Learning to Write with Cooperative Discriminators: Supplementary Material

Ari Holtzman, Jan Buys, Maxwell Forbes, Antoine Bosselut, David Golub, & Yejin Choi
Paul G. Allen School of Computer Science and Engineering
University of Washington
Seattle, WA 98195
{ahai,jbuys,mbforbes,antoineb,golubd,yejin}@cs.washington.edu

Abstract

We describe hyperparameters and details of our model training and evaluation.

3 The Learning Framework

3.1 Base Language Model

We use a 2-layer GRU (Cho et al., 2014) with a hidden size of 1024 for each layer. Following (Inan et al., 2017) we tie the input and output embedding layers' parameters. We use an Adaptive Softmax for the final layer (Grave et al., 2016), which factorizes the prediction of a token into first predicting the probability of k (in our case $k = 3$) clusters of words that partition the vocabulary and then the probability of each word in a given cluster. To regularize we dropout (Srivastava et al., 2014) cells in the output layer of the first layer with probability 0.2. We use mini-batch stochastic gradient descent (SGD) and anneal the learning rate when the validation set performance fails to improve, checking every 1000 batches. Learning rate, annealing rate, and batch size were tuned on the validation set for each dataset. Gradients are backpropagated 35 time steps and clipped to a maximum value of 0.25.

3.2 Cooperative Communication Models

For all the models except the entailment model, training is performed with Adam (Kingma and Ba, 2015) with batch size 64 and learning rate 0.01. The classifier's hidden layer size is 300. Dropout is performed on both the input word embeddings and the non-linear hidden layer before classification with rate 0.5.

Word embeddings are kept fixed during training for the repetition model, but are fine-tuned for all the other models.

3.2.2 Entailment Model

We mostly follow the hyperparameters of Parikh et al. (2016): Word embeddings are projected to a hidden size of 200, which are used throughout the model. Optimization is performed with AdaGrad (Duchi et al., 2011) with initial learning rate 1.0 and batch size 16. Dropout is performed at rate 0.2 on the hidden layers of the 2-layer MLPs in the model.

Our entailment classifier obtains 82% accuracy on the SNLI validation set and 68% accuracy on the MultiNLI validation set.

3.2.3 Relevance Model

The convolutional layer is a one-dimensional convolution with filter size 3 and stride 1; the input sequences are padded such that the input and output lengths are equal.

4 Experiments

4.1 Corpora

For the language model and discriminators we use a vocabulary of 100,000 words – we found empirically that larger vocabularies lead to better generation quality. To train our discriminators and evaluate our models, we use segments of length 10, using the first 5 sentences as context and the second 5 as the reference continuation. For TripAdvisor we use the first 10 sentences of reviews of length at least 10. For the BookCorpus we split books into segments of length 10. We select 20% of each corpus as held-out data (the rest is used for language model training). From the held-out data we select a test set of 2000 examples and two validation sets of 1000 examples each, one of which is used to train the mixture weights of the decoding objective. The rest of the held-out data is used to train the discriminative classifiers.

4.2 Baselines

CACHELM Due to memory constraints, we use a vocabulary size of 50k for CACHELM. Beam search decoding is used, with a beam size 5.

SEQGAN The implementation we used adds a number of modelling extensions to the original SeqGAN. In order to make training tractable, the vocabulary is restricted to 25k words, the maximum sequence length is restricted to 250, Monte Carlo rollouts to length 4, and the discriminator updated once for every 10 generator training steps. Greedy decoding sampling with temperature 0.7 was found to work better than beam search.

SEQ2SEQ Due to memory constraints, we use a vocabulary size of 50k for SEQ2SEQ. Beam search decoding is used, with a beam size 5.

4.3 Evaluation Setup

The forms used on Amazon Mechanical Turk are pictured in Tables 1, 2, 3, and 4.

References

Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the properties of neural machine translation: Encoder–decoder approaches. In *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, pages 103–111.

John Duchi, Elad Hazan, and Yoram Singer. 2011. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(Jul):2121–2159.

Edouard Grave, Armand Joulin, Moustapha Cissé, David Grangier, and Hervé Jégou. 2016. Efficient softmax approximation for gpus. *arXiv preprint arXiv:1609.04309*.

Hakan Inan, Khashayar Khosravi, and Richard Socher. 2017. Tying word vectors and word classifiers: A loss framework for language modeling. In *Proceedings of theInternational Conference on Learning Representations*.

Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of theInternational Conference on Learning Representations*.

Ankur Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. 2016. A decomposable attention model for natural language inference. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 2249–2255. Association for Computational Linguistics.

Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958.

BookCorpus

Instructions

In this task you will be provided with 5 sentences from the middle of a book, and 2 possible continuations.

We'd like you to tell us whether you agree or disagree with the statements below, which compare and contrast the two continuations.

Thank you!

Task

Beginning Text: The cat figure leaned forward and spoke quietly his voice resonating in the room. You were in mortal danger when I chanced upon you in the hospital. The creature who was in that room was an assassin, sent by evil men to kill our queen. Jared looked puzzled and mordalayn paused. Do, do you mean sophie?

Continuation:

Continuation A What is wrong with you?" "I don't know. I didn't know. I don't know." "you have nothing to say.

Continuation B "I 'm sure he didn't mean to kill him." She shook her head. "no, it 's not. It was just a dream." "what happened?"

Question 1: Repetitiveness

Compare the **continuations** based on their **repetitiveness**. A good continuation **DOES NOT** repeat details from earlier in the continuation or from the first 5 sentences.

Repetition Good/Bad Examples (Expand/Collapse)

Beginning

"*...They'd been waiting for a while and Melissa had given in to temptation and was perusing a magazine rack in one of the mini-malls located a few feet from another institution of higher-thought removal; a fast-food chain designed to represent a slow-food eatery.*"

The following continuation is GOOD as it does not repeat previous content.

- "*The waiters and waitresses were all dressed to impress with the appropriate buttons and trinkets placed accordingly. The thought was to put the diner at ease by displaying humorous one-liners which would cause them to feel an attraction for the high-schoolers or employment-challenged and over-order in the process...*"

The following continuation is BAD as it repeats many words and phrases.

- "*The table was set with what looked to be a dining room, a dining table two chairs a table two chairs and a table The chairs were comfortable and comfortable...*"

Statement: Continuation A repeats itself MORE than Continuation B.

○ Strongly Agree ○ Agree ○ They're roughly the same ○ Actually, **Continuation B** repeats itself a little **MORE** ○ **Continuation B** repeats itself **much MORE** than **Continuation A**

Question 2: Contradiction

Compare the **continuations** based on how much they **contradict** themselves or the first 5 sentences. A good continuation should **NOT** contradict itself or the beginning 5 sentences.

Contradiction Good/Bad Examples (Expand/Collapse)

Beginning

"*...On a long stretch of narrow road, the truck stopped abruptly and all the men jumped out, disconnecting the trailer. The process took half a minute, then the truck did a three-point turn back toward the main route north. *"

The following continuation is GOOD as it does not contradict itself or the beginning sentences.

- "*The controllers at Hanscom almost lost contact before they realized the boat was gone. They could not tell that two men had also left the truck under the tree cover...*"

The following continuation is BAD as it contradicts itself as well as details described in the first 5 sentences.

- "*The trees were dense. They were in an open field.*"

Statement: Continuation A contradicts itself MORE than Continuation B.

○ Strongly Agree ○ Agree ○ They're roughly the same ○ Actually, **Continuation B** contradicts itself a little **MORE** ○ **Continuation B** contradicts itself **much MORE** than **Continuation A**

Table 1: The first half of the form for the BookCorpus human evaluation.

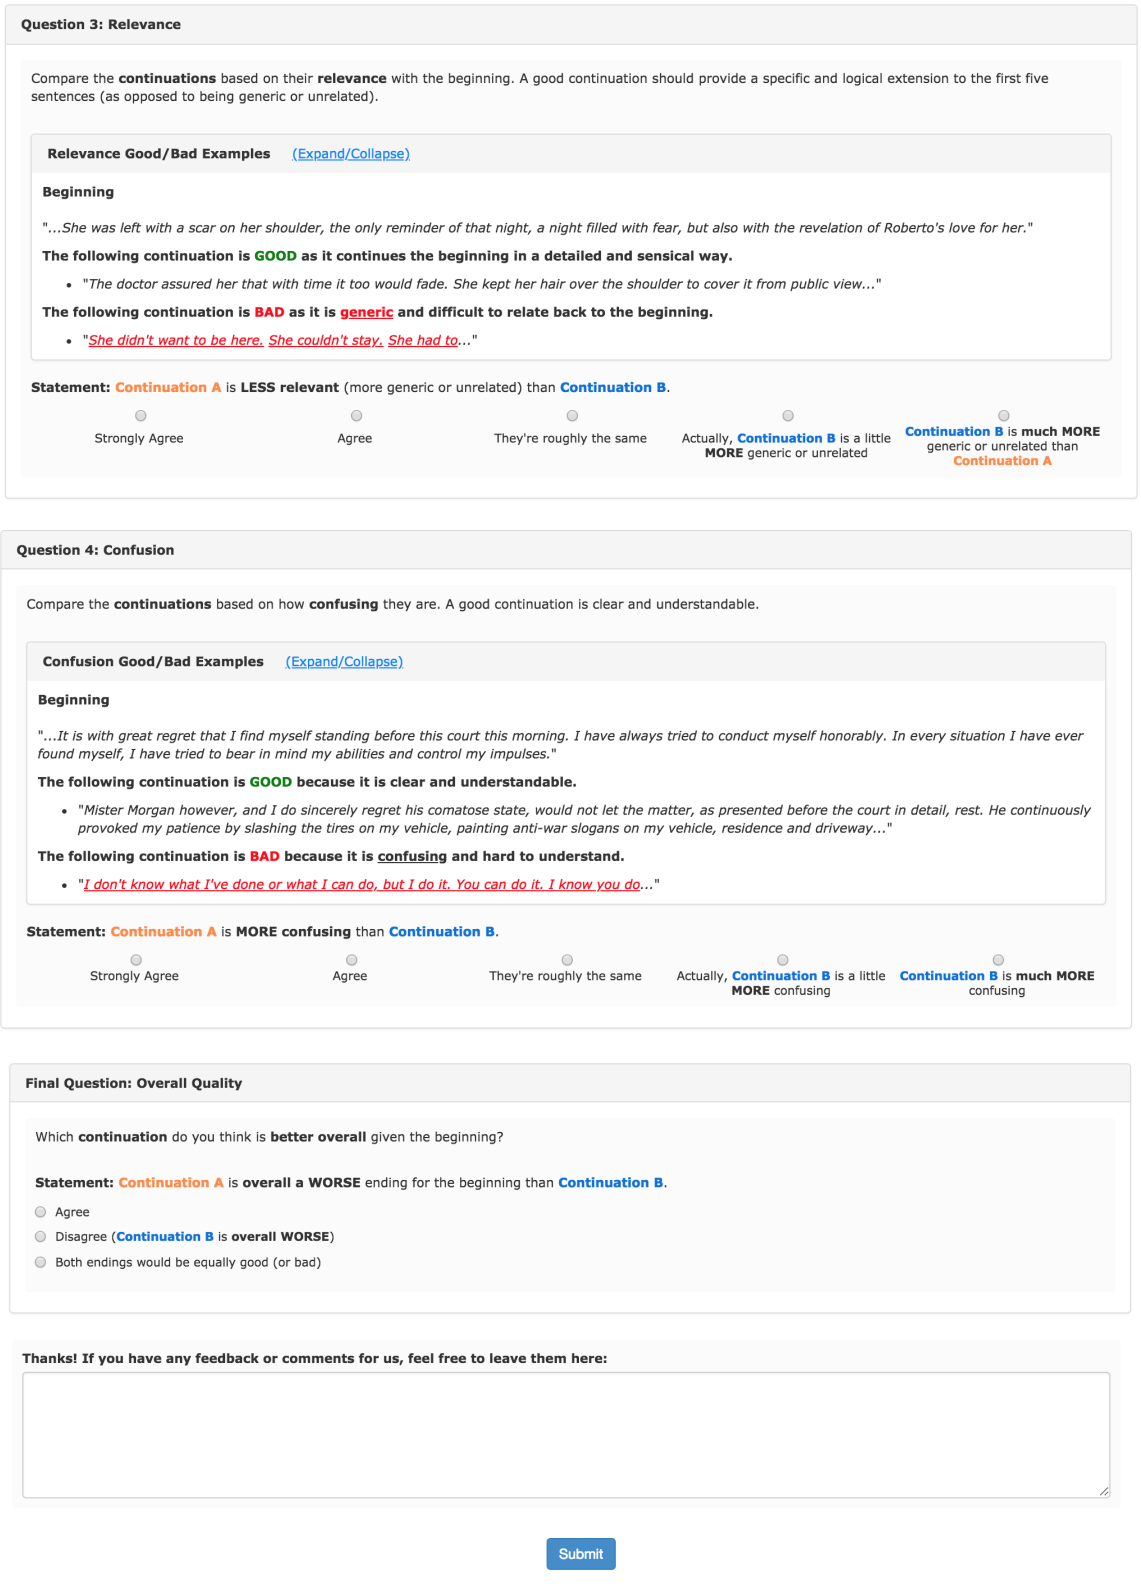

Table 2: The second half of the form for the BookCorpus human evaluation.

TripAdvisor

Instructions

In this task you will be provided with the first 5 sentences of a hotel review, and 2 possible continuations for the review.

We'd like you to tell us whether you agree or disagree with the statements below, which compare and contrast the two continuations.

Thank you!

Task

Beginning Text: We stayed at the hampton inn for a week in mid-aug and thought the hotel was very good. Had requested a room on a high floor that overlooked the pool, and this was exactly what we got. The room was clean and spacious, although surprisingly didn't have either a safe or a fridge. We had arrived in orlando in the middle of a heat wave with temperatures being in the mid 90s and really needed somewhere to cool our drinks, lotion, etc. So we ended up having to order a second ice bucket and use them as a makeshift drinks cooler!

Continuation:

Continuation A The staff at the front desk were very friendly, and helped us with our bags, even though we didn't use them, they were very helpful, and we had no problems. The hotel is in a great location, on the edge of the old town, and only about a 15 minute walk to the heart of the old city, and there is a bus stop right in front of the hotel, and it is just a short bus ride or a cab ride to the old town. All in all, we had a great time, and would definitely go back again, and would recommend this hotel to our friends, and wouldn't hesitate to stay here again. The only thing I would say is that you need to pay for wi-fi

Continuation B We also didn't drink the tap water (there's a tesco nearby! It's about $ 1. 50 to wash & restock, and they'll store your luggage if you're staying for more than one night and don't want to spend an arm and leg). The staff was very helpful in pointing out places to eat, etc. I would probably stay here again because of its location, but don't be fooled by the " free breakfast " - this hotel doesn't offer the " continental breakfast ", with " fresh " fruit juice, " cereal, " or " milk " and " do-it-yourself " waffles. I'm a sucker.

Question 1: Repetitiveness

Compare the **continuations** based on their **repetitiveness**. A good continuation **DOES NOT** repeat details from earlier in the continuation or from the beginning of the review.

Repetition Good/Bad Examples (Expand/Collapse)

Beginning

"*...The pool is very nice with a lazy river.*"

The following continuation is GOOD as it does not repeat previous content.

- "*Our kids had a good time. There's beach volleyball, bikes, golfing, etc....plus you get to enjoy the activities the Ritz has...*"

The following continuation is BAD as it repeats many words and phrases.

- "*The pool is great, and the gym is well equipped. The pool is very nice, and a nice place to relax...*"

Statement: Continuation A repeats itself MORE than Continuation B.

○ Strongly Agree ○ Agree ○ They're roughly the same ○ Actually, **Continuation B** repeats itself a little **MORE** ○ **Continuation B** repeats itself **much MORE** than **Continuation A**

Question 2: Contradiction

Compare the **continuations** based on how much they **contradict** themselves or the beginning of the review. A good continuation should **NOT** contradict itself or the beginning of the review.

Contradiction Good/Bad Examples (Expand/Collapse)

Beginning

"*...I did find the gentleman at the concierge desk a bit cold.*"

The following continuation is GOOD as it does not contradict itself or the beginning of the review.

- "*He told me to wait in line to print boarding passes but when a woman from the AAA convention being held there at the hotel came up he printed hers so she didn't have to wait in line...*"

The following continuation is BAD as it contradicts the beginning of the review.

- "*However, he was very friendly...*"

Statement: Continuation A contradicts itself MORE than Continuation B.

○ Strongly Agree ○ Agree ○ They're roughly the same ○ Actually, **Continuation B** contradicts itself a little **MORE** ○ **Continuation B** contradicts itself **much MORE** than **Continuation A**

Table 3: The first half of the form for the TripAdvisor human evaluation.

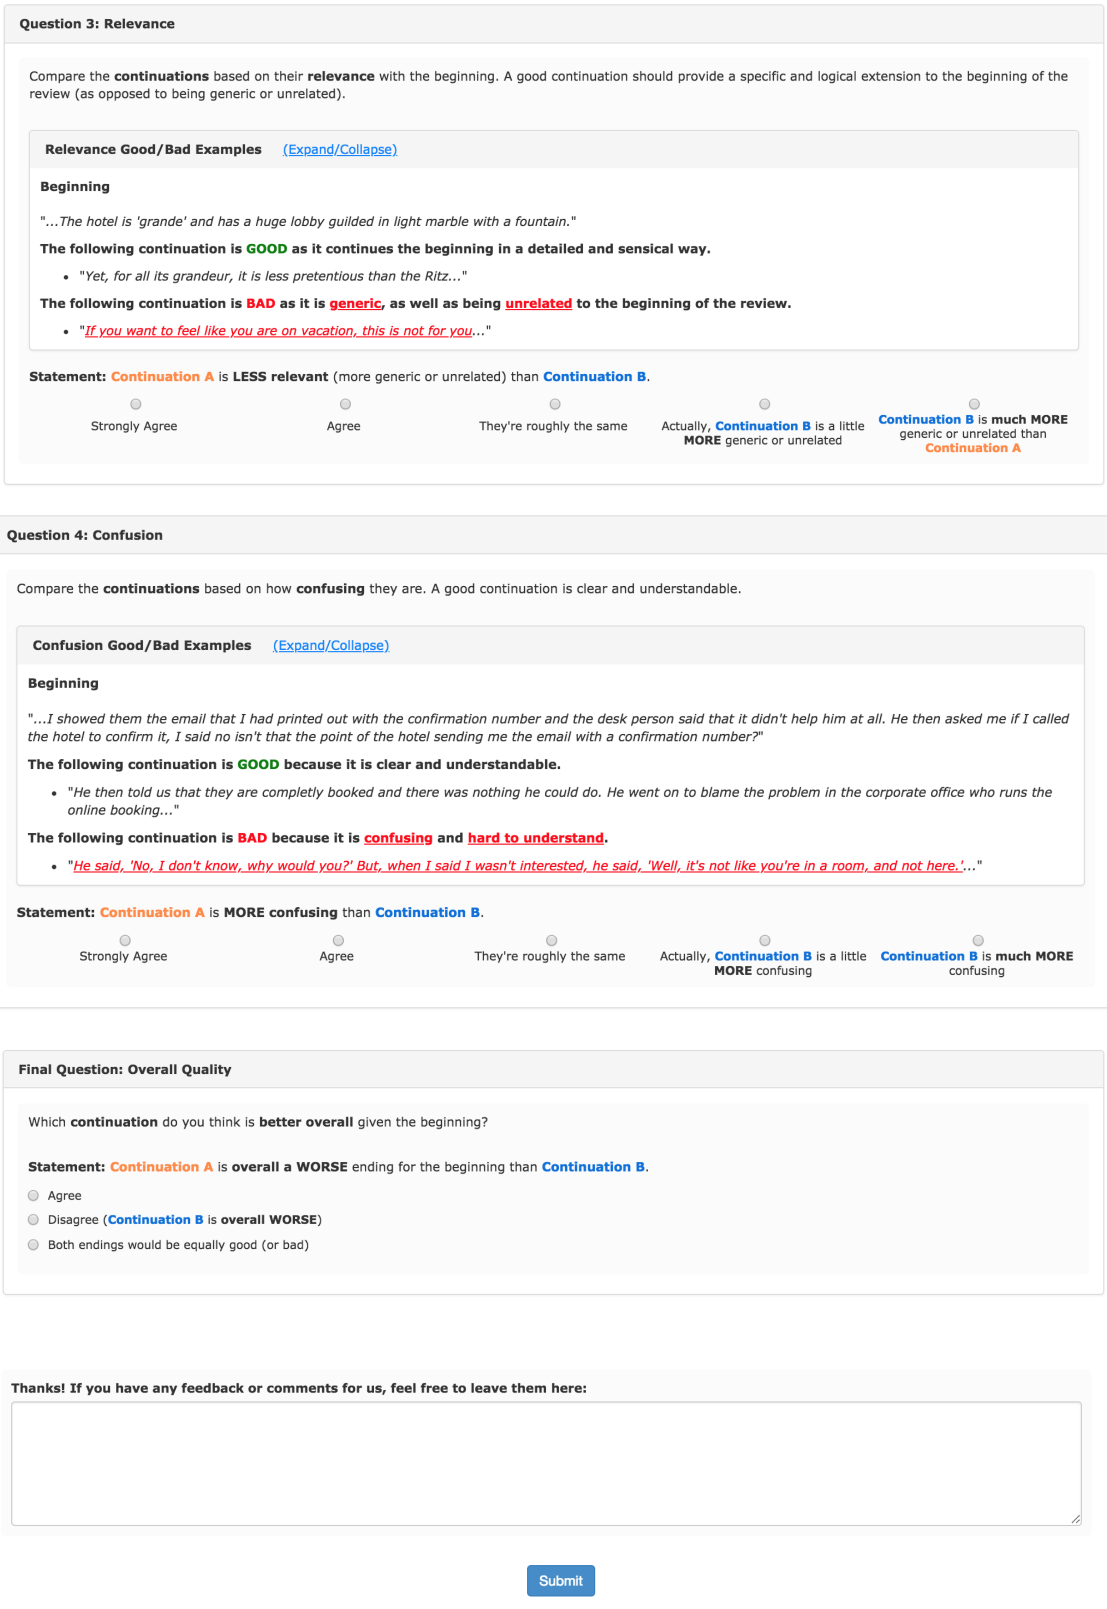

Table 4: The second half of the form for the TripAdvisor human evaluation.