



SememeASR: Boosting Performance of End-to-End Speech Recognition against Domain and Long-Tailed Data Shift with Sememe Semantic Knowledge

Jiaxu Zhu¹, Changhe Song^{1,2,†}, Zhiyong Wu^{1,2,3,*}, Helen Meng³

¹Shenzhen International Graduate School, Tsinghua University, Shenzhen, China

²Peng Cheng Lab, Shenzhen, China

³The Chinese University of Hong Kong, Hong Kong SAR, China

{zhu-jx21, sch19}@mails.tsinghua.edu.cn, zywu@sz.tsinghua.edu.cn, hmmeng@se.cuhk.edu.hk

Abstract

Recently, excellent progress has been made in speech recognition. However, pure data-driven approaches have struggled to solve the problem in domain-mismatch and long-tailed data. Considering that knowledge-driven approaches can help data-driven approaches alleviate their flaws, we introduce sememe-based semantic knowledge information to speech recognition (SememeASR). Sememe, according to the linguistic definition, is the minimum semantic unit in a language and is able to represent the implicit semantic information behind each word very well. Our experiments show that the introduction of sememe information can improve the effectiveness of speech recognition. In addition, our further experiments show that sememe knowledge can improve the model's recognition of long-tailed data and enhance the model's domain generalization ability.

Index Terms: speech recognition, sememe, long-tailed problem, domain generalization

1. Introduction

Automatic Speech Recognition (ASR) is a technology that converts audio into text. In recent years, end-to-end (E2E) ASR has attracted a lot of attention and has made great progress. E2E ASR can convert audio to text using a single network model, greatly simplifying the training and inference process. There are three main types of E2E ASR models: connectionist temporal classification (CTC) [1], recurrent neural network transducer (RNN-T) [2, 3], and attention based encoder-decoder (AED) [4, 5, 6]. E2E ASR models achieve excellent results by leveraging large amounts of training data, which is the so-called pure data-driven approach.

However, pure data-driven approaches suffer poor recognition of long-tailed data and poor domain generalization due to the performance depends entirely on the training data, even though they are extremely characteristic and unevenly distributed. As poor recognition of long-tailed data and weak domain generalization are brought about by the training data itself, introducing external knowledge information can help alleviate this problem. The semantic knowledge information implied behind textual data has become a hot topic of research. As shown in Figure 1, sememe is defined as the minimum semantic unit of languages in linguistics. Sememe knowledge has been widely studied in the field of natural language processing (NLP) [7, 8, 9, 10]. Compared to rough text data, sememe knowledge is more accurate and fundamental and has been refined by professional scholars over many years. The sememe knowledge, capable of representing the semantic information of any word, is more stable and robust and is not affected by the data. A knowledge-driven approach based on sememe knowledge can




Figure 1: Examples of words and sememes. The red blocks are the words, and the green blocks are the set of sememes corresponding to the words above. The next word can be predicted more accurately by combining sememe information.

alleviate the problems of data-driven approaches and improve the effectiveness of models in long-tailed data problems.

As far as we know, we are the first to introduce sememe into ASR. In this paper, we propose a sememe-based semantically enhanced ASR model called SememeASR, which improves the semantic capability of the model by improving the semantic information of the text representation and adding sememe prediction tasks. Inspired by [7, 10], while the traditional data-driven approach is not able to fully exploit the rich semantic information of text, adding the information of the sememe can enrich the semantic representation of text. Our experiments have shown that different methods of adding sememe information to the model can improve the recognition ability and enhances the model's ability to recognize long-tailed data and somewhat enhances the model's domain generalization capability.

2. Methodology

In this section, we review the architecture of the baseline hybrid CTC/AED model in Section 2.1. Then our proposed method will be described in Section 2.2, which aims to apply sememe-based semantic knowledge to improve the ability of the ASR model, thereby improving the recognition ability of long tail data and enhancing the domain generalization ability.

2.1. Hybrid CTC/AED ASR Model

The baseline E2E ASR model we choose in our experiment is similar to the one presented in WeNet [11], which uses both CTC and Attention-based Encoder-Decoder (AED) loss during training to speed convergence and is also a relatively good one among a series of state-of-the-art approaches [12, 13]. As depicted in Figure 2, the hybrid CTC/AED ASR model mainly contains three parts, a *Shared Encoder*, a *CTC Decoder*, and an *Attention Decoder*. The *Shared Encoder* consists of a convolution subsampling layer containing two convolutional layers with stride 2 for downsampling, a linear projection layer, and a

† Equal contribution. * Corresponding author.

positional encoding layer, followed by multiple Conformer [14] encoder layers. The *CTC Decoder* consists of a linear layer and a log softmax layer. The *Attention Decoder* consists of a positional encoding layer, multiple Transformer [15] decoder layers, and a linear projection layer.




Figure 2: Architecture of Hybrid CTC/AED ASR model.

2.2. Proposed Model Architecture

The main idea of our paper is to introduce sememe-based semantic information into ASR simply and effectively. Inspired by [10], we use three simple but effective strategies: 1) adding sememe prediction auxiliary task with sememe loss in a multi-task learning manner, 2) simply adding sememe information to the textual representation, or 3) employing *Sememe Encoder* to improve the semantic representation ability of the text.

2.2.1. Sememe Prediction Task

As depicted in Figure 3, we add the *Sememe Prediction* task with sememe loss after the *Attention Decoder*. A multilabel classification task can be used to build the sememe prediction task, which seeks to predict sememes for the following token. Understanding semantics is closely related to predicting the following token’s sememes, which is frequently easier to learn than directly modeling the likelihood of the next token. Given current contextualized representation \mathbf{g} from Transformer in *Attention Decoder*, we estimate the probability of sememe s associated with next token t as showed in Equation 1:

$$p(t, s) = \sigma(W\mathbf{g} + b) \quad (1)$$

where W and b are the weight and bias associated with sememe s , σ is the sigmoid activation function. We have named the model that uses this approach **SememeASR-SP**.




Figure 3: Architecture of SememeASR-SP.

2.2.2. Semantically Enhanced Text Representation

Traditional data-driven training methods in ASR have difficulty in mining the rich semantic information in the text. Some approaches use BERT [16] to provide semantic information, which helps to increase the semantic information of ASR, but this semantic information is still relatively shallow and poorly interpretable. The introduction of sememe-based semantic information not only provides rich semantic information but also has strong interpretability. As depicted in Figure 4, one simple method to add the sememe information is averaging the corresponding sememe of each token to get the sememe embedding corresponding to the text, and then adding it to the text embedding. We denote original text embedding as $E = (e_1, e_2, \dots, e_i, \dots, e_I)$, sememe embedding as $C =$

$(c_1, c_2, \dots, c_i, \dots, c_I)$ and final semantically enhanced text embedding as $\hat{E} = (\hat{e}_1, \hat{e}_2, \dots, \hat{e}_i, \dots, \hat{e}_I)$, where I is the number of tokens in the text sequence. Formally, we have:

$$c_i = \frac{1}{n_t} \sum_{s \in S(t)} q_s \quad (2)$$

$$\hat{e}_i = c_i + e_i \quad (3)$$

where t represents the i -th token of text sequence, $S(t)$ refers to the sememe set associated with token t , n_t is the number of sememe entries of token t , q_s refers to the embedding of the sememe s . And c_i is formed by averaging the corresponding embeddings of all sememes of token t as shown in Equation 2. The sememe enhanced token embedding \hat{e}_i is thus derived by adding c_i and e_i as showed in Equation 3. Then the semantically enhanced text representation \hat{E} is directly fed into the *Attention Decoder*. We have named the model that uses this approach **SememeASR-SE**.




Figure 4: Architecture of SememeASR-SE.

2.2.3. Sememe Encoder

To make full use of semantic information, we explore better ways to incorporate sememe information and improve text representation. Inspired by [17], a bottom-up and top-down network structure is used to compose the image from high to low resolution to extract stronger semantic features, and the top-down process then raises the resolution to enhance the designated features. Similarly, we use stacked linear layers to achieve a similar effect of the bottom-up and top-down network structure, which we refer to as *Sememe Encoder* as shown in Figure 5.




Figure 5: Architecture of SememeASR-SEP.

We reduce the dimension of the concatenated text and sememe representation to extract stronger semantic features, and then increase the dimension to strengthen the formed features. The dimension changes are shown in Figure 6. In our experiment, the dimension of text embedding and sememe embedding is 256. Dimension after embedding concatenation will be 512 and the *Sememe Encoder* will transform the dimension to 256 which matches the dimension of *Attention Decoder*. We name the model assigned to this method as **SememeASR-SEP**.

2.2.4. SememeASR Loss Function

Just like the WeNet [11], we also combine the CTC and AED losses during training to speed convergence. Furthermore,




Figure 6: Illustration of the dimension change in Sememe Encoder that consists of multiple linear layers.

we calculate the binary cross-entropy loss of sememe prediction as sememe loss L_{SE} to enhance the model’s modeling of sememe-based semantic information. Equation 4 defines the SememeASR objective, L_{CTC}, L_{AED} are the CTC and AED losses respectively, $\lambda \in (0, 1)$ and $\alpha \in (0, 1)$ are hyper-parameters where λ balances the importance of CTC and AED loss while α balances the importance of AED loss and Sememe loss. Finally, the training loss can be represented as follows:

$$Loss = \lambda L_{CTC} + (1 - \lambda)[\alpha L_{AED} + (1 - \alpha)L_{SE}] \quad (4)$$

3. Experiments

3.1. Datasets

In this paper, we train our proposed E2E ASR on public Mandarin datasets Aishell-1 [18]. The Aishell-1 corpus consists of 178 hours of labeled speech collected from 400 speakers. The content of the datasets covers 5 domains including Finance, Science and Technology, Sports, Entertainments, and News. To compare the domain adaptation ability of ASR in the text domain while minimizing the influence of differences in the acoustic environment, we test the trained model on another public Mandarin dataset Aishell-2 [19] that has a similar acoustic environment for sound recording but the corresponding text contents cover different text domains. The Aishell-2 corpus consists of 1000 hours of labeled speech collected from 1991 speakers. The content of Aishell-2 corresponds to the domains of voice commands, digital sequence, places of interest, entertainment, finance, technology, sports, English spellings, and free speaking without specific topics. To further illustrate the robustness of our method, we also conduct evaluation experiment on different domains with WenetSpeech [20], which is a multi-domain Mandarin corpus consisting of high-quality labeled speech but a relatively more complex acoustic environment than Aishell-1.

3.2. Experimental Setup

For all experiments, we use the open-source WeNet toolkit [11] to build both the hybrid CTC/attention baseline and our proposed SememeASR. And we used the default values in the WeNet for the main parameters which have been validated by the WeNet contributor. The input features are 80-dimensional log Mel-filterbank (Fbank) computed on 25ms window with 10ms shift. We use SpecAugment [21] and speed perturb for data augmentation. We choose 4233 characters (including ⟨blank⟩, ⟨unk⟩, ⟨sos/eos⟩ labels) as model units for Aishell-1.

We construct the foundation model using 12 Conformer blocks in the *Shared Encoder* and 6 Transformer blocks in the *Attention Decoder*. We employ $h = 4$ parallel attention heads in both Conformer block and Transformer block. For every layer, we use $d_k = d_v = d_{model}/h = 64$, $d_{ffn} = 2048$. Our proposed SememeASR model adds sememe encoder module and sememe prediction auxiliary task based on the baseline.

We train the model with Adam Optimizer [15] for at most 240 epochs with a batch size of 12. And *learning rate* = 0.002, *warm up* = 25000, and gradient clipping at 5.0. Additionally, during training, we employ the gradient accumulation method, in which the gradients are modified every four batches. Moreover, we employ label smoothing of value $\epsilon_{ls} = 0.1$ and dropout rate of $P_{drop} = 0.1$. We set the weight λ of the CTC branch during joint training to 0.3. Considering that $\alpha=0.3$ achieves better results in our experiment, we choose it as the weight parameter of sememe loss. During joint decoding, we set the CTC-weight λ to 0.5. To avoid overfitting, we averaged the 30 best model parameters in the development dataset.

4. Experimental results

The performance of the models is evaluated based on character error rates (CER) without external language models. Our experimental results are mainly based on the attention-rescore two-step decoding method.

4.1. Results of Different Dataset

We first present the results on the Aishell-1 test dataset. Table 1 compares the CER results of different models. From the results of Aishell-1, we can see our proposed SememeASR model is better than the baseline hybrid CTC/AED model.

Table 1: Comparison of CER on Aishell-1 and Aishell-2

Model	Aishell-1	Aishell-2	Aishell-2
	test	dev	test
Baseline (CTC/AED)	4.56	12.18	12.03
SememeASR-SP	4.53	11.99	12.16
SememeASR-SE	4.59	11.98	11.95
SememeASR-SEP	4.53	11.93	12.03

We also compare the results on the Aishell-2 test and dev datasets, which have a similar acoustic environment with Aishell-1 but cover different text domains. From the results of Aishell-2 in Table 1, we can see that our proposed SememeASR model outperforms the baseline on the new domain data. This indicates the improvement in the domain generalization capability of our proposed model.

To further illustrate the validity of our approach in more difficult text domains and more complex acoustic environments, we conduct further experiments on different domains of WenetSpeech. Experimental results in Table 3 indicate that our method also outperforms the baseline in new domains.

4.2. Results of Long Tail Data

To evaluate the ability of the model to recognize long-tail data, we first counted the long-tail data according to the methodology of [22]. The characters in the bottom 95% of occurrences in the training set were used as long-tail characters. In addition, in order to analyze the impact of long-tail data in more detail, we have divided 10 intervals based on the ratio of the number of long-tail characters in the sentence.

According to the result of Figure 7, the proposed method can improve the recognition of long-tailed data, which will also lead to a little reduction in the recognition of head data. However, as the proportion of long-tailed characters in a sentence increases, our model is less effective than the baseline model. It suggests that in the case of poor recognition, wrong results bring wrong semantic information, which hinders the role of

Table 2: ASR experiments on different decoding methods

Dataset	Model	attention	CTC greedy search	CTC prefix beam search	attention rescoring
Aishell-1 test	Baseline (CTC/AED)	4.86	4.82	4.82	4.56
	SememeASR-SP	4.87	4.91	4.91	4.53
	SememeASR-SE	4.82	5.02	5.02	4.59
	SememeASR-SEP	4.71	4.94	4.94	4.53
Aishell-2 test	Baseline (CTC/AED)	12.53	12.67	12.67	12.03
	SememeASR-SP	12.57	12.79	12.79	12.16
	SememeASR-SE	12.49	12.75	12.75	11.95
	SememeASR-SEP	12.30	12.75	12.75	12.03
Aishell-2 dev	Baseline (CTC/AED)	12.93	12.76	12.75	12.18
	SememeASR-SP	12.43	12.85	12.84	11.99
	SememeASR-SE	12.41	12.66	12.66	11.98
	SememeASR-SEP	12.23	12.78	12.77	11.93

Table 3: Performance on different domains of WenetSpeech

Domain	Baseline	SememeASR		
		SP	SE	SEP
audiobook	15.29	15.47	15.59	15.69
commentary	37.74	37.09	37.00	36.52
documentary	41.04	40.53	40.24	40.36
drama	56.24	54.39	55.68	55.10
interview	38.44	37.92	38.26	38.10
news	31.98	32.15	32.31	31.98
reading	40.56	40.43	40.16	39.01
talk	34.03	33.98	34.29	34.14
variety	57.72	57.21	58.18	57.29
others	34.65	33.60	33.97	33.47




Figure 8: CER of long-tailed data on Aishell-2 test set.




Figure 7: CER of long-tailed data on Aishell-1 test set.

semantic information.

Furthermore, we also evaluate the ability of the model in recognizing long-tail data in the Aishell-2 dataset. As shown in Figure 8, the proposed method can also improve the recognition of long-tail data overall on the new domain data.

4.3. Further Analysis on Different Parts of the Model

To further analyze the effect of sememe information on the different parts of the model, we adopted different decoding approaches for our experiments. Similar to [11], we adopt four decoding methods, **attention**, **CTC greedy search**, **CTC prefix beam search**, **attention rescoring**.

As shown in Table 2, in the **attention** and **attention rescoring** decoding methods, our proposed method performs better than the baseline CTC/AED model. However, at the same time, the performance of the CTC-based decoding approach is infe-

rior to that of the baseline. It indicates that our proposed approach increases the modeling capability of the *Attention Decoder*, but has an impact on the *Shared Encoder*. According to previous studies [23, 24], this is because the *Shared Encoder* part is more correlated with acoustic modeling, while the *Attention Decoder* part is correlated with language modeling. Our approach improves the language modeling capability, but the coupling of acoustic modeling and language modeling makes our proposed method inevitably influential in its acoustic modeling component. And thus the effect of using the CTC decoding approach is degraded.

5. Conclusion and Future Work

In this paper, we introduce sememe knowledge into the E2E ASR model and verify the effectiveness of external semantic knowledge for data-driven models. The proposed SememeASR can improve the recognition of long-tail data and enhance the domain generalization ability of the model. The intention of our work is to validate the effectiveness of sememe information for boosting ASR performance. Therefore, we explore a series of simple but effective model structures. In the future, we will consider optimizing the model structure and exploring different methods to further enhance the role of semantics for ASR.

6. Acknowledgements

This work is supported by Shenzhen Science and Technology Program (WDZC20200818121348001, JCYJ20220818101014 030), the Major Key Project of PCL (PCL2021A06, PCL2022 D01) and AMiner.Shenzhen SciBrain fund.

7. References

- [1] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, “Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks,” in *ACM International Conference on Machine Learning (ICML)*, 2006, pp. 369–376.
- [2] A. Graves, “Sequence transduction with recurrent neural networks,” *arXiv preprint arXiv:1211.3711*, 2012.
- [3] S. Wang, P. Zhou, W. Chen, J. Jia, and L. Xie, “Exploring rnn-transducer for chinese speech recognition,” in *IEEE Asia Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, 2019, pp. 1364–1369.
- [4] J. Chorowski, D. Bahdanau, K. Cho, and Y. Bengio, “End-to-end continuous speech recognition using attention-based recurrent nn: First results,” in *Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2014.
- [5] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *IEEE international conference on acoustics, speech and signal processing (ICASSP)*, 2016, pp. 4960–4964.
- [6] H. Luo, S. Zhang, M. Lei, and L. Xie, “Simplified self-attention for transformer-based end-to-end speech recognition,” in *IEEE Spoken Language Technology Workshop (SLT)*, 2021, pp. 75–81.
- [7] Y. Niu, R. Xie, Z. Liu, and M. Sun, “Improved word representation learning with sememes,” in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2017, pp. 2049–2058.
- [8] Y. Gu, J. Yan, H. Zhu, Z. Liu, R. Xie, M. Sun, F. Lin, and L. Lin, “Language modeling with sparse product of sememe experts,” in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2018, pp. 4642–4651.
- [9] F. Qi, R. Xie, Y. Zang, Z. Liu, and M. Sun, “Sememe knowledge computation: A review of recent advances in application and expansion of sememe knowledge bases,” *Frontiers of Computer Science*, vol. 15, no. 5, pp. 1–11, 2021.
- [10] Y. Zhang, C. Yang, Z. Zhou, and Z. Liu, “Enhancing transformer with sememe knowledge,” in *Proceedings of the 5th Workshop on Representation Learning for NLP*, 2020, pp. 177–184.
- [11] Z. Yao, D. Wu, X. Wang, B. Zhang, F. Yu, C. Yang, Z. Peng, X. Chen, L. Xie, and X. Lei, “Wenet: Production oriented streaming and non-streaming end-to-end speech recognition toolkit,” in *ISCA Conference of the International Speech Communication Association (Interspeech)*, 2021, pp. 4054–4058.
- [12] Y. Peng, S. Dalmia, I. Lane, and S. Watanabe, “Branchformer: Parallel mlp-attention architectures to capture local and global context for speech recognition and understanding,” in *ACM International Conference on Machine Learning (ICML)*, 2022, pp. 17 627–17 643.
- [13] X. Ren, H. Zhu, L. Wei, M. Wu, and J. Hao, “Improving mandarin speech recognition with block-augmented transformer,” *arXiv preprint arXiv:2207.11697*, 2022.
- [14] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, “Conformer: Convolution-augmented Transformer for Speech Recognition,” in *ISCA Conference of the International Speech Communication Association (Interspeech)*, 2020, pp. 5036–5040.
- [15] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2017, pp. 5998–6008.
- [16] Y. Bai, J. Yi, J. Tao, Z. Tian, Z. Wen, and S. Zhang, “Fast end-to-end speech recognition via non-autoregressive models and cross-modal knowledge transferring from bert,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 1897–1911, 2021.
- [17] A. Newell, K. Yang, and J. Deng, “Stacked hourglass networks for human pose estimation,” in *14th European Conference on Computer Vision (ECCV)*. Springer Verlag, 2016, pp. 483–499.
- [18] H. Bu, J. Du, X. Na, B. Wu, and H. Zheng, “Aishell-1: An open-source mandarin speech corpus and a speech recognition baseline,” in *IEEE Conference of the Oriental Chapter of the International Coordinating Committee on Speech Databases and Speech I/O Systems and Assessment (O-COCOSDA)*, 2017, pp. 1–5.
- [19] J. Du, X. Na, X. Liu, and H. Bu, “Aishell-2: Transforming mandarin asr research into industrial scale,” *arXiv preprint arXiv:1808.10583*, 2018.
- [20] B. Zhang, H. Lv, P. Guo, Q. Shao, C. Yang, L. Xie, X. Xu, H. Bu, X. Chen, C. Zeng *et al.*, “Wenetspeech: A 10000+ hours multi-domain mandarin corpus for speech recognition,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 6182–6186.
- [21] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, “SpecAugment: A simple data augmentation method for automatic speech recognition,” in *ISCA Conference of the International Speech Communication Association (Interspeech)*, 2019, pp. 2613–2617.
- [22] K. Deng, G. Cheng, R. Yang, and Y. Yan, “Alleviating asr long-tailed problem by decoupling the learning of representation and classification,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 340–354, 2021.
- [23] P. Wang, T. N. Sainath, and R. J. Weiss, “Multitask training with text data for end-to-end speech recognition,” in *ISCA Conference of the International Speech Communication Association (Interspeech)*, 2021, pp. 2566–2570.
- [24] S. Kim, K. Li, L. Kabela, R. Huang, J. Zhu, O. Kalinli, and D. Le, “Joint audio/text training for transformer rescorer of streaming speech recognition,” in *Findings of the Association for Computational Linguistics: EMNLP*, 2022, pp. 5717–5722.