

Joint Energy-based Model Training for Better Calibrated Natural Language Understanding Models

Tianxing He

MIT

cloudygoose@csail.mit.edu

Bryan McCann*

Salesforce Research

bryan@you.com

Caiming Xiong

Salesforce Research

cxiong@salesforce.com

Ehsan Hosseini-Asl

Salesforce Research

ehosseiniasl@salesforce.com

Abstract

In this work, we explore joint energy-based model (EBM) training during the finetuning of pretrained text encoders (e.g., Roberta) for natural language understanding (NLU) tasks. Our experiments show that EBM training can help the model reach a better calibration that is competitive to strong baselines, with little or no loss in accuracy. We discuss three variants of energy functions (namely *scalar*, *hidden*, and *sharp-hidden*) that can be defined on top of a text encoder, and compare them in experiments. Due to the discreteness of text data, we adopt noise contrastive estimation (NCE) to train the energy-based model. To make NCE training more effective, we train an autoregressive noise model with the masked language model (MLM) objective.

1 Introduction

Calibration refers to how well a classification model’s confidence (reflected by its output posterior probability) aligns with its actual accuracy. As deep learning models achieve amazing accuracy in computer vision (He et al., 2015) or natural language processing (NLP) (Liu et al., 2019; Devlin et al., 2018), more research attention has been drawn to the calibration aspect of these models. As shown by Guo et al. (2017), the high accuracy from deep models does not always lead to better calibration. This motivates an important line of works attempting to achieve a better trade-off between accuracy and calibration.

Most existing calibration methods (Guo et al., 2017; Kumar et al., 2019; Zadrozny and Elkan, 2001) generally rescale the posterior distribution predicted from the classifier after training. Such post-processing methods require a held-out development set with a decent number of samples to be

* Bryan McCann contributed to this work while he was at Salesforce Research.

available. To overcome this constraint, Jung et al. (2020) uses a penalty term to encourage better calibration during training.

In another line of work, Grathwohl et al. (2019) shows that one can jointly train an energy-based model (EBM) during the standard training of a neural classifier. Although calibration is not explicitly addressed during EBM training, the calibration of the resulting model is shown to be greatly improved. Some intuitions of the underlying reasons will be given in Section 2.3. However, the training framework proposed by Grathwohl et al. (2019) is designed for image classifiers, and it can not be readily applied to discrete text data.

In this work, we propose a framework that uses noise contrastive estimation (NCE) to jointly train an energy-based model during the finetuning of pretrained text encoders (e.g., BERT (Devlin et al., 2018) or Roberta (Liu et al., 2019)) for NLU tasks. We compare several variants of energy functions that can be defined on top of the encoder. Our experiments show that the resulting models achieve competitive calibration results comparing to strong baselines, with little or no loss in accuracy.

2 Framework

2.1 Notations and Background

We focus on the finetuning of pretrained text encoder on NLU tasks. We assume samples from the data distribution P_D are in the form of (x, y) pairs, where x usually refers to a single or a pair of sentences, and y refers to the corresponding label. The number of classes are denoted by $|Y|$.

Given input x , we first use a text encoder model (e.g., BERT or Roberta) to encode it and we denote this embedding as $\text{enc}(x)$. For the target classification task, a classifier f_{CLS} , which could be a simple linear transform or a multi-layer perception (MLP), will be applied to $\text{enc}(x)$. We denote the output

logits as $f_{\text{CLS}}(\text{enc}(x))$, whose dimension is equal to the number of possible classes $|Y|$. The y -th logit is denoted by $f_{\text{CLS}}(\text{enc}(x))[y]$. The posterior distribution $P_{\theta}(y|x)$ is obtained by applying a softmax operation to the logits, where θ refers to the parameters in the model.

In standard finetuning, the cross-entropy (CE) loss and gradient based optimizers are used to train the classifier:

$$\mathcal{L}_{\text{CE}} = \mathbb{E}_{(x,y) \sim P_D} (-\log P_{\theta}(y|x)). \quad (1)$$

In the next few sections, we discuss how we define and jointly train an EBM on top of the text encoder.

2.2 Definitions of Energy Function

An energy-based model (LeCun et al., 2006) expresses $P_{\theta}(x)$ as:

$$P_{\theta}(x) = \frac{\exp(-E_{\theta}(x))}{Z}, \quad (2)$$

where Z is the normalization factor, and is usually intractable to compute. We refer to $E_{\theta}(x)$, which returns a scalar value, as *the energy function*. We now define three variants of energy functions.

Variant *scalar*: We introduce another linear layer g_S whose output is a scalar. And we use it to define the energy function:

$$\hat{E}_{\theta}(x) = g_S(\text{enc}(x)). \quad (3)$$

Variant *hidden*: As pointed out by Grathwohl et al. (2019), there’s an EBM “hidden” in every neural classifier with softmax output, and the energy function for x can be derived¹ as:

$$\hat{E}_{\theta}(x) = -\text{LogSumExp}_{y=1}^{|Y|}(f_{\text{CLS}}(\text{enc}(x))[y]). \quad (4)$$

Different from the *scalar* variant, here the energy function directly uses the logits for prediction (visualized in Figure 1). Hence the impact on the model’s classification behavior could be larger.

Variant *sharp-hidden*: The *hidden* variant has a potential weakness that, the correlation between input x and the prediction y is not addressed because the energy is distributed among all the logits. Motivated by this potential issue, we propose the following “sharp” variant:

$$\hat{E}_{\theta}(x) = -\max_y f_{\text{CLS}}(\text{enc}(x))[y]. \quad (5)$$

Note that (5) can be viewed as an approximation to (4), and we find it to work well in practice.

¹Please see Appendix A for the detailed derivation.




Figure 1: Comparison of the *scalar* and the *hidden* variants of energy functions. The modules introduced for EBM are shaded in green.

Finally, for each variant, we define the energy function to be $E_{\theta}(x) = \hat{E}_{\theta}(x) - \log P_N(x)$, where P_N is the *noise distribution* introduced for NCE. We will motivate this design choice below.

2.3 NCE Training

We use noise contrastive estimation (NCE) (Gutmann and Hyvärinen, 2010; Ma and Collins, 2018) to jointly train the energy model. NCE trains the model to discriminate between data samples and noise samples from a given noise distribution P_N . We formulate the NCE loss below:

$$\mathcal{L}_{\text{NCE}} = \mathbb{E}_{x_+ \sim P_D} -\log \frac{\tilde{P}_{\theta}(x_+)}{\tilde{P}_{\theta}(x_+) + K \cdot P_N(x_+)} + K \cdot \mathbb{E}_{x_- \sim P_N} -\log \frac{K \cdot P_N(x_-)}{\tilde{P}_{\theta}(x_-) + K \cdot P_N(x_-)}, \quad (6)$$

where K is the ratio of noise samples. Note that $\tilde{P}_{\theta}(x)$ does not need to be normalized by construction, therefore we set it to be $\tilde{P}_{\theta}(x) = \exp(-E_{\theta}(x))$. In our experiments, we mostly report results with noise ratio $K = 8$, while in some cases we find that a small ratio of $K = 1$ works slightly better. We have also tried with larger ratio such as 16, but the gain is minimal.

If we directly use the formulations of $\hat{E}_{\theta}(x)$ defined in last section as the energy function, the optimization will be difficult because of the $P_N(x)$ terms (which could be of very small value) in the NCE objective. To overcome this issue, we follow Deng et al. (2020) and define $E_{\theta}(x) = \hat{E}_{\theta}(x) - \log P_N(x)$. In this way, the $P_N(x)$ terms are canceled, and the objective is simplified to:

$$\mathcal{L}_{\text{NCE}} = \mathbb{E}_{x_+ \sim P_D} -\log \frac{1}{1 + K \cdot \exp(\hat{E}_{\theta}(x_+))} + K \cdot \mathbb{E}_{x_- \sim P_N} -\log \frac{K}{K + \exp(-\hat{E}_{\theta}(x_-))}. \quad (7)$$

In training, we jointly optimize \mathcal{L}_{CE} and \mathcal{L}_{NCE} with the Adam optimizer (Kingma and Ba, 2014):

$$\mathcal{L}_{\text{joint}} = \mathcal{L}_{\text{CE}} + \mathcal{L}_{\text{NCE}}. \quad (8)$$

Intuitively, joint EBM training makes the model aware of $P(x)$, instead of only focusing on predicting $P(y|x)$ as in standard finetuning. This awareness can potentially help with calibration because the model can be more conservative when it detects the input is out-of-distribution.

2.4 Construction of Noise Distribution

For the choice of noise distribution P_N , in our preliminary trials, we finetune the GPT-2 language model (Radford et al., 2019) with samples from the target training set using the standard LM objective. However during NCE training, we find that the energy model can easily discriminate between data samples and noise samples, which makes training ineffective. To alleviate this issue, we adopt an objective similar² to the masked language model (MLM) loss (Devlin et al., 2018) during the finetuning of the noise model (GPT-2): With a given mask ratio M , we randomly mask part of x , and train the model to complete it:

$$\mathcal{L}_{\text{MLM}} = \mathbb{E}_{x \sim P_D, x^m \sim P_{\text{mask}}(x^m|x;M)} -\log P_N(x|x^m). \quad (9)$$

During noise sample generation, adopting the same mask ratio M , we feed a masked x^m to the LM (x is from the training set), and use the generated sample as the noise sample. In this way, the noise distribution is made closer to the data distribution. In our experiments we set $M = 0.4$. During generation, we use top- k (Fan et al., 2018) sampling with $k = 20$. More details are provided in Appendix B.

3 Experiments

Setting We consider finetuning the Roberta-base model³, on eight GLUE tasks (Wang et al., 2018). We do not include results on STS-B because it is a regression task. To measure calibration error, we follow Jung et al. (2020); Grathwohl et al. (2019) and use the expected calibration error (ECE) metric with B (number of bins) set to 20. To save space, we defer detailed definition of ECE to Appendix C.

For baseline or NCE training, we follow the recommended hyper-parameters (learning rate, batch size, etc.) for Roberta (Liu et al., 2019). Since NCE training requires more computation (because of the noise ratio), we have tried finetuning the baseline with more steps, but we find that gives worse ECE and very little or no improvement on accuracy.

²The difference is that we still train the model to generate the full sentence, instead of only the masked words.

³Our code is based on <https://github.com/huggingface/transformers>.

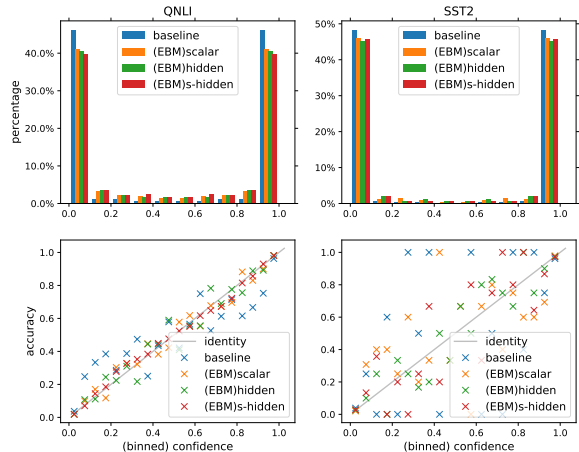


Figure 2: Visualization of calibration on QNLI and SST-2. In the histogram plots, we use 10 bins instead of 20 for better readability. An enlarged version of this figure is provided in Appendix D.

We compare EBM training with three strong baselines for calibration: *posterior calibrated training* (PosCal) (Jung et al., 2020), *temperature scaling* (T-Scal) (Guo et al., 2017), and *scaling-binning calibrator* (Scal-bin) (Kumar et al., 2019). For PosCal and Scal-bin, we use the published code.

Scal-bin and T-Scal require a development set for parameter learning and a test set for evaluation, but for each GLUE task we only have one labeled development set available. Therefore, in this work we treat half of the standard development set as test set, and keep the other half as development set.

Results In Table 1 and Table 2 we compare test-set accuracy⁴ and ECE for different methods on the GLUE tasks. For fair comparison between Scal-bin / T-Scal and EBM training (which does not use the development set), we apply them to the whole training set. We also report their performance when applied to the development set for reference.

In most tasks, all three EBM variants get substantial improvement in ECE with little or no loss in accuracy comparing to the (strong) baseline methods. Moreover, the performance of EBM training is comparable to Scal-bin / T-Scal applied to the development set, while their performance degrades when the development set is not available. Among the three variants, on average, the *sharp-hidden* variant achieves the best accuracy, while the *hidden* variant achieves best calibration. We visualize the calibration error in Figure 2.

⁴For CoLA we report with Matthews correlation coefficient (mcc).

Method	SST-2		MNLI		MNLI(mm)		QNLI		QQP		MRPC		CoLA		Average	
	acc.	ECE	acc.	ECE	acc.	ECE	acc.	ECE	acc.	ECE	acc.	ECE	mcc.	ECE	perf.	ECE
Baseline	.942	.050	.876	.067	.872	.068	.929	.043	.904	.034	.862	.133	.539	.182	.802	.102
Scal-bin(train)	.940	.036	.872	.051	.869	.056	.931	.034	.904	.035	.843	.092	.586	.146	.791	.096
T-Scal(train)	.942	.042	.876	.058	.872	.060	.929	.030	.904	.034	.862	.126	.539	.175	.802	.096
PosCal	.944	.040	.876	.067	.872	.067	.930	.039	.905	.032	.867	.129	.540	.184	.810	.092
(EBM)scalar	.942	.033	.871	.038	.871	.047	.927	.016	.899	.034	.862	.098	.540	.150	.801	.073
(EBM)hidden	.956	.032	.869	.032	.868	.044	.923	.016	.900	.033	.867	.099	.545	.131	.807	.063
(EBM)s-hidden	.947	.038	.875	.027	.872	.031	.930	.016	.900	.032	.862	.089	.563	.133	.815	.069
Scal-bin(dev)	.944	.019	.876	.030	.870	.032	.931	.021	.905	.021	.862	.062	.557	.048	.802	.052
T-Scal(dev)	.942	.037	.876	.024	.872	.026	.929	.018	.904	.026	.862	.126	.539	.109	.802	.072

Table 1: Test-set accuracy and ECE results for different methods on GLUE tasks. “s-hidden” refers to the *sharp-hidden* variant. The leading zeros are omitted to save space. Note that T-Scal and Scal-bin are applied to the training set or the development set, respectively. Due to space constraint, results on RTE and WNLI are deferred to Table 2. The average value is compute on all nine test sets. For each task, the method that achieves best calibration without using the development set are shown in bold.

Method	RTE		WNLI	
	acc.	ECE	acc.	ECE
Baseline	.724	.279	.571	.058
Scal-bin(train)	.717	.271	.457	.144
T-Scal(train)	.724	.275	.571	.063
PosCal	.789	.206	.571	.060
(EBM)scalar	.753	.207	.542	.033
(EBM)hidden	.797	.148	.542	.036
(EBM)s-hidden	.811	.182	.571	.073
Scal-bin(dev)	.731	.042	.542	.189
T-Scal(dev)	.724	.235	.571	.046

Table 2: (Following Table 1) Main results on RTE and WNLI.




Figure 3: (QNLI) Left: How ECE changes during training. Right: The trade-off between accuracy and ECE for checkpoints (every 500 iterations) during training.




Figure 4: The entropy of the posterior ($P_\theta(\cdot|x)$) versus energy value $\hat{E}_\theta(x)$ for SST-2 test-set samples.

Text: *when the film ended, i felt tired and drained and wanted to lie on my own deathbed.* Label: 1
 $\hat{E}_\theta(x)$: -9.37 Baseline: (.999, .001) \rightarrow EBM: (.998, .002)

Text: *sit through this one, you won't need a magic watch to stop time; your dvd player will do it for you.* Label: 1
 $\hat{E}_\theta(x)$: -7.57 Baseline: (.006, .994) \rightarrow EBM: (.345, .655)

Table 3: The change of the model’s confidence (posterior distribution) for low and high-energy data samples of SST-2. The EBM variant shown is *sharp-hidden*. We also provide QNLI examples in Appendix D.

In Figure 3, we plot how test-set ECE changes during training. It is shown as the training reaches the high-accuracy area, the calibration for baseline model becomes worse, while EBM training is able to reach a better trade-off between accuracy and calibration.

How does the model get better calibration? In Figure 4, we compute and plot the energy value $\hat{E}_\theta(x)$ versus the entropy of the posterior distribution $\mathcal{H}(P_\theta(\cdot|x)) = \sum_{y=1}^{|Y|} -P_\theta(y|x) \log P_\theta(y|x)$, for samples in the SST-2 test set. It is shown that models trained with the *hidden* and *sharp-hidden* variants tend to assign more conservative predictions (reflected by higher entropy) for higher-energy (less likely) samples. We suspect this is due to the strong coupling between the energy function and the classification logits. We provide concrete examples in Table 3. However, we need to mention that we do not observe this interesting trend (Figure 4) in all datasets (e.g., QNLI).

4 Related Works

Finally, we review applications of NCE or energy-based models in the NLP literature. Due to its self-normalizing property, NCE training has been used for faster inference (Mnih and Teh, 2012; Chen et al., 2015; Labeau and Allauzen, 2018) of auto-

regressive language models. It has also been used in an attempt to train a sentence-level bi-directional LM (He et al., 2016).

More closely related to our work, Deng et al. (2020) adopts NCE to train an EBM defined on top of a text encoder (the *scalar* variant), and uses it to improve language generation. EBM has also been recently used in non-autoregressive machine translation (Tu et al., 2020).

5 Conclusion

In this work, we explore joint EBM training during the finetuning of pretrained text encoders with noise contrastive estimation. We find that joint EBM training can greatly improve the calibration of NLU models, with little or no loss on accuracy.

References

- Xie Chen, Xunying Liu, Mark J. F. Gales, and Philip C. Woodland. 2015. [Recurrent neural network language model training with noise contrastive estimation for speech recognition](#). In *ICASSP*, pages 5411–5415. IEEE.
- Yuntian Deng, Anton Bakhtin, Myle Ott, Arthur Szlam, and Marc’Aurelio Ranzato. 2020. [Residual energy-based models for text generation](#). In *International Conference on Learning Representations*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: pre-training of deep bidirectional transformers for language understanding](#). *CoRR*, abs/1810.04805.
- Angela Fan, Mike Lewis, and Yann Dauphin. 2018. [Hierarchical neural story generation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898, Melbourne, Australia. Association for Computational Linguistics.
- Will Grathwohl, Kuan-Chieh Wang, Jörn-Henrik Jacobsen, David Duvenaud, Mohammad Norouzi, and Kevin Swersky. 2019. [Your classifier is secretly an energy based model and you should treat it like one](#).
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. 2017. [On calibration of modern neural networks](#). volume 70 of *Proceedings of Machine Learning Research*, pages 1321–1330, International Convention Centre, Sydney, Australia. PMLR.
- Michael Gutmann and Aapo Hyvärinen. 2010. [Noise-contrastive estimation: A new estimation principle for unnormalized statistical models](#). In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pages 297–304, Chia Laguna Resort, Sardinia, Italy. PMLR.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. [Deep residual learning for image recognition](#). *arXiv preprint arXiv:1512.03385*.
- Tianxing He, Yu Zhang, Jasha Droppo, and Kai Yu. 2016. [On training bi-directional neural network language model with noise contrastive estimation](#). *CoRR*, abs/1602.06064.
- Taehee Jung, Dongyeop Kang, Hua Cheng, Lucas Mentch, and Thomas Schaaf. 2020. [Posterior calibrated training on sentence classification tasks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2723–2730, Online. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). Cite arxiv:1412.6980Comment: Published as a conference paper at the 3rd International Conference for Learning Representations, San Diego, 2015.
- Ananya Kumar, Percy S Liang, and Tengyu Ma. 2019. [Verified uncertainty calibration](#). In *Advances in Neural Information Processing Systems 32*, pages 3792–3803. Curran Associates, Inc.
- Matthieu Labeau and Alexandre Allauzen. 2018. [Learning with noise-contrastive estimation: Easing training by learning to scale](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3090–3101, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Yann LeCun, Sumit Chopra, Raia Hadsell, Fu Jie Huang, and et al. 2006. [A tutorial on energy-based learning](#). In *PREDICTING STRUCTURED DATA*. MIT Press.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized BERT pretraining approach](#). *CoRR*, abs/1907.11692.
- Zhuang Ma and Michael Collins. 2018. [Noise contrastive estimation and negative sampling for conditional models: Consistency and statistical efficiency](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3698–3707, Brussels, Belgium. Association for Computational Linguistics.
- Andriy Mnih and Yee Whye Teh. 2012. [A fast and simple algorithm for training neural probabilistic language models](#). In *Proceedings of the 29th International Conference on Machine Learning, ICML’12*, page 419–426, Madison, WI, USA. Omnipress.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#).

Lifu Tu, Richard Yuanzhe Pang, Sam Wiseman, and Kevin Gimpel. 2020. [ENGINE: Energy-based inference networks for non-autoregressive machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2819–2826, Online. Association for Computational Linguistics.

Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.

Bianca Zadrozny and Charles Elkan. 2001. Obtaining calibrated probability estimates from decision trees and naive bayesian classifiers. In *In Proceedings of the Eighteenth International Conference on Machine Learning*, pages 609–616. Morgan Kaufmann.

Appendices

A Derivation of the *hidden* Variant

Remember from Section 2.1, the posterior distribution is obtained from a softmax operation on the logits, in other words:

$$P_{\theta}(y|x) \propto \exp(f_{\text{CLS}}(\text{enc}(x))[y]). \quad (10)$$

Without changing any parameters, one can reuse the logits to define an energy based model of the joint distribution of data point x and labels y via:

$$P_{\theta}(x, y) = \frac{\exp(f_{\text{CLS}}(\text{enc}(x))[y])}{Z(\theta)}, \quad (11)$$

where $Z(\theta)$ is the normalizing factor. Note that Equation 11 is **consistent** with Equation 10 in the sense that Equation 10 is a direct consequence of Equation 11.

Now by marginalizing out y , we get:

$$P_{\theta}(x) = \frac{\sum_{y=1}^{|Y|} \exp(f_{\text{CLS}}(\text{enc}(x))[y])}{Z(\theta)}, \quad (12)$$

which is equivalent to

$$P_{\theta}(x) = \frac{\exp(-E_{\theta}(x))}{Z(\theta)}, \quad (13)$$

where

$$E_{\theta}(x) = -\text{LogSumExp}_y(f_{\text{CLS}}(\text{enc}(x))[y]). \quad (14)$$

For more intuition behind this derivation we refer readers to Grathwohl et al. (2019).

B Details About the Noise Distribution

We show some examples of generated noise samples and the masking in Table 4. Note that the masks could be applied to a consecutive span of words (Masking is applied to each token independently with probability M).

Input: absolutely and completely <M> (ridiculous)
Gen: absolutely and completely hilarious
Input: <M> (as a) young <M> (woman) of great charm, <M> (generosity) and diplomacy
Gen: of a young man with a great charm, wit and diplomacy

Table 4: Example of generated noise samples on SST-2. The original words that are masked are also shown.

Another possible way to get noise samples is that we can sample from BERT or Roberta with masked

input. However, due to the nature of masked language modeling and the architecture of BERT / Roberta, the sampled tokens will be independent of each other, which could result in unnatural noise samples. That is why we choose to utilize an autoregressive LM (e.g., GPT-2).

C Definition of ECE

Given an input sample x , for each label y , we say that the model predicts that x belongs to label y with confidence $P_{\theta}(y|x)$. Assuming the test-set contains n samples, we will have $n \times |Y|$ predictions.

ECE first partitions all predictions into B equally-spaced bins by its confidence. Following Jung et al. (2020); Grathwohl et al. (2019), we set $B = 20$, which means the width of each bin is 0.05. For example, the first bin contains all predictions that have confidence in the range of $[0, 0.05)$. Then for each bin ECE computes how the average of confidence is different from its actual accuracy:

$$\text{ECE} = \frac{1}{|Y|} \sum_{y=1}^{|Y|} \sum_{b=1}^B \frac{|B_{yb}|}{n} |\text{acc}(B_{yb}) - \text{conf}(B_{yb})|, \quad (15)$$

where n is the number of samples in the test set, and $\text{acc}(B_{yb})$ is simply the ratio of samples (x) whose true label is indeed y in B_{yb} .

D Auxiliary Results and Examples

Examples of the model’s confidence for low and high-energy data samples in QNLI are shown in Table 5.

The histogram of energy values $\hat{E}_{\theta}(x)$ for samples in the test set of QNLI and SST-2 are shown in Figure 5.

In Figure 6, we provide an enlarged version of Figure 2.

Text: *Q: What city north of New York was settled by Huguenots? A: Huguenot immigrants did not disperse or settle in different parts of the country, but rather, formed three societies or congregations; one in the city of New York, another 21 miles north of New York in a town which they named New Rochelle, and a third further upstate in New Paltz.* Label: 1
 $\hat{E}_\theta(x)$: -8.48 Baseline: (.997, .003) \rightarrow EBM: (.995, .005)

Text: *Q: What is the source of oxygen production through electrocatalytic means? A: A similar method is the electrocatalytic O2 evolution from oxides and oxoacids.* Label: 1
 $\hat{E}_\theta(x)$: 4.22 Baseline: (.252, .748) \rightarrow EBM: (.472, .527)

Table 5: The change of the model’s confidence (posterior distribution) for low and high-energy data samples in the test set of QNLI. The EBM variant shown is *sharp-hidden*.




Figure 5: The histogram of energy values $\hat{E}_\theta(x)$ for samples in the test set of QNLI and SST-2.

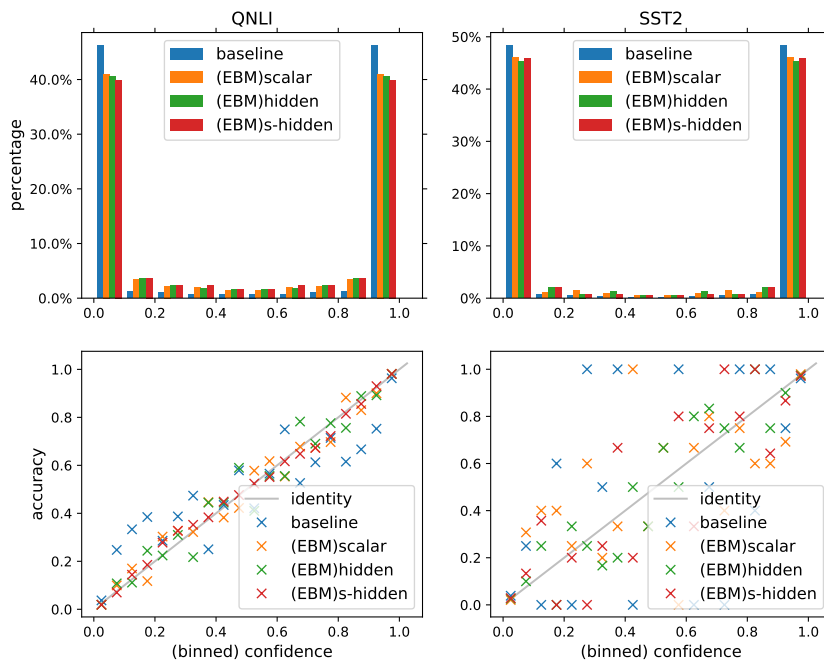


Figure 6: Visualization of calibration on QNLI and SST-2. Enlarged version.