



How Does Pretraining Improve Discourse-Aware Translation?

Zhihong Huang^{1*} Longyue Wang^{2*} Siyou Liu Derek F. Wong¹

¹NLP²CT Lab, University of Macau

²Tencent AI Lab

nlp2ct.zhihong@gmail.com, vinnylywang@tencent.com, helen.liu103@gmail.com,
derekfw@um.edu.mo

Abstract

Pretrained language models (PLMs) have produced substantial improvements in discourse-aware neural machine translation (NMT), for example, improved coherence in spoken language translation. However, the underlying reasons for their strong performance have not been well explained. To bridge this gap, we introduce a probing task to interpret the ability of PLMs to capture discourse relation knowledge. We validate three state-of-the-art PLMs across encoder-, decoder-, and encoder-decoder-based models. The analysis shows that (1) the ability of PLMs on discourse modelling varies from architecture and layer; (2) discourse elements in a text lead to different learning difficulties for PLMs. Besides, we investigate the effects of different PLMs on spoken language translation. Through experiments on IWSLT2017 Chinese-English dataset, we empirically reveal that NMT models initialized from different layers of PLMs exhibit the same trends with the probing task. Our findings are instructive to understand how and when discourse knowledge in PLMs should work for downstream tasks.

Index Terms: spoken Language, discourse, pretrained language models, machine translation, linguistic probing

1. Introduction

Translating spoken language is a significantly challenging task due to its inherent characteristics such as irregular expressions and discourse properties [1, 2, 3]. In recent years, discourse-aware neural machine translation (NMT) has performed better by initializing the Transformer-based [4] models with pretrained language models (PLMs) in encoder [5], decoder [6], both [7] or themselves [8]. The common assumption is that NMT models can utilize rich knowledge from PLMs to tackle complex discourse phenomena [9]. For example, some works found that better-translated results often contain more connective words [8], which can be classified as explicit with the non-tree-structure shallow discourse relations [10]. Table 1 shows an example of discourse-aware translation.

However, it is still unclear how discourse knowledge is embedded in PLMs, and when PLMs are leveraged in discourse-level NMT. Towards the understanding of PLMs, probing tasks are exploited to provide fine-grained analysis of model ability [11]. Related works either probed different PLMs about tree-structure discourse knowledge based on rhetorical structure theory (RST) [12, 13], or probed multilingual PLMs about part of the discourse relations [14]. Therefore, a more comprehensive study on the effects of discourse knowledge in PLMs on NMT is needed, particularly for spoken language translation.

* Longyue Wang and Zhihong Huang contributed equally to this work.

Table 1: An example of discourse-aware translation from IWSLT2017 dataset. The *token* is an explicit connective word while © is implicit that is invisible to models. “Inp.” and “Out.” represent the Chinese input and English translation, respectively. As seen, a coherent translation should accurately transfer discourse relation between sentences from source to target language (e.g., *words*).

Inp.	有一个志愿者在他呆的村庄里发明了这个小玩意儿 ©就是把废纸压缩成块状用来燃烧 但是它烧得很慢
Out.	This is one such volunteer and this is a device that he had built in the village where he worked. And the idea was that you could take waste paper; you could compress it and make briquettes that could be used for fuel. But this device was very slow.

To bridge this gap, we propose a method to probe the ability of advanced PLMs (i.e., BERT [15], BART [16], and GPT-2 [17]) to capture discourse knowledge. Analysis results on PDTB dataset [18] demonstrate that encoder-decoder-based PLMs perform best especially on higher layers (except for GPT-2). About the translation tasks, we leverage PLMs to discourse-aware NMT by utilizing part/all of their parameters. Experiment results on the Chinese-English IWSLT2017 dataset [19] show that (i) PLMs in the source language help more than that in the target language; (ii) NMT prefers PLMs with the same architecture (i.e., BART); (iii) NMT initialized with the single discourse-aware layer can achieve a close performance to using all layers; (iv) translation performance exhibits the same trends with the probing task at the layer level.

2. Methodology

2.1. Probing Discourse Knowledge

Our probing tasks mainly focus on the shallow discourse relation in a sentence with two semantic arguments [20], rather than consider the RST relations in several sentences. The shallow discourse relations can be characterized into five types: (1) *Explicit* relation means that the connective words in the sentence are visible; (2) *Implicit* relation means that the sentence has no connective words but can be annotated manually; when the sentence has no connective word but shows a discourse relation by its expressions or entities, it contains (3) *AltLex* or (4) *EntRel* relation; (5) *NoRel* relation means there is no discourse relation in the sentence. Further, *Explicit*, *Implicit*, and *AltLex* relations




Figure 1: The framework of our proposed probing task. The input is two adjacent sentences in a document, which are split into two spans (i.e., Argument 1 and 2) with explicit/implicit connective words. Second, we extract corresponding embeddings by feeding the input to a pretrained model. Third, the embeddings are used to train an MLP model for learning to classify the discourse relation between Argument 1 and 2. The classification accuracy is used to reflect the ability of PLM on modelling discourse.

can be annotated with sense in three levels: class, type, and subtype. We list three examples below, whose words underlined are connective words, highlighted in italics are argument-1, and in boldface are argument-2.

- (1) *It was a far safer deal for lenders since **NWA had a healthier cash flow and more collateral on hand.***
- (2) *Some have raised their cash positions to record levels. (Implicit: BECAUSE) **High cash positions help buffer a fund when the market falls.***
- (3) *Ms. Bartlett's previous work, which earned her an international reputation in the non-horticultural art world, often took gardens as its nominal subject. (AltLex) **Mayhap this metaphorical connection made the BPC Fine Arts Committee think she had a literal green thumb.***

In the sentence (2), the manual-annotation connective word BECAUSE has an *Implicit* relation, a *Contingency* class, a *Cause* type, and a *Reason* subtype. This kind of fine-grained sense is the discourse relation label of the sentences in our tasks dataset.

According to the shallow discourse relation, we propose two probing tasks to assess PLMs' ability to encode the discourse relations in different dimensional views: (1) To interpret the overall ability of PLMs, we probe the whole sentence embeddings in the complete dataset, including the four discourse relations. In particular, we adopt two forms of BERT representation to extract the sentences: *[CLS]* embeddings and average pooling embeddings of all tokens. (2) To determine if the PLMs distribute the discourse knowledge in different linguistic elements, we individually probe the embeddings of connective words and sentiment arguments after obtaining the whole-sentence embeddings in the explicit data. The whole-sentence embeddings of implicit and AltLex data are also probed to compare with the explicit data.

The overview of the probing tasks' structure is shown in Figure 1. First, we extract the representations of the whole sentences with the frozen-parameters PLMs. Then we feed the different combinations of embeddings into two-layer MLPs (Multilayer Perceptron) to train probing models with the parameters updated. The probing models are classifiers, and the labels are fine-grained sense relations of the input sentence. All layers of three PLMs (i.e., BERT, GPT-2, and BART) are probed separately. We evaluated the classification accuracy of the probing

Table 2: Data statistics of the datasets on discourse probing and machine translation tasks. We calculate the sizes of training, validation, and testing sets in terms of sentence number. The number in brackets denotes the size of instances with explicit discourse relation.

Task	Probing	Translation
Dataset	PDTB2.0 [18]	IWSLT2017 [19]
Train	32,535 (18,459)	231,266
Valid	1,436 (812)	879
Test	1,928 (1,090)	6,046

models. Specifically, we treat BART's encoder layers as its first to sixth model layers, and decoder layers as its seventh to twelfth model layers. The layer with a higher accuracy means that it is more capable of embedding the discourse knowledge.

2.2. Discourse-Aware NMT with Pretraining

Inspired by the work of Rothe et al.[6], we adopt the following strategies to leverage the PLMs to Transformer-based NMT models: (1) For BERT models, we initialize the encoder with the PLMs and randomly initialize the decoder. (2) For GPT-2 models, we initialize the decoder with the PLMs and randomly initialize the encoder. (3) For BART models, we use them directly as a sequence-to-sequence model. The three models are trained on Chinese data. We also exploit Chinese and English versions for all three types of PLMs, and multilingual-BERT to investigate whether discourse knowledge in the source or target language is more significant to discourse-aware translation. For a fair comparison, all models have similar size of parameters.

We employ a document-level NMT model for the translation task, which can utilize context in the source language. Following [21], we use one previous source sentence in the document as context information when translating each current sentence. Taking Table 1 for instance, the first Chinese sentence can be used as context for translating the second one. When the translation model can make better use of context information, that is, the translation of the current text contains more complete discourse knowledge, then its quality will be better. We consider such

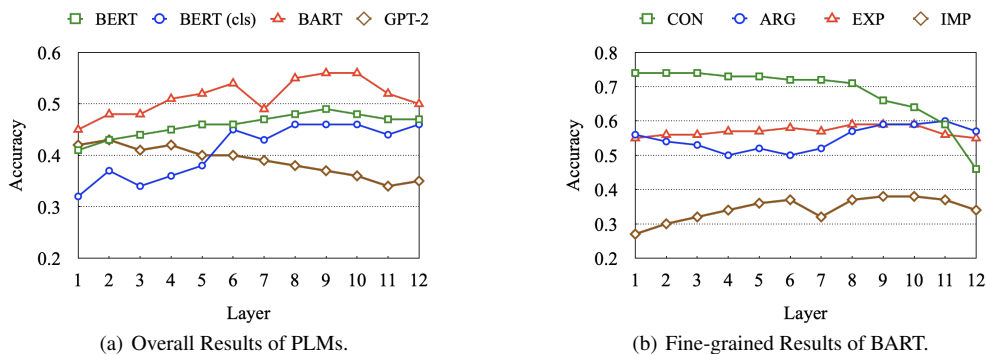


Figure 2: Results of probing tasks for evaluating discourse properties embedded in different layers of representations in different pretrained models. (a) shows the overall results of three PLMs on the complete dataset. (b) shows the fine-grained results of BART by feeding parts of embeddings into the probing model. Note that, “CON” means only connective words while “ARG” indicates two arguments without connective words. We also report results on parts of data with explicit (“EXP”), and implicit (“IMP”) discourse relations, respectively, including connective words.

a translation task to be a discourse-aware NMT. Note that the sentence number of translation input is also consistent with that in the probing task, making the results of the two experiments comparable. We use BLEU [22], TER [23], and METEOR [24] scores to measure the performance of the translation systems.

Besides, in a fine-grained view, we propose a method to investigate how PLMs perform in each layer when they are utilized for NMT. We train special discourse-aware NMT tasks by only updating the parameters of a specific layer each time, while the other layers’ are frozen. Based on the results of the probing tasks, we select the first, the middle, the last layer (i.e., layers 1, 6, 12), and the discourse-aware as the specific layer. We wonder if the performance of NMT models with a single updated layer is related to the probing results. In this additional task, we only leverage the English-version PLMs, and keep the rest of the experiment settings the same as the normal discourse-aware translation task.

3. Experiments

3.1. Experimental Setup

We summarize all data used in experiments in Table 2. For probing tasks, we conduct experiments on PDTB2.0 [18], which only contains English data. We simplify the discourse relation labels from 35 to 19 based on the strategy from [25]. *EntRel* and *NoRel* are also included as individual labels to consider as many shallow discourse phenomena as possible.

For discourse-aware NMT, we conduct experiments on IWSLT2017 Chinese-English dataset. IWSLT2017 is generated from TED talk, which is a spoken language dataset and has coherent sentences. According to our translation tasks, the adjacent sentences in the dataset can be combined as translation units for their discourse relation. To make the model understand which sentence is context, a break token *[SEP]* is inserted between two source sentences of every unit during the training. We use dev2010 for development and tst2010-2013 for testing from IWSLT dataset. We train all NMT models up to 200K steps with 16 batch sizes. Length penalty is set to 1. Adam [26] is used to optimize parameters with the $2e-5$ learning rate. The beam size of decoding is 4.

3.2. Results of Probing Task

In Figure 2(a), we present the first probing task performance for three basic models generated from the 12 layers. For the accuracy of all models’ layers under 0.6, it seems that PLMs have a weaker capture ability in discourse than other linguistic knowledge [27]. It’s clear that BART, as an encoder-decoder model, is the best among the three models. BART has a process of gradual increase in discourse knowledge in both all encoder layers and decoder layers 7-10. In BERT, the average pooling of embeddings can extract more discourse knowledge than the *[CLS]* embeddings. Both BART and BERT have their discourse-aware layer in the ninth layer and have a significant decline after the discourse-aware layer. Different from the other two PLMs, GPT-2 has its strongest discourse-aware capability in the first layer and then continually decreases in the after layers.

In Figure 2(b), we show the results of BART performed in the second probing task. The curves of EXP and IMP both perform the same trend as BART (all data) in Figure 2(a). This indicates that the relative capability between the different layers of PLM is the same in different kinds of discourse manifestations. We can see that the accuracy of implicit data is much lower than that of the various embeddings of explicit data. Without connective words, pretrained language models can not embed discourse knowledge well. As the number of encoder layers increases, the discourse knowledge declines in both CON and ARG. But as the number of decoder layers increases, the discourse knowledge turns to rise in ARG, while that declines sharply in CON. CON contains larger discourse knowledge than EXP and ARG before layer 11 but is overtaken in layer 11 by ARG.

3.3. Fine-grained Analysis

As observed, the embeddings of connective words within most layers of PLMs encompass discourse knowledge. In Figure 3, we further investigate the effects of three pretrained models on three types of discourse relations. As seen, the accuracies of Implicit (IMP) and AltLex (ALT) relations are comparable and notably lower than that of the Explicit (EXP) relation. Based on these consistent phenomena observed across the three models, we can reaffirm the conclusion drawn in Section 3.2 that connective words serve as essential elements for PLMs in comprehending discourse relations. As demonstrated in [8], mBART




Figure 3: Results of probing tasks for evaluating different discourse relations, including explicit (“EXP”), implicit (“IMP”), and AltLex (“ALT”) discourse relation, embedded in different layers of representations in different pretrained models. We label the results of all data, including five discourse relations, as “ALL”

Table 3: Results of Chinese→English translation task in terms of different evaluation metrics. We compare three pretrained models in different languages, including English (“EN”), Chinese (“ZH”), or multilingual (“Multi”).

Lang.	Model	BLEU [↑]	TER [↓]	METEOR [↑]
EN	BERT	5.67	83.52	0.21
	BART	6.53	83.96	0.25
	GPT-2	4.60	86.07	0.19
ZH	BERT	6.72	81.74	0.23
	BART	7.60	78.57	0.27
	GPT-2	4.64	86.18	0.19
Multi	BERT	12.90	75.82	0.32

exhibits a preference for translating implicit discourse relations into explicit connective words, which in turn enhances the performance of document-level translation. This finding underscores the significant role of explicit discourse knowledge in PLMs in facilitating discourse-aware translation.

3.4. Results of Translation Task

Table 3 shows the overall performance of PLMs leveraged in NMT models. Except for English BART’s TER score which is weaker than that of English BERT, BART performs best on both language versions and all. GPT-2 is the worst. As BART is an encoder-decoder model, we consider that document-level NMT task prefers PLMs with the same architecture. Besides, all three Chinese PLMs get better scores than the English PLMs. The multilingual-BERT even performs a huge improvement over all other PLMs. We consider that a PLM with the source language discourse properties will perform better than the target language.

Table 4 shows the fine-grained results of the translation tasks. According to all three scores, the discourse-aware layer of PLMs revealed in probing tasks still performs better than the other layers in NMT tasks. During the experiments, we found that only updating one layer’s parameters can save around 13% of training time, and the performance of the discourse-aware layer is close to that of all layers. The conclusions from Table 3 and Table 4 show the same trend with the probing tasks.

Table 4: Fine-grained results of translation task in Table 3. We use only a specific layer of PLMs for initializing NMT models.

Layer	1	6	9	12	ALL
<i>BLEU[↑]</i>					
BERT	4.86	4.89	5.04	4.84	5.67
BART	5.27	5.46	6.03	5.66	6.53
GPT-2	4.52	3.69	3.51	3.37	4.60
<i>TER[↓]</i>					
BERT	85.12	84.97	84.92	84.92	83.52
BART	85.84	86.28	84.33	84.68	83.96
GPT-2	85.55	88.35	88.67	88.75	86.07
<i>METEOR[↑]</i>					
BERT	0.2002	0.2011	0.2019	0.2011	0.2150
BART	0.2287	0.2215	0.2424	0.2374	0.2479
GPT-2	0.1932	0.1690	0.1676	0.1650	0.1916

4. Conclusion and Future Work

This paper analyzes discourse knowledge in pretrained language models for spoken language translation. The coincident trend in probing tasks and machine translation tasks reveals how the discourse knowledge changes in different layers of different pretrained language models when the model is leveraged in NMT. Our work complements existing probing tasks about discourse knowledge in PLMs and provides a fine-grained interpretive perspective for the application of PLMs in document-level NMT.

In future work, we intend to delve deeper into the challenges of discourse-aware MT. Our plan is to leverage large language models (LLMs), which have shown great potential in dealing with complex tasks, to tackle this particular challenge [28, 29]. Moreover, we plan to evaluate our approach on more discourse-aware tasks [30, 31, 32, 33].

5. Acknowledgements

This work was supported in part by the Science and Technology Development Fund, Macau SAR (Grant Nos. FDCT/060/2022/AFJ, FDCT/0070/2022/AMJ) and the Multi-year Research Grant from the University of Macau (Grant No. MYRG2020-00054-FST). This work was performed in part at SICCC which is supported by SKL-IOTSC, and HPCC supported by ICTO of the University of Macau.

6. References

- [1] L. Wang, Z. Tu, A. Way, and Q. Liu, “Exploiting cross-sentence context for neural machine translation,” in *EMNLP*, 2017.
- [2] L. Wang, “Discourse-aware neural machine translation,” Ph.D. dissertation, Ph. D. thesis, Dublin City University, Dublin, Ireland, 2019.
- [3] L. Wang, S. Liu, M. Xu, L. Song, S. Shi, and Z. Tu, “A survey on zero pronoun translation,” in *ACL*, 2023.
- [4] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” *NeurIPS*, 2017.
- [5] Z. Guo and M. Le Nguyen, “Document-level neural machine translation using bert as context encoder,” in *AAACL: Student Research Workshop*, 2020.
- [6] S. Rothe, S. Narayan, and A. Severyn, “Leveraging pre-trained checkpoints for sequence generation tasks,” *TACL*, 2020.
- [7] L. Wang, Z. Tu, X. Wang, L. Ding, L. Ding, and S. Shi, “Tencent AI lab machine translation systems for WMT20 chat translation task,” in *Proceedings of the Fifth Conference on Machine Translation*, 2020.
- [8] Y. Liu, J. Gu, N. Goyal, X. Li, S. Edunov, M. Ghazvininejad, M. Lewis, and L. Zettlemoyer, “Multilingual denoising pre-training for neural machine translation,” *TACL*, vol. 8, 2020.
- [9] R. R. Agrawal, M. Turchi, and M. Negri, “Contextual handling in neural machine translation: Look behind, ahead and on both sides,” in *EAMT*, 2018.
- [10] R. Prasad, B. Webber, and A. Joshi, “Reflections on the penn discourse treebank, comparable corpora, and complementary annotation,” *Computational Linguistics*, vol. 40, no. 4, 2014.
- [11] N. F. Liu, M. Gardner, Y. Belinkov, M. E. Peters, and N. A. Smith, “Linguistic knowledge and transferability of contextual representations,” 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)*, 2019, pp. 1073–1094.
- [12] L. Carlson, D. Marcu, and M. E. Okurowski, “Building a discourse-tagged corpus in the framework of rhetorical structure theory,” in *Current and new directions in discourse and dialogue*. Springer, 2003.
- [13] F. Koto, J. H. Lau, and T. Baldwin, “Discourse probing of pre-trained language models,” in *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2021, pp. 3849–3864.
- [14] M. Kurfali and R. Östling, “Probing multilingual language models for discourse,” in *Proceedings of the 6th Workshop on Representation Learning for NLP (RepL4NLP-2021)*, 2021, pp. 8–19.
- [15] J. D. M.-W. C. Kenton and L. K. Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of NAACL-HLT*, 2019, pp. 4171–4186.
- [16] M. Lewis, Y. Liu, N. Goyal, M. Ghazvininejad, A. Mohamed, O. Levy, V. Stoyanov, and L. Zettlemoyer, “Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension,” in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020, pp. 7871–7880.
- [17] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever *et al.*, “Language models are unsupervised multitask learners,” *OpenAI blog*, vol. 1, no. 8, p. 9, 2019.
- [18] R. Prasad, N. Dinesh, A. Lee, E. Miltsakaki, L. Robaldo, A. Joshi, and B. Webber, “The penn discourse treebank 2.0,” in *LREC*, 2008.
- [19] M. Cettolo, M. Federico, L. Bentivogli, N. Jan, S. Sebastian, S. Katsutho, Y. Koichiro, and F. Christian, “Overview of the iwslt 2017 evaluation campaign,” in *International Workshop on Spoken Language Translation*, 2017.
- [20] M. Palmer, D. Gildea, and P. Kingsbury, “The proposition bank: An annotated corpus of semantic roles,” *Computational linguistics*, vol. 31, no. 1, 2005.
- [21] J. Tiedemann and Y. Scherrer, “Neural machine translation with extended context,” *DiscoMT 2017*, p. 82, 2017.
- [22] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, “Bleu: a method for automatic evaluation of machine translation,” in *ACL*, 2002.
- [23] M. Snover, B. Dorr, R. Schwartz, L. Micciulla, and J. Makhoul, “A study of translation edit rate with targeted human annotation,” in *Proceedings of the 7th Conference of the Association for Machine Translation in the Americas: Technical Papers*, 2006.
- [24] S. Banerjee and A. Lavie, “Meteor: An automatic metric for mt evaluation with improved correlation with human judgments,” in *Proceedings of the ACL workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, 2005.
- [25] N. Xue, H. T. Ng, S. Pradhan, R. Prasad, C. Bryant, and A. Rutherford, “The conll-2015 shared task on shallow discourse parsing,” in *Proceedings of the Nineteenth Conference on Computational Natural Language Learning-Shared Task*, 2015.
- [26] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [27] I. Vulić, E. M. Ponti, R. Litschko, G. Glavaš, and A. Korhonen, “Probing pretrained language models for lexical semantics,” in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020, pp. 7222–7240.
- [28] L. Wang, C. Lyu, T. Ji, Z. Zhang, D. Yu, S. Shi, and Z. Tu, “Document-level machine translation with large language models,” *arXiv preprint arXiv:2304.02210*, 2023.
- [29] C. Lyu, J. Xu, and L. Wang, “New trends in machine translation using large language models: Case examples with chatgpt,” *arXiv preprint arXiv:2305.01181*, 2023.
- [30] L. Wang, Z. Tu, X. Zhang, H. Li, A. Way, and Q. Liu, “A novel approach for dropped pronoun translation,” in *NAACL*, 2016.
- [31] L. Wang, Z. Tu, A. Way, and Q. Liu, “Learning to jointly translate and predict dropped pronouns with a shared reconstruction mechanism,” in *EMNLP*, 2018.
- [32] L. Wang, Z. Tu, X. Wang, and S. Shi, “One model to learn both: Zero pronoun prediction and translation,” in *EMNLP-IJCNLP*, 2019.
- [33] L. Wang, M. Xu, D. F. Wong, H. Liu, L. Song, L. S. Chao, S. Shi, and Z. Tu, “GuoFeng: A benchmark for zero pronoun recovery and translation,” in *EMNLP*, 2022.