

Learning to Rank for Plausible Plausibility

Zhongyang Li^{†‡*} Tongfei Chen[‡] Benjamin Van Durme[‡]

[†] Harbin Institute of Technology

[‡] Johns Hopkins University

z yli@ir.hit.edu.cn, {tongfei,vandurme}@cs.jhu.edu

Abstract

Researchers illustrate improvements in contextual encoding strategies via resultant performance on a battery of shared Natural Language Understanding (NLU) tasks. Many of these tasks are of a categorical prediction variety: given a conditioning context (e.g., an NLI *premise*), provide a label based on an associated prompt (e.g., an NLI *hypothesis*). The categorical nature of these tasks has led to common use of a cross entropy log-loss objective during training. We suggest this loss is intuitively wrong when applied to *plausibility* tasks, where the prompt by design is neither categorically *entailed* nor *contradictory* given the context. Log-loss naturally drives models to assign scores near 0.0 or 1.0, in contrast to our proposed use of a margin-based loss. Following a discussion of our intuition, we describe a confirmation study based on an extreme, synthetically curated task derived from MultiNLI. We find that a margin-based loss leads to a more plausible model of plausibility. Finally, we illustrate improvements on the Choice Of Plausible Alternative (COPA) task through this change in loss.

1 Introduction

Contextualized encoders such as GPT (Radford et al., 2018) and BERT (Devlin et al., 2019) have led to improvements on various structurally similar Natural Language Understanding (NLU) tasks such as variants of Natural Language Inference (NLI). Such tasks model the conditional interpretation of a sentence (e.g., an NLI *hypothesis*) based on some other context (usually some other sentence, e.g., an NLI *premise*). The structural similarity of these tasks points to a structurally similar modeling approach: (1) concatenate the conditioning context (premise) to a sentence to be interpreted, (2)

*This work was done while the first author was visiting Johns Hopkins University.

p	<i>I just stopped where I was</i>	
h_E	<i>I stopped in my tracks</i>	✓
h_N	<i>I stopped running right were I was</i>	
h_N	<i>I stopped running right were I was</i>	✓
h_C	<i>I continued on my way</i>	

Figure 1: COPA-like pairs may be constructed from datasets such as MultiNLI, where a *premise* and two *hypotheses* are presented, where the correct – most plausible – item depends on the competing hypothesis.




Figure 2: Dev set score distribution on COPA-pairs derived from MNLI, after training with cross entropy log-loss and margin-loss. Margin-loss leads to a more intuitively *plausible* encoding of *Neutral* statements.

read this pair using a contextualized encoder, then (3) employ the resultant representation to support classification under the label set of the task. NLI datasets employ a categorical label scheme (*Entailment*, *Neutral*, *Contradiction*) which has led to the use of a cross-entropy log-loss objective at training time: learn to maximize the probability of the correct label, and thereby minimize the probability of the competing labels.

We suggest that this approach is intuitively problematic when applied to a task such as COPA (Choice Of Plausible Alternative) by Roemmele et al. (2011), where one is provided with a premise and two or more alternatives, and the model must select the most sensible hypothesis, with respect to the premise *and the other options*. As compared

to NLI datasets, COPA was designed to have alternatives that are neither strictly true nor false in context: a procedure that maximizes the probability of the correct item at training time, thereby minimizing the probability of the other alternative(s), will seemingly learn to *misread* future examples.

We argue that COPA-style tasks should intuitively be approached as *learning to rank* problems (Borges et al., 2005; Cao et al., 2007), where an encoder on competing items is trained to assign *relatively* higher or lower scores to candidates, rather than maximizing or minimizing probabilities. In the following we investigate three datasets, beginning with a constructed COPA-style variant of MultiNLI (Williams et al., 2018, later MNLI), designed to be adversarial (see Figure 1). Results on this dataset support our intuition (see Figure 2). We then construct a second synthetic dataset based on JOCI (Zhang et al., 2017), which employed a finer label set than NLI, and a margin-based approach strictly outperforms log-loss in this case. Finally, we demonstrate state-of-the-art on COPA, showing that a BERT-based model trained with margin-loss significantly outperforms a log-loss alternative.

2 Background

A series of efforts have considered COPA: by causality estimation through pointwise mutual information (Gordon et al., 2011) or data-driven methods (Luo et al., 2016; Sasaki et al., 2017), or through a pre-trained language model (Radford et al., 2018, GPT).¹

Under the Johns Hopkins Ordinal Commonsense Inference (JOCI) dataset (Zhang et al., 2017), instead of selecting which hypothesis is the most plausible, a model is expected to directly assign ordinal 5-level Likert scale judgments (from *impossible* to *very likely*). If taking an ordinal interpretation of NLI, this can be viewed as a 5-way variant of the 3-way labels used in SNLI (Bowman et al., 2015) and MNLI (Williams et al., 2018).

In this paper, we recast MNLI and JOCI as COPA-style plausibility tasks by sampling and constructing (p, h, h') triples from these two datasets. Each premise-hypothesis pair (p, h) is labeled with different levels of plausibility $y_{p,h}$.²

¹ As reported in <https://blog.openai.com/language-unsupervised/>.

² For MNLI, *entailment* > *neutral* > *contradiction*; for JOCI, *very likely* > *likely* > *plausible* > *technically possible* > *impossible*.

3 Models

In models based on GPT and BERT for plausibility or NLI, similar neural architectures have been employed. The premise p and hypothesis h are concatenated into a sequence with a special delimiter token, along with a special sentinel token CLS inserted as the token for feature extraction:

$$\begin{aligned} \text{BERT: } & [\text{CLS}; p; \text{SEP}; h; \text{SEP}] \\ \text{GPT: } & [\text{BOS}; p; \text{EOS}; h; \text{CLS}] \end{aligned}$$

The concatenated string is passed into the BERT or GPT encoder. One takes the encoded vector of the CLS state as the feature vector extracted from the (p, h) pair. Given the feature vector, a dense layer is stacked upon to get the final score $F(p, h)$, where $F: \mathcal{P} \times \mathcal{H} \rightarrow \mathbb{R}$ is the model.

Cross entropy loss The model is trained to maximize the probability of the correct candidate, normalized over all candidates in the set (leading to a cross entropy log-loss between the posterior distribution of the scores and the true labels):

$$P(h_i | p) = \frac{\exp F(p, h_i)}{\sum_{j=1}^N \exp F(p, h_j)}. \quad (1)$$

Margin-based loss As we have argued before, the cross entropy loss employed in Equation 1 is problematic. Instead we propose to use the following margin-based triplet loss (Weston and Watkins, 1999; Chechik et al., 2010; Li et al., 2018):

$$L = \frac{1}{N} \sum_{h > h'} \max\{0, \xi - F(p, h) + F(p, h')\}, \quad (2)$$

where N is the number of pairs of hypotheses where the first is more plausible than the second under the given premise p ; $h > h'$ means that h ranks higher than (i.e., is more plausible than) h' under premise p ; and ξ is a margin hyperparameter denoting the desired scores difference between these two hypotheses.

4 Recasting Datasets

We consider three datasets: MNLI, JOCI, and COPA. These are all cast as plausibility datasets, into a format comprising (p, h, h') triples, where h is more plausible than h' under the context of premise p .

Dataset	Train	Eval
MNLI ₁	410k	dev: 8.2k
MNLI ₂	142k	dev: 130k
JOCI ₁	8.7k	dev: 3.0k
JOCI ₂	2.3k	dev: 1.9k
COPA	500	test: 500

Table 1: Statistics of various plausibility datasets. All numbers are numbers of (p, h, h') triplets.

MNLI In MNLI, each premise p is paired with 3 hypotheses. We cast the label on each hypothesis as a relative plausibility judgment, where *entailment* $>$ *neutral* $>$ *contradiction* (we label them as 2, 1, and 0). We construct two 2-choice plausibility tasks from MNLI:

$$\text{MNLI}_1 = \{(p, h, h') \mid y_{p,h} > y_{p,h'}\}$$

$$\text{MNLI}_2 = \{(p, h, h') \mid (y_{p,h}, y_{p,h'}) \in \{(2, 1), (1, 0)\}\}$$

MNLI₁ comprises all pairs labeled with 2/1, 2/0, or 1/0; whereas MNLI₂ removes the presumably easier 2/0 pairs. For MNLI₁, the training set is constructed from the original MNLI training dataset, and the dev set for MNLI₁ is derived from the original MNLI matched dev dataset. For MNLI₂, all of the examples in our training and dev sets is taken from the original MNLI training dataset, hence the same premise exists in both training and dev. This is by our adversarial design: each neutral hypothesis appears either as the preferred (beating contradiction), or dispreferred alternative (beaten by entailment), which is flipped at evaluation time.

JOCI In JOCI, every inference pair is labeled with their ordinal inference Likert-scale labels 5, 4, 3, 2, or 1. Similar to MNLI, we cast these to 2-choice problems under the following conditions:

$$\text{JOCI}_1 = \{(p, h, h') \mid y_{p,h} > y_{p,h'} \geq 3\}$$

$$\text{JOCI}_2 = \{(p, h, h') \mid (y_{p,h}, y_{p,h'}) \in \{(5, 4), (4, 3)\}\}$$

We ignore inference pairs with scores below 3, aiming for sets akin to COPA, where even the dispreferred option is still often semi-plausible.

COPA We label alternatives as 1 (the more plausible one) and 0 (otherwise). The original dev set in COPA is used as the training set.

Table 1 shows the statistics of these datasets.

5 Experiments and Analyses

Setup We fine-tune the BERT-BASE-UNCASED (Devlin et al., 2019) using our proposed margin-

Dataset	Log loss	Margin loss
MNLI ₁	93.6	93.4
MNLI ₂	87.9	87.9
JOCI ₁	86.6	86.9
JOCI ₂	76.6	78.0

Table 2: Results on recast MNLI and JOCI.

based loss, and perform hyperparameter search on the margin parameter ξ .

For the recast MNLI and JOCI datasets, the margin hyperparameter $\xi = 0.2$. Since COPA does not have a training set, we use the original dev set as the training set, and perform 10-fold cross validation to find the best hyperparameter $\xi = 0.37$. We employ the Adam optimizer (Kingma and Ba, 2014) with initial learning rate $\eta = 3 \times 10^{-5}$, fine-tune for at most 3 epochs and use early-stopping to select the best model.

Results on Recast MNLI and JOCI Table 2 shows results on the recast MNLI and JOCI datasets. We find that for the two synthetic MNLI datasets, margin-loss performs similarly to cross entropy log-loss. Shifting to the JOCI datasets, with less extreme (contradiction / entailed) hypotheses, especially in the adversarial JOCI₂ variant, margin-loss outperforms log-loss.

Though log-loss and margin-loss give close quantitative results on predicting the more plausible (p, h) pairs, they do so in different ways, confirming our intuition. From Figure 3 we find that the log-loss always predicts the more plausible (p, h) pair with very high probabilities close to 1, and predicts the less plausible (p, h) pair with very low probabilities close to 0. Figure 3, showing a per-premise normalized score distribution from margin-loss, is more reasonable and explainable: hypotheses with different plausibility are distributed hierarchically between 0 and 1.

Method	Acc (%)
PMI (Jabeen et al., 2014)	58.8
PMI_EX (Gordon et al., 2011)	65.4
CS (Luo et al., 2016)	70.2
CS_MWP (Sasaki et al., 2017)	71.2
BERT _{log} (ours)	73.4
BERT _{margin} (ours)	75.4

Table 3: Experimental results on COPA test set.

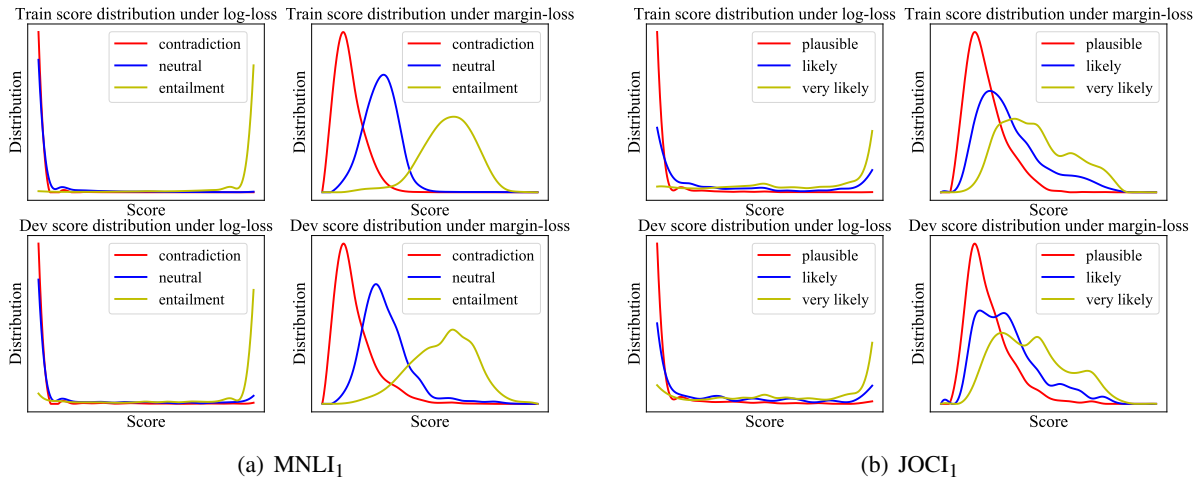


Figure 3: Train and dev score distribution after training with a cross entropy log-loss and a margin-loss.

Dataset	Premise	Hypotheses	Gold	Log	Margin
MNLI ₁	(1) <i>I just stopped where I was.</i>	(a) <i>I stopped in my tracks</i>	2	0.919	0.568
		(b) <i>I stopped running right where I was.</i>	1	0.0807	0.358
		(c) <i>I continued on my way.</i>	0	1.71×10^{-8}	0.0739
MNLI ₁	(2) <i>An organization’s activities, core processes and resources must be aligned to support its mission and help it achieve its goals.</i>	(a) <i>An organization is successful if its activities, resources and goals align.</i>	2	0.505	0.555
		(b) <i>Achieving organizational goals reflects a change in core processes.</i>	1	0.495	0.257
		(c) <i>A company’s mission can be realized even without the alignment of resources.</i>	0	3.48×10^{-5}	0.187
JOCI ₁	(3) <i>A few people and cars out on their daily commute on a rainy day.</i>	(a) <i>The commute is a journey.</i>	5	0.994	0.473
		(b) <i>The commute is bad.</i>	4	5.79×10^{-3}	0.230
		(c) <i>The commute becomes difficult.</i>	3	1.28×10^{-3}	0.157
JOCI ₁	(4) <i>Cheerleaders in red uniforms perform a lift stunt.</i>	(a) <i>The stunt is a feat.</i>	5	0.508	0.304
		(b) <i>The stunt is no fluke.</i>	4	0.486	0.279
		(c) <i>The stunt is dangerous.</i>	3	2.72×10^{-4}	0.166
		(d) <i>The stunt is remarkable.</i>	3	4.13×10^{-3}	0.153
		(e) <i>The stunt backfires.</i>	3	2.36×10^{-4}	0.107
COPA	(5) <i>She jumped off the diving board.</i> (5') <i>She ran on the pool deck.</i>	(a) <i>The girl landed in the pool.</i>	1	0.972	0.520
			0	0.028	0.480
COPA	(6) <i>The student knew the answer to the question.</i>	(a) <i>He raised his hand.</i>	1	0.982	0.738
		(b) <i>He goofed off.</i>	0	0.018	0.262

Table 4: Examples of premises and their corresponding hypotheses in various plausibility datasets, with gold labels and scores given by the log-loss and margin-loss trained models.

Results on COPA Table 3 shows our results on COPA. Compared with previous state-of-the-art knowledge-driven baseline methods, a BERT model trained with a log-loss achieves better performance. When training the BERT model with a margin-loss instead of a log-loss, our method gets the new state-of-the-art result on the established COPA splits, with an accuracy of 75.4%.³

Analyses Table 4 shows some examples from the MNLI₁, JOCI₁ and COPA datasets, with scores

³ We exclude a blog-posted GPT result, which comes without experimental conditions and is not reproducible.

normalized with respect to all hypotheses given a specific premise.

For the premise (1) from MNLI₁, log-loss results in a very high score (0.919) for the entailment hypothesis (1a), while assigning a low score (0.0807) for the neutral hypothesis (1b), and an extremely low score (1.71×10^{-8}) for the contradiction hypothesis (1c). Though the log-loss can achieve high accuracy by making these extreme prediction scores, we argue these scores are unintuitive. For the premise (2) from MNLI₁, log-loss again gives a very high score (0.505) for the hypothesis (2a).

But it also gives a high score (0.495) for the neutral hypothesis (2b). The contradiction hypothesis (2c) still gets an extremely low score (3.48×10^{-5}).

These are the two ways for the log-loss approach to make predictions with high accuracy: always giving very high score for the entailment hypothesis and low score for the contradiction hypothesis, but giving either very high or very low score for the neutral hypothesis. In contrast, the margin-loss gives more intuitive scores for these two examples. Also, we get similar observations from the JOCI₁ examples (3) and (4).

Example (5) from COPA is asking for a more plausible *cause* premise for the *effect* hypothesis. Here, each of the two candidate premises (5) and (5') is a possible answer. The log-loss gives very high (0.972) and very low (0.028) scores for the two candidate premises, which is unreasonable. Whereas the margin-loss gives much more rational ranking scores for them (0.52 and 0.48). For example (6), which is asking for a more likely effect hypothesis for the cause premise, margin-loss still gets more reasonable prediction scores than the log-loss.

Our qualitative analysis is related to the concept of *calibration* in statistics: are these resulting scores close to their class membership probabilities? Our intuitive qualitative results might be thought as a type of calibration for the plausibility task (more “reliable” scores) instead of the more common multi-class classification (Zadrozny and Elkan, 2002; Hastie and Tibshirani, 1998; Niculescu-Mizil and Caruana, 2005).

6 Conclusion

In this paper, we propose that margin-loss in contrast to log-loss is a more plausible training objective for COPA-style *plausibility* tasks. Through adversarial construction we illustrated that a log-loss approach can be driven to encode plausible statements (Neutral hypotheses in NLI) as either extremely likely or unlikely, which was highlighted in contrasting figures of per-premise normalized hypothesis scores. This intuition was shown to lead to a new state-of-the-art in the original COPA task, based on a margin-based loss.

Acknowledgements

This work was partially sponsored by the China Scholarship Council. It was also supported in part by DARPA AIDA. The authors thank the reviewers

for their helpful comments.

References

- Samuel R Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. 2015. A large annotated corpus for learning natural language inference. In *Proc. EMNLP*, pages 632–642.
- Christopher Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Gregory N Hullender. 2005. Learning to rank using gradient descent. In *Proc. ICML*, pages 89–96.
- Zhe Cao, Tao Qin, Tie-Yan Liu, Ming-Feng Tsai, and Hang Li. 2007. Learning to rank: from pairwise approach to listwise approach. In *Proc. ICML*, pages 129–136. ACM.
- Gal Chechik, Varun Sharma, Uri Shalit, and Samy Bengio. 2010. Large scale online learning of image similarity through ranking. *J. Mach. Learn. Res.*, 11(3):1109–1135.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proc. NAACL*.
- Andrew S Gordon, Cosmin A Bejan, and Kenji Sagae. 2011. Commonsense causal reasoning using millions of personal stories. In *Proc. AAAI*.
- Trevor Hastie and Robert Tibshirani. 1998. Classification by pairwise coupling. In *Proc. NeurIPS*, pages 507–513.
- Shahida Jabeen, Xiaoying Gao, and Peter Andrae. 2014. Using asymmetric associations for commonsense causality detection. In *Pacific Rim International Conference on Artificial Intelligence*, pages 877–883. Springer.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Zhongyang Li, Xiao Ding, and Ting Liu. 2018. Constructing narrative event evolutionary graph for script event prediction. In *Proc. IJCAI*, pages 4201–4207. AAAI Press.
- Zhiyi Luo, Yuchen Sha, Kenny Q Zhu, Seung-won Hwang, and Zhongyuan Wang. 2016. Commonsense causal reasoning between short texts. In *Fifteenth International Conference on the Principles of Knowledge Representation and Reasoning*.
- Alexandru Niculescu-Mizil and Rich Caruana. 2005. Predicting good probabilities with supervised learning. In *Proc. ICML*, pages 625–632. ACM.

- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. URL <https://s3-us-west-2.amazonaws.com/openai-assets/research-covers/languageunsupervised/language-understanding-paper.pdf>.
- Melissa Roemmele, Cosmin Adrian Bejan, and Andrew S Gordon. 2011. Choice of plausible alternatives: An evaluation of commonsense causal reasoning. In *2011 AAAI Spring Symposium Series*.
- Shota Sasaki, Sho Takase, Naoya Inoue, Naoaki Okazaki, and Kentaro Inui. 2017. Handling multiword expressions in causality estimation. In *IWCS 12th International Conference on Computational Semantics Short papers*.
- Jason Weston and Chris Watkins. 1999. Support vector machines for multi-class pattern recognition. In *ESANN 1999, 7th European Symposium on Artificial Neural Networks, Bruges, Belgium, April 21-23, 1999, Proceedings*, pages 219–224.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *Proc. NAACL*, pages 1112–1122. Association for Computational Linguistics.
- Bianca Zadrozny and Charles Elkan. 2002. Transforming classifier scores into accurate multiclass probability estimates. In *Proc. KDD*, pages 694–699. ACM.
- Sheng Zhang, Rachel Rudinger, Kevin Duh, and Benjamin Van Durme. 2017. Ordinal common-sense inference. *Trans. ACL*, 5(1):379–395.