



# The Emotion is Not One-hot Encoding: Learning with Grayscale Label for Emotion Recognition in Conversation

Joosung Lee

Kakao Enterprise Corp., South Korea

rung.joo@kakaointerprise.com

## Abstract

In emotion recognition in conversation (ERC), the emotion of the current utterance is predicted by considering the previous context, which can be utilized in many natural language processing tasks. Although multiple emotions can coexist in a given sentence, most previous approaches take the perspective of a classification task to predict only a given label. However, it is expensive and difficult to label the emotion of a sentence with confidence or multi-label. In this paper, we automatically construct a grayscale label considering the correlation between emotions and use it for learning. That is, instead of using a given label as a one-hot encoding, we construct a grayscale label by measuring scores for different emotions. We introduce several methods for constructing grayscale labels and confirm that each method improves the emotion recognition performance. Our method is simple, effective, and universally applicable to previous systems. The experiments show a significant improvement in the performance of baselines.

## 1. Introduction

As interest in interactive applications (e.g. chatbots) increases, emotion recognition in conversation becomes more important. Emotions are additional information that can better understand the speaker's state of conversation, and this is used to design an empathic dialogue system [1, 2, 3]. Emotion also helps provide personalized results, such as social media opinion mining [4] and recommendation systems [5].

In Emotion Recognition in Sentence (ERS) studies, a grayscale label is constructed because one utterance can have multiple emotions. [6] introduces a grayscale label using emotion lexicons such as NRC [7] and Emotinet [8]. [9] introduces grayscale labels with pre-trained word embeddings. However, since these methods do not consider utterance, there is a limitation that one-hot encoding is mapped to one grayscale label. [10] trains the model with grayscale labels through the basic predictor. However, it depends on the basic predictor with poor performance and cannot handle the noise.

In emotion recognition in conversation (ERC), an utterance has a more sensitive distribution of emotions than in ERS because emotions must be recognized in consideration of both context and utterance. However, most ERC datasets are labeled for only one emotion, and the latest models are trained with labels of one-hot encoding. That is, when the utterance "yeah, i do" contains both *joy* and *neutral* emotions depending on the context, learning with only one emotion can cause an error. Our paper focuses on solving the limitation that a model is trained with only one emotion, but it is very expensive for humans to re-label multiple emotions in every utterance. Therefore, we construct grayscale labels in several automatic ways. We propose the following methods for constructing grayscale labels: 1) **Category** 2) **Word-Embedding** 3) **Self** 4) **Self-Adjust** 5) **Future-**

**Self.** "Category" and "Word-embedding" are mapping one-hot encodings to grayscale label (i.e. soft-label encoding), which means that training samples with the same emotion are always mapped to the same grayscale label. "Self"-type methods utilize a pre-trained self-teacher-model to construct grayscale according to utterances, which are mapped to different grayscale labels regardless of ground-truth emotion. To the best of our knowledge, our work is the first attempt to create multiple grayscale labels in ERC.

We construct grayscale labels to four ERC datasets: IEMO-CAP, dailydialog, MELD, and EmoryNLP. Experimental results show that RoBERTa-large [11] achieves competitive performance with only grayscale labels, which is simple but effective. We also show that grayscale labels can be used in the latest approaches, Psychological [12], CoG-BART [13], DAG-ERC [14] and CoMPM [15].

## 2. Proposed Approach

### 2.1. Overview

Figure 1 is an overview of our approach to recognizing the emotion of the  $k$ -th utterance  $u_k$ . First, we construct grayscale labels on our training samples in various ways. Then, the model is trained on both the one-hot encoding and the grayscale label. The input is concatenated with the previous context and the current utterance, and special tokens indicate the speaker by prepending each utterance. Also, we prepend the  $[CLS]$  token to the input in RoBERTa. The model is trained to predict emotions with the  $[CLS]$  token.

### 2.2. Construction of Grayscale Label

We propose 5 ways to automatically construct grayscale labels. The "category" method constructs a grayscale label heuristically. The "word-embedding" method constructs a grayscale label using word-embedding. Other methods (i.e. self-methods) construct grayscale using logit of the self-model (i.e. teacher-model).

#### 2.2.1. Category

[6] generates grayscale labels as a strategy based on linguistic resources, which uses words from each sentence and ground-truth together. Inspired by this, using only the ground-truth, we simply divide emotions into sentiment categories and score emotions differently for each category. Categories are shown in Table 2. The scores ( $s_i$ ) of emotions ( $e_i$ ) are as follows:

$$s_i = \begin{cases} 1 & \text{if } e_i = e_{gt} \\ 0.5 & \text{if } \text{category}(e_i) = \text{category}(e_{gt}) \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

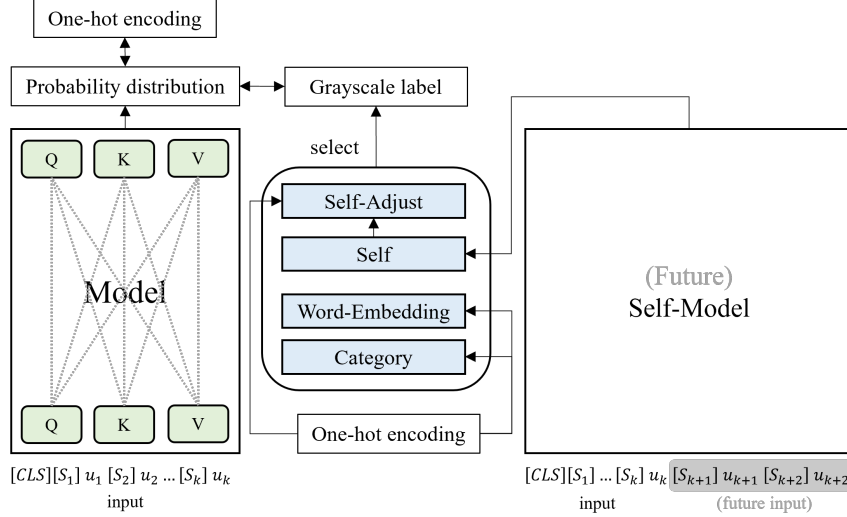


Figure 1: An overview of our framework based on RoBERTa. For each input, a grayscale label is constructed, and the model is trained so that the distance between the probability distribution of emotion and the distribution of the grayscale label is close.

where  $e_{gt}$  is the the ground-truth emotion. In the category method, emotions included in the same sentiment give positive scores to each other. We calculate the grayscale label  $\mathbf{g} = \{g_1, g_2, \dots, g_k\}$  ( $k$  = the number of class) by normalizing  $s$  as follows:

$$g_i = \frac{s_i}{\sum_{j=1}^k s_j} \quad (2)$$

### 2.2.2. Word-Embedding

Inspired by [9], we calculate the similarity between emotions using FastText [16], a publicly released word-embedding<sup>1</sup>. Word-embedding refers to the latent representation of a word. Therefore, the similarity between words with similar meanings is high, and the similarity between words with different meanings is low. The score for each emotion is calculated using cosine similarity as follows:

$$s_i = \max\left(\frac{w_i \cdot w_{gt}}{\|w_i\| \|w_{gt}\|}, 0\right) \quad (3)$$

where  $w_{gt}$  is the embedding vector of the ground-truth emotion.

We construct the grayscale label  $\mathbf{g}$  in two steps. First, if the similarity score is negative, it is changed to 0. Because a negative cosine similarity value indicates a less relevant emotion, the  $s_i$  of the unrelated emotion is changed to 0 before linear normalizing. Then, the grayscale label is calculated through normalizing as in Equation 2. The spelling for each emotion is explained in Table 2.

### 2.2.3. Self

The "category" and "word-embedding" methods have a limitation in that grayscale labels are calculated only in the relation to emotions without considering utterances. Also, the word-embedding method has a disadvantage in that the embedding can be similar even if the meaning of the word is different (e.g *night* and *day*). So we propose the self-method, which is similar to [10] in ERS. The self-grayscale label is constructed from the

<sup>1</sup>We used FastText, but other public word-embeddings are fine. A more advanced method is to use word-embedding learned from data such as emotion lexicon.

emotion of the utterance predicted by the self-model (fine-tuned model), which is a kind of distillation [17]. Since distillation is not our goal, we use a self-model instead of exploring a new teacher-model.

First, the teacher-model is trained as a classification task using one-hot encoding. The logit of the pre-trained teacher-model becomes the score, and then the grayscale label is calculated through the softmax as follows:

$$g_i = \frac{e^{s_i}}{\sum_{j=1}^k e^{s_j}} \quad (4)$$

When the student-model is trained, the parameters of the teacher-model are fixed, where the student-model is the final model trained with ERC data (RoBERTa). Each training sample has a different grayscale label according to the utterance regardless of the ground-truth emotion.

### 2.2.4. Self-Adjust

The self-method depends on the performance of the self-model. So, for example, if the one-hot encoding of the ground-truth is (1, 0, 0, 0) and the self-grayscale label is (0.3, 0.4, 0.2, 0.1), there is a risk that the model is confused about the best emotion. To alleviate this problem, we propose a self-adjust-method to adjust the self-grayscale label. Figure 2 is an example of the self-adjust-method.

If ground-truth emotion is different from the maximum probability of the self-grayscale label, we adjust as follows: 1) The self-adjust-grayscale value corresponding to the gold emotion is adjusted to 0.5. 2) Other emotions have a value divided by 0.5 as much as the distribution of self-grayscale labels. The adjusting function is calculated as follows:

$$g'_i = \begin{cases} 0.5 & \text{if } i = gt \\ \frac{0.5g_i}{1-g_{gt}} & \text{otherwise} \end{cases} \quad (5)$$

where  $gt$  is the index corresponding to the ground-truth emotion, and  $g'_i$  is the adjusted grayscale label.

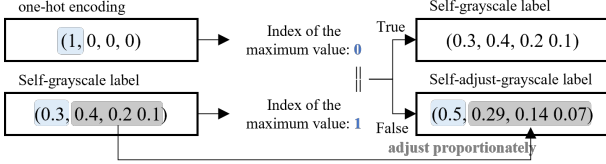


Figure 2: The example of constructing a self-adjust-grayscale label.

Dataset	dialogues			utterance		
	train	dev	test	train	dev	test
IEMOCAP	108	12	31	5163	647	1623
DailyDialog	11118	1000	1000	87170	8069	7740
MELD	1038	114	280	9989	1109	2610
EmoryNLP	713	99	85	9934	1344	1328

Table 1: Statistics for 4 datasets.

### 2.2.5. Future-Self

We improve the performance with an additional strategy when a self-model is trained. In ERC research, it is common to predict the emotion of an utterance by considering only the past context. If the model captures information in a future context, it cannot predict emotions in real-time during inference. A non-real-time emotion recognition system has the disadvantage of being difficult to be utilized in other natural language processing tasks. However, it can be helpful in the training phase because the future context has information about the current emotion that affects the listener’s response.

Therefore, we propose a method to improve the performance of the self-model by using the future context as input. The input is concatenated with the utterances of two future turns, and this model is called the future-self-model.

### 2.3. Loss

We train the model on both a one-hot encoding and a grayscale label. The one-hot encoding is used to predict the top one of emotions, and the grayscale label is used to predict the probability distribution reflecting the distance between emotions. Both tasks use cross-entropy loss:

$$L_b = -\frac{1}{N} \sum_{n=1}^N \sum_{i=1}^k o_i \log(p_i) \quad (6)$$

$$L_g = -\frac{1}{N} \sum_{n=1}^N \sum_{i=1}^k g_i \log(p_i) \quad (7)$$

$$L = L_b + \alpha L_g \quad (8)$$

where  $N$  is the number of training data,  $k$  is the number of emotion classes,  $o_i$  is a one-hot encoding,  $g_i$  is a grayscale label, and  $p_i$  is the emotion probability predicted by the model.  $\alpha$  is described in Section 3.4.

## 3. Experiments

In the “category” and “word-embedding” methods, the model is jointly trained with a one-hot encoding and a grayscale label as Equation 8. In the self-methods, first, the teacher-model is trained as a classification task using one-hot encoding. Then,

Dataset	category		
	positive	negative	neutral
IEMOCAP	excited, happy	angry, frustrated, sad	neutral
DailyDialog	happy	anger, disgust, fear, sad	neutral, surprise
MELD	joy	anger, disgust, fear, sad	neutral, surprise
EmoryNLP	joy, peaceful, powerful	mad, sad, scared	neutral

Table 2: The sentiment category of the emotion corresponding to each dataset

the parameters of the teacher-model are frozen and the student-model is jointly trained using one-hot encoding and grayscale labels as Equation 8.

### 3.1. Training Setup

We use the pre-trained model from the huggingface library<sup>2</sup>. The optimizer is AdamW and the learning rate is 1e-6 as an initial value. The learning rate scheduler used for training is *get\_linear\_schedule\_with\_warmup*, and the maximum value of 10 is used for the gradient clipping. We select the model with the best performance on the validation set. All experiments are conducted on one A100 GPU.

### 3.2. Dataset and Evaluation

We experiment on four datasets: IEMOCAP [18], DailyDialog [19], MELD [20], and EmoryNLP [21]. Table 1 shows the statistics of the data. DailyDialog uses 7 classes for training, but we measure Macro-F1 for only 6 classes excluding neutral. Other datasets are evaluated with weighted average F1.

In IEMOCAP, the emotional inventory (6) is given as “happy, sad, angry, excited, frustrated and neutral”. In DailyDialog, the emotional inventory (7) is given as “anger, disgust, fear, joy, surprise, sadness and neutral”. In MELD, the emotional inventory (7) is given as “anger, disgust, sadness, joy, surprise, fear and neutrality”. In EmoryNLP, the emotional inventory (7) is given as “joyful, peaceful, powerful, scared, mad, sad and neutral”. Table 2 shows the sentiment category of each emotion, and the corresponding word-embedding is used.

### 3.3. Results and Discussion

Table 3 shows our experimental results. First, our proposed RoBERTa has slightly improved performance by using special tokens compared to the RoBERTa proposed in previous studies. RoBERTa trained with grayscale labels generally improves performance. However, the performance doesn’t improve according to the combination of data and grayscale type (i.e. EmoryNLP+(W or S)). In particular, +S negatively affects RoBERTa in EmoryNLP. We infer that +S contains a lot of noise due to the insufficient performance of the self-model in EmoryNLP. +SA proposed to alleviate the noise problem is effective in improving the performance. In RoBERTa, we observe that self-methods are generally superior to +C and +W. +C and +W are very simple and effective, but there is a limit that does not consider utterance. However, the self-methods considering dialogue further improve performance because they are more fine-grained grayscale labels. Because future-self-RoBERTa is inferior to original RoBERTa in IEMOCAP, RoBERTa+FSA was not tested. In other datasets, +FSA outperforms +SA because the future-self-model constructs fine-grained grayscale labels more than the self-model.

We apply +S and +SA, which were effective in RoBERTa,

<sup>2</sup><https://github.com/huggingface/transformers>

Models	IEMOCAP	DailyDialog		MELD	EmoryNLP
	W-Avg F1	Macro F1	Micro F1	W-Avg F1	W-Avg F1
RoBERTa (ours)	63.17	51.17	58.63	64.79	36.3
+C (category)	64.83	53.44	58.91	65.27	36.97
+W (word-embedding)	64.85	51.77	60.15	65.48	36.05
+S (self)	64.78	53.89	60.67	65.71	29.09
+SA (self-adjust)	<b>65.85</b>	53.75	59.73	66	37.69
+FSA (future-self-adjust)	-	<b>55.84</b>	<b>61.67</b>	<b>66.49</b>	<b>38</b>
Psychological [12]	63.37	52.38	60.54	64.38	37.32
+S	<b>65.43</b>	53.63	59.98	<b>65.24</b>	<b>38.96</b>
+SA	64.92	<b>55</b>	<b>60.79</b>	65.04	38.33
DAG-ERC [14]	68.03	-	59.25	<b>63.63</b>	39.02
+S	<b>68.57</b>	-	<b>59.4</b>	63.55	<b>40.23</b>
+SA	68.43	-	59.33	63.57	39.85
CoG-BART [13]	63.8	53.99	55.52	64.77	36.48
+S	64.55	<b>57.14</b>	<b>57.15</b>	64.26	36.97
+SA	<b>64.77</b>	56.07	56.57	<b>65.28</b>	<b>37.42</b>
CoMPM [15]	66.33	52.46	60.41	65.53	38.56
+S	67.5	<b>52.88</b>	<b>61.1</b>	66.19	39.05
+SA	<b>67.52</b>	52.59	60.66	<b>66.5</b>	<b>39.95</b>

Table 3: Our results are the average of three runs. The performance of the comparison system is the result of our re-experiment with the published code. Bold text indicates the best performance in each part.

to other models. In other models, unlike RoBERTa, +FSA is not used because future inputs cannot simply be combined with the models. +S and +SA give a positive signal to the performance of the other models. In addition, unlike RoBERTa, +S also improves model performance on EmoryNLP. Analysis of this should be explored, but we assume that (teacher-) comparative systems learned about EmoryNLP do not have false biases, unlike RoBERTa. Therefore, the proposed grayscale labels can improve the average performance regardless of the model structure.

Instead of focusing on modeling, our research focuses on constructing a grayscale label with the distribution of emotions contained in the utterance. We achieve high performance competitive with the state-of-the-art models just by combining grayscale with RoBERTa. We also show performance improvements by combining grayscale with the original state-of-the-art models. In other words, grayscale labels are easy to effectively combine with other models. However, since grayscale labels have different effects depending on the combination of model and data, it is difficult to select the optimal grayscale label and  $\alpha$  for the model through several experiments. Exploring  $\alpha$  is shown in the next section.

### 3.4. Performance According to $\alpha$

This section demonstrates that the optimal  $\alpha$  is determined by a combination of model structure, data, and grayscale type. When RoBERTa is trained with MELD, the performance according to  $\alpha$  for the most effective +FSA and least effective +C is measured and the results are shown in Figure 3. Since +FSA and +C have the largest difference in effect, it is considered intuitive to compare these two methods.

As  $\alpha$  increases, RoBERTa is trained with more emphasis on the distribution of grayscale labels.  $\alpha$  up to the threshold improves the performance, but when it becomes larger than the threshold, the performance decreases. The threshold depends on the type of data and grayscale. We focus on the effect of grayscale labels, not to find the optimal  $\alpha$  for each combination. Therefore,  $\alpha$  is used as a fixed value (=1) in our experiments.

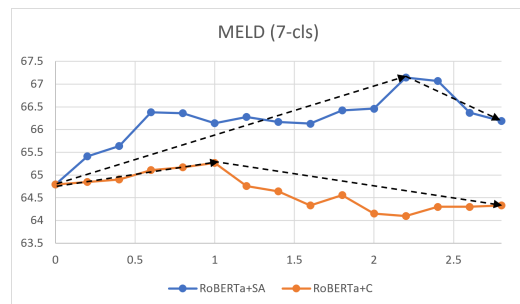


Figure 3: Performance of RoBERTa+SA and RoBERTa+C according to  $\alpha$  in MELD.

We confirm that  $\alpha$  shows similar effects if it is not too small, which is verified through multiple runs. Our future work is to find the optimal  $\alpha$  according to the learning environment.

## 4. Conclusion

We introduce a novel approach in ERC, which automatically constructs different types of grayscale labels taking into account correlations between emotions. The "category" and "word-embedding" methods construct a grayscale label based on the ground-truth emotion, while the other methods construct a grayscale label based on the dialogue. However, the self-model reflecting the dialogue may have incorrect information with ground-truth because it leverages the pre-trained model to generate grayscale labels. Therefore we introduce enhanced methods such as SA and FSA. We show competitive performance by learning RoBERTa as a grayscale label. It can also improve the performance of state-of-the-art models.

The different types of grayscale labels we proposed have different effects depending on the data and model structure. Also, finding the optimal  $\alpha$  in the loss Equation 8 should be determined through several experiments. We will explore further in this related future study.

## 5. References

- [1] S. Poria, N. Majumder, R. Mihalcea, and E. Hovy, "Emotion recognition in conversation: Research challenges, datasets, and recent advances," *IEEE Access*, vol. 7, pp. 100 943–100 953, 2019.
- [2] B. Li, Z. Wang, H. Liu, Y. Jiang, Q. Du, T. Xiao, H. Wang, and J. Zhu, "Shallow-to-deep training for neural machine translation," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Online: Association for Computational Linguistics, Nov. 2020, pp. 995–1005. [Online]. Available: <https://aclanthology.org/2020.emnlp-main.72>
- [3] Z. Lin, P. Xu, G. I. Winata, F. B. Siddique, Z. Liu, J. Shin, and P. Fung, "Caire: An end-to-end empathetic chatbot," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 09, pp. 13 622–13 623, Apr. 2020. [Online]. Available: <https://ojs.aaai.org/index.php/AAAI/article/view/7098>
- [4] A. Chatterjee, K. N. Narahari, M. Joshi, and P. Agrawal, "SemEval-2019 task 3: EmoContext contextual emotion detection in text," in *Proceedings of the 13th International Workshop on Semantic Evaluation*. Minneapolis, Minnesota, USA: Association for Computational Linguistics, Jun. 2019, pp. 39–48. [Online]. Available: <https://aclanthology.org/S19-2005>
- [5] H. Costa and L. Macedo, "Emotion-based recommender system for overcoming the problem of information overload," in *Highlights on Practical Applications of Agents and Multi-Agent Systems*, J. M. Corchado, J. Bajo, J. Kozlak, P. Pawlewski, J. M. Molina, V. Julian, R. A. Silveira, R. Unland, and S. Giroux, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 178–189.
- [6] Y. Zhang, J. Fu, D. She, Y. Zhang, S. Wang, and J. Yang, "Text emotion distribution learning via multi-task convolutional neural network," in *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*. International Joint Conferences on Artificial Intelligence Organization, 7 2018, pp. 4595–4601. [Online]. Available: <https://doi.org/10.24963/ijcai.2018/639>
- [7] S. M. Mohammad and P. D. Turney, "Nrc emotion lexicon," *National Research Council, Canada*, vol. 2, 2013.
- [8] S. Poria, A. Gelbukh, E. Cambria, A. Hussain, and G.-B. Huang, "Emosentencespace: A novel framework for affective commonsense reasoning," *Knowledge-Based Systems*, vol. 69, pp. 108–123, 2014.
- [9] Z. Li, H. Xie, G. Cheng, and Q. Li, "Word-level emotion distribution with two schemas for short text emotion classification," *Knowledge-Based Systems*, vol. 227, p. 107163, 2021.
- [10] B. Guo, S. Han, X. Han, H. Huang, and T. Lu, "Label confusion learning to enhance text classification models," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 14, pp. 12 929–12 936, May 2021. [Online]. Available: <https://ojs.aaai.org/index.php/AAAI/article/view/17529>
- [11] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "Roberta: A robustly optimized BERT pretraining approach," *CoRR*, vol. abs/1907.11692, 2019. [Online]. Available: <http://arxiv.org/abs/1907.11692>
- [12] J. Li, Z. Lin, P. Fu, and W. Wang, "Past, present, and future: Conversational emotion recognition through structural modeling of psychological knowledge," in *Findings of the Association for Computational Linguistics: EMNLP 2021*. Punta Cana, Dominican Republic: Association for Computational Linguistics, Nov. 2021, pp. 1204–1214. [Online]. Available: <https://aclanthology.org/2021.findings-emnlp.104>
- [13] S. Li, H. Yan, and X. Qiu, "Contrast and generation make bart a good dialogue emotion recognizer," *Proceedings of the AAAI Conference on Artificial Intelligence*, 2022.
- [14] W. Shen, S. Wu, Y. Yang, and X. Quan, "Directed acyclic graph network for conversational emotion recognition," in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Online: Association for Computational Linguistics, Aug. 2021, pp. 1551–1560. [Online]. Available: <https://aclanthology.org/2021.acl-long.123>
- [15] J. Lee and W. Lee, "Compm: Context modeling with speaker's pre-trained memory tracking for emotion recognition in conversation," in *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2022.
- [16] T. Mikolov, E. Grave, P. Bojanowski, C. Puhersch, and A. Joulin, "Advances in pre-training distributed word representations," in *Proceedings of the International Conference on Language Resources and Evaluation (LREC 2018)*, 2018.
- [17] J. Gou, B. Yu, S. J. Maybank, and D. Tao, "Knowledge distillation: A survey," *International Journal of Computer Vision*, vol. 129, no. 6, pp. 1789–1819, Jun 2021. [Online]. Available: <https://doi.org/10.1007/s11263-021-01453-z>
- [18] 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," *Lang. Resour. Evaluation*, vol. 42, no. 4, pp. 335–359, 2008. [Online]. Available: <http://dblp.uni-trier.de/db/journals/lre/lre42.html>
- [19] Y. Li, H. Su, X. Shen, W. Li, Z. Cao, and S. Niu, "DailyDialog: A manually labelled multi-turn dialogue dataset," in *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Taipei, Taiwan: Asian Federation of Natural Language Processing, Nov. 2017, pp. 986–995. [Online]. Available: <https://www.aclweb.org/anthology/I17-1099>
- [20] S. Poria, D. Hazarika, N. Majumder, G. Naik, E. Cambria, and R. Mihalcea, "MELD: A multimodal multi-party dataset for emotion recognition in conversations," in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy: Association for Computational Linguistics, Jul. 2019, pp. 527–536. [Online]. Available: <https://www.aclweb.org/anthology/P19-1050>
- [21] S. M. Zahiri and J. D. Choi, "Emotion detection on TV show transcripts with sequence-based convolutional neural networks," in *The Workshops of the The Thirty-Second AAAI Conference on Artificial Intelligence, New Orleans, Louisiana, USA, February 2-7, 2018*, ser. AAAI Workshops, vol. WS-18. AAAI Press, 2018, pp. 44–52. [Online]. Available: <https://aaai.org/ocs/index.php/WS/AAAIW18/paper/view/16434>