE-SCALE DATASET AND BENCHMARK FOR SPEECH EMOTION RECOGNITION

Weiquan Fan<sup>1</sup> Xiangmin Xu<sup>1</sup> Xiaofen Xing<sup>1</sup> Weidong Chen<sup>1</sup> Dongyan Huang<sup>2</sup>

\* School of Electronic and Information Engineering, South China University of Technology, China

† UBTECH Robotics Corp, China

## ABSTRACT

Speech emotion recognition is a vital contributor to the next generation of human-computer interaction (HCI). However, current existing small-scale databases have limited the development of related research. In this paper, we present LSSSED, a challenging large-scale english speech emotion dataset, which has data collected from 820 subjects to simulate real-world distribution. In addition, we release some pre-trained models based on LSSSED, which can not only promote the development of speech emotion recognition, but can also be transferred to related downstream tasks such as mental health analysis where data is extremely difficult to collect. Finally, our experiments show the necessity of large-scale datasets and the effectiveness of pre-trained models. The dataset will be released on <https://github.com/tobefans/LSSSED>.

**Index Terms**— speech emotion recognition, dataset, pre-trained model, deep learning

## 1. INTRODUCTION

Speech emotion recognition (SER) is a necessary part of the human-computer interaction system. Although emotion itself is very abstract, it still has some obvious intonation characteristics. Intuitively, sad voices are generally low-pitched and slow while happy voices are usually the opposite. Up to now, many algorithms have emerged for existing dataset.

There are lots of researches carried out on SER. In [1], Schuller et al. applies continuous Hidden Markov Models (HMM) to introduce SER with a self-collected speech corpus. Since 2004, some standardized speech emotion databases have been released. AIBO[2], EMODB[3], ENTERFACE[4], RML[5], IEMOCAP[6], AFEW[7], and MELD[8]. Among them, IEMOCAP[6] and MELD[8] are the databases with the most data. IEMOCAP[6] collects 7,433 sentences (13 hours and 40 minutes in total) spoken by 10 people. MELD[8] contains 13,708 sentences (about 12 hours) from 407 people. In [9], decision tree is utilized to mitigate error propagation on

AIBO[2] and IEMOCAP[6]. In [10], RBM is applied to learn discriminatory features on EMODB[3] and ENTERFACE[4].

With the rapid development of deep learning, Zhang et al. [11] utilizes DCNN to bridge the affective gap in speech signals on EMODB[3], RML[5], ENTERFACE[4]. At the same time, Satt et al. [12] presents a system based on an end-to-end LSTM-CNN with raw spectrograms on IEMOCAP[6]. Recently, Yeh et al. [13] proposes a dialogical emotion decoding algorithm to consecutively decode the emotion states of each utterance on IEMOCAP[6] and MELD[8].

Although there have been certain level of progression on SER, there is still a potentially serious overfitting problem, which may limit the development of SER. As shown in [14–17], even if a high accuracy is achieved on a certain database, their performance may be poor when transferring to another database. This is because the existing databases are generally small in scale, resulting in insufficient diversity, which is far from the real-world scenarios thus leading to the tendency of model overfitting. Therefore, a large-scale emotion dataset that can more comprehensively represent the real distribution is urgently needed to improve the generalization of existing algorithms.

Generally speaking, transfer learning can to a certain extent improve the performance of an algorithm. Boigne et al. [18] points out task-related transfer learning of recognizing emotions on small datasets. For emotion recognition related task, a good pre-trained model is urgent since data collection is very difficult. Taking the depression detection task as an example, there are only about a hundred subjects at most till date. In our opinion, the pre-trained model from the SER task is more suitable for detecting depression, since it is more inclined to obtain acoustic features while the model from ASR task is prone to extract linguistic features.

In this paper, we present LSSSED, a challenging large-scale english dataset for speech emotion recognition. It contains 147,025 sentences (206 hours and 25 minutes in total) spoken by 820 people. Based on our dataset, we can simulate a more comprehensive and rich data distribution of real-world scenarios so that deep neural networks can better model their distribution. Furthermore, since there is currently no non-semantic large-scale pre-training model, we release some pre-trained models with speech emotion recognition task.

Corresponding author. Email: xfxing@scut.edu.cn

Thanks to Datatang for support.

Thanks to UBTECH-SCUT Joint Research Lab for support.

**Table 1.** Comparison to existing public speech emotion datasets.

| Corpus       | Published Time | Language | Speakers | Naturalness | h:mm    | Number of sentences | Classes |
|--------------|----------------|----------|----------|-------------|---------|---------------------|---------|
| AIBO[2]      | 2004           | Multiple | 51       | Natural     | 9:20    | 48401               | 11      |
| EMODB[3]     | 2005           | German   | 10       | Acted       | 0:22    | 494                 | 7       |
| ENTERFACE[4] | 2006           | English  | 43       | Acted       | 1:00    | 1170                | 6       |
| RML[5]       | 2008           | Multiple | 8        | Acted       | 0:42    | 500                 | 6       |
| IEMOCAP[6]   | 2008           | English  | 10       | Both        | 12:00   | 7433                | 10      |
| AFEW[7]      | 2012           | Multiple | 330      | Acted       | unknown | 1426                | 7       |
| MELD[8]      | 2018           | English  | 407      | Acted       | 13:40   | 13708               | 7       |
| LSSSED       | 2020           | English  | 820      | Natural     | 206:25  | 147025              | 11      |

## 2. LSSSED

In this section, we introduce our dataset, LSSSED in details. LSSSED collects a total of 147,025 utterances from 820 subjects, with an average duration of 5.05s. As shown in Table 1, the data volume of LSSSED is very large, and its total duration (over 200 hours) can reach dozens of times than existing databases.

### 2.1. Collection and Labeling

The subjects that participate in the experiment are widely distributed with representations from both genders and variety of age groups. Each subject would be recorded in one or several emotional videos sessions in an indoor lab environment with a camera pointing at him or her. In the video, the subject is induced by random questions as their utterances are associated with an emotional label. The total length of a video is about 10-20 minutes.


The utterances in each video dialogue are annotated by a professional annotation team. Each utterance is annotated with the corresponding emotion label, including anger, happiness, sadness, disappointment, boredom, disgust, excitement, fear, surprise, normal, and other. Note that some utterances in the video contain two or more emotions. In addition, each utterance is also annotated with auxiliary information, including the gender and age of the subject.

### 2.2. Data Distribution

As mentioned above, our database covers various groups of people. Table 2 shows the conditional and joint distribution of the ages and genders. In LSSSED, the gender distribution is relatively balanced. The age distribution however has fewer elderly people.

**Table 2.** Data distribution for gender and age.

|               | Young | Middle-aged | Old | Total |
|---------------|-------|-------------|-----|-------|
| <b>Female</b> | 253   | 167         | 65  | 485   |
| <b>Male</b>   | 155   | 141         | 39  | 335   |
| <b>Total</b>  | 408   | 308         | 104 | 820   |

**Fig. 1.** Distribution of data for each emotion labels

A pie chart of the distribution of data for emotion labels is shown in Fig 1. Since the subjects speak in a spontaneous environment, the more common neutral samples accounted for a larger proportion. Next is happy, sad, disappointed, excited, and angry samples respectively. The samples of these six categories account for 81% of the total sample. Then, the samples of boring, disgusting, fearful, and surprised are fewer, accounting for only 6%. In addition, 13% of other uncommon samples can be used for tasks to distinguish whether they are common emotions.

In order to standardize future training benchmarks, we divided our LSSSED dataset into training and test sets. Specifically, we first shuffle the order of all samples, then set 20% of the samples as the test set, and the rest as the training set. It should be noted that we ensure the distribution of each emotion category in the training set and test set are the same or at least similar. Table 3 shows the specific distribution of data for emotion labels in the training set and test set respectively.

### 2.3. Preprocessing and feature extraction

After obtaining the videos, we then convert them into audio signals at a sampling rate of 16kHz. According to the start time and end time of each utterance, we cut out 147,025 audio utterances. For each sentence, we use spectral subtraction algorithm [19] to perform audio denoising. It subtracts noise on the short-time spectrum and then restores the audio. Next,

**Table 3.** Distribution of data for emotion labels in the subsets.

|                     | Training      | Test         | Total         |
|---------------------|---------------|--------------|---------------|
| <b>Angry</b>        | 1298          | 5192         | 6490          |
| <b>Neutral</b>      | 12092         | 48369        | 60461         |
| <b>Happy</b>        | 5406          | 21627        | 27033         |
| <b>Sad</b>          | 2410          | 9641         | 12051         |
| <b>Disappointed</b> | 1781          | 7124         | 8905          |
| <b>Bored</b>        | 583           | 2333         | 2916          |
| <b>Disgusted</b>    | 636           | 2543         | 3179          |
| <b>Excited</b>      | 1046          | 4182         | 5228          |
| <b>Surprise</b>     | 331           | 1325         | 1656          |
| <b>Fear</b>         | 126           | 502          | 628           |
| <b>Other</b>        | 3696          | 14782        | 18478         |
| <b>Total</b>        | <b>117620</b> | <b>29405</b> | <b>147025</b> |

we increase the audio volume by a factor of 2 to make the sound louder.

After preprocessing, we perform STFT with Hann window length of 1024 points and the window shift of 512 points. A square operation follows to obtain the power spectrum. The power spectrum is then passed through a triangular filter bank with 128 Mel-scales to simulate the human auditory perception system.

#### 2.4. Pre-trained Models

We firstly select VGG [20] and ResNet [21] for pre-training, which are useful in many scenarios. VGG builds a unified and simple structure to deepen the network, while ResNet proposes residual learning to ease the training procedure.

In order to better adapt to the specificity of speech, we propose PyResNet, an improved model of ResNet [21]. Due to the sufficient amount of data, PyResNet is based on ResNet50, ResNet101 or ResNet152.

Specifically, the second convolution layer in each layer of ResNet is replaced with a pyramid convolution [22] that can capture multi-scale information to solve the problem of uncertain time position of valid speech information. In addition, we replaced the GAP layer with average pooling layer only in the time dimension to make the model insensitive to time and preserve the frequency information.

### 3. DATASET EXPERIMENTS

#### 3.1. Effectiveness of LSSSED

Although the current algorithms have achieved good results on many small-scale datasets, pre-trained models often cannot be well generalized to other datasets. This triggered our thinking about the scale of databases resulting in the collection and building of a large amount of database, that can be informative enough to train a model with good generalization.

In order to verify the effectiveness of different-scale datasets, we calculate the performance degradation based on ResNet152 as shown in Table 4. The metrics we choose are weighted accuracy (WA, the overall classification accuracy) and unweighted accuracy (UA, the averaged accuracy of each category). As indicated, the performance degradation is very large when the model trained from small-scale IEMOCAP[6] is tested on large-scale LSSSED, while it is less when tested from large-scale to small-scale. This demonstrates the effectiveness of LSSSED, since it simulates the real-world distribution.

**Table 4.** Performance degradation when testing in the target database compared with the source database (training in the source database).

| Source  | Target  | -WA          | -UA          |
|---------|---------|--------------|--------------|
| IEMOCAP | LSSSED  | 0.596        | 0.342        |
| LSSSED  | IEMOCAP | <b>0.119</b> | <b>0.071</b> |

#### 3.2. Speech Emotion Recognition Benchmark

In the experiments, all algorithms use the training set and test set from LSSSED. Our models are iterated for 60 epochs with batch size of 256 through the SGD optimizer with a weight decay of 0.001. The learning rate (initialized to 0.01) drops to 10% of the original every 20 epochs. Consistent with the current mainstream SER experiments, we use four emotion categories, including angry, neutral, happy and sad.

The results are shown in table 5. This shows that the performance of existing algorithms on large-scale LSSSED is not satisfactory. More importantly, the accuracy (weighted and unweighted) of these algorithms is even lower than that of the basic VGG and ResNet models. In addition, it is worth noting that our PyResNet achieves better results than the basic backbone models. This demonstrates that the improvement based on pyramid convolution is effective on large-scale database. Since these algorithms are not excellent in overall performance on large-scale databases, it should be indicated that LSSSED still has great challenges which means that speech emotion recognition is still a long way from being perfectly widely applicable.

Confusion matrices of both MTS-3 branches and PyResNet that use ResNet152 as the backbone is shown below. Although they all use multi-scale convolution kernels, the former uses multi-scale kernels derived from one kernel, while the latter directly uses multiple different kernels with more powerful modeling capabilities. As shown in Figure 2, we can observe that neutral samples have a high probability of being correctly predicted, which is also the most common emotion. But we should also note that both models have a prediction bias problem for the neutral class. We speculate that this is because each individual has different neutral standards. In our future work, we will also take into account the

**Table 5.** The performance of different methods on LSSSED.

| Algorithm         | Backbone  | WA           | UA           |
|-------------------|-----------|--------------|--------------|
| FCN-Attention[23] | ALEXNet   | 0.570        | 0.250        |
| MTS-3branches[24] | ALEXNet   | 0.570        | 0.250        |
| MTS-5branches[24] | ALEXNet   | 0.570        | 0.250        |
| MTS-3branches[24] | ResNet152 | 0.585        | 0.296        |
| MTS-5branches[24] | ResNet152 | 0.582        | 0.311        |
| ADV-Real[25]      | VGG16     | 0.570        | 0.250        |
| ADV-Fake[25]      | VGG16     | 0.570        | 0.250        |
| ADV-Real[25]      | ResNet152 | 0.548        | 0.381        |
| ADV-Fake[25]      | ResNet152 | 0.453        | 0.339        |
| VGG[20]           | VGG11     | 0.595        | 0.337        |
| VGG[20]           | VGG13     | 0.604        | 0.393        |
| VGG[20]           | VGG16     | 0.585        | 0.313        |
| VGG[20]           | VGG19     | 0.585        | 0.370        |
| ResNet[21]        | ResNet18  | 0.594        | 0.382        |
| ResNet[21]        | ResNet34  | 0.598        | 0.355        |
| ResNet[21]        | ResNet50  | 0.587        | 0.377        |
| ResNet[21]        | ResNet101 | 0.592        | 0.332        |
| ResNet[21]        | ResNet152 | 0.601        | 0.396        |
| PyResNet          | ResNet50  | <b>0.615</b> | <b>0.420</b> |
| PyResNet          | ResNet101 | <b>0.616</b> | <b>0.428</b> |
| PyResNet          | ResNet152 | <b>0.624</b> | <b>0.429</b> |

resting (neutral) state of each individual. In comparison, our PyResNet has a significant improvement in the angry, happy, and sad categories which are less predictable.

|         | Angry | Neutral | Happy  | Sad   |
|---------|-------|---------|--------|-------|
| Angry   | 5.48% | 72.47%  | 20.12% | 1.93% |
| Neutral | 0.08% | 95.17%  | 4.12%  | 0.63% |
| Happy   | 0.24% | 85.04%  | 13.42% | 1.30% |
| Sad     | 0.17% | 91.31%  | 3.99%  | 4.53% |

(a) The confusion matrix of MTS-3 branches

|         | Angry  | Neutral | Happy  | Sad    |
|---------|--------|---------|--------|--------|
| Angry   | 32.43% | 47.23%  | 16.56% | 3.78%  |
| Neutral | 1.45%  | 86.85%  | 9.54%  | 2.17%  |
| Happy   | 2.77%  | 59.85%  | 35.18% | 2.20%  |
| Sad     | 2.12%  | 71.04%  | 9.83%  | 17.01% |

(b) The confusion matrix of our PyResNet

**Fig. 2.** Confusion matrices of different algorithms.

### 3.3. Pre-trained Model in Downstream Task

With the above pre-trained models, we want to further explore its applicability to downstream tasks. We choose speech-based depression detection as our downstream task. Due to the high professional requirements, it is very difficult to collect data on patients with depression. This leads to the current unsatisfactory effect of automatic depression detection. It is therefore a natural idea to use a pre-trained model with sufficient prior knowledge to improve the detection accuracy. These series of experiments are carried out on the DAIC-WOZ depression database, which is a subset of the Distress Analysis Interview Corpus (DAIC) [26].

We choose SER task and ASR task for transfer. For SER, we use the pre-trained PyResNet with ResNet152 as a backbone. For ASR, we use ESPNet [27], which is an end-to-end encoder-decoder structure network.

**Table 6.** The performance of different pre-trained models on DAIC-WOZ.

| Algorithm      | WA           | UA           |
|----------------|--------------|--------------|
| ESPNet (ASR)   | 0.657        | 0.500        |
| PyResNet (SER) | <b>0.714</b> | <b>0.583</b> |

As Table 6 shows, the performance of transfer based on SER is better than that based on ASR. This is because the features extracted by ASR are bias towards semantics while the features extracted by SER are bias towards acoustics. Depression detection pays more attention to acoustic features which has smaller gaps with SER tasks.

Moreover, we also considered the differences in bandwidth between SER and ASR when framing. ASR generally uses a narrow window length of about 25ms. This means that it pays more attention to changes in time and has a higher time resolution. For SER, we use a wide window length of about 65ms, which means that the frequency information in each frame is richer and the frequency resolution is higher. Since a high frequency resolution is conducive to extracting acoustic features, the SER pre-trained model may be a better choice for some downstream tasks such as depression detection.

## 4. CONCLUSION

In this work, we present LSSSED, a challenging large-scale english database for speech emotion recognition that can simulate real distribution. We point out that existing algorithms tend to overfit small-scale databases and thus cannot be well generalized to real scenes. Furthermore, we release some pre-trained models based on LSSSED. These models can not only promote the development of SER, but can also be transferred to similar downstream tasks like mental health analysis where data is extremely difficult to collect.

## 5. ACKNOWLEDGEMENT

The work is supported in part by Key-Area Research and Development Program of Guangdong Province, China, under grant 2019B010154003; in part by the National Natural Science Foundation of China under Grant U1801262,61972163; in part by the Fundamental Research Funds for the Central Universities, under grant 2019PY21,2019MS028, in part by Guangdong Provincial Key Laboratory under grant 2020B121202011, and in part by the Guangzhou Key Laboratory of Body Data Science under Grant201605030011.

## References

- [1] B. Schuller, G. Rigoll, and M. Lang, "Hidden markov model-based speech emotion recognition," in *2003 IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003. Proceedings.(ICASSP'03)*. IEEE, 2003, vol. 2, pp. II–1.
- [2] A. Batliner, C. Hacker, S. Steidl, E. Nöth, S. D'Arcy, M. J. Russell, and M. Wong, "'you stupid tin box'-children interacting with the aibo robot: A cross-linguistic emotional speech corpus.," in *Lrec*, 2004.
- [3] F. Burkhardt, A. Paeschke, M. Rolfes, W. F. Sendlmeier, and B. Weiss, "A database of german emotional speech," in *Ninth European Conference on Speech Communication and Technology*, 2005.
- [4] O. Martin, I. Kotsia, B. Macq, and I. Pitas, "The interface'05 audio-visual emotion database," in *22nd International Conference on Data Engineering Workshops (ICDEW'06)*. IEEE, 2006, pp. 8–8.
- [5] Y. Wang and L. Guan, "Recognizing human emotional state from audiovisual signals," *IEEE transactions on multimedia*, vol. 10, no. 5, pp. 936–946, 2008.
- [6] C. Busso, M. Bulut, C.-C. Lee, A. Kazemzadeh, E. Mower, S. Kim, J. N. Chang, S. Lee, and S. S. Narayanan, "Iemocap: Interactive emotional dyadic motion capture database," *Language resources and evaluation*, vol. 42, no. 4, pp. 335, 2008.
- [7] A. Dhall, R. Goecke, S. Lucey, and T. Gedeon, "Collecting large, richly annotated facial-expression databases from movies," *IEEE multimedia*, , no. 3, pp. 34–41, 2012.
- [8] S. Poria, D. Hazarika, N. Majumder, G. Naik, E. Cambria, and R. Mihalcea, "Meld: A multimodal multi-party dataset for emotion recognition in conversations," *arXiv preprint arXiv:1810.02508*, 2018.
- [9] C.-C. Lee, E. Mower, C. Busso, S. Lee, and S. Narayanan, "Emotion recognition using a hierarchical binary decision tree approach," *Speech Communication*, vol. 53, no. 9-10, pp. 1162–1171, 2011.
- [10] A. Stuhlsatz, C. Meyer, F. Eyben, T. Zielke, G. Meier, and B. Schuller, "Deep neural networks for acoustic emotion recognition: Raising the benchmarks," in *2011 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2011, pp. 5688–5691.
- [11] S. Zhang, S. Zhang, T. Huang, and W. Gao, "Speech emotion recognition using deep convolutional neural network and discriminant temporal pyramid matching," *IEEE Transactions on Multimedia*, vol. 20, no. 6, pp. 1576–1590, 2017.
- [12] A. Satt, S. Rozenberg, and R. Hoory, "Efficient emotion recognition from speech using deep learning on spectrograms.," in *Interspeech*, 2017, pp. 1089–1093.
- [13] S.-L. Yeh, Y.-S. Lin, and C.-C. Lee, "A dialogical emotion decoder for speech motion recognition in spoken dialog," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6479–6483.
- [14] Y. Zong, W. Zheng, T. Zhang, and X. Huang, "Cross-corpus speech emotion recognition based on domain-adaptive least-squares regression," *IEEE signal processing letters*, vol. 23, no. 5, pp. 585–589, 2016.
- [15] M. Abdelwahab and C. Busso, "Domain adversarial for acoustic emotion recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 12, pp. 2423–2435, 2018.
- [16] N. Liu, Y. Zong, B. Zhang, L. Liu, J. Chen, G. Zhao, and J. Zhu, "Unsupervised cross-corpus speech emotion recognition using domain-adaptive subspace learning," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5144–5148.
- [17] H. Luo and J. Han, "Nonnegative matrix factorization based transfer subspace learning for cross-corpus speech emotion recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2047–2060, 2020.
- [18] J. Boigne, B. Liyanage, and T. Östrem, "Recognizing more emotions with less data using self-supervised transfer learning," 2020.
- [19] N. Upadhyay and A. Karmakar, "Speech enhancement using spectral subtraction-type algorithms: A comparison and simulation study," *Procedia Computer Science*, vol. 54, pp. 574–584, 2015.
- [20] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [21] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [22] I. C. Duta, L. Liu, F. Zhu, and L. Shao, "Pyramidal convolution: Rethinking convolutional neural networks for visual recognition," *arXiv preprint arXiv:2006.11538*, 2020.
- [23] Y. Zhang, J. Du, Z. Wang, J. Zhang, and Y. Tu, "Attention based fully convolutional network for speech emotion recognition," in *2018 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2018, pp. 1771–1775.
- [24] E. Guizzo, T. Weyde, and J. B. Leveson, "Multi-time-scale convolution for emotion recognition from speech audio signals," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6489–6493.
- [25] Z. Ren, A. Baird, J. Han, Z. Zhang, and B. Schuller, "Generating and protecting against adversarial attacks for deep speech-based emotion recognition models," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7184–7188.
- [26] J. Gratch, R. Artstein, G. M. Lucas, G. Stratou, S. Scherer, A. Nazarian, R. Wood, J. Boberg, D. DeVault, S. Marsella, et al., "The distress analysis interview corpus of human and computer interviews.," in *LREC*, 2014, pp. 3123–3128.
- [27] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. E. Y. Soplin, J. Heymann, M. Wiesner, N. Chen, et al., "Espnet: End-to-end speech processing toolkit," *arXiv preprint arXiv:1804.00015*, 2018.