Better Chinese Sentence Segmentation with Reinforcement Learning

Srivatsan Srinivasan
DeepMind
vatsan@deepmind.com

Chris Dyer
DeepMind
cdyer@deepmind.com

Abstract

A long-standing challenge in Chinese–English machine translation is that sentence boundaries are ambiguous in Chinese orthography, but inferring good splits is necessary for obtaining high quality translations. To solve this, we use reinforcement learning to train a segmentation policy that splits Chinese texts into segments that can be independently translated so as to maximise the overall translation quality. We compare to a variety of segmentation strategies and find that our approach improves the baseline BLEU score on the WMT2020 Chinese–English news translation task by +0.3 BLEU overall and improves the score on input segments that contain more than 60 words by +3 BLEU.

1 Introduction

Machine translation systems typically operate on sentence-like units, where sentences are translated independently of each other (Vaswani et al., 2017; Bahdanau et al., 2015; Koehn et al., 2003; Brown et al., 1993), in some cases with additional conditioning on a representation of adjacent sentences to improve coherence (Miculicich et al., 2018; Zhang et al., 2018). While many pairs of languages use similar orthographic conventions to designate sentence boundaries, English and Chinese diverge considerably: complete sentences in Chinese may be terminated either unambiguously with a full stop (。) or ambiguously with a comma. See Figure 1 for an example.

This divergence poses a challenge for Chinese–English translation systems since they must either be able to cope with potentially long, multi-sentence inputs (i.e., translating any text that falls between unambiguous sentence-ending punctuation) or, alternatively, they must be able to determine which comma occurrences terminate complete sentences that can be translated independently and which do not.

习近平主席2015年9月宣布成立的"中国-联合国和平与发展基金"，将支持非洲维和、反恐、实现可持续发展作为重点，迄今已开展34个涉非项目，成为中国、非洲、联合国三方合作的新平台。

The "China-United Nations Peace and Development Fund" announced by President Xi Jinping in September 2015 focuses on supporting African peacekeeping, counter-terrorism, and achieving sustainable development. So far, 34 projects involving Africa have been carried out, becoming a tripartite cooperation between China, Africa and the United Nations.

Figure 1: An example taken from the WMT2020 test set that shows a single source Chinese segment is translated into two separate English sentences. The highlighted comma separates the two corresponding complete sentences in the Chinese text, whereas the other two commas are sentence-internal boundaries.

Being able to directly accommodate long, multi-sentence inputs has clear appeal. However, in practice, the training data available for translation models is dominated by the relatively short (sub)sentence pairs that are preferentially recovered by standard approaches to sentence alignment (Tiedemann, 2011; Gale and Church, 1993), and as a result of the natural distribution of sentence lengths. Unfortunately, generalisation from training on short sequences to testing on long sequences continues to be an unsolved problem even in otherwise well-performing translation models (Lake and Baroni, 2018; Koehn and Knowles, 2017). Rather than addressing the length generalisation problem directly, in this paper we side-step it by learning to make decisions about segmentation so as to maximise the performance of an unreliable machine translation system that operates optimally only on shorter segments of input Chinese text.

While numerous text segmentation techniques designed to improve machine translation have been proposed over the years (§5), these have typically

Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021, pages 293–302

been based on heuristics that capture linguistic or statistical notions about what "minimal translatable units" consist of. In Chinese, robustly identifying such units is particularly challenging, on account of the lack of overt tense and frequent argument dropping which deprive an annotator of important clues (Huang, 1984). In contrast, we formalise the segmentation problem as a series of classification decisions about whether or not to split at candidate segmentation boundaries and use reinforcement learning to train the segmentation policy to optimise the aggregate BLEU score that results from translating the resulting segments with a *particular* translation system. Our approach is therefore robust to the idiosyncrasies of the underlying translation system, it is capable of discovering a policy that deviates from perhaps unreliable intuitions about minimal translatable units, and it can easily be retrained as the translation system improves.

Experiments indicate that the proposed approach outperforms a baseline that carries out no sentence splitting other than at unambiguous points, a classification approach based on linguistic criteria, and a heuristic system used in prior work. Overall, we improve the BLEU score on the WMT2020 Chinese–English news translation task by 0.3 BLEU, but for segments consisting of more than 60 words, the BLEU score increases by 3 BLEU.

2 Problem setup

We setup segmentation problem as a Markov decision problem (MDP) whose state, actions and rewards are characterized as follows. Every example in the training dataset is treated as a new episode and the objective is to maximise the sentence level BLEU score of the translation.

- The action set $\mathcal{A} = \{\text{SPLIT}, \text{CONTINUE}\}$.
- The state representation at time t, $\mathbf{s}_t = [\boldsymbol{\phi}(x)_{1,\ldots,n}]; [\mathbf{p}^{(t)}_{1,\ldots,n}]$ where n is the length of the input sequence $\boldsymbol{x}$ in words, $\boldsymbol{\phi}(x_i)$ is a vector encoding of ith token in context, and $\mathbf{p}^{(t)}_i$ is a record of previous decisions taken by the classifier about split decisions. The decision state for each token at timestep t can take on 4 possible discrete values: **no punctuation** (no actions are taken on these), **undecided punctuation** (punctuation on which an action still needs to be taken), **un-split** punctuation and **split** punctuation. At each timestep t, an action is taken at the next immediate undecided punctuation from the left and the state update involves appropriately updating the punctuation marker (to unsplit or split) corresponding to that position. The episode is considered terminal when there are no unattended punctuation markers in the sentence.
- For our reward, we use $r_t = \text{BLEU}(\tau(\mathbf{s}_{t+1}), \boldsymbol{y}^*) - \text{BLEU}(\tau(\mathbf{s}_t), \boldsymbol{y}^*)$, the marginal difference added to BLEU score based on the current action decision similar to Wu et al. (2018). τ represents the translation of the source inputs constructed from the state definition, where we split the sentence according to the punctuation markers in $\mathbf{s}_t$, translate each segment independently and recombine the best translation of each segment. An action of no split yields 0 rewards as segments remain identical and splitting would produce positive or negative rewards depending on the improvement/degradation to the quality of the overall sentence translation. Marginal BLEU sums to the sentence level BLEU obtained on the full sequence of translations, but provides a denser reward signal which makes policy learning more efficient.

Network and Learning Algorithm. We learn optimal segmentation policies using a distributed deep RL algorithm, IMPALA (Espeholt et al., 2018) which provides data-efficient stable learning at high throughput by combining decoupled acting and learning. At each time-step, our policy receives a state $\mathbf{s}_t$ and defines a distribution over the discrete action set $\mathcal{A}$ – i.e. $\pi(s_t) : \mathbb{R}^n \rightarrow \mathcal{A}\{\text{SPLIT}, \text{CONTINUE}\}$. Our algorithm also employs a value network that learns the value of a state $V(s_t) : \mathbb{R}^n \rightarrow \mathbb{R}$, which is used to regularize the policy network. In this work, we use a transformer encoder with self-attention layers as an observation encoder and then apply a feed-forward classifier on this encoded observation to learn the policy's action distribution. More details on the network architecture can be found in Appendix A.

3 Experiments

Our experiments are carried out on the WMT2020 Chinese–English news translation task, which is subjected to the same pre-processing steps as described in (Yu et al., 2020). Figure 2 shows the distribution of the maximum number of possible segments that we could split for each given example in our test dataset if a model were to split on every available punctuation (comma and full-stop

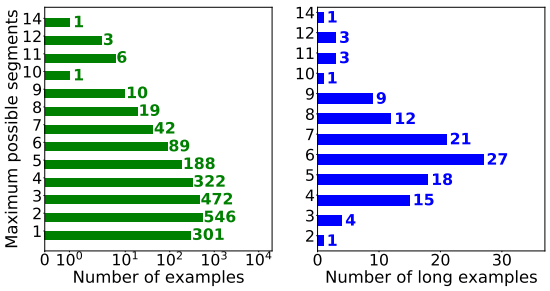

Figure 2: Distribution of maximum possible number of segments (if a model were to split on every available punctuation) for all the examples in the entire WMT Chinese–English test dataset and on long examples (source length ≥60 words) only. Note that the test dataset contains 2000 examples in total and 115 long examples.

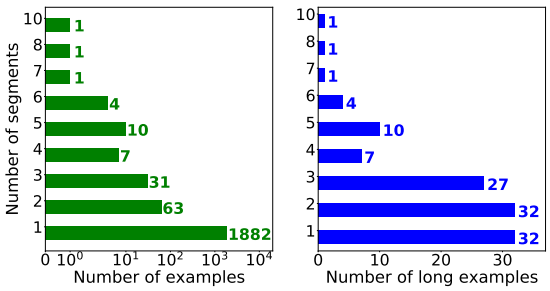

Figure 3: Distribution of number of segments split by our model RLSEGMENT on the entire WMT20 Chinese–English test dataset and on long examples (source length ≥60 words) only. Out of the 115 long examples, the model leaves 32 (28%) of them unsegmented and segments the rest into two or more segments.

in our case). We report case-sensitive BLEU, as computed with `sacrebleu` (Post, 2018). All the model/training details such as datasets, model architectures and hyperparameters pertaining to the baseline models are listed in Appendix B. We compare RLSEGMENT, our proposed solution's performance to six other baselines which help highlight different trade-offs involved in segmentation decisions for machine translation.

- NOSPLIT – Our key baseline that we wish to improve – a strategy of not doing any splits on the source inputs beyond unambiguous full stops.
- ALLSPLIT – An aggressive segmentation policy where we segment the source on every possible comma and full-stop.
- ORACLE – To compute the oracle score, we translate possible splits of a given source sentence and select the one that maximizes the example-level BLEU score. This benchmark is the upper limit of any segmentation policy, given a translation model. It is quite expensive to compute since it requires decoding all possible segmentations of every sequence.
- ORACLESUP – Oracle segmentation decisions from the training corpus are used to setup a supervised binary classification task with an architecture similar to RLSEGMENT's policy network.
- COMMACLASS – Using the syntactic patterns from the Penn Chinese Treebank 6.0, this system builds a comma classifier to disambiguate terminal and non-terminal commas similar to (Xue and Yang, 2011). This uses a transformer encoder followed by a positional feed-forward network to classify every comma.

Model	full BLEU	full BP	≥ 60 BLEU	≥ 60 BP
NOSPLIT (Baseline) [2000]	31.89	89.42	25.73	78.17
ALLSPLIT [6233]	29.45	93.18	27.88	94.43
ORACLESUP [3107]	31.57	90.89	27.66	88.02
COMMACLASS [2569]	31.82	89.98	26.93	81.16
HEURISTIC [2071]	32.02	89.59	27.05	80.80
RLSEGMENT (Ours) [2228]	**32.21**	**91.34**	**29.03**	**88.09**
ORACLE [3513]	*36.36*	*96.19*	*34.71*	*93.45*

Table 1: BLEU and brevity penalty scores, both on the corpus and long sentences only (source length ≥ 60 words) on the test dataset of WMT20 Chinese–English. Reported in square brackets are the number of segments independently translated under each policy.

- HEURISTIC – The uses a combination of predictions from COMMACLASS together with heuristic length constraints that only split long inputs (> 60 words) on terminal commas suggested by the model and terminal punctuations, and only if the resulting segments are not too short (> 10 words).

As discussed above, there is no standard segmentation of Chinese texts into sentences, and therefore all the "supervised" approaches—including our baselines ORACLESUP and HEURISTIC, construct their own training data to train a classifier. RLSEGMENT on the other hand, requires only a translation system and a corpus of standard parallel data for training and learns without any hand-engineered constraints on the model itself, thus presenting a generic solution that can scale across languages and system variants.

Source (showing segmentation candidate split points and RLSEGMENT policy decision in grey)

中新网深圳10月1日电(郑小红徐殿伟刘姝梦)10月1日为内地国庆"黄金周"首日，也是香港国庆公众假期的最后一天，许多内地居民利用假期出国境游玩、购物，也有许多来深的香港居民当日返港，深圳各口岸出现客流"爆棚"的情况，罗湖、深圳湾口岸新启用22条自助通道，提升口岸通关效率，便利旅客高效快捷通关。

NOSPLIT *Shenzhen, Oct. 1 (Zheng Xiaohong Xu Dianwei Liu Jiameng) Oct. 1 is the first day of the Mainland National Day "Golden Week" and the last day of the Hong Kong National Day Public Holiday. Many mainland residents used the holiday to visit and shop abroad. Many Hong Kong residents who came to Shenzen returned to Hong Kong on the same day. There are passenger flow "burst shed" at all ports.*

RLSEGMENT *Shenzhen, Oct. 1 (Zheng Xiaohong Xu Dianwei Liu Jiameng) Oct. 1 is the first day of the Mainland National Day "Golden Week" and the last day of the Hong Kong National Day Public Holiday. Many mainland residents used their holidays to visit and shop abroad. Many Hong Kong residents who came to Shenzhen returned to Hong Kong on the same day. There are passenger flow "burst shed" at all ports in Shenzhen. 22 new self-service channels opened at Lo Wu and Shenzhen Bay ports to improve the customs clearance efficiency at ports, facilitate efficient and rapid customs clearance of passengers.*

Figure 4: A example translation where segmentation with RLSEGMENT mitigates premature truncation in our system; material dropped by the baseline system is highlighted in grey.

4 Results

Figure 3 shows the distribution of segments proposed by our model RLSEGMENT. We see that a lot of the long sentences (source length ≥ 60 words) are split into two or more independent segments to mitigate premature truncation seen in transformer models with these long sentences. Table 1 compares different segmentation policies when translating the WMT20 Chinese–English news test dataset. While the BLEU scores indicate the quality of translations on the entire corpus, we also report BLEU scores on sources longer than 60 words as a metric to show the performance of these models on longer sentences where standard transformers tend to produce translations that are too short. In both the cases, we also report the brevity penalty (BP), a component of BLEU to show the impact on the overall length of the translation.

We see that our proposed segmentation policy, RLSEGMENT improves both the BLEU scores and brevity penalties as compared to the baseline translation case NOSPLIT. Specifically, the RL model improves BLEU scores on long sentences by 3+ BLEU points and BP on those sentences by about 9+ points. This shows that our model, via smart segmentation, suffers less because of premature truncation of long translations as compared to the baseline—a common problem (Meister et al., 2020; Koehn and Knowles, 2017). While segmentation of long sentences at appropriate punctuations helps performance, segmentation at all punctuations is expected to hurt performance as it is highly likely to produce extremely small segments which lose a lot of necessary source context when individually translated. This is demonstrated by poorer BLEU score of the ALLSPLIT baseline, even though it achieves good BP scores both on the corpus and long translations. Compared to supervised baselines trained on syntactic data such as COMMACLASS and HEURISTIC, our model performs competitively on both BLEU and BP without any supervised data for segmentation or hand-engineered length constraints. In Figure 4, we see an example where RLSEGMENT mitigates premature truncation of the resultant translation. In this example, although input (and resulting English translation) consists of three sentences separated by commas, the segmentation policy has only chosen to split at only *one* position, having learned that the underlying translation system is capable of translating some two-sentence inputs. This example thus illustrates the practicality of learning a segmentation policy based on the abilities of the underlying translation system, not just on the basis of normative notions of translatable units. (More examples can be found in Appendix C.1, C.2)

While our model does better than the baselines, there is a sufficient performance gap to the oracle BLEU scores (because of the different data/length characteristics between training and test time) that could be achieved via "perfect" segmentation, demonstrating the value for further research into better segmentation strategies. However, we also note that the RLSEGMENT outperforms ORACLESUP—especially on long sentences. We suspect this has to do with the relative scarcity of such examples in the training data—while a supervised learner can happily ignore those rare cases at little cost in terms of cross-entropy, they have an out-sized impact on BLEU, and therefore the RL learner is sensitive to them.

Finally, it is important to note that while RLSEG-

MENT improves BLEU at a corpus level, there exist cases where individual translation examples (Appendix C.3) are worse because of inappropriate segmentation.

5 Related Work

The segmentation of long texts and sentences into segments suitable for translation has been a recurring topic in machine translation research (Tien and Minh, 2019; Pouget-Abadie et al., 2014; Goh and Sumita, 2011; Doi and Sumita, 2003); however, we are the first to apply reinforcement learning to solve the problem. A related problem to the segmentation problem occurs in automated simultaneous interpretation, where the system must produce translations as quickly as possible, but it is necessary to wait until sufficient context has been received before an accurate translation can be produced. Grissom II et al. (2014) used an RL approach, targeting a reward that balances translation quality with translation latency.

Chinese comma disambiguation has likewise been studied. However, without exception these have sought to predict normative notions of what constitutes a complete clause or elementary discourse unit (Xu and Li, 2013; Xue and Yang, 2011; Jin et al., 2004), on the basis of syntactic annotations in the Chinese Treebank (Xue et al., 2005). In contrast, our solution is directly targeted at developing a segmentation strategy that results in a good downstream translation, rather than conforming to any single normative notion of what constitutes a complete sentence.

6 Conclusion

In this work, we have addressed a key challenge in Chinese-English machine translation : the ambiguity of English-like sentence boundaries in Chinese, resulting in long Chinese sentence data for machine translation tasks. Our solution casts the Chinese sentence segmentation problem into a sequential decision making problem and then uses Reinforcement Learning to learn an optimal segmentation policy to maximize the BLEU scores of the eventual translation from the independent segment translations. Our solution does not require any paired training data for segmentation and is able to learn an optimal strategy purely from paired machine translation data. Our model is able to outperform a baseline translation strategy that segments on only unambiguous full-stops by 0.3 BLEU at a corpus level and by 3 BLEU on a sub-corpus comprising only source sentences longer than 60 words.

Acknowledgements

We would like to thank Laurent Sartran, Wojciech Stokowiec, Lei Yu and Wang Ling for helpful discussions and guidance during different stages of this work in terms of experimentation, presentation and software development.

References

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proc. ICLR*.

Peter F. Brown, Stephen A. Della Pietra, Vincent J. Della Pietra, and Robert L. Mercer. 1993. The mathematics of statistical machine translation: Parameter estimation. *Computational Linguistics*, 19(2).

Takao Doi and Eiichiro Sumita. 2003. Input sentence splitting and translating. In *Proceedings of the HLT/-NAACL Workshop on Building and Using Parallel Texts*.

Lasse Espeholt, Hubert Soyer, Rémi Munos, Karen Simonyan, Volodymyr Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim Harley, Iain Dunning, Shane Legg, and Koray Kavukcuoglu. 2018. IMPALA: scalable distributed deep-RL with importance weighted actor-learner architectures. In *Proc. ICML*.

William A. Gale and Kenneth W. Church. 1993. A program for aligning sentences in bilingual corpora. *Computational Linguistics*, 19(1).

Chooi-Ling Goh and Eiichiro Sumita. 2011. Splitting long input sentences for phrase-based statistical machine translation. In *Proceedings of the Association of Natural Language Processing*.

Alvin Grissom II, He He, Jordan Boyd-Graber, John Morgan, and Hal Daumé III. 2014. Don't until the final verb wait: Reinforcement learning for simultaneous machine translation. In *Proc. EMNLP*.

C. T. James Huang. 1984. On the distribution and reference of empty pronouns. *Linguistic Inquiry*, 15:531–574.

Mei xun Jin, Mi-Young Kim, Dongil Kim, and Jong-Hyeok Lee. 2004. Segmentation of Chinese long sentences using commas. In *Proceedings of the Third SIGHAN Workshop on Chinese Language Processing*.

Philipp Koehn and Rebecca Knowles. 2017. Six challenges for neural machine translation. In *Proceedings of the First Workshop on Neural Machine Translation*.

Philipp Koehn, Franz J. Och, and Daniel Marcu. 2003. Statistical phrase-based translation. In *Proc. NAACL*.

Brendan Lake and Marco Baroni. 2018. Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks. In *Proc. ICML*.

Clara Meister, Ryan Cotterell, and Tim Vieira. 2020. If beam search is the answer, what was the question? In *Proc. EMNLP*.

Lesly Miculicich, Dhananjay Ram, Nikolaos Pappas, and James Henderson. 2018. Document-level neural machine translation with hierarchical attention networks. In *Proc. EMNLP*.

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, Brussels, Belgium. Association for Computational Linguistics.

Jean Pouget-Abadie, Dzmitry Bahdanau, Bart van Merriënboer, Kyunghyun Cho, and Yoshua Bengio. 2014. Overcoming the curse of sentence length for neural machine translation using automatic segmentation. In *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*.

Noam Shazeer. 2019. Fast transformer decoding: One write-head is all you need. *CoRR*, abs/1911.02150.

Jörg Tiedemann. 2011. *Bitext Alignment*. Synthesis Lectures on Human Language Technologies. Morgan & Claypool.

Ha Nguyen Tien and Huyen Nguyen Thi Minh. 2019. Long sentence preprocessing in neural machine translation. In *IEEE-RIVF International Conference on Computing and Communication Technologies*.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proc. NeurIPS*.

Lijun Wu, Fei Tian, Tao Qin, Jianhuang Lai, and Tie-Yan Liu. 2018. A study of reinforcement learning for neural machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3612–3621, Brussels, Belgium. Association for Computational Linguistics.

Shengqin Xu and Peifeng Li. 2013. Recognizing Chinese elementary discourse unit on comma. In *Proceedings of the 2013 International Conference on Asian Language Processing*.

Nianwen Xue, Fei Xia, Fu-Dong Chiou, and Martha Palmer. 2005. The Penn Chinese TreeBank: Phrase structure annotation of a large corpus. *Natural Language Engineering*, 11(2).

Nianwen Xue and Yaqin Yang. 2011. Chinese sentence segmentation as comma classification. In *Proc. ACL*.

Lei Yu, Laurent Sartran, Po-Sen Huang, Wojciech Stokowiec, Domenic Donato, Srivatsan Srinivasan, Alek Andreev, Wang Ling, Sona Mokra, Agustin Dal Lago, Yotam Doron, Susannah Young, Phil Blunsom, and Chris Dyer. 2020. The DeepMind Chinese–English document translation system at WMT2020. In *Proceedings of the Fifth Conference on Machine Translation*, pages 326–337, Online. Association for Computational Linguistics.

Jiacheng Zhang, Huanbo Luan, Maosong Sun, Feifei Zhai, Jingfang Xu, Min Zhang, and Yang Liu. 2018. Improving the transformer translation model with document-level context. In *Proc. EMNLP*.

A Data and Model details for RLSEGMENT

A.1 Dataset

We use the training, validation and test datasets from the WMT2020 Chinese-English constrained data provided to shared task participants. All the examples are aligned paired translations with sentence-like units provided as part of the dataset. Pre-processing of text is done in exactly the same methodology as described in Section 3 of (Yu et al., 2020). We also use the same `sentencepiece` tokenizer described in Section 3 of that work to convert the text into integer tokens. In total, we operate on a dataset with 18 million training examples, 1619 validation examples and 2000 test examples.

A.2 Translation Model

For all experiments and baselines, we use the same transformer model to translate any source input (Vaswani et al., 2017). The model used has 6 encoder layers and 2 decoder layers with 8 encoder and decoder attention heads each. The model uses a feed-forward layer of size 2048 and employs Multi-Query Attention layers (Shazeer, 2019) for faster inference. It uses different dropout values for different components: 0.1 for the attention, 0.05 in the feed-forward network, and finally 0.3 after the sub-layer. The other hyperparameters and learning rate schedules of the model are similar to those described in Section 4.1 of (Yu et al., 2020).

A.3 RL Models

The policy and value networks share common encoding layers, with different feed-forward networks on top of the shared encoder layer. The shared encoder consists of 2 stacked self-attention layers with 8 attention heads each and feed-forwards of size 512, with 0.1 as the rate for all attention, feed-forward and sub-layer dropouts. Sequence lengths for training were restricted to 280 tokens (maximum sentence length in the validation set is 277). The policy network applies a feed-forward network of sizes 256-2 on the outputs of the shared encoder and value network applies a feed-forward network of sizes 256-1. Adam optimizer with learning rate=0.0002, b1=0.0 and b2=0.99 was used for training. For the IMPALA style loss function (Espeholt et al., 2018), the weights of policy, baseline and entropy loss are set as 1.0, 0.5 and 0.0005 respectively. Key hyperparameters such as the weights of losses, learning rates and model sizes were tuned (a single trial for each hyperparameter configuration) using BLEU scores on the WMT20 Chinese–English validation dataset.

Bounds for hyperparameters - Number of attention layers [2,8], Number of attention heads [2,8], Size of feed-forward network [128, 2048], Learning rates [0.01, 0.00001], Weights of losses [0.0001, 1].

Compute and other details - For both inference and learning, we use 2x2 slices (4 TPUs) of Google TPU v2 with 2 cores per device (8 cores were split into 2 cores used for inference and 6 for learning). 512 actors on CPU were run in parallel to generate transition data for the distributed learner to learn from. Learning was done on 70 million episodes with a batch size of 256 per core and the entire experiment had an average run-time of approximately 6 hours.

B Data and Model details for BASELINES

All baseline models use the same inference model as described in Appendix A.2. All baselines are evaluated on the WMT20 Chinese–English test data as described in Appendix A.1. The baselines NOSPLIT, ALLSPLIT and ORACLE do not require any training data and are directly employed during test time.

B.1 COMMACLASS and HEURISTIC

Dataset - The comma classifier used in these baselines is trained on Chinese Treebank data prepared in the same format as described in (Xue and Yang, 2011).

Model - Both these baselines rely on a comma classifier model. In our experiments, this model uses a standard Transformer style encoder with a feed-forward classifier for Chinese comma disambiguation. It uses the encoder layers from a pre-trained transformer (which was trained on WMT20 Chinese–English translation task) since the Chinese Treebank dataset alone is small for training large models. The encoder of this model has 16 heads, 9 layers and a feed-forward size of 4096. On top of this encoder, a feed-forward classifier of size 512-128-2 is used for classification. The pre-trained encoders are further fine-tuned during classification. We use Adam Optimizer with a custom

learning rate schedule similar to (Vaswani et al., 2017).

Bounds for hyperparameters - Number of attention layers [2,16], Number of attention heads [2,16], Size of feed-forward network [128, 4096], Learning rates [0.01, 0.00001].

B.2 ORACLESUP

Dataset - The ORACLESUP baseline uses WMT2020 Chinese-English data described in Appendix A.1. From this dataset, supervised labels were generated by splitting an example into its all possible splits on punctuations and then choosing the split set which optimized the BLEU score at an individual example level.

Model - This model uses the same network as the policy network described in Appendix A.3, but with 6 attention layers and 8 attention heads each. The model uses a learning rate schedule and Adam optimizer, similar to (Vaswani et al., 2017). The classes (No split:split - 1:2.2) were assigned different weights in the loss function to account for class imbalance in the data.

Bounds for hyperparameters - Number of attention layers [2,8], Number of attention heads [2,8], Size of feed-forward network [128, 2048], Learning rates [0.01, 0.00001], Weight for the split class [1,5].

C Example Model Outputs

The splits in the source sentence have been highlighted.

C.1 RLSEGMENT addresses truncation in long source sentences.

Source

中国人民银行授权中国外汇交易中心公布, 2018 年1 月18 日银行间外汇市场人民币汇率中间价为: 1 美元对人民币6.4401 元, 1 欧元对人民币7.8482 元, 100 日元对人民币5.7854 元, 1 港元对人民币0.82399 元, 1 英镑对人民币8.9017 元, 1 澳大利亚元对人民币5.1302 元, 1 新西兰元对人民币4.6779 元, 1 新加坡元对人民币4.8546 元, 1 瑞士法郎对人民币6.6731 元, 1 加拿大元对人民币5.1747 元, 人民币1 元对0.61399 马来西亚林吉特, 人民币1 元对8.8352 俄罗斯卢布, 人民币1 元对1.9128 南非兰特, 人民币1 元对165.95 韩元, 人民币1 元对0.57030 阿联酋迪拉姆, 人民币1 元对0.58233 沙特里亚尔, 人民币1 元对39.3106 匈牙利福林, 人民币1 元对0.53093 波兰兹罗提, 人民币1 元对0.9487 丹麦克朗, 人民币1 元对1.2501 瑞典克朗, 人民币1 元对1.2240 挪威克朗,

Ground Truth

the People 's Bank of China has authorized the China Foreign Exchange Trade System to release the interbank foreign exchange market central parity rates of the RMB : USD 1 to RMB 6.4401 , EUR1 to RMB 7.8482 , JPY 100 to RMB 5.7854 , HKD1 to RMB 0.82399 , GBP 1 to RMB 8.9017 , AUD1 to RMB 5.1302 , NZD1 to RMB 4.6779 , SGD 1 to RMB 4.8546 , CHF1 to RMB 6.6731 , CAD1 to RMB 5.1747 , RMB 1 to MYR 0.61399 , RMB 1 to RUB8.8352 , RMB 1 to ZAR1.9128 , RMB 1 to KRW 165.95 , RMB 1 to AED0.57030 , RMB 1 to SAR0.58233 , RMB 1 to HUF39.3106 , RMB 1 to PLN0.53093 , RMB 1 to DKK0.9487 , RMB 1 to SEK1.2501 , RMB 1 to NOK1.2240 , RMB 1 to TRY0.59303 , RMB 1 to MXN2.9071 .

NOSPLIT Translation

the central parity rate of RMB in the interbank foreign exchange market was : US$ 1 to RMB 6.4401 , 1 euro to RMB 7.8482 , 100 yen to RMB 5.7854 , 1 Hong Kong dollars to RMB 0.82399 , 1 British pound to RMB 8.9017 , 1 Australian dollar to RMB 5.1302 , 1 NZ dollar versus RMB 4.6779 ,

RLSEGMENT Translation

the People 's Bank of China authorized the China Foreign Exchange Trading Center to announce that the median exchange rate of RMB in the interbank foreign exchange market on January 18 , 2018 was : US$ 1 to RMB 6.4401 , 1 euro to RMB 7.8482 , 100 yen to RMB 5.7854 , 1 Hong Kong dollar

to RMB 0.82399 , 1 pound to RMB 8.9017 , 1 Australian dollar to RMB 5.1302 , 1 NZ dollar versus RMB 4.6779 , 1 Singapore dollar to RMB 4.8546 and 1 Swiss franc to RMB 6.6731 , 1 Canadian dollar to RMB 5.1747 , 1 yuan to 0.61399 Malaysian ringgit , 1 yuan to 8.8352 Russian rubles , RMB 1 to 1.9128 South African rand and RMB 1 to 165.95 won , RMB 1 to 0.57030 UAE dirham , RMB 1 to 0.58233 Saudi SAR , RMB 1 to 39.3106 forint , RMB 1 to 0.53093 Polish zloty , RMB 1 to DKr 0.9487, RMB 1 to 0.59303 Turkish lire , RMB 1 to 2.9071 Mexican pesos.

Source

根据广东省委省政府部署, 广东广电网络积极实施新数字家庭4K 行动计划, 2017 年12 月发布了全国首个支持IPv6 、全国广电首款支持4 路真4K 并发的新数字家庭网关终端产品- - - “ U 点家庭服务器” (搭配AVS2 标准的4K 盒子、讯飞语音遥控器) , 加快建设以4K 超高清视频、下一代互联网、有线无线融合网络应用为标志的新数字家庭示范区, 积极参与广东省4K 电视网络应用示范区建设。

Ground Truth

according to the deployment of the Guangdong Provincial Party Committee and the provincial government , Guangdong Cable Network actively implemented the new digital home 4K action plan . in December 2017 , it released the first new digital home gateway terminal product in China that supports IPv6 and the national radio and television industry supporting the 4 - way real 4K concurrency . – “ U - Point Home Server ” (with AVS2 standard 4K box , iFlytek voice remote control) to accelerate the construction of a new digital home demonstration area marked by 4K Ultra HD video , next - generation Internet , wired and wireless integrated network applications , and actively participate in the construction of the 4K TV networks application demonstration area in Guangdong Province .

NoSplit Translation

according to the deployment of the provincial government of Guangdong Province , Guangdong Radio and Television Network actively implemented the new digital family 4K action plan . in December 2017 , it released the first new digital family gateway terminal product in the country to support IPv6 and the first national radio and television support 4zhen 4K concurrence - “ U - point family server ” (with AVS2 standard 4K box and message flying voice remote control) to accelerate the construction of

RLSegment Translation

according to the deployment of the provincial government of Guangdong Province , the Guangdong Radio and Television Network actively implemented the new digital family 4K action plan . in December 2017 , it released the first new digital family gateway terminal product in the country to support IPv6 and the first national radio and television to support 4 Zhen 4K concurrently - “ U - point family server ” (with AVS2 standard 4K box and message flying voice remote control) . it accelerate the construction of a new digital family demonstration area marked by 4K ultra - high - clearing video , next - generation Internet, wired and wireless fusion network applications, actively participate in the construction of 4K TV network application demonstration zone in Guangdong Province .

C.2 RLSegment does not split on all punctuations.

Source

此前, 为防范台风“ 潭美” , 截至28 日11 时统计, 福建已撤离268 艘渔船, 转移渔

排上老弱妇幼人员2282 人。

RLSEGMENT Translation

previously , in order to prevent Typhoon “ Tanmei ” , as of 11:00 on the 28th , Fujian had evacuated 268 fishing boats and transferred 2,282 elderly , weak women and children

Source

除上述便民措施外, 卫生局辖下各卫生中心及卫生站自推出计划后, 已延长周一至五的流感疫苗接种服务时间至晚上8 时, 截至9 月30 日下午4 时, 已为4,158 人接种, 与去年同期的3,714 人多。

RLSEGMENT Translation

in addition to the above convenience measures , health centers and health stations under the Health Bureau have extended the influenza vaccination service from Monday to Friday to 8 pm . as of 4 pm on September 30 , 4,158 people had been vaccinated , up from 3,714 in the same period last year .

C.3 Bad segmentation decisions produced by RLSEGMENT

Source

虽然海风剧烈, 但法国ABM 队的舵手皮埃尔· 伊维斯· 杜兰德在赛前见面会上诉说了自己与比赛举办地青岛的温情往事: 5 年前的国际极限帆船系列赛自己意外受伤, 肌肉断裂, 多亏当地赛事医疗团队的处理才让自己还能在5 年后的今天站在比赛场地。

Ground Truth

though the sea wind blew violently , the helmsman of French ABM Team Pierre - Yves Durand told his warm memories with Qingdao at the pre - game meeting that he was accidentally injured in the Extreme Sailing Series 5 years ago , resulting in the muscular rupture , and fortunately , the local race medical team helped him deal with it so that he could stand here to take part in the race after 5 years .

RLSEGMENT Translation

despite the intense sea breeze , Pierre Ives Durand , the helmsman of the French ABM team , told his warm past with Qingdao , the venue of the competition, that he was injured accidentally in the International Extreme Sailing Series five years ago . caused muscle fracture , thanks to the handling of the local event medical team , he can still stand on the venue today five years later .

Source

“ 喀山” 号、“ 新西伯利亚” 号、“ 克拉斯诺亚尔斯克” 号、“ 阿尔汉格尔斯克” 号多用途潜艇均为“ 白蜡树- M ” 改进型, 排水量为13800 吨, 下潜深度520 米, 人员编制64 人, 水下航速31 节, 将携带水雷、鱼雷和“ 口径” 及“ 缟玛瑙” 巡航导弹。

Ground Truth

the multipurpose submarines of “ Kazan ” , “ Novosibirsk ” , “ Krasnoyarsk ” and “ Arkhangelsk ” are all of “ Fraxinus - M ” improved type , with displacement of 13,800 tons , depth of 520 m , staffing of 64 , submerged speed of 31 knots . all the submarines will carry mines , torpedoes and “ Kalibr ” and “ Onyx ” cruise missiles .

RLSEGMENT Translation

Kazan , Nova Scotia , Krasnoyarsk and Arkhangelsk are all improved models of White wax tree - M , with a displacement of 13,800 tons and a submersible depth of 520 meters . staffing 64 , underwater speed of 31 knots. They will carry mines , torpedoes and “ caliber ” and “ agate ” cruise missiles .