

Autocorrect in the Process of Translation — Multi-task Learning Improves Dialogue Machine Translation

Tao Wang^{1,2}, Chengqi Zhao¹, Mingxuan Wang¹, Lei Li¹, Deyi Xiong^{3*}

¹ByteDance AI Lab

²School of Computer Science and Technology, Soochow University, Suzhou, China

³College of Intelligence and Computing, Tianjin University, Tianjin, China

{wangtao.960826, zhaochengqi.d, wangmingxuan.89}@bytedance.com

{lilei.02}@bytedance.com

dyxiong@tju.edu.cn

Abstract

Automatic translation of dialogue texts is a much needed demand in many real life scenarios. However, current neural machine translation systems usually deliver unsatisfying translation results of dialogue texts. In this paper, we conduct a deep analysis of a dialogue corpus and summarize three major issues on dialogue translation, including pronoun dropping (ProDrop), punctuation dropping (PunDrop), and typos (DialTypo). In response to these challenges, we propose a joint learning method to identify omission and typo in the process of translating, and utilize context to translate dialogue utterances. To properly evaluate the performance, we propose a manually annotated dataset with 1,931 Chinese-English parallel utterances from 300 dialogues as a benchmark testbed for dialogue translation. Our experiments show that the proposed method improves translation quality by 3.2 BLEU over the baselines. It also elevates the recovery rate of omitted pronouns from 26.09% to 47.16%. The code and dataset are publicly available at <https://github.com/rgwt123/DialogueMT>.

1 Introduction

Remarkable progress has been made in Neural Machine Translation (NMT) (Bahdanau et al., 2015; Wu et al., 2016; Lin et al., 2020; Liu et al., 2020) in recent years, which has been widely applied in everyday life. A typical scenario for such application is translating dialogue texts, in particular the record of group chats or movie subtitles, which helps people of different languages understand cross-language chat and improve their comprehension capabilities.

However, traditional NMT models translate texts in a sentence-by-sentence manner and focus on the formal text input, such as WMT news translation

(1)	Nancy怎么了? [她] _{drop} 是不是哭了啊。
MT	What happened to Nancy? Did you cry?
REF	What happened to Nancy? Did she cry?
(2)	Nancy怎么了[?] _{drop} 是不是哭了啊。
MT	Did Nancy cry?
REF	What happened to Nancy? Did she cry?
(3)	Nancy怎么[乐] _{typo} ?
MT	How happy is Nancy?
REF	What happened to Nancy?

Table 1: Examples of ProDrop (1), PunDrop (2) and DialTypo (3). MT is translation results from Google Translate while REF is references.

(Barrault et al., 2020), while the translation of dialogue must take the meaning of context and the input noise into account. Table 1 shows examples of dialogue fragment in Chinese and their translation in English. Example (1) demonstrates that the omission in traditional translation (e.g., dropped pronouns in Chinese) leads to inaccurate translation results.

Despite its vast potential application, efforts of exploration into dialogue translation are far from enough. Existing works (Wang et al., 2016; Maruf et al., 2018) focus on either extracting dialogues from parallel corpora, such as OpenSubtitles (Lison et al., 2019), or leveraging speaker information for integrating dialogue context into neural models. Also, the lack of both training data and benchmark test set makes current dialogue translation models far from satisfying and need to be further improved.

In this paper, we try to alleviate the aforementioned challenges in dialogue translation. We first analyze a fraction of a dialogue corpus and summarize three critical issues in dialogue translation, including ProDrop, PunDrop, and DialTypo. Then we design a Multi-Task Learning (MTLDIAL) approach that learns to self-correct sentences in the process of translating. The model’s

*Corresponding author.

encoder part automatically learns how to de-noise the noise input via explicit supervisory signals provided by additional contextual labeling. We also propose three strong baselines for dialogue translation, including repair (REPAIRDIAL) and robust (ROBUSTDIAL) model. To alleviate the challenges arising from the scarcity of dialogue data, we use sub-documents in the bilingual parallel corpus to enable the model to learn from cross-sentence context.

Additionally as for evaluation, the most commonly used BLEU metric (Papineni et al., 2001) for NMT is not good enough to provide a deep look into the translation quality in such a scenario. Thus, we build a Chinese-English test set containing sentences with the issues in `ProDrop`, `PunDrop` and `DialTypo`, attached with the human translation and annotation. Finally, we get a test set of 300 dialogues with 1,931 parallel sentences.

The main contributions of this paper are as follows: a) We analyze three challenges `ProDrop`, `PunDrop` and `DialTypo`, which greatly impact the understanding and translation of a dialogue. b) We propose a contextual multi-task learning method to tackle the analyzed challenges. c) We create a Chinese-English test set specifically containing those problems and conduct experiments to evaluate proposed method on this test set.

2 Analysis on Dialogue Translation

There were already some manual analyses of translation errors, especially in the field of discourse translation. Voita et al. (2019) study English-Russian translation and find three main challenges for discourse translation: deixis, ellipsis, and lexical cohesion. For Chinese-English translation, tense consistency, connective mismatch, and content-heavy sentences are the most common issues (Li et al., 2014).

Different from previous works, we mainly analyze the specific phenomena in dialogue translation. We begin with a study on a bilingual dialogue corpus (Wang et al., 2018).¹ We translate source sentences into the target language at sentence level and compare translation results with reference at dialogue level. Around 1,000 dialogues are evaluated, and the results are reported in Table 2. From the statistic, we observe two persistent dialogue translation problems: pronoun dropping (`ProDrop`), punctuation drop-

Types of phenomena	Frequency
Correct	88.1%
<code>ProDrop</code>	4.3%
<code>PunDrop</code>	3.2%
Incorrect segmentation	2.4%
Other translation errors	2.0%

Table 2: Manual evaluation of dialogue samples.

ping(`PunDrop`). The phenomenon is consistent with the issue we collect in practical Instant Messaging (IM) chat scenarios, except for typos since the analyzed dialogue corpus has been proofread to remove typos.

2.1 Pronoun Dropping

Pronouns are frequently omitted in pro-drop languages (Huang, 1989), such as Chinese, Japanese, Korean, Vietnamese, and Slavic languages. Such phenomenon are more frequent in dialogue, where the interlocutors are both aware of what’s omitted in the context. However, when translating a pro-drop language into a non-pro-drop language (e.g., English)², it is hard to translate those omitted pronouns, resulting in grammatical errors or semantic inaccuracies in the target language. The first conversation in Table 1 is an example.

2.2 Punctuation Dropping

In dialogue scenarios, such as IM software, punctuation is often omitted and users tend to segment sentences with spaces. The problem becomes much serious in languages with no spaces, such as Chinese, Japanese, Korean, and Thai. Table 1 shows this phenomenon in Example (2).

2.3 Dialogue Typos

Typo repairing is another fundamental but very challenging practical problem. In dialogue translation, typos or misspellings are very common, which dramatically undermine the quality of translation output produced by machine translation. Table 1 shows this phenomenon in Example (3).

3 Approach to NMTDIAL

This section aims to propose a unified framework that facilitates NMT to correct noisy inputs in dialogue neural machine translation (NMTDIAL). The framework includes three different methods, which are REPAIRDIAL, ROBUSTDIAL and MTL-DIAL.

¹<https://github.com/longyuewangdcu/tvsu>

²https://en.wikipedia.org/wiki/Pro-drop_language

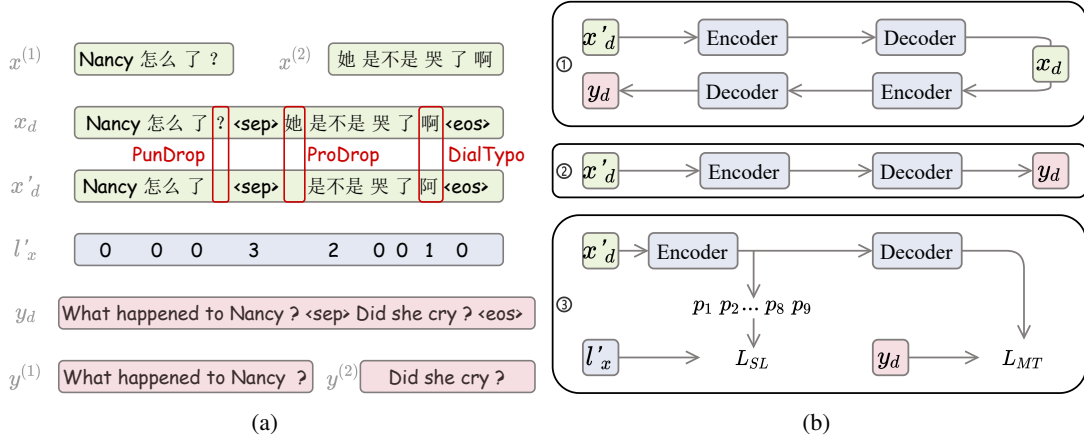


Figure 1: Overall diagram of NMTDIAL. (a) demonstrates the process of data generation, and (b) displays the three proposed methods. ①/②/③ represent REPAIRDIAL, ROBUSTDIAL and MTLDIAL respectively.

3.1 Contextual Perturbation Example Generation

The most challenging problem for NMTDIAL is the data distribution gap between training and inference stage, where the training data are clean sentence-level pairs while the test data are noisy dialogue-level conversations.

To bridge the distribution gap, the first step is to generate perturbation examples based on training instances. The data generation mainly consists of two steps. The first step is to obtain sub-documents with cross-sentence context, and the second step is to generate examples with word perturbations within sub-documents. Figure 1a shows a complete process.

Cross-sentence Context It is difficult to acquire dialog-level parallel training data. As an alternative approach, we use parallel document data to catch dependencies across sentences.

Formally, let $x_d = \{x^{(1)}, x^{(2)}, \dots, x^{(M)}\}$ be a source-language document containing M source sentences. And $y_d = \{y^{(1)}, y^{(2)}, \dots, y^{(M)}\}$ is the corresponding target-language document containing the same number of sentences as that of the source document. To get more context information, we randomly sample consecutive sub-document pairs (x_d, y_d) of N sentences (i.e., snippet pairs from aligned documents). We set $N \in [1, 10]$ in this paper.

We use a special token $\langle \text{sep} \rangle$ as the separator to concatenate sentences into a parallel sub-document $\{(x_d, y_d)\}$, as shown in Figure 1a.

Contextual Perturbation We then consider generating perturbation example x'_d from x_d with respect to sub-document context. For ProDrop,

PunDrop and DialTypo, we build a Chinese pronoun table T_{ProDrop} , a common punctuation table T_{PunDrop} and a Chinese homophone table T_{DialTypo} respectively.

For ProDrop and PunDrop, we traverse source sentences of x_d , discard pronouns/punctuation in these sentences with a probability of 30% and record deletion positions with corresponding labels (see details below); to construct a typo, we choose a word with a probability of 1%, of which 80% is replaced with one of its homophones according to T_{DialTypo} and 20% is replaced with another random word. We determine these percentages by observing the generated perturbation data. For annotation labels, we tag correct words with 0, words of DialTypo with 1, ProDrop words with 2 and PunDrop words with 3.

Finally we get x_d , x'_d and their corresponding label sequences ℓ_x, ℓ'_x . ℓ_x is a sequence of all 0s.

3.2 NMTDIAL Base Models

With the created training data, we first introduce two methods for NMTDIAL as our strong baselines, which will be elaborated here for model comparison.

REPAIRDIAL A natural way for NMTDIAL is to train a dialog repair model to transform dialogue inputs into forms that an ordinary NMT system can deal with. REPAIRDIAL involves training a repair model to transform x'_d to x_d and a clean translation model that translates x_d to y_d . As a pipeline method, REPAIRDIAL may suffer from error propagation.

ROBUSTDIAL We extend the robust NMT

(Cheng et al., 2018) to dialogue-level translation. Specifically, we take both the original (x_d, y_d) and the perturbed (x'_d, y_d) bilingual pairs as training instances. So the model is more resilient on dialogue translation. During the inference stage, the robust model directly translates raw inputs into the target language.

3.3 MTLDIAL

ROBUSTDIAL has the potential to handle translation problems caused by noisy dialogue inputs. However, the internal mechanism is rather implicit and in a black box. Therefore, the improvement is limited, and it is not easy to analyze the improvement. To address this issue, we introduce a context-aware multi-task learning method MTLDIAL for NMTDIAL.

As shown in ③ of Figure 1b, the only difference is that we have a contextual labeling module based on the encoder. We denote the final layer output of the Transformer encoder as H . For each token h_i in $H = (h_1, h_2, \dots, h_m)$, the probability of contextual labeling is defined as:

$$P(p_i = j|X) = \text{softmax}(W \cdot h_i + b)[j] \quad (1)$$

where $X = (x_1, x_2, \dots, x_m)$ is the input sequence, $P(p_i = j|X)$ is the conditional probability that token x_i is labeled as j ($j \in 0, 1, 2, 3$ as defined above).

Here we make the labeling module as simple as possible, so that the Transformer encoder can behave like BERT (Devlin et al., 2019), learning more information related to perturbation and guiding the decoder to find desirable translations.

During the training phrase, the model takes $(x_d, x'_d, \ell_x, \ell'_x, y_d)$ as the training data. The learning process is driven by optimizing two objectives, corresponding to sequence labeling as auxiliary loss (\mathcal{L}_{SL}) and machine translation as the primary loss (\mathcal{L}_{MT}) in a multi-task learning framework.

$$\mathcal{L}_{SL} = -\log(P(\ell_x|x_d) + P(\ell'_x|x'_d)) \quad (2)$$

$$\mathcal{L}_{MT} = -\log(P(y_d|x_d) + P(y_d|x'_d)) \quad (3)$$

The two objective are linearly combined as the overall objective in learning.

$$\mathcal{L} = \mathcal{L}_{MT} + \lambda \cdot \mathcal{L}_{SL} \quad (4)$$

λ is coefficient. During experiments, we set as follows according the best practice:

$$\lambda = \max(1.0 - \frac{\text{update_num}}{10^5}, 0.2) \quad (5)$$

where update_num is the number of updating steps during training.

We introduce multi-task learning for two reasons: 1) The labeling performance reflects the model’s understanding of sentences containing the mentioned phenomena. 2) Contextual Labeling can be seen as a pre-training process based on the BERT-like model, and explicit guidance can enable the encoder to learn more about the information we annotate.

3.4 Modeling Dialogue Context

The modes for exploring dialogue context during decoding can be divided into offline and online. For the offline setting, all sentences in a dialogue are concatenated one by one with $\langle \text{sep} \rangle$. The concatenated sequence is translated, and the target translation for each sentence can be easily detected according to the separator $\langle \text{sep} \rangle$.

The offline mode can be used for dialogue translation where the entire source dialogue has already been available before translation (e.g., movie subtitles). However, we continuously get new source sentences for online chat and need to generate corresponding translations immediately. We refer to this mode as the online setting.

We experiment with two online methods. One is *online-cut* where the current sentence is concatenated to the previous context with the separator $\langle \text{sep} \rangle$. The trained NMTDIAL model then translates the concatenated sequence and the last target segment is used as the translation for the current source sentence. The other is *online-fd*. *Online-fd* is a force decoding method. It forces the decoder to use translated history and continues decoding instead of re-translating the entire concatenated sequence. *Online-fd* brings more consistent translation.

4 Experiments

4.1 Test Set

For better evaluation of NMTDIAL, we create a Chinese-English test set covering all issues discussed above based on the corpus we analyze in the second section. Statistics on the built test set are displayed in Table 3. Building such a test set is hard and time-consuming as we need to perform manual selection, translation and annotation.

As for translation quality evaluation, we use other metrics in addition to BLEU. For `PunDrop` and `DialTypo`, we evaluate BLEU scores on sen-

Item	Count
#dialogues	300
#sentence pairs	1,931
#total tokens	19,155/15,976
#average tokens	9.92/8.27
#ProDrop	299
#PunDrop	542
#DialTypo	203

Table 3: Statistics on the test set. “/” denote numbers in Chinese and English separately.

tences containing missing punctuation or typos according to the annotation information. As for ProDrop, we evaluate the translation quality by the percentage of correctly recovering and translating the dropped pronouns.

4.2 Settings

We adopt the Chinese-English corpus from WMT2020³, with about 48M sentence pairs, as our bilingual training data D . We select newstest2019 as the development set. After splicing, we get D_{doc} with 1.2M pairs and corresponding perturbed dataset D' and D'_{doc} with 48M and 1.2M pairs respectively.

We use byte pair encoding compression algorithm (BPE) (Sennrich et al., 2016) to process all these data and limit the number of merge operations to a maximum of 30K. In our studies, all translation models are Transformer-big, including 6 layers for both encoders and decoders, 1024 dimensions for model, 4096 dimensions for FFN layers and 16 heads for attention.

During training, we use label smoothing = 0.1 (Szegedy et al., 2016), attention dropout = 0.1 and dropout (Hinton et al., 2012) with a rate of 0.3 for all other layers. We use Adam (Kingma and Ba, 2015) to train the NMT models. β_1 and β_2 of Adam are set to 0.9 and 0.98, the learning rate is set to 0.0005, and gradient norm 5. The models are trained with a batch size of 32,000 tokens on 8 Tesla V100 GPUs during training. During decoding, we employ beam search algorithm and set the beam size to 5. We use sacrebleu (Post, 2018) to calculate uncased BLEU-4 (Papineni et al., 2001).

4.3 Results of Offline Setting

The offline mode aims at using the entire source dialogue for translation. We experiment with all the methods in the offline setting, and the results

³This corpus includes News Commentary, Wiki Titles, UN Parallel Corpus, CCMT Corpus, WikiMatrix and Back-translated news.

Methods	Overall BLEU	Details		
		ProDrop	PunDrop	DialTypo
BASE	32.7	26.09%	28.2	24.0
REPAIRDIAL	34.0	29.77%	31.2	27.4
ROBUSTDIAL	34.1	45.48%	33.0	28.8
MTLDIAL	35.9	47.16%	34.3	28.7
GOLD+BASE	36.8	97.32%	34.6	36.8

Table 4: Experiment results on our constructed dialogue translation test set in offline setting. The GOLD+BASE represents translations of completely correct inputs (without ProDrop, PunDrop or DialTypo) using BASE model, which is used to show the oracle results with Transformer on the test set.

Methods	Overall BLEU	Details		
		ProDrop	PunDrop	DialTypo
BASE	32.8(+0.1)	19.06%(-7.03%)	28.1(-0.1)	22.3(-1.7)
REPAIRDIAL	33.8(-0.2)	24.75%(-5.02%)	32.0(+0.8)	28.3(+0.9)
ROBUSTDIAL	34.2(+0.1)	36.79%(-8.69%)	32.7(-0.3)	28.9(-0.5)
MTLDIAL	35.3(-0.6)	34.78%(-12.38%)	34.3(-0.0)	28.6(-0.1)
GOLD+BASE	37.1(+0.3)	96.66%(-0.66%)	35.3(+0.7)	35.9(-0.9)

Table 5: Results on our constructed dialogue translation test set in online setting at the sentence level.

are shown in Table 4. BASE is a Transformer-big model trained with D and D_{doc} . GOLD+BASE represents the oracle result on this test set. We can see that MTL DIAL has achieved the best results, reducing the gap between $test_{wrong}$ and $test_{gold}$ from 4.1 to 0.9. Compared with ROBUSTDIAL and MTL DIAL, REPAIRDIAL performs relatively poorly. We believe that this is due to the error propagation caused by the pipeline.

From the specific indicators, we can draw the following conclusions: 1) DialTypo has a very obvious impact on BLEU, and the gap between BASE and GOLD+BASE is more than 12 points; 2) The recovery of ProDrop is a relatively difficult task. Although compared with BASE, the current best result of 47.16% has been greatly improved, but is still far away from the golden result 97.32%; 3) PunDrop seems to be a relatively easy task for each method to address.

4.4 Results of Online Setting

The online mode only makes use of previous context during translation. An extreme situation of online setting is that there is no context, that is, sentence-level translation. We show the results of all the methods on the test set at the sentence level in Table 5. Despite the lack of context, our approaches can still bring general benefits. We find that ProDrop relies heavily on context, especially for MTL DIAL, where the absence of context results in a 12.38% drop in performance. This is in line

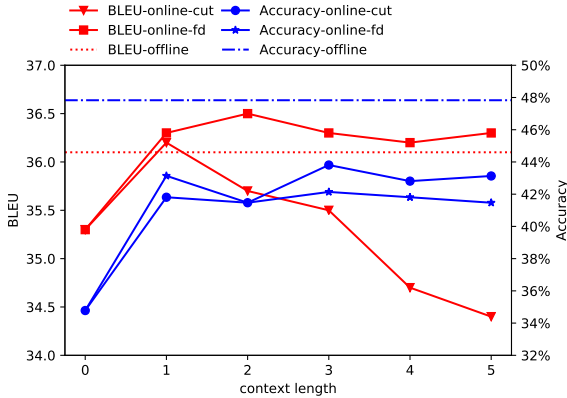


Figure 2: Overall BLEU and ProDrop recovery performance (Accuracy) of MTLdIAL with different context length. Dash lines are the offline results.

Data		Precision	Recall	F1
validation	ProDrop	61.3	48.7	54.3
	PunDrop	80.0	63.6	70.9
	DialTypo	85.3	64.2	73.2
test	ProDrop	48.6	32.2	38.8
	PunDrop	96.6	87.9	92.1
	DialTypo	83.3	31.0	45.2

Table 6: Labeling performance on the validation/test set.

with our expectations, as in many cases machine translation system heavily depends on context to fulfill the dropped pronouns.

We further experiment on how context lengths can affect NMTdIAL. The results are shown in Figure 2. In the *online-cut* setting, we can see that using previous few sentences as context may improve overall BLEU score, but continuously adding more preceding texts will lead to a continuous decline. *Online-fd* performs well because using historical translation records to continue decoding can bring more consistent translation results. For the recovery accuracy of ProDrop, *online-cut* is better than *online-fd* in contrast, because forced decoding may cause wrong pronoun transmission.

5 Analysis

5.1 Labeling Performance

To better understand how our proposed MTLdIAL make sense, we calculate the labeling performance on both validation and test set. Table 6 shows the overall performance. The validation set follows the same processing progress of training data, while the test set is the real dialogue data set built manually.

The proposed model obtains 54.3% F1 score on the validation set for ProDrop, 70.9% for

	zh	en
(1)	艾丽最近怎么样了 她已经不在我的律所了 什么(她/she)为什么走了 (她/she)开了自己的律所	What's going on with Ellie? She is no longer in my law firm. What, why are you/is she going? I open my/She opens her own law firm.
(2)	琼斯(我/I)问你件事	Jones asked/ I want to ask you something.
(3)	他上次帮我私下搞 (他/He)差点工作都丢了	He helped me out in private last time. I/He nearly lost my/his job.

Table 7: Examples of ProDrop recovery errors.




Figure 3: ProDrop recovery performance of BASE and contextual MTLdIAL. Total means the total number of occurrence of corresponding pronouns in the test set. We ignore pronouns with a total occurrence number less than 5.

PunDrop, and 73.2% for DialTypo. When testing on the real test data, the performance on ProDrop has declined a lot because of the difference between synthetic training/validation data and real test data. Especially noteworthy is the fact that F1 score of DialTypo drops the most, reaching 26%, because of its low recall. It may be due to the considerable difference between the typos generated by our automatic method and the actual distribution.

5.2 Effects of Pronoun Correcting

We further explore the auto-correction of specific pronouns. As shown in Figure 3, we can find that pronouns such as I/you, which occur mostly in the corpus, generally have a higher recovery success rate. We believe this is due to the data imbalance. Compared with BASE, MTLdIAL has a much better performance. While ProDrop recovery accuracy has been improved, it still has not achieved 50%. The most common error is that the model does not capture any context or captures previous inappropriate context. We summarize frequently-occurring recovery errors in Table 7.

6 Related Work

Our work is related with both dialogue translation and robust training.

Dialogue Translation

There has been some work on building bilingual dialogue data sets for the translation task in recent years. Wang et al. (2016) propose a novel approach to automatically construct parallel discourse corpus for dialogue machine translation and release around 100K parallel discourse data with manual speaker and dialogue boundary annotation. Maruf et al. (2018) propose the task of translating Bilingual Multi-Speaker Conversations. They introduce datasets extracted from Europarl and Opensubtitles and explore how to exploit both source and target-side conversation histories. Bawden et al. (2019) present a new English-French test set for evaluating of Machine Translation (MT) for informal, written bilingual dialogue. Recently WMT2020 has also proposed a new shared task - machine translation for chats,⁴ focusing on bilingual customer support chats (Farajian et al., 2020).

Robust Training

Neural models have been usually affected by noisy issues. Many efforts (Li et al., 2017; Sperber et al., 2017; Vaibhav et al., 2019; Yang et al., 2020) focus on data augmentation to alleviate the problem by adding synthetic noise to the training set. However, generating noise has always been a challenge, as natural noise is always more diversified than artificially constructed noise (Belinkov and Bisk, 2018; Anastasopoulos, 2019; Anastasopoulos et al., 2019).

7 Conclusions

In this paper, we manually analyze challenges in dialogue translation and detect three main problems. In order to tackle these issues, we propose a multi-task learning method with contextual labeling. For deep evaluation, we construct dialogues with translation and detailed annotations as a benchmark test set. Our proposed model achieves substantial improvements over the baselines. What is more, we further analyze the performance of contextual labeling and pronoun recovery errors.

Acknowledgments

We thank the bilingual speakers for test set construction, and the anonymous reviewers for suggestions. Deyi Xiong is partially supported by the Natural Science Foundation of Tianjin (Grant No. 19JCZDJC31400) and the Royal Society (London) (NAF\R1\180122).

⁴<http://www.statmt.org/wmt20/chat-task.html>

References

- Antonios Anastasopoulos. 2019. An analysis of source-side grammatical errors in nmt. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 213–223.
- Antonios Anastasopoulos, Alison Lui, Toan Q Nguyen, and David Chiang. 2019. Neural machine translation of text from non-native speakers. In *NAACL-HLT (1)*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations, ICLR 2015*.
- Loïc Barrault, Magdalena Biesialska, Ondřej Bojar, Marta R Costa-jussà, Christian Federmann, Yvette Graham, Roman Grundkiewicz, Barry Haddow, Matthias Huck, Eric Joanis, et al. 2020. Findings of the 2020 conference on machine translation (wmt20). In *Proceedings of the Fifth Conference on Machine Translation*, pages 1–55.
- Rachel Bawden, Sophie Rosset, Thomas Lavergne, and Eric Bilinski. 2019. Diabla: A corpus of bilingual spontaneous written dialogues for machine translation. *arXiv preprint arXiv:1905.13354*.
- Yonatan Belinkov and Yonatan Bisk. 2018. Synthetic and natural noise both break neural machine translation. In *International Conference on Learning Representations*.
- Yong Cheng, Zhaopeng Tu, Fandong Meng, Junjie Zhai, and Yang Liu. 2018. Towards robust neural machine translation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 1756–1766. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT (1)*.
- M Amin Farajian, António V Lopes, André FT Martins, Sameen Maruf, and Gholamreza Haffari. 2020. Findings of the wmt 2020 shared task on chat translation. In *Proceedings of the Fifth Conference on Machine Translation*, pages 65–75.
- Geoffrey E. Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan R. Salakhutdinov. 2012. Improving neural networks by preventing co-adaptation of feature detectors. *CoRR*, abs/1207.0580.
- CT James Huang. 1989. Pro-drop in chinese: A generalized control theory. In *The null subject parameter*, pages 185–214. Springer.

- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of ICLR*.
- Junyi Jessy Li, Marine Carpuat, and Ani Nenkova. 2014. Assessing the discourse factors that influence the quality of machine translation. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 283–288.
- Yitong Li, Trevor Cohn, and Timothy Baldwin. 2017. Robust training under linguistic adversity. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 21–27.
- Zehui Lin, Xiao Pan, Mingxuan Wang, Xipeng Qiu, Jiangtao Feng, Hao Zhou, and Lei Li. 2020. Pre-training multilingual neural machine translation by leveraging alignment information. *arXiv preprint arXiv:2010.03142*.
- Pierre Lison, Jörg Tiedemann, Milen Kouylekov, et al. 2019. Open subtitles 2018: Statistical rescoring of sentence alignments in large, noisy parallel corpora. In *LREC 2018, Eleventh International Conference on Language Resources and Evaluation*. European Language Resources Association (ELRA).
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. Multilingual denoising pre-training for neural machine translation. *Transactions of the Association for Computational Linguistics*, 8:726–742.
- Sameen Maruf, André FT Martins, and Gholamreza Haffari. 2018. Contextual neural model for translating bilingual multi-speaker conversations. *WMT 2018*, page 101.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2001. Bleu: a method for automatic evaluation of machine translation. In *ACL*.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Belgium, Brussels. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1715–1725.
- Matthias Sperber, Jan Niehues, and Alex Waibel. 2017. Toward robust neural machine translation for noisy input sequences. In *International Workshop on Spoken Language Translation (IWSLT)*.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826.
- Vaibhav Vaibhav, Sumeet Singh, Craig Stewart, and Graham Neubig. 2019. Improving robustness of machine translation with synthetic noise. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1916–1920.
- Elena Voita, Rico Sennrich, and Ivan Titov. 2019. When a good translation is wrong in context: Context-aware machine translation improves on deixis, ellipsis, and lexical cohesion. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1198–1212.
- Longyue Wang, Zhaopeng Tu, Shuming Shi, Tong Zhang, Yvette Graham, and Qun Liu. 2018. Translating pro-drop languages with reconstruction models. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Longyue Wang, Xiaojun Zhang, Zhaopeng Tu, Andy Way, and Qun Liu. 2016. Automatic construction of discourse corpora for dialogue translation.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- Jiacheng Yang, Mingxuan Wang, Hao Zhou, Chengqi Zhao, Weinan Zhang, Yong Yu, and Lei Li. 2020. Towards making the most of bert in neural machine translation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 9378–9385.