



# Are you sure? Analysing Uncertainty Quantification Approaches for Real-world Speech Emotion Recognition

Oliver Schrüfer<sup>1</sup>, Manuel Milling<sup>2</sup>, Felix Burkhardt<sup>1</sup>, Florian Eyben<sup>1</sup>, Björn Schuller<sup>1,2,3</sup>

<sup>1</sup>audEERING GmbH, Germany

<sup>2</sup>CHI – Chair of Health Informatics, MRI, Technical University of Munich, Germany

<sup>3</sup>GLAM – the Group on Language, Audio, & Music, Imperial College London, London, UK

oschruefer@audEERING.com, manuel.milling@tum.de, [fburkhardt, fe, bs]@audEERING.com

## Abstract

Uncertainty Quantification (UQ) is an important building block for the reliable use of neural networks in real-world scenarios, as it can be a useful tool in identifying faulty predictions. Speech emotion recognition (SER) models can suffer from particularly many sources of uncertainty, such as the ambiguity of emotions, Out-of-Distribution (OOD) data or, in general, poor recording conditions. Reliable UQ methods are thus of particular interest as in many SER applications no prediction is better than a faulty prediction. While the effects of label ambiguity on uncertainty are well documented in the literature, we focus our work on an evaluation of UQ methods for SER under common challenges in real-world application, such as corrupted signals, and the absence of speech. We show that simple UQ methods can already give an indication of the uncertainty of a prediction and that training with additional OOD data can greatly improve the identification of such signals.

**Index Terms:** Uncertainty Quantification, Speech Emotion Recognition, Prior Networks, EDL, Out-of-Distribution

## 1. Introduction

Intelligent audio analysis is becoming more and more important in our lives, whether in voice assistant systems or assistive robots, e. g., in the care sector. While the field of automatic speech recognition (ASR) has seen major advances through Deep Learning (DL), enabling robust real-world applications, speech emotion recognition (SER) still sees some major challenges after more than two decades of research [1]. A particular challenge hereby is the ambiguity, subjectivity [2] or even untypical expression [3] of emotions and thus the lack of ground truth labels[4]. Even though a rough intuition of emotions can be considered common sense, a wide range of emotion definitions and an almost arbitrary granularity of emotions make it already challenging to compare models between datasets [5]. A common approach to mitigate these challenges, which we also apply in this work, is to model only a subset of the available emotions, which tend to be available in most datasets and have a high agreement amongst annotators.

Still, robust cross-corpus SER models are hard to develop and they struggle when confronted with OOD data. Beyond this, the paralinguistic information necessary to detect emotions is easily corrupted by real-world challenges, such as poor recording conditions or background noise [6].

In line with the subjectivity of emotions, many application scenarios of emotional context do not require a pin-point accurate estimation of emotion. Instead, they rather follow the philosophy of no (or a neutral) emotion prediction being preferable to a faulty emotion prediction. An expressive example of this is the frequent discussion of confusion matrices when analysing

SER performance, where for instance mistaking happiness for anger is considered worse than the confusion with neutral. SER systems could thus benefit vastly from knowing when an emotion prediction is accurate and when it is not. A key approach to put this idea into practice is the quantification of uncertainty.

Uncertainties have been studied extensively in the literature and can usually be divided into two or three different categories [7, 8]: aleatoric uncertainty, caused by uncertainty in the data, on the one hand and epistemic uncertainty (model uncertainty), as well as distributional uncertainty on the other hand. In the context of SER, aleatoric uncertainty arises for instance through the ambiguity of the emotions themselves or through corrupted data. Epistemic uncertainty arises from the model itself and describes the inability of the model to precisely model the underlying distribution. One common reason for this is a lack of variability in the training data [9], for instance unbalanced gender or age distribution or a mismatch of the language between training data and real-world application. Distributional uncertainty arises, as the name suggests, from OOD samples. In the case of SER, this could be emotions which were not included in the training set or recordings that do not contain speech at all, e. g., music, silence or ambient sounds. This is also often referred to as Out-of-Domain samples [10, 11], which for the sake of simplicity we will summarise with the mathematical term Out-of-Distribution under the abbreviation OOD.

Popular methods aiming to quantify these effects of uncertainty on a sample level can generally be split into four different types: single deterministic methods, Bayesian methods, ensemble methods and test-time augmentation methods [8]. The latter three all rely on multiple forward passes through a network for an uncertainty estimation and thus add significant computational cost at inference time, making efficient deployment challenging. Single deterministic methods, however, only add minimal computational overhead and are thus promising candidates for real-time real-world uncertainty estimation in SER applications. Apart from being used directly by an application, UQ can also help to leverage more sophisticated training methods, such as active [12] or weakly-supervised learning.

In this work, we therefore aim to investigate the behaviour of several common (mostly single deterministic) UQ methods under simulated challenges of real-world SER scenarios, in particular when confronted with OOD data and different noise levels.

To our knowledge, this is the first systematic investigation to evaluate the most prominent UQ methods on multiple SER datasets under realistic challenges of real-world application.

## 2. Uncertainty Tests

In order to obtain a good picture of our SER models with respect to different contexts, we measure the models' behaviour with

respect to four different uncertainty tests.

### 2.1. Rater Agreement

A rater agreement test is a straight-forward approach to evaluate aleatoric uncertainty in the context of SER [13]. It calculates the sample-level correlation between the uncertainty of the model and the level of agreement of different annotators during the labelling process. We thus obtain an impression of whether the sources of uncertainty are similar for a SER model and human annotators.

### 2.2. Unknown Emotions

For uncertainty quantification in classification tasks, one of the main questions is, how the UQ method reacts to data points from classes not included during the training, in general aiming for higher uncertainties for unknown classes. Accordingly, we will test our SER models for samples with emotion labels that are unknown to the model. It is important to note, however, that the expressiveness of this test suffers from limitations in the SER case, as the ambiguity and potential overlaps of emotions may not be properly accounted for in the datasets[14].

### 2.3. Non Speech Data

A more clearly interpretable test is to expose the model to OOD data, which does not contain speech at all, such as “background music” or environmental sounds. In this case, it would be desirable for the UQ measures to have high scores, in order to avoid incorrect predictions. Most real-world SER applications apply Voice Activity Detection to prevent this type of faulty predictions, thus adding additional computational resources.

### 2.4. Corrupted Signals

Our final test aims to simulate bad recording quality, one of the major challenges of real-world applications. For that purpose, we augment data with varying levels of white noise or environmental sounds and investigate the interplay of model performance and uncertainty predictions for increasing levels of noise.

## 3. Uncertainty Quantification

With respect to the above discussed tests, we analyse four UQ methods, three of which belong to the category of single deterministic UQ methods, with only Monte Carlo (MC) Dropout needing several forward passes through the model.

### 3.1. Entropy

One of the easiest ways to get insights about the uncertainty of a categorical task is to calculate the entropy [15] over logits. The entropy is calculated as  $H = -\sum_{i=1}^N p_i \log(p_i)$  for  $N$  classes, where  $p_i$  is the probability to predict class  $i$ . We consider this method our baseline as it can be calculated a posteriori for each classification model trained with a standard softmax layer and (weighted) cross-entropy loss.

### 3.2. MC Dropout

Many State-of-the-art (SOTA) models like PANNs [16] or wav2vec2.0 [17] contain dropout layers to prevent overfitting during training. This again enables an out-of-the-box use of our second baseline, MC dropout [18], which works by reactivating the dropout layers during inference to obtain multiple outputs for the same input sample. The UQ is then calculated as the




Figure 1: Simplexes of Dirichlet Distributions for a 3-class problem with sharp (a,b) and flat (c) distributions

variance over the different outputs. For classification tasks, it is alternatively possible to use the mean probability per class to calculate the entropy, as in section 3.1.

### 3.3. Evidential Deep Learning

Evidential Deep Learning (EDL) is a more sophisticated approach of UQ and formulates learning as an evidence acquisition process [19, 20], where each of the training samples adds to an evidential distribution. It is based on the Dempster-Shafer theory of belief functions [21]. For application, Sensoy et al. [19] specify non-negative belief masses  $b_k$  for all  $K$  classes and an uncertainty mass  $u$ , which all together add up to one  $u + \sum_{k=1}^K b_k = 1$ . The model outputs  $K$  values, describing the evidence per class  $e_k$ . The evidences are thereby modelled by a Dirichlet distribution with the parameters  $\alpha = (\alpha_1, \dots, \alpha_K)$ , as  $\alpha_k = e_k + 1$ , which are then used to calculate the uncertainty and belief masses.

### 3.4. Prior Networks

Another approach specifically tailored to detect OOD samples are Prior Networks (PN)s [22]. Similar to EDL, PNs formulate the ground truth via an underlying Dirichlet distribution (see Figure 1). This further allows for a distinction between data uncertainty and OOD uncertainty. By defining sharp distributions towards the respective classes, the network is trained for its classification task (Figure 1 a). By specifying a flat distribution (Figure 1 c) it is possible to train on OOD data without the need for an additional OOD class. For a 4-class-problem for instance, one would use, instead of the one hot encoding  $(1, 0, 0, 0)$ , a Dirichlet distribution with  $\alpha = (101, 1, 1, 1)$  and for OOD data  $\alpha = (1, 1, 1, 1)$  with the Kullback-Leibler (KL) divergence as the loss function. For PNs it is possible to differentiate between data uncertainty and OOD data by calculating the precision (sharpness) of the distribution. For the sake of simplicity, however, we will only refer to the combined uncertainty in this paper, which is once more calculated via the entropy.

## 4. Experiments

### 4.1. Model Architecture

We run the experiments for a SOTA wav2vec2.0 pre-trained model for SER as described in Wagner et al. [23] followed by a classification head. We use the first 11 transformer layers of the pre-trained ‘facebook/wav2vec2-large-robust’ model [24] followed by 2 dense layers with a  $\tanh$  activation function as classification head, which adds up to 153M parameters.

### 4.2. Datasets

We combine three datasets for categorical emotional speech, the two English datasets CRowd-sourced Emotional Multimodal Ac-

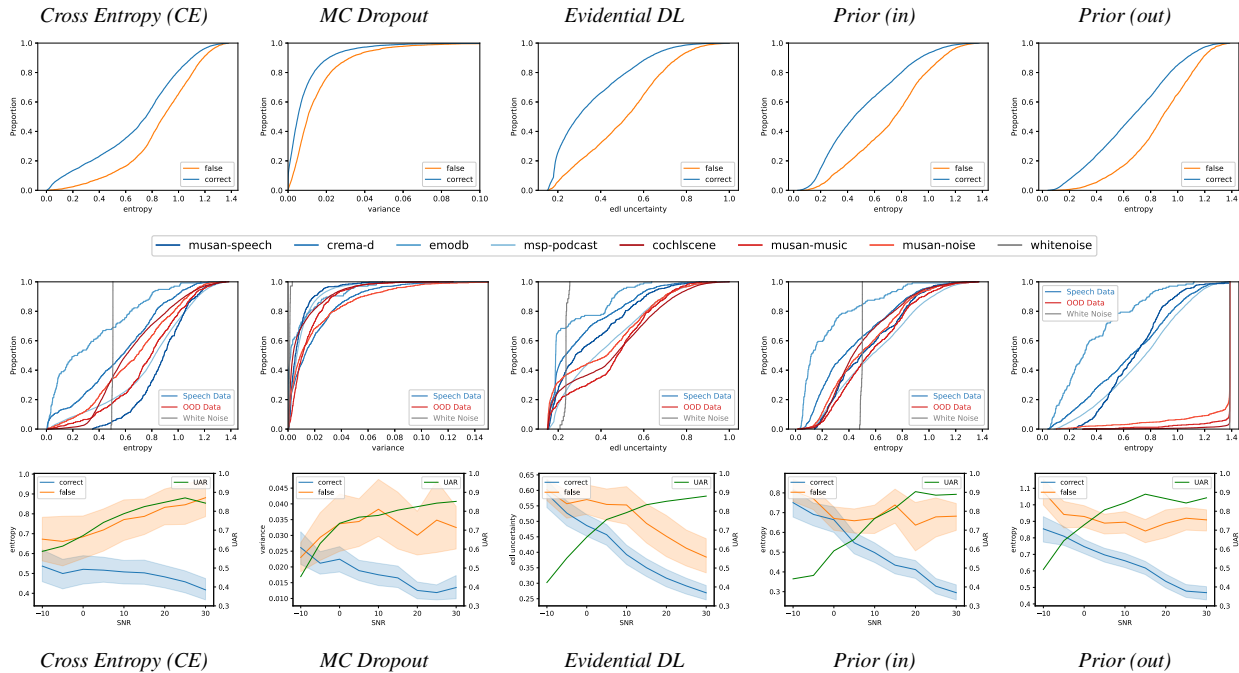


Figure 2: Overview of the performance of the 5 UQ methods on different tests. Higher values on the x-axis (and y-axis in the bottom row) represent higher uncertainty in all cases. The first row shows Empirical Cumulative Distribution Function (CDF) plots for correct and wrong predictions on the test data of all 3 SER datasets. The middle row shows CDF plots of the uncertainty on speech (blue) and non-speech datasets (red + grey). The bottom row shows the mean uncertainty for EmoDB test data augmented with noise for different Signal to Noise Ratio (SNR) levels with 95% confidence interval (left axis) and Unweighted Average Recall (UAR) (right axis) for different SNRs.

tors Dataset (CREMA-D) [25], MSP-Podcast [26] (displaying natural emotional expressions), and the Berlin Database of Emotional Speech (EmoDB) [27] for German acted emotional speech. All the corpora provide a rater agreement per utterance. To train a common model for all datasets we choose the subset of the “basic-4” emotions (anger, happiness, sadness, and neutral) for all datasets, as is also done in [28][23]. Excluding the samples with other emotions, the train set is around 45 h long and contains 28 945 utterances in total, divided into 60% neutral, 26% happiness, 8% anger, and 6% sadness. Beyond, this allows us to treat left out and less frequently occurring emotions as OOD data. To test the models when confronted with non-speech data set, we pick cochlscene [29] which contains over 200 h environmental sounds from 13 different acoustic scenes and the over 100 h long MUSIC, Speech, And Noise Corpus (MUSAN) [30].

### 4.3. Training Procedure

We train four different models to evaluate the different UQ methods without any further fine-tuning or adaptation<sup>1</sup>: The first model, trained with standard weighted cross-entropy loss is used for the two baselines entropy and MC Dropout UQ. For MC Dropout, we run 20 forward passes during inference. The second model is trained with EDL loss<sup>2</sup> in order to test the corresponding EDL UQ. Finally, we train two models for PNs [22] with the implementation of Rylan Schaeffer et al.<sup>3</sup>, the first one being only trained on the three SER datasets (PN(in)), while the second

one also includes cochlscene as OOD training data (PN(out)).

We train all models with a learning rate of 1e-4 and without hyperparameter tuning. It takes around 5 epochs and around 2 h on an Nvidia RTX A4000 GPU for the model to converge. We use the default, speaker-independent, train, development, and test splits of the datasets.

### 4.4. Correctness

The performances on the test sets are listed in Table 1. For MC Dropout, we use the mean over all forward passes. All models have comparable performance.

Table 1: Unweighted Average Recall (UAR) and Acc on test sets of  $E = \text{EmoDB}$ ,  $C = \text{CREMA-D}$ ,  $M = \text{MSP test 1}$ ,  $CE = \text{Cross Entropy}$ ,  $MC = \text{Monte Carlo Dropout}$ ,  $EDL = \text{Evidential Deep Learning}$ ,  $PN(\text{in and out})$  for Prior Networks trained on speech (in) and with additional OOD data (out)

		CE	MC	EDL	PN(in)	PN(out)
E	Acc	.875	.875	.875	.919	.904
	UAR	.843	.843	.843	.898	.880
C	Acc	.715	.728	.792	.828	.766
	UAR	.677	.704	.665	.745	.711
M	Acc	.635	.639	.662	.668	.691
	UAR	.609	.603	.576	.561	.558

In order to test, whether higher uncertainty correlates with a larger number of misclassifications, we visualise the uncertainty

<sup>1</sup><https://github.com/audeering/ser-uncertainty-quantification>

<sup>2</sup><https://github.com/dougbrion/pytorch-classification-uncertainty/>

<sup>3</sup><https://github.com/RylanSchaeffer/HarvardAM207-prior-networks>

discrepancy between correct and false predictions with the help of CDF plots, similar to [31] (see first row of Figure 2). Ideally, we would expect a considerably larger area under the curve (AUC) for the correct predictions versus the false predictions in the CDF graph. While to a certain extent true for all UQ, in particular, our baselines CE and MC show very small differences, thus apparently limiting their reliability to detect in-domain (ID) misclassifications.

#### 4.5. Correlation with rater agreement

We further determine the Pearson Correlation Coefficient (PCC) between the UQ predictions for all models and the utterance-level rater agreement over all three SER datasets. The latter can be calculated as the ratio percentage of raters that chose the majority class as their annotation (see Table 2).

Table 2: PCC: Correlation between rater agreement and uncertainty, lower is better

	CE	MC	EDL	PN(in)	PN(out)
EmoDB	-.236	-.165	-.153	-.186	-.259
Crema-d	-.120	-.174	-.265	-.217	-.238
MSP-1	-.128	-.082	-.119	-.144	-.194

Even though a weak negative correlation can be observed in all cases, indicating that lower rater agreement also leads to higher uncertainty, the effect is minimal. A reason for the weak correlation might be that the model only has to distinguish between 4 emotions, while the raters had a larger amount of emotions to discriminate. Beyond, even samples with very high annotator disagreement are presented with a “ground truth” emotion during training, thus potentially not sensitising the model for rater disagreements.

#### 4.6. Out-of-Distribution Emotions

Table 3 shows the mean uncertainty for emotions included during training vs the emotions omitted during training on the EmoDB test data. As expected, all methods have greater uncertainty for unseen emotions than for known ones.

Table 3: Mean Uncertainty for EmoDB emotions included in the training (anger, happiness, neutral, sadness) vs boredom, disgust, fear (out)

	CE	MC	EDL	PN(in)	PN(out)
in	.333	.010	.256	.239	.374
out	.850	.023	.387	.725	.813

#### 4.7. Out-of-Domain Data

In order to test how the UQ behave under non-speech OOD data, we compare the CDF plots of the SER datasets to those of the other audio databases cochlsce (acoustic ambient noise), musan (speech/music/noise), and additionally, artificially generated “white noise” (45 samples from near silence with gradually increasing Gaussian white noise) in the second row of Figure 2. As the latter type of data is completely unfit for the SER task, we expect lower AUC for their graphs (red/grey) compared to those of the speech datasets (blue). At first glance, it stands out that PN(out) has by far the largest separation between ID and OOD data, as the speech and noise datasets are clearly separated in

the CDF plot. A small separation can also be observed for EDL, while any of the other models’ UQs fail to produce a visible distinction. It is important to note at this point that the PN(out) model is the only one that has been trained on some OOD due to its design allowing for this in a straightforward manner. However, it has only seen OOD data in the form of ambient noise from cochlsce, and it manages to generalise well to other types of OOD data. Regarding the individual databases, EmoDB has almost always the lowest uncertainty, which aligns with models showing the highest performance on it.

#### 4.8. Sensor Uncertainty

Our final test aims to simulate sensor uncertainty by augmenting the EmoDB dataset with different levels of white noise, starting with an Signal to Noise Ratio (SNR) of 30 (original utterance is 30 dB louder than the noise) and increasing noise levels up to an SNR of -10. The last row of Figure 2 shows the mean uncertainty for correct and false predictions over the different SNRs. The plot also contains the UAR, which, as expected, gets progressively worse as the noise level increases. For an SNR of 30 dB, the uncertainties for the correct and false predictions are clearly separated for all UQ methods, even more prominently than in the top row, as we only evaluate on the ‘easiest’ dataset EmoDB.

For increasing noise levels, we would expect an increasing uncertainty for the correct predictions as more and more samples are correctly classified by accident, which we can observe to varying degrees for all methods. However, with (Cross) Entropy and MC Dropout, the uncertainty of the misclassified samples also decreases, which makes it impossible to reliably distinguish between false and correct predictions based on the UQ at this point. For EDL however, we see exactly the opposite behaviour, as all predictions become more uncertain with decreasing SNR, which is a promising behaviour in order to prevent faulty classifications based on UQ. Still, the fact that the uncertainty increases much quicker than UAR indicates that the UQ is too sensitive to noise. Interestingly, for EDL we see higher uncertainties for speech samples with added noise than for pure white noise in the middle row of Figure 2. The most ideal behaviour is shown by the two PNs with the uncertainty of false predictions initially remaining stable, but the uncertainty of correct predictions steadily increasing.

## 5. Conclusions and Outlook

We tested the behaviour of different UQ methods under realistic sources of uncertainty for SER applications, such as OOD emotions, OOD data or corrupted signals. We find that entropy-based UQ predictions on models with a standard cross-entropy training pipeline can already give good intuition on whether an emotion prediction can be considered trustworthy, thus reducing the importance of additional gatekeepers such as SNR models or voice activity detection systems. Nevertheless, the reliability of UQ can greatly be increased by incorporating the UQ already in the network design, in particular through Prior Networks and by exposing the model to OOD data during training. In this regard, further improvement might be expected through methods that artificially create the OOD data [32] or calibration [33]. However, a closer investigation about the separability of data uncertainty and OOD data through PNs and other UQs in the context of SER seems necessary.

## 6. Acknowledgements

Manuel Milling and Björn Schuller are with MCML – Munich Center for Machine Learning, Germany, Björn Schuller is with MDSI – Munich Data Science Institute, Germany

This research has been partly funded by the European SHIFT (MetamorphoSiS of cultural Heritage Into augmented hypermedia assets For enhanced accessibiliTy and inclusion) project (Grant Agreement number: 101060660).

## 7. References

- [1] B. W. Schuller, “Speech emotion recognition: Two decades in a nutshell, benchmarks, and ongoing trends,” *Communications of the ACM*, vol. 61, no. 5, pp. 90–99, 2018.
- [2] G. Sahu, “Multimodal speech emotion recognition and ambiguity resolution,” *arXiv e-prints*, pp. arXiv–1904, 2019.
- [3] M. Milling, A. Baird, K. D. Bartl-Pokorny, S. Liu, A. M. Alcorn, J. Shen, T. Tavassoli, E. Ainger, E. Pellicano, M. Pantic *et al.*, “Evaluating the impact of voice activity detection on speech emotion recognition for autistic children,” *Frontiers in Computer Science*, vol. 4, p. 837269, 2022.
- [4] G. Rizos and B. W. Schuller, “Average jane, where art thou?—recent avenues in efficient machine learning under subjectivity uncertainty,” in *Information Processing and Management of Uncertainty in Knowledge-Based Systems: 18th International Conference, IPMU 2020, Lisbon, Portugal, June 15–19, 2020, Proceedings, Part I 18*. Springer, 2020, pp. 42–55.
- [5] M. Gerczuk, S. Amiriparian, S. Ottl, and B. W. Schuller, “Emonet: A transfer learning framework for multi-corpus speech emotion recognition,” *IEEE Transactions on Affective Computing*, vol. 14, no. 2, pp. 1472–1487, 2021.
- [6] M. Jaiswal and E. Mower Provost, “Best practices for noise-based augmentation to improve the performance of emotion recognition in the wild,” *arXiv e-prints*, pp. arXiv–2104, 2021.
- [7] M. Abdar, F. Pourpanah, S. Hussain, D. Rezazadegan, L. Liu, M. Ghavamzadeh, P. Fieguth, X. Cao, A. Khosravi, U. R. Acharya *et al.*, “A review of uncertainty quantification in deep learning: Techniques, applications and challenges,” *Information Fusion*, vol. 76, pp. 243–297, 2021.
- [8] J. Gawlikowski, C. R. N. Tassi, M. Ali, J. Lee, M. Humt, J. Feng, A. Kruspe, R. Triebel, P. Jung, R. Roscher *et al.*, “A survey of uncertainty in deep neural networks,” *Artificial Intelligence Review*, vol. 56, no. Suppl 1, pp. 1513–1589, 2023.
- [9] E. Hüllermeier and W. Waegeman, “Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods,” *Machine Learning*, vol. 110, pp. 457–506, 2021.
- [10] A. R. Tan, S. Urata, S. Goldman, J. C. Dietschreit, and R. Gómez-Bombarelli, “Single-model uncertainty quantification in neural network potentials does not consistently outperform model ensembles,” *npj Computational Materials*, vol. 9, no. 1, p. 225, 2023.
- [11] J. Liu, Z. Lin, S. Padhy, D. Tran, T. Bedrax Weiss, and B. Lakshminarayanan, “Simple and principled uncertainty estimation with deterministic deep learning via distance awareness,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 7498–7512, 2020.
- [12] Z. Zhang, J. Deng, E. Marchi, and B. Schuller, “Active learning by label uncertainty for acoustic emotion recognition,” in *Proceedings INTERSPEECH 2013, 14th Annual Conference of the International Speech Communication Association, Lyon, France, 2013*.
- [13] J. Han, Z. Zhang, M. Schmitt, M. Pantic, and B. Schuller, “From hard to soft: Towards more human-like emotion recognition by modelling the perception uncertainty,” in *Proceedings of the 25th ACM international conference on Multimedia*, 2017, pp. 890–897.
- [14] N. Amir, R. Rubinstein, A. Shlomov, and G. Diamond, “Comparing categorical and dimensional ratings of emotional speech,” in *2015 11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*, vol. 5. IEEE, 2015, pp. 1–5.
- [15] C. E. Shannon, “A mathematical theory of communication,” *The Bell system technical journal*, vol. 27, no. 3, pp. 379–423, 1948.
- [16] Q. Kong, Y. Cao, T. Iqbal, Y. Wang, W. Wang, and M. D. Plumbley, “Panns: Large-scale pretrained audio neural networks for audio pattern recognition,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2880–2894, 2020.
- [17] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” *Advances in neural information processing systems*, vol. 33, pp. 12 449–12 460, 2020.
- [18] Y. Gal and Z. Ghahramani, “Dropout as a bayesian approximation: Representing model uncertainty in deep learning,” in *international conference on machine learning*. PMLR, 2016, pp. 1050–1059.
- [19] M. Sensoy, L. Kaplan, and M. Kandemir, “Evidential deep learning to quantify classification uncertainty,” *Advances in neural information processing systems*, vol. 31, 2018.
- [20] A. Amini, W. Schwarting, A. Soleimany, and D. Rus, “Deep evidential regression,” in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 14 927–14 937. [Online]. Available: <https://proceedings.neurips.cc/paper/2020/file/aab085461de182608ee9f607f3f7d18f-Paper.pdf>
- [21] L. Liu and R. R. Yager, “Classic works of the dempster-shafer theory of belief functions: An introduction,” *Classic works of the Dempster-Shafer theory of belief functions*, pp. 1–34, 2008.
- [22] A. Malinin and M. Gales, “Predictive uncertainty estimation via prior networks,” *Advances in neural information processing systems*, vol. 31, 2018.
- [23] J. Wagner, A. Triantafyllopoulos, H. Wierstorf, M. Schmitt, F. Burkhardt, F. Eyben, and B. W. Schuller, “Dawn of the transformer era in speech emotion recognition: closing the valence gap,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- [24] W.-N. Hsu, A. Sriram, A. Baevski, T. Likhomanenko, Q. Xu, V. Pratap, J. Kahn, A. Lee, R. Collobert, G. Synnaeve *et al.*, “Robust wav2vec 2.0: Analyzing domain shift in self-supervised pre-training,” *arXiv e-prints*, pp. arXiv–2104, 2021.
- [25] 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.
- [26] R. Lotfian and C. Busso, “Building naturalistic emotionally balanced speech corpus by retrieving emotional speech from existing podcast recordings,” *IEEE Transactions on Affective Computing*, vol. 10, no. 4, pp. 471–483, 2017.
- [27] F. Burkhardt, A. Paeschke, M. Rolfes, W. F. Sendlmeier, B. Weiss *et al.*, “A database of german emotional speech.” in *Interspeech*, vol. 5, 2005, pp. 1517–1520.
- [28] W. Wu, C. Zhang, X. Wu, and P. C. Woodland, “Estimating the uncertainty in emotion class labels with utterance-specific dirichlet priors,” *IEEE Transactions on Affective Computing*, 2022.
- [29] I.-Y. Jeong and J. Park, “Cochlscene: Acquisition of acoustic scene data using crowdsourcing,” in *2022 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2022, pp. 17–21.
- [30] D. Snyder, G. Chen, and D. Povey, “Musan: A music, speech, and noise corpus,” *arXiv e-prints*, pp. arXiv–1510, 2015.
- [31] M. Sensoy, L. Kaplan, F. Cerutti, and M. Saleki, “Uncertainty-aware deep classifiers using generative models,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 04, 2020, pp. 5620–5627.
- [32] J. Van Amersfoort, L. Smith, Y. W. Teh, and Y. Gal, “Uncertainty estimation using a single deep deterministic neural network,” in *International conference on machine learning*. PMLR, 2020, pp. 9690–9700.
- [33] M. P. Naeini, G. Cooper, and M. Hauskrecht, “Obtaining well calibrated probabilities using bayesian binning,” in *Proceedings of the AAAI conference on artificial intelligence*, vol. 29, no. 1, 2015.