



Self-paced ensemble learning for speech and audio classification

Nicolae-Cătălin Ristea¹, Radu Tudor Ionescu²

¹University Politehnica of Bucharest, Romania

²University of Bucharest, Romania

r.catalin196@yahoo.ro, raducu.ionescu@gmail.com

Abstract

Combining multiple machine learning models into an ensemble is known to provide superior performance levels compared to the individual components forming the ensemble. This is because models can complement each other in taking better decisions. Instead of just combining the models, we propose a self-paced ensemble learning scheme in which models learn from each other over several iterations. During the self-paced learning process based on pseudo-labeling, in addition to improving the individual models, our ensemble also gains knowledge about the target domain. To demonstrate the generality of our self-paced ensemble learning (SPEL) scheme, we conduct experiments on three audio tasks. Our empirical results indicate that SPEL significantly outperforms the baseline ensemble models. We also show that applying self-paced learning on individual models is less effective, illustrating the idea that models in the ensemble actually learn from each other.

Index Terms: ensemble learning, self-paced learning, speech emotion recognition, tropical species detection, mask detection.

1. Introduction

State-of-the-art machine learning models are trained using supervision in the form of labeled data samples provided by an oracle (a domain expert or an ordinary human that knows how to label the samples). However, effective human learners have the ability to guide themselves through the learning activities, considering increasingly difficult concepts at their own pace, without the help of an oracle (a teacher). Indeed, most of the times, students choose what, how, when and how long to study, implicitly using a self-paced curriculum. These self-regulated aspects of learning influence the future results of the learner. This cognitive learning paradigm inspired researchers to study and propose machine learning models based on self-paced learning (SPL) [1, 2]. We underline that SPL can be seen as a particular case of curriculum learning [3, 4], in which the difficulty of the examples is estimated by the model itself, hence the self-pacing. A variety of SPL schemes have been designed for different computer vision [2, 5] and other pattern recognition tasks [6, 7], with demonstrated improvements over the standard supervised learning paradigm in specific scenarios. In signal processing, there are only a handful of approaches based on SPL [8].

In the most straightforward SPL approach, the model uses its own confidence scores to determine what are the next examples to learn. However, if the examples come from a different (target) domain and their ground-truth labels are not available, the model can only use the labels predicted by itself (pseudo-labels). Therefore, in the context of unsupervised domain adaptation, we encounter a paradox: there is little new information to be learned from pseudo-labeled examples that are selected by the model with high confidence. To overcome the paradox, some methods [9, 10] use an external difficulty predictor. In this

work, we conjecture that useful information can be learned simply by training multiple models together. To this end, we propose to jointly train an ensemble of models that learn from each other through SPL. More precisely, an individual model from the ensemble will benefit from the pseudo-labels and the confidence scores provided by the whole ensemble, thus being able to take useful information from the other models included in the ensemble. To the best of our knowledge, we are the first to propose self-paced ensemble learning (SPEL) in the audio domain. By combining self-paced learning and ensemble learning into a single framework, there are two essential benefits: higher effectiveness and increased generality. To demonstrate the benefits of SPEL, we conduct experiments on three audio benchmark tasks, namely speech emotion recognition on CREMA-D [11], speech mask detection on MSC [12] and bird and frog species detection on RCSAD [13]. Our empirical results indicate that SPEL provides superior results than the baseline ensembles and the state-of-the-art methods. In addition, we present ablation results on CREMA-D showing that employing self-paced learning on an individual model is less effective, indicating that models in an ensemble can effectively learn from each other.

2. Related Work

Ensemble learning. On the one hand, traditional machine learning models may fail to attain the desired performance level when dealing with high-dimensional, noisy or imbalanced data. The reason behind this is that such models belong to restricted hypothesis classes, thus not being able to discover complex patterns revealing the underlying structure of the data. On the other hand, deep neural networks can model such complex patterns when sufficient data is available, but such models lack stability, mainly due to the random initialization of parameters and the stochastic optimization process. In both cases, ensemble learning [14, 15] offers improved effectiveness and stability by integrating multiple models into a single framework. There are several ensemble learning methods in literature, ranging from majority voting [16] and Bayesian model averaging [17] to bagging [18], boosting [19] and classifier stacking [20]. As other ensemble learning methods applied in the audio domain [21, 22], we employ Bayesian model averaging in our method. Different from the related methods, we adapt our ensemble to the target domain using self-paced learning, attaining superior results.

Self-paced learning. The vast majority of SPL approaches proposed in literature are focused on computer vision topics, the technique being proved effective in many vision tasks [23, 24]. Sangineto et al. [23] proposed a new protocol based on the self-paced learning paradigm, where the main idea is to iteratively select a subset of images and bounding boxes that are the most reliable, and use them for training. The proposed method outperforms related weakly-supervised object detectors, attaining state-of-the-art results. Zhou et al. [24] designed a model based

on an SPL regime performing person re-identification, which usually suffers from noisy samples with background clutter and mutual occlusion, making it extremely difficult to distinguish different individuals across disjoint camera views. Their solution seems to overpass these issues, attaining superior performance in comparison with state-of-the-art approaches.

Different from these methods [23, 24], we employ an algorithm that combines the advantages of ensemble learning and SPL, resulting in a more effective and stable approach.

Self-paced ensemble learning. Pi et al. [25] proposed a self-paced boost learning method that incorporates an adaptive easy-to-hard pace in the boosting process. The method is demonstrated on generic tasks using weak learners such as Logistic Regression or Support Vector Machines. Liu et al. [26] introduced a self-paced ensemble method to cope with the class imbalance problem. The method is based on combining self-paced learning with undersampling to generate a robust ensemble. As Pi et al. [25], they demonstrated the utility of the proposed method on ensembles of weak learners. Since deep neural networks require large data sets during training, the undersampling approach of Liu et al. [26] is not suitable for deep models. In contrast, we propose an ensemble of deep models and apply SPL for unsupervised domain adaptation.

SPL has also been applied on clustering methods. For example, Zhou et al. [7] proposed a novel self-paced clustering ensemble that gradually adds instances, from easy to difficult, into the ensemble learning. Their approach overcomes the existing problems of ensemble clustering methods, which usually exploit all data samples to learn a consensus clustering result, ignoring the adverse effects caused by difficult instances. Ghasedi et al. [27] proposed a different clustering approach, in which a balanced self-paced learning algorithm improves the training process of deep generative adversarial clustering networks, tackling the drawbacks of training deep clustering models in an unsupervised manner. Unlike these methods [7, 27], we focus on supervised classification and detection problems.

Wang et al. [28] used SPL in a very particular scenario, focusing on a Random Forest ensemble applied specifically to lung cancer prognosis. Different from Wang et al. [28], we consider ensembles composed of various deep neural architectures applied on a broad variety of audio tasks.

Unlike the aforementioned methods [7, 25, 26, 27, 28], we employ an ensemble of deep neural networks that are (i) trained on the source domain using standard supervision, then (ii) fine-tuned on the target domain using self-paced learning with pseudo-labeling, leading to a more robust ensemble with superior results on multiple audio classification and detection tasks. To the best of our knowledge, there is no prior work based on self-paced ensemble learning in the audio domain.

3. Method

Data representation. We first transform each audio sample into a 2D representation in order to be able to employ state-of-the-art convolutional neural networks (CNNs) [29, 30] based on 2D convolutions. To this end, we compute the discrete Short Time Fourier Transform (STFT), as follows:

$$STFT\{x[n]\}(m, k) = \sum_{n=-\infty}^{\infty} x[n] \cdot w[n-mR] \cdot e^{-j \frac{2\pi}{N_x} kn}, \quad (1)$$

where $x[n]$ is the discrete input signal, $w[n]$ is a window function (in our approach, Hann), N_x is the STFT length and R is the hop (step) size [31]. Afterwards, we compute the spectrograms as the squared magnitude of the STFT, convert the result-

ing values to a logarithmic scale (decibels) and normalize them to the interval $[-1, 1]$, generating a single-channel grayscale image. Moreover, the frequency bins were mapped onto the Mel scale with 256 Mel bands.

Baseline model ensemble. As our learning approach, we consider an ensemble of residual neural networks [29, 32] that are combined through Bayesian model averaging. We employ residual nets because such architectures are known to produce state-of-the-art results on a wide variety of tasks, mainly due to their very high depth. In order to significantly increase the depth over the previous models [30, 33], He et al. [29] introduced skip connections to eliminate vanishing gradient problems in training very deep neural models, providing alternative pathways for the gradients during back-propagation. We opted for different residual architectures for each data set, adjusting each architecture and ensemble to match the complexity of the task and the number of data samples. Depending on the data set, our ensemble is typically formed of four or five residual nets.

Self-paced ensemble learning. In self-paced learning, machine learning models learn at their own pace, taking into consideration the samples with high confidence predictions first. This technique is robust in solving problems with noisy labels [34]. In our work, we consider an unsupervised domain adaptation scenario, in which the ground-truth labels of the target domain samples used for SPL are unknown. Hence, the model needs to learn from its own labels, called pseudo-labels. We propose an automatic self-paced ensemble learning technique that iteratively enriches the training data with self-annotated examples from the target domain. The proposed approach assumes no overhead at testing time, the network architectures being identical (we only update the models' weights). The steps required by our framework are formally described in Algorithm 1. The algorithm is divided into three stages. The first two stages are about training the ensemble, while the third stage is about evaluating the ensemble on the final test set.

In the first stage (steps 1-6), we employ the conventional training procedure based on stochastic gradient descent with mini-batches. Each neural network f_i which is part of the ensemble is optimized until we reach an optimal convergence point. At the current iteration t , the first step (step 4) is to randomly sample a mini-batch of labeled examples from the training data set (X, T) . Then, at step 5, the weights θ_i corresponding to the network f_i are updated in the negative direction of the gradient of the loss function $\nabla \mathcal{L}$, the update step being controlled through the learning rate η .

In the second stage (steps 7-18), we employ our self-paced ensemble learning procedure for a number of k iterations, where k is a pre-established hyperparameter of our method. At each iteration, we start by computing the predicted labels $P_{\bar{Y}}$ and the confidence scores $C_{\bar{Y}}$ provided by the ensemble (step 8) for the unlabeled set \bar{Y} , which belongs to the target domain. Next, at step 9, we sort the set of samples \bar{Y} and the associated pseudo-labels $P_{\bar{Y}}$ with respect to $C_{\bar{Y}}$, in decreasing order of the confidence scores. Afterwards, we select $m \cdot j$ samples (step 10) and the corresponding pseudo-labels (step 11), which are added (steps 12 and 13) to the new training set (X', T') . We underline that at each iteration j , the number of extra samples is m , where m is another pre-established hyperparameter of our method. As the confidence scores of the previously included samples may change over time, we re-evaluate them at each iteration j , regenerating X' and T' . Hence, at iteration j , we add $m \cdot j$ pseudo-labeled samples. Upon forming the new training set (X', T') , we continue the training of the individual models, applying stochastic gradient descent with mini-batches

Algorithm 1 Self-paced ensemble learning (SPEL)

Input: (X, T) - a training set of samples and associated labels from the source domain; Y - a test set from the target domain; \bar{Y} - an unlabeled set from the target domain; n - the number of models in the ensemble; k - the number of SPEL steps; m - the number of extra examples added at a SPEL step; η - a learning rate; \mathcal{L} - a loss function.

Notations: f_i - the i -th model in the ensemble; θ_i - the weights of the i -th model; avg - a function that computes Bayesian model averaging; $sort$ - a function that jointly sorts the input set; $\mathcal{N}(0, \Sigma)$ - the normal distribution of mean 0 and standard deviation Σ ; $\mathcal{U}(S)$ - the uniform distribution over the set S ; $S[i : j]$ - subset of elements $\{s_i, s_{i+1}, \dots, s_j\}$ from S .

Initialization: $\theta_i^{(0)} \sim \mathcal{N}(0, \Sigma), \forall i \in \{1, \dots, n\}$

Output: P_Y - the predictions for the test set Y .

Stage 1: Pre-training of individual models

```
1: for  $i \leftarrow 1$  to  $n$  do
2:    $t \leftarrow 0$ 
3:   while converge criterion not met do
4:      $X^{(t)}, T^{(t)} \leftarrow \text{mini-batch} \sim \mathcal{U}((X, T))$ 
5:      $\theta_i^{(t+1)} = \theta_i^{(t)} - \eta^{(t)} \nabla \mathcal{L}(\theta_i^{(t)}, X^{(t)}, T^{(t)})$ 
6:      $t \leftarrow t + 1$ 
```

Stage 2: Self-paced ensemble learning

```
7: for  $j \leftarrow 1$  to  $k$  do
8:    $P_{\bar{Y}}, C_{\bar{Y}} \leftarrow avg(f_1(\theta_1^{(t)}, \bar{Y}), \dots, f_n(\theta_n^{(t)}, \bar{Y}))$ 
9:    $\bar{Y}', P_{\bar{Y}'} \leftarrow sort(\bar{Y}, P_{\bar{Y}})$  with respect to  $C_{\bar{Y}}$ 
10:   $\bar{X}^{(j)} \leftarrow \bar{Y}'[1 : m \cdot j]$ 
11:   $\bar{T}^{(j)} \leftarrow P_{\bar{Y}'}[1 : m \cdot j]$ 
12:   $X' \leftarrow X \cup \bar{X}^{(j)}$ 
13:   $T' \leftarrow T \cup \bar{T}^{(j)}$ 
14:  for  $i \leftarrow 1$  to  $n$  do
15:    while converge criterion not met do
16:       $X^{(t)}, T^{(t)} \leftarrow \text{mini-batch} \sim \mathcal{U}((X', T'))$ 
17:       $\theta_i^{(t+1)} = \theta_i^{(t)} - \eta^{(t)} \nabla \mathcal{L}(\theta_i^{(t)}, X^{(t)}, T^{(t)})$ 
18:       $t \leftarrow t + 1$ 
```

Stage 3: Prediction

```
19:  $P_Y \leftarrow avg(f_1(\theta_1^{(t)}, Y), \dots, f_n(\theta_n^{(t)}, Y))$ 
```

in steps 14-18. We underline that, by eliminating the second training stage, we fall back on the baseline ensemble learning.

In the third and final stage (step 19), we apply our ensemble based on Bayesian model averaging on the test set Y , obtaining the final predictions P_Y . We note that the proposed self-paced ensemble learning algorithm can be applied on top of any type of ensemble and it does not involve any additional overhead at test time.

4. Experiments

4.1. Data sets

CREMA-D. The CREMA-D multi-modal database [11] contains 7,442 videoclips of 91 actors (48 male and 43 female) with different ethnic backgrounds. The actors were asked to simulate particular emotions while producing, with different intonations, 12 particular sentences that evoke the target emotions. There are six emotion categories: neutral, happy, anger, disgust, fear and sad. In our experiments, we consider only the audio modal-

ity as source of information. We split the audio samples into 70% for training, 15% for validation and 15% for testing.

RCSAD. The Rainforest Connection Species Audio Detection (RCSAD) data set [13] is released by Rainforest Connection (RFCx) and is based on acoustic data collected while monitoring the ecosystem soundscape at tropical locations. RCSAD contains 4,727 audio files for training and 1,992 audio files for testing. We kept 20% of the training set for validation. The data set contains 24 classes representing bird and frog species from tropical regions. An audio sample may contain one or more species. Therefore, the goal is to solve a multi-label classification task. The data set is released in a recent Kaggle competition, implying that the only way to evaluate a model is by making a submission on Kaggle, i.e. the test labels are private.

MSC. The Mask Augsburg Speech Corpus (MSC) is provided by the ComParE organizers [12]. It comprises recordings of 32 German native speakers, with or without wearing surgical masks. Each data sample is a recording of 1 second at a sampling rate of 16 KHz. The data set is partitioned into a training set of 10,895 samples, a development set of 14,647 samples and a test set of 11,012 samples. Since the test labels are private, we report results on the development set, keeping 20% of the training set for validation.

4.2. Evaluation setup

Performance measures. For the CREMA-D data set, we report the classification accuracy. For the RCSAD data set, the Kaggle competition organizers decided to rank participants based on the weighted label-ranking average precision (WLRAP). Therefore, we report our performance in terms of this measure. Regarding the MSC data set, we report the unweighted average recall (UAR), this being the official measure in the ComParE challenge [12].

Baselines. On all data sets, we consider as baseline the model ensemble trained without the second stage described in Algorithm 1. For the CREMA-D data set, we added as baselines two state-of-the-art methods [35, 36] that reported the accuracy on the audio modality. For RCSAD, we refer to our ranking in the Kaggle competition, as there are no prior publications reporting results on this data set. For both CREMA-D and RCSAD, we also report the results of the best model in the baseline ensemble. For the MSC data set, we considered the official competition baseline as well as the approach proposed in [37], the latter approach coinciding with our baseline model ensemble.

Data preprocessing and tuning. We first standardized all audio clips to a fixed dimension. For CREMA-D, we padded or clipped the audio files to 4 seconds, while for RCSAD, we extracted an 8 second window centered on each training detection and we applied the sliding window algorithm at test time, keeping the maximum probabilities. The length of the MSC audio files is already established by the organizers at 1 second. When applying the STFT, we used $N_x = 1024$, $R = 64$ and a window size of 512. Each utterance is represented as a Mel spectrogram of $1 \times L \times 256$ components, where 256 is the number of Mel bins and L represents the number of time bins, which is 512 for CREMA-D and 618 for RCSAD. Exceptionally, on the MSC data set, we performed the preprocessing proposed in [37], which involves only STFT.

While our ensembles are all based on residual nets, we used slight different architectures on the three data sets. For CREMA-D, we employed an ensemble of five ResNet-18 models [29] with leaky ADA activations [38] which increase the model's capability to classify non-linearly separable data, such as emotions from voice recordings. For RCSAD, we consid-

Table 1: Results of our SPEL approach versus various baseline and state-of-the-art methods on CREMA-D. Significantly better self-paced learning results are marked with ‡, using a paired McNemar’s test [40] at the significance level 0.01.

Method	Accuracy
Shukla et al. [35]	55.01
He et al. [36]	58.71
ResNet-18 + leaky ADA [38]	65.37
ResNet-18 + leaky ADA [38] + SPL	66.04 [‡]
5×ResNet-18 + leaky ADA [38]	66.54
5×ResNet-18 + leaky ADA [38] + SPEL	68.12 [‡]

Table 2: Results of our SPEL approach versus several baselines on RCSAD. Significantly better self-paced learning results are marked with ‡, using a paired McNemar’s test [40] at the significance level 0.01.

Method	WLRAP	Kaggle Rank
ResNeSt-50	0.862	521 / 1143
5×ResNeSt-50	0.900	98 / 1143
5×ResNeSt-50 + SPEL	0.906 [‡]	84 / 1143
5×ResNeSt-50 + SPEL + PP	0.943 [‡]	28 / 1143

ered an ensemble of five ResNeSt-50 models [32]. We modified the classification layer of the ResNeSt-50 network to predict 24 classes and added a sigmoid activation function, such that the network is able to output the probability of each class independently of the other classes. For MSC, the ensemble comprises a ResNet-18, a ResNet-34, a ResNet-50 and a ResNet-101, being identical to the ensemble proposed by Ristea et al. [37].

All models are optimized with Adam [39]. For CREMA-D, we trained the ResNet-18 models for 70 epochs on mini-batches of 16 samples using a learning rate of $5 \cdot 10^{-4}$. For RCSAD, we set the learning rate to 10^{-3} and trained the models for 30 epochs on mini-batches of 16 samples. For MSC, we used the same hyperparameters as Ristea et al. [37]. Our SPEL framework, has two additional hyperparameters. For the parameter m , we considered values in the set $\{50, 100, 150, 200\}$. For the parameter k , we considered all possible values subject to $m \cdot k \leq 1,000$, i.e. we added at most 1,000 pseudo-labeled samples during our second training stage. On CREMA-D, we obtained optimal results with $m = 50$ and $k = 3$. On the other two data sets, we obtained optimal results with $m = 150$ and $k = 4$. All hyperparameters are tuned on the validation sets.

4.3. Results

In Tables 1, 2 and 3, we present the results obtained on three different data sets by our SPEL approach against various baseline methods. We observe that SPEL attains the highest improvements on CREMA-D (see Table 1), where it outperforms all state-of-the-art results with an accuracy of 68.12%. The improvement of 1.58% over the corresponding baseline ensemble is statistically significant according to a paired McNemar’s test [40] performed at the significance level 0.01. Additionally, we empirically demonstrate that our SPEL approach brings a higher gain (of 1.58%) than the gain (of 0.67%) of SPL applied on top of a single model. This indicates that the models in the ensemble are able to boost each others’ accuracy.

On RCSAD (see Table 2), the ensemble model trained with SPEL surpasses the baseline ensemble by a statistically

Table 3: Results of our SPEL approach versus the official competition baseline and the ensemble proposed by Ristea et al. [37] on MSC. Significantly better self-paced learning results are marked with ‡, using a paired McNemar’s test [40] at the significance level 0.01.

Method	UAR
DeepSpectrum [12]	63.4
Ristea et al. [37]	72.2
Ristea et al. [37] + SPEL	72.5 [‡]

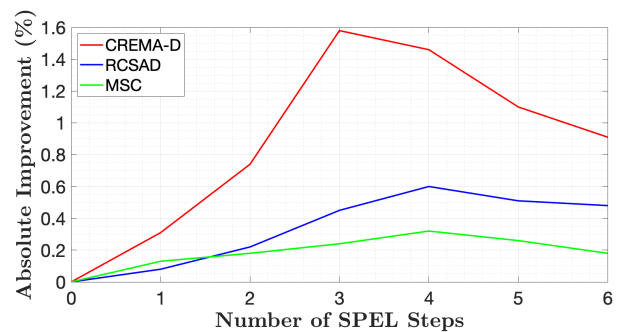


Figure 1: The absolute improvements brought by our SPEL algorithm at various steps, with respect to the corresponding baseline model ensemble. Best viewed in color.

significant margin of 0.6%. We obtained an additional post-competition performance boost (of 3.7%), when we learned that a cheap post-processing (PP)¹ of the probabilities is useful.

In Table 3, we observe that the model proposed by Ristea et al. [37] gains a performance boost of 0.3% when SPEL is applied. Even if the magnitude of the performance gain seems small, a paired McNemar’s test [40] performed at the significance level 0.01 reveals that the difference is significant.

In Figure 1, we show the absolute performance gains brought by our approach with respect to the number of SPEL steps k . We observe that the generic behavior of SPEL is to increase the performance until a certain point (typically, after 3 to 5 steps), and afterwards, the performance slowly declines. We note that SPEL brings in pseudo-labeled examples during training, which means that a small percentage of the newly added sample have incorrect labels. The proportion of incorrectly labeled samples increases with the number of SPEL steps. Hence, when the signal-to-noise ratio of the labels becomes unfavorable, the performance starts to decline.

5. Conclusion

In this paper, we presented a self-paced ensemble learning algorithm in which models learn from each other over several iterations. To demonstrate the generality of our SPEL scheme, we conducted comprehensive experiments with neural architectures of various depths on three audio tasks, where we obtained statistically significant performance gains. In future work, we aim to study increasing the regularization when the ensemble’s performance starts to decline.

Acknowledgements. Work supported by a grant of the Romanian Ministry of Education and Research, CNCS - UEFISCDI, project no. PN-III-P1-1.1-TE-2019-0235, within PNCDI III.

¹<https://www.kaggle.com/c/rfcx-species-audio-detection/discussion/220389>

6. References

- [1] M. P. Kumar, B. Packer, and D. Koller, "Self-Paced Learning for Latent Variable Models," in *Proceedings of NIPS*, vol. 23, 2010, pp. 1189–1197.
- [2] L. Jiang, D. Meng, Q. Zhao, S. Shan, and A. Hauptmann, "Self-paced curriculum learning," in *Proceedings of AAAI*, 2015, pp. 2694–2700.
- [3] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in *Proceedings of ICML*, 2009, pp. 41–48.
- [4] P. Soviany, R. T. Ionescu, P. Rota, and N. Sebe, "Curriculum learning: A survey," *arXiv preprint arXiv:2101.10382*, 2021.
- [5] W. Lin, J. Gao, Q. Wang, and X. Li, "Pixel-Level Self-Paced Learning For Super-Resolution," in *Proceedings of ICASSP*, 2020, pp. 2538–2542.
- [6] W. Zheng, X. Zhu, G. Wen, Y. Zhu, H. Yu, and J. Gan, "Unsupervised feature selection by self-paced learning regularization," *Pattern Recognition Letters*, vol. 132, pp. 4–11, 2020.
- [7] P. Zhou, L. Du, X. Liu, Y.-D. Shen, M. Fan, and X. Li, "Self-Paced Clustering Ensemble," *IEEE Transactions on Neural Networks and Learning Systems*, 2020.
- [8] R. Shang, Y. Yuan, L. Jiao, Y. Meng, and A. M. Ghalamzan, "A self-paced learning algorithm for change detection in synthetic aperture radar images," *Signal Processing*, vol. 142, pp. 375–387, 2018.
- [9] P. Soviany, C. Ardei, R. T. Ionescu, and M. Leordeanu, "Image Difficulty Curriculum for Generative Adversarial Networks (CUGAN)," in *Proceedings of WACV*, 2020, pp. 3463–3472.
- [10] P. Soviany, R. T. Ionescu, P. Rota, and N. Sebe, "Curriculum self-paced learning for cross-domain object detection," *Computer Vision and Image Understanding*, vol. 204, p. 103166, 2021.
- [11] H. Cao, D. G. Cooper, M. K. Keutmann, R. C. Gur, A. Nenkova, and R. Verma, "CREMA-D: Crowd-sourced emotional multimodal actors dataset," *IEEE Transactions on Affective Computing*, vol. 5, no. 4, pp. 377–390, 2014.
- [12] B. W. Schuller, A. Batliner, C. Bergler, E.-M. Messner, A. Hamilton, S. Amiriparian, A. Baird, G. Rizos, M. Schmitt, L. Stappen, H. Baumeister, A. D. MacIntyre, and S. Hantke, "The INTERSPEECH 2020 Computational Paralinguistics Challenge: Elderly Emotion, Breathing & Masks," in *Proceedings of INTERSPEECH*, 2020, pp. 2042–2046.
- [13] "Rainforest Connection Species Audio Detection," <https://www.kaggle.com/c/rfcx-species-audio-detection/>, accessed: 2021-03-21.
- [14] X. Dong, Z. Yu, W. Cao, Y. Shi, and Q. Ma, "A survey on ensemble learning," *Frontiers of Computer Science*, vol. 14, no. 2, pp. 241–258, 2020.
- [15] X. Zhang, D. Povey, and S. Khudanpur, "A Diversity-Penalizing Ensemble Training Method for Deep Learning," in *Proceedings of INTERSPEECH*, 2015, pp. 3590–3594.
- [16] K. Raza, "Improving the prediction accuracy of heart disease with ensemble learning and majority voting rule," in *U-Healthcare Monitoring Systems*, 2019, pp. 179–196.
- [17] E. Fersini, E. Messina, and F. A. Pozzi, "Sentiment analysis: Bayesian ensemble learning," *Decision Support Systems*, vol. 68, pp. 26–38, 2014.
- [18] R. Bryll, R. Gutierrez-Osuna, and F. Quek, "Attribute bagging: improving accuracy of classifier ensembles by using random feature subsets," *Pattern Recognition*, vol. 36, no. 6, pp. 1291–1302, 2003.
- [19] Z. Jianqiang, "Combing semantic and prior polarity features for boosting twitter sentiment analysis using ensemble learning," in *Proceedings of DSC*, 2016, pp. 709–714.
- [20] R. Sikora and O. H. Al-laymoun, "A Modified Stacking Ensemble Machine Learning Algorithm Using Genetic Algorithms," *Journal of International Technology and Information Management*, vol. 23, no. 1, 2014.
- [21] A. Elowsson and A. Friberg, "Predicting the perception of performed dynamics in music audio with ensemble learning," *The Journal of the Acoustical Society of America*, vol. 141, no. 3, pp. 2224–2242, 2017.
- [22] S. Zahid, F. Hussain, M. Rashid, M. H. Yousaf, and H. A. Habib, "Optimized audio classification and segmentation algorithm by using ensemble methods," *Mathematical Problems in Engineering*, vol. 2015, 2015.
- [23] E. Sanginetto, M. Nabi, D. Culibrk, and N. Sebe, "Self paced deep learning for weakly supervised object detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 3, pp. 712–725, 2018.
- [24] S. Zhou, J. Wang, D. Meng, X. Xin, Y. Li, Y. Gong, and N. Zheng, "Deep self-paced learning for person re-identification," *Pattern Recognition*, vol. 76, pp. 739–751, 2018.
- [25] T. Pi, X. Li, Z. Zhang, D. Meng, F. Wu, J. Xiao, and Y. Zhuang, "Self-paced boost learning for classification," in *Proceedings of IJCAI*, 2016, pp. 1932–1938.
- [26] Z. Liu, W. Cao, Z. Gao, J. Bian, H. Chen, Y. Chang, and T.-Y. Liu, "Self-paced Ensemble for Highly Imbalanced Massive Data Classification," in *Proceedings of ICDE*, 2020, pp. 841–852.
- [27] K. Ghasedi, X. Wang, C. Deng, and H. Huang, "Balanced self-paced learning for generative adversarial clustering network," in *Proceedings of CVPR*, 2019, pp. 4391–4400.
- [28] Q. Wang, Y. Zhou, W. Ding, Z. Zhang, K. Muhammad, and Z. Cao, "Random Forest with Self-Paced Bootstrap Learning in Lung Cancer Prognosis," *ACM Transactions on Multimedia Computing, Communications, and Applications*, vol. 16, no. 1s, 2020.
- [29] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *Proceedings of CVPR*, 2016, pp. 770–778.
- [30] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," in *Proceedings of NIPS*, 2012, pp. 1097–1105.
- [31] J. B. Allen and L. R. Rabiner, "A unified approach to short-time Fourier analysis and synthesis," *Proceedings of the IEEE*, vol. 65, no. 11, pp. 1558–1564, 1977.
- [32] H. Zhang, C. Wu, Z. Zhang, Y. Zhu, Z. Zhang, H. Lin, Y. Sun, T. He, J. Mueller, R. Manmatha, M. Li, and A. Smola, "ResNeSt: Split-Attention Networks," *arXiv preprint arXiv:2004.08955*, 2020.
- [33] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," in *Proceedings of ICLR*, 2014.
- [34] D. Meng, Q. Zhao, and L. Jiang, "A theoretical understanding of self-paced learning," *Information Sciences*, vol. 414, pp. 319–328, 2017.
- [35] A. Shukla, K. Vougioukas, P. Ma, S. Petridis, and M. Pantic, "Visually guided self-supervised learning of speech representations," in *Proceedings of ICASSP*, 2020, pp. 6299–6303.
- [36] G. He, X. Liu, F. Fan, and J. You, "Image2Audio: Facilitating Semi-supervised Audio Emotion Recognition with Facial Expression Image," in *Proceedings of CVPR Workshops*, 2020, pp. 912–913.
- [37] N. Ristea and R. T. Ionescu, "Are you Wearing a Mask? Improving Mask Detection from Speech Using Augmentation by Cycle-Consistent GANs," in *Proceedings of INTERSPEECH*, 2020, pp. 2102–2106.
- [38] M.-I. Georgescu, R. T. Ionescu, N.-C. Ristea, and N. Sebe, "Non-linear Neurons with Human-like Apical Dendrite Activations," *arXiv preprint arXiv:2003.03229*, 2020.
- [39] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proceedings of ICLR*, 2015.
- [40] T. G. Dietterich, "Approximate Statistical Tests for Comparing Supervised Classification Learning Algorithms," *Neural Computation*, vol. 10, no. 7, pp. 1895–1923, 1998.