

CopyNE: Better Contextual ASR by Copying Named Entities

Anonymous ACL submission

Abstract

End-to-end automatic speech recognition (ASR) systems have made significant progress in general scenarios. However, it remains challenging to transcribe contextual named entities (NEs) in the contextual ASR scenario. Previous approaches have attempted to address this by utilizing the NE dictionary. These approaches treat entities as individual tokens and generate them token-by-token, which may result in incomplete transcriptions of entities. In this paper, we treat entities as indivisible wholes and introduce the idea of copying into ASR. We design a systematic mechanism called CopyNE, which can copy entities from the NE dictionary. By copying all tokens of an entity at once, we can reduce errors during entity transcription, ensuring the completeness of the entity. Experiments demonstrate that CopyNE consistently improves the accuracy of transcribing entities compared to previous approaches. Even when based on the strong Whisper, CopyNE still achieves notable improvements.

1 Introduction

End-to-end automatic speech recognition (ASR) systems have achieved impressive performance in general scenarios (Chan et al., 2016; Rao et al., 2017; Gulati et al., 2020; Boulianne, 2022). However, in the contextual ASR scenario where speech often contains numerous contextual entities, it remains a challenge for ASR systems to get accurate transcriptions. For instance, when utilizing personal voice assistants like Siri or Alexa, it is common to encounter contextual entities such as personal names, place names, and organization names. ASR models trained solely on speech-text data often struggle to transcribe these personalized entities due to their infrequent occurrence in the training set. Since contextual entities always cover a wealth of semantic information. It is important to improve the accuracy of transcribing entities for downstream natural language processing (NLP)

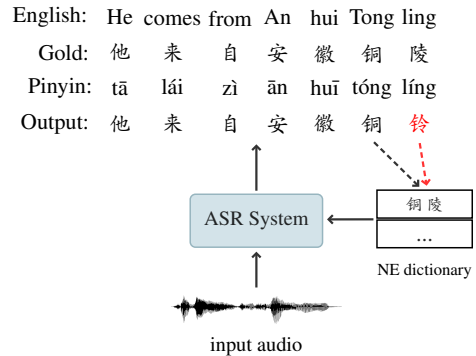


Figure 1: An example with homophonic errors. Pinyin is the Mandarin pronunciation of each token. The red text indicates the wrongly predicted token.

tasks such as information retrieval and spoken language understanding (Ganesan et al., 2021; Wu et al., 2022).

Recently, researchers have started leveraging the information of textual modality as additional contextual knowledge to help contextual ASR. The most typical approach, premised on the assumption that entities are already known before, use a contextual named entity (NE) dictionary as contextual knowledge (Chen et al., 2019; Alon et al., 2019; Jain et al., 2020; Han et al., 2021; Huber et al., 2021; Fu et al., 2023). Two representative approaches are “contextual listen, attend and spell” (CLAS) (Pundak et al., 2018) and contextual bias attention (CBA) (Zhang and Zhou, 2022). CLAS employs the knowledge of the dictionary to aid the prediction of each token. They use dictionary representation as extra inputs for token prediction in the decoder. The decoder attends to each entity, and the dictionary representation is an aggregated representation of all entities, weighted by the attention scores. CBA extends CLAS and uses an extra training loss. The loss explicitly makes use of the attention scores and force the model attend to a proper entity in the dictionary if the token to be predicted is related with the entity.

Previous methods have achieved considerable improvements, especially in transcribing entities. However, they all treat entities as individual tokens. These models utilize contextual knowledge to aid in predicting independent tokens without considering the role of these tokens in constituting a complete entity. In other words, a multi-token entity is broken into isolated tokens during decoding. We argue that this is problematic. For instance, model may erroneously generate the subsequent tokens of an entity, despite correctly producing the preceding tokens. As shown in Figure 1, when transcribing the speech “他来自安徽铜陵” (He comes from Anhui Tongling), an incorrect output of “他来自安徽铜铃” (He comes from Anhui copper bell) is obtained. Despite the model’s awareness of the location entity “铜陵” in the NE dictionary and its accurate prediction of the first token “铜”, it mistakenly transcribes “陵” (ling) as “铃” (bell) during the token-level prediction process. This occurs because the model predicts tokens independently, neglecting the integrity of the token span as a complete entity. Furthermore, “陵” and “铃” share the same pronunciation “líng”, with “铃” being a more frequently occurring token in the training data. Consequently, the model tends to generate the wrong token “铃”.

In this paper, we propose a new approach for contextual ASR called CopyNE. Unlike previous approaches, we view entities as indivisible wholes. To the best of our knowledge, we are the first to introduce the idea of copying into ASR. We design a systematic and effective mechanism to copy entities from a dictionary. Specifically, CopyNE uses a copy loss that guides the model to copy the correct entity from the dictionary. During inference, our CopyNE has the flexibility to either predict a token from the token vocabulary or copy an entity from the NE dictionary at each decoding step. By copying multiple tokens simultaneously, we can alleviate errors within the entity, thus ensuring the token span as a complete entity.

Experiments on Chinese Aishell (Bu et al., 2017), ST-cmds¹, and English Eng (Yadav et al., 2020) datasets show that our CopyNE achieves significant improvements across all scenarios, particularly in the contextual ASR scenario. Compared to previous methods using dictionary, CopyNE achieves relative reductions in CER of 13.5% and 20.8% on Aishell and ST-cmds in the contextual scenario. Notably, CopyNE shows more remark-

able improvements when it comes to transcribing entities, with relative reductions of 55.4% and 53.9% in the NE-CER metric on Aishell and ST-cmds. Moreover, when based on Whisper (Radford et al., 2022) and evaluated in its domain of expertise, Eng dataset, CopyNE still achieves an impressive 6.4% and 16.8% relative reductions in WER and NE-WER. We will release our codes, configurations, and models at <https://github.com/>.

2 The CTC-Transformer Model

In this work, we build our proposed approach on the end-to-end CTC-Transformer model, since it is the most widely used and achieves competitive performance in the ASR field (Hori et al., 2017; Kim et al., 2017; Miao et al., 2020; Omachi et al., 2021; Gong et al., 2022). However, it is worth noting that our idea can be applied to other ASR approaches as well.

The CTC-Transformer is built upon the seq-to-seq Transformer (Vaswani et al., 2017), with a connectionist temporal classification (CTC) layer added after the audio encoder. As shown in Figure 2, it takes a sequence of acoustic frames $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_T)$ as input and generates the corresponding transcription text $\mathbf{y} = (y_1, \dots, y_U)$ as output. The model consists of two main components: an encoder and a decoder. First, the encoder encodes the acoustic frames \mathbf{X} into hidden states $\mathbf{H} = (\mathbf{h}_1, \dots, \mathbf{h}_T)$. Then, the decoder predicts the target sequence \mathbf{y} in an auto-regressive manner. At each decoding step u , the decoder predicts the next target token y_{u+1} based on the encoder’s output \mathbf{H} and the previously predicted tokens $y_{\leq u} = (y_1, \dots, y_u)$. This process is expressed as follows:

$$\mathbf{H} = \text{AudioEncoder}(\mathbf{X}) \quad (1)$$

$$\mathbf{d}_u = \text{Decoder}(y_{\leq u}, \mathbf{H}) \quad (2)$$

$$P(y_{u+1}|y_{\leq u}) = \text{softmax}(\mathbf{W}\mathbf{d}_u + \mathbf{b}) \quad (3)$$

Here, $\mathbf{d}_u \in \mathbb{R}^d$ denotes the hidden state at step u , and $P(y_{u+1}|y_{\leq u})$ is the posterior distribution of predicting token y_{u+1} . $\mathbf{W} \in \mathbb{R}^{|\mathcal{V}| \times d}$ and $\mathbf{b} \in \mathbb{R}^{|\mathcal{V}|}$ are learned parameters, where \mathcal{V} is the token vocabulary, and $|\mathcal{V}|$ is the size of the vocabulary.

The loss of Transformer, $\mathcal{L}_{trans}(\mathbf{y})$, comes from minimizing the negative log probability of \mathbf{y} .

$$\mathcal{L}_{trans}(\mathbf{y}) = - \sum_{u=0}^{U-1} \log P(y_{u+1}|y_{\leq u}) \quad (4)$$

¹<https://www.openslr.org/38/>




Figure 2: The CTC-Transformer model.

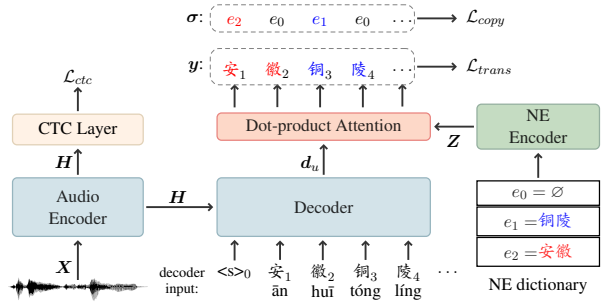


Figure 3: Our CopyNE model.

As commonly used in previous works, the CTC loss is also applied here. CTC aligns each acoustic frame with a token from left to right. For a given target sequence y , there may be multiple valid alignments. The CTC loss is derived from maximizing the sum of these valid alignments, and has been proved to be able to enhance the representational capacity of the audio encoder (Kim et al., 2017). Finally, the overall loss is the a weighted sum of the $\mathcal{L}_{trans}(y)$ and $\mathcal{L}_{ctc}(y)$, as follows:

$$\mathcal{L}(y) = \lambda \mathcal{L}_{trans}(y) + (1 - \lambda) \mathcal{L}_{ctc}(y) \quad (5)$$

where λ is a hyper-parameter that determines the relative weight of each loss term.

In inference, the model selects the most probable transcription using beam search as follows:

$$\hat{y} = \arg \max_y \left(\sum_u \log P(y_{u+1} | y_{\leq u}) \right) \quad (6)$$

Here, there are many ways to use scores in decoding, such as combining CTC scores and Transformer scores as in training, or using CTC-prefix beam search followed by re-scoring with Transformer to select the optimal result. To compare with most previous works, we use the simplest decoding strategy, as shown in Equation 6.

3 Our CopyNE Model

This section describes our proposed CopyNE model. The basic idea is that the model incorporates a contextual NE dictionary as external knowledge and can choose to directly copy NEs from the dictionary. We design a systematic framework to implement the idea. During training, a copy loss is designed to encourage the model to copy a corresponding entity from the dictionary when the following tokens appears as entity in the dictionary. During inference, at each generation step, the model can either predict a single token from token vocabulary or directly copy a entity from the given dictionary.

3.1 The Model Framework

Figure 3 illustrates the framework of our CopyNE model, which shares the same encoder as the CTC-Transformer model, but with a distinct decoder. In the decoder, we introduce an extra NE encoder that takes the NE dictionary as input and encodes it into NE representations. Then, we use a dot-product attention module to compute copy probabilities based on the obtained NE representations, which are then aggregated to form the overall dictionary (Dict) representation. The decoder can not only utilize copy probabilities to select entities for copying but also leverage the Dict representation to aid in predicting the next token.

NE Representation. We denote an NE dictionary as $E = (e_0, e_1, \dots, e_N)$. We use $e_0 = \emptyset$ as a pseudo entity to handle the case where the text to be transcribed has no relation to any entity and the model should not copy any entity at current step.

For each entity e_i , we apply a multi-layer LSTM as the NE encoder to encode the token sequence and use the last hidden state of the NE encoder as the entity representation. It is a popular practice in previous contextual ASR works.

$$z_i = \text{LSTM}(e_i) \quad (7)$$

After that, we get entity representations $Z = (z_0, z_1, \dots, z_N)$, where $Z \in \mathbb{R}^{N \times d}$.

Copy Probability. Once the NE representations are obtained, the copy probability is computed by a dot-product attention mechanism. It is used to determine which entity to copy. First, we compute the attention score $a_u^{e_i}$ for entity e_i at step u as follows:

$$a_u^{e_i} = \frac{(\mathbf{W}_q d_u)^\top (\mathbf{W}_k z_i)}{\sqrt{d}} \quad (8)$$

where $\mathbf{W}_q, \mathbf{W}_k \in \mathbb{R}^{d_a \times d}$ are two learned parameters. d_a denotes the dimension of the atten-

tion. After that we obtain the attention probability $P_c(e_i|y_{\leq u})$ for entity e_i by softmax.

$$P_c(e_i|y_{\leq u}) = \frac{\exp(a_u^{e_i})}{\sum_{e_j \in E} \exp(a_u^{e_j})} \quad (9)$$

Here, $P_c(e_i|y_{\leq u})$ not only represents the attention probability of e_i but also naturally serves as the copy probability for the entity. During inference, we use the copy probabilities to select the entities for copying.

Dict Representation. With copy (attention) probabilities, we can obtain the Dict representation \mathbf{r}_u at decoding step u . It is used to help the prediction of subsequent tokens. Specifically, $\mathbf{r}_u \in \mathbb{R}^d$ is computed by weighted summing the entity representations with the copy (attention) probabilities.

$$\mathbf{r}_u = \sum_{e_i \in E} P_c(e_i|y_{\leq u}) \mathbf{z}_i \quad (10)$$

Dict-enhanced Prediction. Finally, we get the overall Dict representation and copy probabilities. Following Pundak et al. (2018), the Dict representation is applied to help the generation of the next token. So Equation 3 is extended as follows:

$$P(y_{u+1}|y_{\leq u}, E) = \text{softmax}(\mathbf{W} [\mathbf{d}_u, \mathbf{r}_u] + \mathbf{b}) \quad (11)$$

3.2 Training

During training, to guide the model in selecting correct entities from the NE dictionary for copying, we introduce an additional copy loss \mathcal{L}_{copy} . First, based on the ground truth transcription \mathbf{y} and the NE dictionary, we construct a copy target σ_{u+1} for each decoding step u , telling the model whether to copy an entity from the dictionary or not, and which one to copy. Then we compute the copy loss \mathcal{L}_{copy} according to the copy target σ_{u+1} and the copy probability $P_c(\sigma_{u+1}|y_{\leq u})$.

The Computation of Copy Loss. Provided that we have an NE dictionary E^b , we construct a copy target, denoted as σ_{u+1} , for decoding step u . In order to build the copy target, we perform maximum matching on the transcription text \mathbf{y} from left to right based on the dictionary E^b . If the token span $\mathbf{y}_{i,j} = (y_i, \dots, y_j)$ matches the k -th entity e_k in E^b , then we set the copy target $\sigma_i = e_k$, and $\sigma_{i+1 \sim j} = \emptyset$. This indicates that the model can copy the k -th entity from the dictionary at decoding step $i - 1$, but cannot copy any entity from

decoding step i to $j - 1$. When it comes to a span of length 1, i.e., $i = j$, during the left-to-right maximum matching process, we also set σ_i to \emptyset^2 .

For example, in the instance shown in Figure 3, the span “安徽” (An hui) matches the second entity in the dictionary, and the span “铜陵” (Tong ling) matches the first entity in the dictionary. This means that at steps 0 and 2, the model can choose to copy the second and first entities from the dictionary, respectively. Therefore, $\sigma_1 = e_2$ and $\sigma_3 = e_1$, while $\sigma_2 = \emptyset$ and $\sigma_4 = \emptyset$.

After constructing all the copy targets $\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_U)$, we can compute the copy loss as follows:

$$\mathcal{L}_{copy}(\boldsymbol{\sigma}) = - \sum_{u=0}^{U-1} \log P_c(\sigma_{u+1}|y_{\leq u}) \quad (12)$$

where $P_c(\sigma_{u+1}|y_{\leq u})$ is the copy probability computed in Equation 9, meaning the probability of copying entity σ_{u+1} at decoding step u . It is worth noting that the copy loss and the bias loss in CBA have fundamental differences. The bias loss in CBA provides information to each token, including tokens within entities, about which entity to attend to. In contrast, our copy loss solely instructs the model to copy the entity from the dictionary at the first token of the entity.

Finally, the loss in our CopyNE model is formed as follows:

$$\mathcal{L} = \lambda \mathcal{L}_{trans}(\mathbf{y}) + (1 - \lambda) \mathcal{L}_{ctc}(\mathbf{y}) + \mathcal{L}_{copy}(\boldsymbol{\sigma}) \quad (13)$$

Dictionary Construction. To construct the copy target and compute the copy loss, we should first build a contextual NE dictionary for training. Provided that the entities have been labeled in the dataset, we build a NE dictionary E^b for each data batch following previous works.

Firstly, to construct E^b , we extract all entities in the instances of this batch and add them to the dictionary. For instances that do not contain any entities, in order to ensure an adequate number of positive examples, we randomly select one or two substrings of length 2 or 3 from the transcription and include them in the dictionary as pseudo entities. In order to improve the ability of copying the correct entity from a wide range of entities, we also extract additional negative entities from the

²Please note that in this paper, we primarily focus on entities with a length greater than 1, and therefore only retain such entities in our dictionary.

training set. We analyze the influence of the quantity of negative entities on the model. Due to page constraints, we have included this section in §C.

3.3 Inference

During inference, our CopyNE model generates the transcription in an auto-regressive manner. However, unlike previous token-level approaches, our model has the flexibility to predict either a token from the vocabulary or an entity from the NE dictionary. In the latter case, the model can produce multiple tokens at one step. By copying the tokens of an entity at once, our CopyNE model can avoid errors that occur when predicting multiple tokens separately. As shown in Figure 3, our CopyNE model can directly copy the two entities “安徽” and “铜陵” from the dictionary.

Specifically, at step u , our prediction is based on both the model’s probability for a token v , i.e., $P(v|\hat{y}_{\leq u}, E)$, and the copy probability for an entity e , i.e., $P_c(e|\hat{y}_{\leq u})$. The former represents the probability of predicting a token v from the token vocabulary, while the latter is normalized on all entities, originally indicating the attention probability over entity e , which can be naturally interpreted as the probability of copying entity e from the dictionary. To consider both probabilities on the same scale, we devise an elegant decoding strategy by taking use of the copy probability of \emptyset , i.e., $P_c(\emptyset|\hat{y}_{\leq u})$, and re-normalize the probabilities to create an unified searching space Q .

$$Q(i|\hat{y}_{\leq u}) = \begin{cases} P_c(\emptyset|\hat{y}_{\leq u})P(i|\hat{y}_{\leq u}, E), & i \in \mathcal{V} \\ P_c(i|\hat{y}_{\leq u}), & i \in E, i \neq \emptyset \end{cases} \quad (14)$$

Here, to ensure the sum of the probabilities of all elements is 1, we use $P_c(\emptyset|\hat{y}_{\leq u})$ as a prior probability, representing the probability of the text to be transcribed has no relation with the entities in the dictionary and the text should be generated from the token vocabulary. If the element is from the token vocabulary \mathcal{V} , we obtain the probability by multiplying the prior probability and the model’s probability for the token. Otherwise, we use the copy probability directly.

However, in our experiments, we observe that the model occasionally selects irrelevant entities for copying. To enhance the quality of copying, we introduce a confidence threshold γ during decoding to filter out low-confidence copies. Specifically, we set $P_c(i|\hat{y}_{\leq u}) = 0, i \in E, i \neq \emptyset$, and $P_c(\emptyset|\hat{y}_{\leq u}) = 1$ when $\max\{P_c(i|\hat{y}_{\leq u})|i \in E, i \neq$

$\emptyset\} < \gamma$. This means that if the model’s maximum copy probability over real entities is less than γ , it is prevented from copying entities from the dictionary and instead generates tokens from the token vocabulary. In section 4.2, we discuss the influence of the γ in detail.

Finally, we use beam search to select the best element at each step to form the final prediction³.

$$\hat{y} = \arg \max_y \left(\sum_u \log Q(i|y_{\leq u}) \right) \quad (15)$$

4 Experiments

4.1 Experimental Setup

Datasets. Experiments on Chinese Mandarin are conducted on two widely used datasets, Aishell (Bu et al., 2017) and ST-cmds⁴. We use the Eng dataset released by Yadav et al. (2020) to perform experiments on English. Furthermore, to compare the performance of different methods in contextual ASR scenarios where speeches contain entities, we extract instances containing entities from the dev and test sets, forming the corresponding “*-NE” datasets. Detailed introduction about the datasets can be found in §A.

NE Dictionary. Aishell and ST-cmds were initially released without entity annotations. In contrast, the Eng dataset was simultaneously released with audio, transcribed text, and corresponding entity annotations. Chen et al. (2022) further annotated entities in each transcription for Aishell. So, in our experiments with Aishell and Eng, we use the related entities to build the NE dictionary. For ST-cmds, we use HanLP⁵ to get three types of entities: person, location, and organization. The dictionary is build with these predictions.

Evaluation Metrics. Character error rate (CER) and word error rate (WER) are used to assess the overall performance of models in Mandarin and English ASR tasks. In this paper, to evaluate the model’s entity transcription accuracy, we also employ NE-CER and NE-WER metrics (Han et al., 2021). We align the predicted hypothesis and reference using the minimum edit distance algorithm, and subsequently calculate NE-C(W)ER by measuring the C(W)ER between the entity text in the reference and its counterpart in the hypothesis.

³It has to be noted that the partially predicted \hat{y} is still encoded at token-level.

⁴<https://www.openslr.org/38/>

⁵<https://github.com/hankcs/HanLP>

Model	Aishell		Aishell-NE		ST-cmds		ST-cmds-NE	
	Dev	Test	Dev	Test	Dev	Test	Dev	Test
Joint CTC-Transformer	6.12	6.70	7.36	9.00	10.63	10.56	13.67	13.63
CLAS (Pundak et al., 2018)	6.04	6.72	7.06	8.73	10.10	10.09	12.64	12.85
CBA (Zhang and Zhou, 2022)	6.11	6.56	6.73	8.00	10.73	10.72	12.69	12.43
Our CopyNE	5.59	6.35	5.36	6.92	9.76	9.89	9.90	9.84
Whisper	5.28	5.97	6.32	7.68	9.14	9.06	12.22	12.35
+ CLAS	4.50	5.23	5.30	6.72	9.10	9.20	12.12	12.25
+ CBA	5.41	6.06	6.13	7.60	7.96	7.87	10.03	9.96
+ Our CopyNE	4.71	5.10	5.40	6.42	7.35	7.19	8.98	8.40

Table 1: CER (%) on the Chinese datasets in general scenarios (Aishell, ST-cmds) and contextual scenarios (Aishell-NE, ST-cmds-NE).

Parameter Setting. The parameter setting in our work is the same as that in most previous works, and the detailed descriptions can be found in §B. To ensure a fair comparison with prior works, we carefully reproduced the CLAS (Pundak et al., 2018) and CBA (Zhang and Zhou, 2022) using the same structure as our model, including the same audio encoder, decoder, and NE encoder, among others. Moreover, to verify the effectiveness of our approach on pre-trained large models, we also conducted experiments on OpenAI Whisper (Radford et al., 2022). Specifically, we use the Whisper model as our transformer encoder and decoder. We choose seeds randomly to run models for 3 times and report the average results.

4.2 Results




Figure 4: Effect of the Confidence Threshold γ .

Analysis about γ . We first investigate the influence of the copy threshold γ during inference. Figure 4 illustrates how the CER changed on the Aishell dev and Aishell-NE dev with different γ values. Our findings reveal that when the threshold is low, the CER is high, indicating that copying results in more errors when the model copies entities with low confidence. As we increase the γ ,

Model	Aishell-NE		ST-cmds-NE	
	Dev	Test	Dev	Test
Joint CTC-Transformer	11.64	14.03	21.63	21.41
CLAS (Pundak et al., 2018)	11.24	13.12	19.70	20.10
CBA (Zhang and Zhou, 2022)	7.78	9.44	15.72	15.92
Our CopyNE	3.00	4.21	7.60	7.34
Whisper	10.31	12.24	21.30	21.83
+ CLAS	8.97	11.64	20.82	21.07
+ CBA	9.13	11.79	15.91	15.41
+ Our CopyNE	6.74	8.79	11.93	11.29

Table 2: NE-CER (%) on the Chinese datasets.

the CER decreases, indicating improved reliability of our CopyNE when the model had higher confidence. However, when the threshold becomes too high (above 0.9), the model has fewer opportunities to choose to copy entities, resulting in a higher CER. This happens because it becomes more difficult for the model to trigger the copy mechanism. So, we set γ to 0.9 for all experiments and discussions to enhance the robustness of our model.

Results on Chinese. Table 1 and 2 show the CER and NE-CER of different models on the Chinese dataset. In Table 1, we note that while our primary focus is on improving entity transcription accuracy, we also achieve significant improvements in overall text transcription. Without Whisper, our CopyNE model outperforms the previous CBA approach with a 3.2% relative CER reduction on the Aishell Test and 7.7% on the ST-cmds Test. In contextual ASR scenarios, the improvements are even more pronounced, with a 13.5% relative CER reduction on the Aishell-NE Test and 20.8% on the ST-cmds-NE Test. Even with the powerful Whisper, our CopyNE consistently excels, especially on the ST-cmds dataset, with relative reductions of 8.6% and 15.7% on the two test sets, respectively. Additionally, we observed that CLAS performs

Model	Eng. _w		Eng-NE _w		Eng-NE _{NW}	
	Dev	Test	Dev	Test	Dev	Test
Whisper	8.54	8.73	8.53	8.71	28.16	26.61
+ CLAS	7.90	8.28	7.86	8.31	27.23	26.55
+ CBA	9.17	9.52	9.21	9.49	30.01	30.82
+ CopyNE	7.47	7.85	7.42	7.78	23.29	22.09

Table 3: Results on the English datasets. W and NW denote WER and NE-WER respectively.

well on Aishell, closely matching CopyNE, but its performance on ST-cmds is comparatively weaker, sometimes even worse than the Whisper baseline, a reverse pattern also seen with CBA. In contrast, CopyNE consistently performs well across different datasets, demonstrating its better adaptability.

In this paper, our main goal is to improve the accuracy of transcribing entities in contextual ASR scenarios. From the results presented in Table 2, it is evident that our approach exhibits significant improvements in entity transcription compared to previous methods. When not using Whisper, our CopyNE model achieves an impressive relative NE-CER reduction of 55.4% on the Aishell-NE Test and 53.9% on the ST-cmds-NE Test. Even based on the powerful Whisper model, our CopyNE continues to achieve remarkable improvements, with a relative NE-CER reduction of 25.4% and 26.7% on the two test sets. This demonstrates that copying entities from the dictionary significantly improves the accuracy of transcribing entities.

Results on English. Whisper (Radford et al., 2022) has already shown strong performance in English, so we directly use it as our baseline for experiments on English. As seen in Table 3, CopyNE still outperforms other methods, achieving a 5.2% relative WER reduction compared to CLAS in the general scenario on the Eng test dataset. In contextual scenarios, CopyNE demonstrates 6.4% relative WER reductions and 16.8% relative NE-WER reductions. Additionally, we observed that CBA lags behind the Whisper baseline. We suspect that this might be due to its approach of encouraging the model to generate entity tokens by modifying Whisper’s output logits, which can disrupt the model’s overall probability distribution, especially given Whisper’s strong fit on English data. On the contrary, our CopyNE is more stable.

4.3 Impact of Noisy NEs

In this paper, our primary focus lies on the contextual ASR scenario where the majority of NEs are already known. Following previous works (Pundak

Noisy NEs	Dev		Test	
	CER	NE-CER	CER	NE-CER
0	5.36	3.00	6.92	4.21
2k	5.61	3.50	7.12	4.82
4k	5.85	3.92	7.39	5.18
6k	6.02	4.09	7.66	5.56
8k	6.15	4.27	7.68	5.81

Table 4: The impact of noisy NEs on CopyNE.

	Transcriptions	Dictionary
English	A company in Yangluo...	
Gold	阳逻的一家公司...	阳逻 (Yangluo)
CLAS	扬罗的一家公司...	
CBA	阳罗的一家公司...	
CopyNE	[阳逻]的一家公司...	杨丙卿 (Yang Bingqing)
English	Yang Bingqing served as manager...	
Gold	杨丙卿担任经理...	冈山 (Gangshan)
CLAS	杨澄清担任经理...	
CBA	杨炳卿担任经理...	
CopyNE	[杨丙卿]担任经理...	桃太郎体育馆 (Taotailang gym)
English	The Taotailang gym in Gangshan.	
Gold	冈山的桃太郎体育馆	
CLAS	冈山的淘汰郎体育馆	
CBA	刚山的淘汰狼体育馆	
CopyNE	[冈山]的[桃太郎体育馆]	

Table 5: Generations of different models. Red text indicates errors, while text enclosed in square brackets represents entities that were copied from the dictionary.

et al., 2018; Alon et al., 2019; Han et al., 2021), we report the main results based on exact NE dictionaries from the test sets. However, in some real-world settings, the dictionary may contain NEs that are not present in the test set. To analyze the performance of CopyNE in this setting, we randomly extract entities from the training set that are not included in the test set as noisy NEs and add them to the original dictionary. Table 4 displays the impact of these noisy entities on the Aishell-NE dataset.

It can be observed that the introduction of noisy NEs indeed leads to a reduction in the model’s performance. However, even when expanding the dictionary size to 10k (the original exact NE dictionary was around 2k in size), the average error rate increased by only about 1%. With 8k noisy NEs, our CopyNE still outperforms CLAS and CBA, even though they only use the exact dictionary.

4.4 Qualitative Analysis

CopyNE demonstrates significant improvements. To gain further insight into CopyNE’s performance, we conduct a qualitative analysis of its generations. Table 5 shows examples of transcriptions from different ASR models. We can see that in the second example, where CBA successfully identified the correct person entity “杨丙卿” from the dictionary

and produced a transcription that is close to gold, it still made a mistake by transcribing “炳” instead of “丙” due to the same pronunciation (bǐng). In contrast, CopyNE can copy all the tokens of the entity from the NE dictionary. For example, in the third example “冈山的桃太郎体育馆” (The Taotailang gym in Gangshan), CopyNE directly copies the location entity “冈山” (Gangshan) and the organization entity “桃太郎体育馆” (Taotailang gym), achieving a completely correct transcription.

5 Related Works

Contextual ASR. Researchers have explored various approaches to help models in the contextual ASR scenario, with the primary approaches being the utilization of external dictionaries and language models (LMs). Among them, using an external dictionary is the most straightforward method. CLAS (Pundak et al., 2018) was the first to introduce the use of dictionary information to aid in prediction. Alon et al. (2019) extend CLAS by adding phonetically similar alternative terms to the dictionary as challenging negative examples, aiming to improve the model’s ability to distinguish entities with similar pronunciations. Huber et al. (2021) propose to utilize the representation of a single entry in the dictionary that is most relevant to the current decoding status. Fu et al. (2023) propose to apply the character-based NE encoder to better capture generalizable acoustic features useful for transcribing rare entities. Different from our CopyNE, these methods all treat entities as individual tokens, which may result in incomplete NE transcriptions.

LMs trained on large-scale text data can learn rich linguistic and contextual knowledge, and thus can be used to assist contextual ASR. There are typically two approaches to leverage LMs for contextual ASR. The first approach involves using the LM to constrain the output of the decoder during ASR decoding, resulting in text that aligns better with the target language style (Chorowski and Jaitly, 2016; Sriram et al., 2017; Zhao et al., 2019a). The second approach is multi-modal pre-training. Researchers have explored joint pre-training of speech and text models, aiming to leverage information from both modalities, and have achieved promising results (Chung et al., 2021; Ao et al., 2022; Zhang et al., 2022). However, compared to using contextual dictionaries, models that rely on LMs tend to have much more parameters, which means that training and deploying require more time and

computational resources.

The Copy Mechanism. The copy mechanism can be traced back to the pointer network (Vinyals et al., 2015), which can predict output sequences from the input. The copy mechanism (Gu et al., 2016) extends the pointer network by enabling the model to generate sequences that are not present in the input. According to the source of copying, it can be divided into copying from input text, from document, and from external dictionary.

Copying from Input Text. The copy mechanism is most commonly used to copy text from input text. For instance, in text summarization tasks, it is common to employ the copy mechanism to copy keywords from the input text (Cheng and Lapata, 2016; Xu et al., 2020). In grammatical error correction tasks, where only a small portion requires correction, the copy mechanism is used to copy the correct text from the input text (Zhao et al., 2019b).

Copying from Document. In addition to copying from input text, the copy mechanism can be employed to copy text from other texts when the input text is not available. Lan et al. (2023) introduced the copy mechanism in decoder-only language models, where text fragments are selected from a vast amount of documents to generate the target text.

Copying from External Dictionary. In this paper, we introduce a systematic framework, encompassing both training and decoding, which seamlessly integrates the process of copying from an external dictionary to aid in generation. We believe that our framework can also be applied to other generation tasks.

6 Conclusion

In this paper, we consider entities as indivisible elements and introduce a copy mechanism into ASR for the first time to assist in transcribing entities. We devise a systematic copy framework that can copy all the tokens of an entity from the NE dictionary at once, preserving the token span as a complete entity. Our approach demonstrates substantial improvements on both English and Chinese datasets. In summary, CopyNE represents a significant advancement in contextual ASR, providing a promising direction in this field.

633 Limitations

634 From our experiments, we have found that an ex-
635 cessive number of noisy entities can impact the
636 performance. As part of our future work, we intend
637 to explore methods for dynamically filtering out
638 interfering entities from the dictionary during the
639 decoding process.

640 References

641 Uri Alon, Golan Pundak, and Tara N Sainath. 2019.
642 Contextual speech recognition with difficult negative
643 training examples. In *2019 IEEE International Con-
644 ference on Acoustics, Speech and Signal Processing
645 (ICASSP)*, pages 6440–6444.

646 Junyi Ao, Rui Wang, Long Zhou, Chengyi Wang, Shuo
647 Ren, Yu Wu, Shujie Liu, Tom Ko, Qing Li, Yu Zhang,
648 Zhihua Wei, Yao Qian, Jinyu Li, and Furu Wei.
649 2022. [SpeechT5: Unified-modal encoder-decoder
650 pre-training for spoken language processing](#). In *Pro-
651 ceedings of the 60th Annual Meeting of the Asso-
652 ciation for Computational Linguistics*, pages 5723–
653 5738.

654 Gilles Boulianne. 2022. [Phoneme transcription of en-
655 dangered languages: an evaluation of recent ASR
656 architectures in the single speaker scenario](#). In *Find-
657 ings of the Association for Computational Linguistics:
658 ACL 2022*, pages 2301–2308.

659 Hui Bu, Jiayu Du, Xingyu Na, Bengu Wu, and Hao
660 Zheng. 2017. Aishell-1: An open-source mandarin
661 speech corpus and a speech recognition baseline. In
662 *2017 20th Conference of the Oriental Chapter of
663 the International Coordinating Committee on Speech
664 Databases and Speech I/O Systems and Assessment
665 (O-COCOSDA)*, pages 1–5.

666 William Chan, Navdeep Jaitly, Quoc Le, and Oriol
667 Vinyals. 2016. [Listen, attend and spell: A neural
668 network for large vocabulary conversational speech
669 recognition](#). In *2016 IEEE International Confer-
670 ence on Acoustics, Speech and Signal Processing
671 (ICASSP)*, pages 4960–4964.

672 Boli Chen, Guangwei Xu, Xiaobin Wang, Pengjun Xie,
673 Meishan Zhang, and Fei Huang. 2022. Aishell-ner:
674 Named entity recognition from chinese speech. In
675 *2022 IEEE International Conference on Acoustics,
676 Speech and Signal Processing (ICASSP)*, pages 8352–
677 8356.

678 Zhehuai Chen, Mahaveer Jain, Yongqiang Wang,
679 Michael L Seltzer, and Christian Fuegen. 2019. Joint
680 grapheme and phoneme embeddings for contextual
681 end-to-end asr. In *2019 Conference of the Inter-
682 national Speech Communication Association (Inter-
683 speech)*, pages 3490–3494.

684 Jianpeng Cheng and Mirella Lapata. 2016. [Neural sum-
685 marization by extracting sentences and words](#). In

*Proceedings of the 54th Annual Meeting of the As-
sociation for Computational Linguistics*, pages 484–
494.

Jan Chorowski and Navdeep Jaitly. 2016. Towards
better decoding and language model integration
in sequence to sequence models. *arXiv preprint
arXiv:1612.02695*.

Yu-An Chung, Chenguang Zhu, and Michael Zeng.
2021. [SPLAT: Speech-language joint pre-training
for spoken language understanding](#). In *Proceedings
of the 2021 Conference of the North American Chap-
ter of the Association for Computational Linguistics:
Human Language Technologies*, pages 1897–1907.

Xuandi Fu, Kanthashree Mysore Sathyendra, Ankur
Gandhe, Jing Liu, Grant P. Strimel, Ross McGowan,
and Athanasios Mouchtaris. 2023. [Robust acoustic
and semantic contextual biasing in neural transducers
for speech recognition](#). In *2023 IEEE International
Conference on Acoustics, Speech and Signal Process-
ing (ICASSP)*, pages 1–5.

Karthik Ganesan, Pakhi Bamdev, Jaivarsan B, Amresh
Venugopal, and Abhinav Tushar. 2021. [N-best ASR
transformer: Enhancing SLU performance using mul-
tiple ASR hypotheses](#). In *Proceedings of the 59th
Annual Meeting of the Association for Computational
Linguistics and the 11th International Joint Confer-
ence on Natural Language Processing*, pages 93–98.

Zhuo Gong, Daisuke Saito, Sheng Li, Hisashi Kawai,
and Nobuaki Minematsu. 2022. [Can we train a lan-
guage model inside an end-to-end ASR model? -
investigating effective implicit language modeling](#).
In *Proceedings of the Second Workshop on When
Creative AI Meets Conversational AI*, pages 42–47.

Jiatao Gu, Zhengdong Lu, Hang Li, and Victor O.K. Li.
2016. [Incorporating copying mechanism in sequence-
to-sequence learning](#). In *Proceedings of the 54th
Annual Meeting of the Association for Computational
Linguistics*, pages 1631–1640.

Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki
Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo
Wang, Zhengdong Zhang, Yonghui Wu, et al.
2020. [Conformer: Convolution-augmented trans-
former for speech recognition](#). *arXiv preprint
arXiv:2005.08100*.

Minglun Han, Linhao Dong, Shiyu Zhou, and Bo Xu.
2021. [Cif-based collaborative decoding for end-to-
end contextual speech recognition](#). In *2021 IEEE
International Conference on Acoustics, Speech and
Signal Processing (ICASSP)*, pages 6528–6532.

Takaaki Hori, Shinji Watanabe, and John Hershey. 2017.
[Joint CTC/attention decoding for end-to-end speech
recognition](#). In *Proceedings of the 55th Annual Meet-
ing of the Association for Computational Linguistics*,
pages 518–529.

686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739

740	Christian Huber, Juan Hussain, Sebastian Stüker, and Alexander Waibel. 2021. Instant one-shot word-learning for context-specific neural sequence-to-sequence speech recognition. In <i>2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)</i> , pages 1–7. IEEE.	<i>and Evaluation (LREC'12)</i> , pages 125–129, Istanbul, Turkey. European Language Resources Association (ELRA).	796 797 798
746	Mahaveer Jain, Gil Keren, Jay Mahadeokar, Geoffrey Zweig, Florian Metze, and Yatharth Saraf. 2020. Contextual rnn-t for open domain asr. <i>arXiv preprint arXiv:2006.03411</i> .	Anuroop Sriram, Heewoo Jun, Sanjeev Satheesh, and Adam Coates. 2017. Cold fusion: Training seq2seq models together with language models. <i>arXiv preprint arXiv:1708.06426</i> .	799 800 801 802
750	Suyoun Kim, Takaaki Hori, and Shinji Watanabe. 2017. Joint ctc-attention based end-to-end speech recognition using multi-task learning. In <i>2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)</i> , pages 4835–4839.	Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. <i>Advances in neural information processing systems</i> , 30.	803 804 805 806 807
755	Tian Lan, Deng Cai, Yan Wang, Heyan Huang, and Xian-Ling Mao. 2023. <i>Copy is all you need</i> . In <i>The Eleventh International Conference on Learning Representations</i> .	Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. <i>Pointer networks</i> . In <i>Advances in Neural Information Processing Systems</i> , volume 28.	808 809 810
759	Haoran Miao, Gaofeng Cheng, Changfeng Gao, Pengyuan Zhang, and Yonghong Yan. 2020. <i>Transformer-based online ctc/attention end-to-end speech recognition architecture</i> . In <i>2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)</i> , pages 6084–6088.	Tongtong Wu, Guitao Wang, Jinming Zhao, Zhaoran Liu, Guilin Qi, Yuan-Fang Li, and Gholamreza Haf-fari. 2022. <i>Towards relation extraction from speech</i> . In <i>Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing</i> , pages 10751–10762.	811 812 813 814 815 816
765	Motoi Omachi, Yuya Fujita, Shinji Watanabe, and Matthew Wiesner. 2021. <i>End-to-end ASR to jointly predict transcriptions and linguistic annotations</i> . In <i>Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies</i> , pages 1861–1871.	Song Xu, Haoran Li, Peng Yuan, Youzheng Wu, Xi-aodong He, and Bowen Zhou. 2020. <i>Self-attention guided copy mechanism for abstractive summarization</i> . In <i>Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics</i> , pages 1355–1362.	817 818 819 820 821 822
772	Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. <i>Librispeech: An asr corpus based on public domain audio books</i> . In <i>2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)</i> , pages 5206–5210.	Hemant Yadav, Sreyan Ghosh, Yi Yu, and Rajiv Ratn Shah. 2020. <i>End-to-end named entity recognition from english speech</i> . In <i>Interspeech 2020, 21st Annual Conference of the International Speech Communication Association, Virtual Event, Shanghai, China, 25-29 October 2020</i> , pages 4268–4272. ISCA.	823 824 825 826 827 828
777	Golan Pundak, Tara N Sainath, Rohit Prabhavalkar, Anjuli Kannan, and Ding Zhao. 2018. Deep context: end-to-end contextual speech recognition. In <i>2018 IEEE Spoken Language Technology Workshop (SLT)</i> , pages 418–425.	Zhengyi Zhang and Pan Zhou. 2022. End-to-end contextual asr based on posterior distribution adaptation for hybrid ctc/attention system. <i>arXiv preprint arXiv:2202.09003</i> .	829 830 831 832
782	Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2022. Robust speech recognition via large-scale weak supervision. <i>arXiv preprint arXiv:2212.04356</i> .	Ziqiang Zhang, Long Zhou, Junyi Ao, Shujie Liu, Lirong Dai, Jinyu Li, and Furu Wei. 2022. <i>SpeechUT: Bridging speech and text with hidden-unit for encoder-decoder based speech-text pre-training</i> . In <i>Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing</i> , pages 1663–1676.	833 834 835 836 837 838 839
786	Kanishka Rao, Haşim Sak, and Rohit Prabhavalkar. 2017. Exploring architectures, data and units for streaming end-to-end speech recognition with rnn-transducer. In <i>2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)</i> , pages 193–199.	Ding Zhao, Tara N Sainath, David Rybach, Pat Rondon, Deepti Bhatia, Bo Li, and Ruoming Pang. 2019a. Shallow-fusion end-to-end contextual biasing. In <i>2019 Conference of the International Speech Communication Association (Interspeech)</i> , pages 1418–1422.	840 841 842 843 844 845
792	Anthony Rousseau, Paul Deléglise, and Yannick Estève. 2012. <i>TED-LIUM: an automatic speech recognition dedicated corpus</i> . In <i>Proceedings of the Eighth International Conference on Language Resources</i>	Wei Zhao, Liang Wang, Kewei Shen, Ruoyu Jia, and Jingming Liu. 2019b. <i>Improving grammatical error correction via pre-training a copy-augmented architecture with unlabeled data</i> . In <i>Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies</i> , pages 156–165.	846 847 848 849 850 851 852

Appendices

A Datasets

Aishell (Bu et al., 2017) and ST-cmds⁶ are two widely used Chinese Mandarin datasets. Aishell contains about 150 hours of speech. ST-cmds was built based on commonly used online chatting and user command speeches, which contains about 110 hours of speech. For the English dataset, we utilize the portion of data that has been manually annotated with entities by Yadav et al. (2020), which comprises approximately 150 hours. The Eng dataset is built by extracting content from well-known English datasets, including Librispeech (Panayotov et al., 2015), CommonVoice⁷, Tedlium (Rousseau et al., 2012), and Voxforge⁸.

Table 6 shows the detailed statistics of the datasets used in our experiments. “Sent” means the number of instances. “NE” is the number of different named entities in the dataset and also the size of the contextual entity dictionary used during inference.

Dataset	Train		Dev		Test	
	Sent	NE	Sent	NE	Sent	NE
Aishell	119919	14241	14326	2194	7176	1186
Aishell-NE	119919	14241	4949	2194	2244	1186
ST-cmds	82080	17376	10260	3029	10260	3124
ST-cmds-NE	82080	17376	3285	3029	3241	3124
Eng	64570	11858	3100	2568	3100	2508
Eng-NE	64570	11858	2677	2568	2690	2508

Table 6: Statistics of the used datasets.

B Parameter Settings

We use 80-dimensional log-mel acoustic features with 25ms frame window and 10ms frame shift. The log-mel features are first fed into a 2D convolutional layer for downsampling and mapped to 256 dimensions before being inputted into the Audio Encoder. Both the audio encoder and decoder consist of 6 Transformer layers with 4 attention heads each. The NE Encoder is composed of three LSTM layers, with the input being a randomly initialized 256-dimensional embedding vector and the hidden size being 512. The relative weight λ in Equation 13 is set to 0.7. The experiments are conducted on two NVIDIA A100 GPUs.

⁶<https://www.openslr.org/38/>

⁷https://en.wikipedia.org/wiki/Common_Voice

⁸<https://en.wikipedia.org/wiki/VoxForge>

In addition, for the experiments on Whisper, we use Whisper-small model⁹, which includes 12 transformer layers in both its encoder and decoder, and is pre-trained on a total of 680,000 hours of multi-lingual and multi-task data. We replace our audio encoder and decoder with the Whisper model and fine-tune the parameters of the entire model on our training set for a maximum of 20 epochs. During fine-tuning, the initial learning rate for the Whisper model’s parameters is set to 1e-5, while the learning rate for other parameters is set to 1e-3 with 10,000 warm-up steps. During inference, we used beam search with a beam size of 5 and 10 for models with and without Whisper, respectively.

C Influence of Negative Entities in Training

During training, we construct an NE dictionary for each batch. To enhance the model’s ability of copying correct entities, we sample additional negative examples.

Suppose the dictionary already contains m entities, either real entities or pseudo sub-string entities. We sample $\beta \cdot m$ entities as negative examples from the training set. We utilize the parameter β to control the number of negative examples. Thus, we get the final dictionary for this batch which contains a total of $(\beta + 1) \cdot m$ entities. As shown in Figure 5, adding 1 or 2 times the number of negative samples can reduce transcription errors. Specifically, when $\beta = 2$, the CER and NE-CER decreased by 0.42% and 0.44% compared to $\beta = 0$. However, as β continues to increase, the error rate started to rise. We think that this is due to the presence of excessive noise. This causes the model to excessively focus on the negative samples, thus affecting its ability to accurately copy entities. Therefore, we set β to 2 during training.

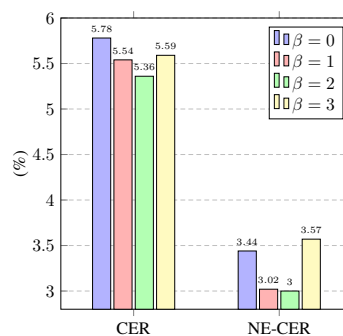


Figure 5: Effect of β .

⁹<https://huggingface.co/openai/whisper-small>