

When in Doubt: Improving Classification Performance with Alternating Normalization

Menglin Jia^{1,2} Austin Reiter² Ser-Nam Lim² Yoav Artzi¹ Claire Cardie¹

¹Cornell University ²Facebook AI

{mj493, yya5, claire}@cornell.edu, {areiter, sernamlim}@fb.com

Abstract

We introduce Classification with Alternating Normalization (CAN), a non-parametric post-processing step for classification. CAN improves classification accuracy for challenging examples by re-adjusting their predicted class probability distribution using the predicted class distributions of high-confidence validation examples. CAN is easily applicable to any probabilistic classifier, with minimal computation overhead. We analyze the properties of CAN using simulated experiments, and empirically demonstrate its effectiveness across a diverse set of classification tasks¹.

1 Introduction

Classification is core to NLP, and many language problems can be effectively addressed as supervised classification tasks. However, even the most effective classifier can suffer when given examples to classify that are close to its decision boundary. The reasons for such failures vary, and include lack of training data coverage, limited representation expressivity, or over-fitting the training data. Despite significant progress, including using pre-trained models (Devlin et al., 2019) to address these issues, every classifier has its weak spots, and some examples will be hard to classify correctly.

In this paper, we study a simple non-parameterized post-processing step to improve classifier accuracy on difficult examples. At the core of our approach is using *Alternating Normalization* (AN; Sinkhorn and Knopp, 1967) to re-adjust the prediction of low-confidence examples using the predicted class distributions of a *reference set* of high-confidence validation examples.

Our process, Classification with Alternating Normalization (CAN), is applicable to any classifier that generates a distribution over target classes. We first

¹Our code for the work are open sourced at github.com/kmnp/can.

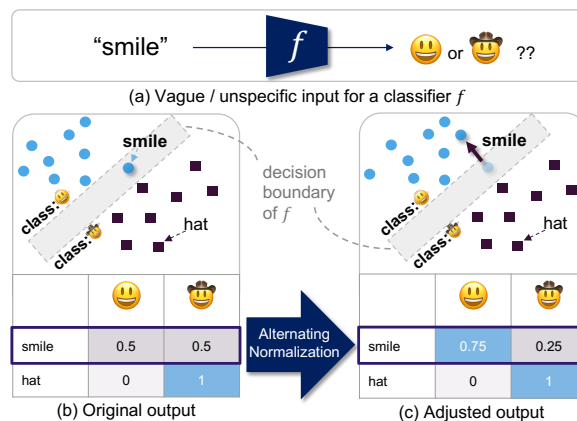


Figure 1: An illustration of Classification with Alternating Normalization (CAN). Given an example “smile”, the classifier f assigns equal probability to both emojis (b). CAN re-scales the class probability distribution and produces a less ambiguous prediction (c).¹

identify challenging examples and a disambiguating reference set using the ambiguity level of the predicted class probability distributions. Then we perform a series of normalizations, alternating between normalizing across examples for each class and for each example across classes. For example, in Figure 1, we classify an input example (“smile”) to one of two labels (😊 and 😐). The example sits at the decision boundary between the two target labels, and is completely ambiguous. The reference set includes a single example (“hat”), which the classifier can resolve with high confidence. Here, we use a single alternating normalization step, including normalizing across rows (examples) and columns (target labels), which disambiguates this simple example to classify it correctly.

We study CAN on randomly generated matrices (Section 3) and evaluate it on several text classification tasks (Section 4). In general, we find that CAN is most effective when the original predictions are of higher ambiguity. Our experiments also suggest

¹We use a classical pragmatic reasoning example for our illustration, highlighting our inspiration in the Rational Speech Act (RSA; Frank and Goodman, 2012) model, which we discuss in Section 6.

CAN is more beneficial for tasks with many labels. All code will be released upon publication.

2 Classification with Alternating Normalization

We modify the output of a probabilistic classifier for an ambiguous input example at test time using alternating normalization with respect to an additional set of examples to re-adjust the input example’s predicted probabilities. Let \mathcal{X} and \mathcal{Y} be random variables for input and labels, respectively. Given a *challenging* example $x \in \mathcal{X}$, and a probability distribution $\mathbb{P}(\mathcal{Y} = y | \mathcal{X} = x)$ over a set of m classes produced by a probabilistic classifier, we compute $\mathbb{P}'(y|x)$ as an adjusted distribution in three steps, the second of which is an iterative process.

Our method requires a reference set of examples. We use the confident portion of the validation data commonly used in standard evaluation settings. Given n confident examples, we create a row stochastic matrix $A_0 \in \mathbb{R}^{n \times m}$ by concatenating the predicted probability distributions. We describe the proposed method below.

Step 1: Identify Hard Examples Let $\mathbf{b}_0 \in \mathbb{R}^m$ be the source prediction, *i.e.*, the predicted class probability distribution for a challenging example x . We identify examples that can benefit from CAN by computing the *ambiguity level* of its class probability distribution \mathbf{b}_0 . With a selected example, we construct a matrix $L_0 \in \mathbb{R}^{(n+1) \times m}$ by concatenating A_0 and \mathbf{b}_0 along the rows:³

$$L_0 = \begin{bmatrix} A_0 \\ \mathbf{b}_0^T \end{bmatrix}. \quad (1)$$

A common way to identify challenging examples is by measuring the entropy of their predicted class distribution $\mathcal{H}(\mathbf{b}_0) = -\sum_i \mathbf{b}_{0,i} \log_m(\mathbf{b}_{0,i})$. The higher the value of $\mathcal{H}(\mathbf{b}_0)$ is, the more uniform the distribution \mathbf{v} is, which indicates higher ambiguity level. We observe $\mathcal{H}(\mathbf{b}_0)$ may not be ideal to capture ambiguity well for our classification purpose, which is concerned mainly with high probability events. For example, consider two distributions $\mathbf{b}_0^1 = [0.5 \ 0.25 \ 0.25]$ and $\mathbf{b}_0^2 = [0.5 \ 0.5 \ 0]$, for which $\mathcal{H}(\mathbf{b}_0^1) > \mathcal{H}(\mathbf{b}_0^2)$. However, \mathbf{b}_0^2 expresses a more uncertain classification result.

Instead, we select challenging examples based on the entropy near the peak of the distribution. We define top- k -entropy to focus on the top of the distribution. Let $\mathcal{T} : \mathbb{R}^m \rightarrow \mathbb{R}^k$ be the top- k

³ \mathbf{b}_0^T denotes the transpose of \mathbf{b}_0 .

operator. The top- k -entropy is:

$$\mathcal{H}_{\text{top-}k}(\mathbf{b}_0) = \mathcal{H}(\mathcal{T}(\mathbf{b}_0, k)). \quad (2)$$

We use a base of m so that $0 < \mathcal{H}_{\text{top-}k}(\mathbf{b}_0) \leq 1$.

Given a scalar threshold $0 \leq \tau \leq 1$ and the number of classes k_{\max} , the ambiguity level of a probability distribution \mathbf{b}_0 is larger than τ , if for any $k \in [2, k_{\max}]$, the top- k -entropy of \mathbf{b}_0 is greater than τ .

Step 2: One Iteration of AN Each iteration d of AN normalizes L_0 , first across its rows (row norm) and then along its columns (column norm). Let $\mathcal{D} : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times n}$ turn a vector \mathbf{v} into a diagonal matrix, and let \mathbf{e} be a vector of ones.

Step 2.1: Row Norm The *row normalization* of L_{d-1} is:

$$\Lambda_S = \mathcal{D}((L_{d-1}^\alpha)^T \mathbf{e}) \quad (3)$$

$$S_d = L_{d-1}^\alpha \Lambda_S^{-1}, \quad (4)$$

where $\alpha > 0$, L_{d-1}^α is the matrix exponentiation of L_{d-1} , and Λ_S^{-1} is the inverse of Λ_S . The diagonal entries of $\Lambda_S \in \mathbb{R}^{m \times m}$ represent the column sums of L_{d-1}^α , so that S_d is column stochastic. The parameter α controls the rate of convergence of \mathbf{b}_0 to a high confidence state.

Step 2.2: Column Norm The *column normalization* step is:

$$\Lambda_L = \mathcal{D}(S_d \Lambda_q \mathbf{e}) \quad (5)$$

$$L_d = \Lambda_L^{-1} S_d \Lambda_q, \quad (6)$$

where $\Lambda_q \in \mathbb{R}^{m \times m}$ is a diagonal matrix that represents the class priors, which we approximate using the training class distribution. The diagonal entries of $\Lambda_L \in \mathbb{R}^{(n+1) \times (n+1)}$ are the row sums of $S_d \Lambda_q$ so that L_d is row stochastic.⁴

Step 3: Re-adjusted Output Extraction Let L_d be the resulting matrix after d steps of Step 2:

$$L_d = \begin{bmatrix} A_d \\ \mathbf{b}_d^T \end{bmatrix}. \quad (7)$$

We keep \mathbf{b}_d as the re-adjusted class probability distribution $\mathbb{P}'(y|x)$, and discard A_d .

3 Simulations on Random Matrices

We study the effect of the ambiguity level of the source prediction and the reference set using Monte Carlo simulations. We randomly generate A_0 , \mathbf{b}_0 , and Λ_q to evaluate the expected performance

⁴Each normalization step (Steps 2.1–2.2) takes $\mathcal{O}(mn)$ because the matrices $\Lambda_S, \Lambda_L, \Lambda_q$ are diagonal.

Datasets	# Classes	Method	Marco F1		Mirco F1	
			val	test	val	test
Ultrafine Entity Typing	10331	BASELINE (Multitask; Choi et al., 2018)	31.32	31.98	27.92	28.80
		CAN	($\uparrow 2.15$) 33.47	($\uparrow 1.71$) 33.69	($\uparrow 2.51$) 30.43	($\uparrow 1.89$) 30.69
		BASELINE (Denoised; Onoe and Durrett, 2019)	40.07	40.22	37.88	37.87
		CAN	($\uparrow 0.34$) 40.41	($\uparrow 0.53$) 40.75	($\uparrow 0.59$) 38.47	($\uparrow 0.84$) 38.71
DialogRE	36	BASELINE (BERT; Yu et al., 2020)	35.89	35.76	59.44	57.93
		CAN	($\uparrow 0.91$) 36.80	($\uparrow 0.70$) 36.45	($\uparrow 0.16$) 59.60	($\uparrow 0.34$) 58.27
		BASELINE (BERTs; Yu et al., 2020)	40.58	39.45	62.18	59.52
		CAN	($\uparrow 0.83$) 41.41	($\uparrow 0.68$) 40.13	($\uparrow 0.33$) 62.51	($\uparrow 0.29$) 59.81

Table 1: Performance on the Ultrafine Entity Typing and DialogRE tasks.




Figure 2: Averaged expected accuracy gain ($\delta(\mathbf{b}_0, \mathbf{b}_1)$) and classification accuracy gain ($\Delta(B_0, B_1)$). Original prediction \mathbf{b}_0 with high ambiguity level yields higher expected accuracy gain.

change after each iteration of AN ($\mathbf{b}_0 \rightarrow \mathbf{b}_1$) as a function of the ambiguity level of \mathbf{b}_0 and A_0 .

Setup The ambiguity levels are grouped into 4 intervals: $\{[0, 0.25), [0.25, 0.5), [0.5, 0.75), [0.75, 1]\}$. Given the number of classes m and an ambiguity interval, we randomly generate $A_0 \in \mathbb{R}^{(m-1) \times m} \Lambda_q$, and $B_0 \in \mathbb{R}^{n \times m}$ independently,⁵ where each row of B_0 represents a randomly generated \mathbf{b}_0 . For each interval and each $m \in \{2, \dots, 10, 20, \dots, 100\}$, we randomly generate A_0 , B_0 , and Λ_q 200 times with $n = 100$. We compute each \mathbf{b}_1 separately using CAN with $\alpha \in \{0.1, \dots, 0.9, 1, \dots, 9\}$, and construct B_1 .

Evaluation Metrics We define two metrics to evaluate the expected classifier performance change: (a) $\delta(\mathbf{b}_0, \mathbf{b}_1)$ measures the expected performance gain of \mathbf{b}_1 w.r.t. \mathbf{b}_0 ; and (b) $\Delta(B_0, B_1)$ measures the accuracy gain (i.e., of the arg max) of a set of input examples B_1 w.r.t. B_0 .⁶ We explore how $\delta(\mathbf{b}_0, \mathbf{b}_1)$ and $\Delta(B_0, B_1)$ change as a function of the ambiguity level of B_0 and A_0 .

Effect of Ambiguity Level Figure 2 shows the averaged $\delta(\mathbf{b}_0, \mathbf{b}_1)$ and $\Delta(B_0, B_1)$ across all ma-

⁵Similar to Yuan et al. (2018), we set L/S as square matrices for simplicity. In practice, L/S do not need to be square, as shown in Section 4.

⁶Appendix A.2 provides formal definitions.

trix sizes, simulations, and values of α . We observe that (a) the expected accuracy gains are positive across all ambiguity levels of A_0 and \mathbf{b}_0 ; (b) CAN tends to improve the performance of \mathbf{b}_0 with high ambiguity level, especially using a reference set with low top- k entropy; and (c) the performance is robust to the ambiguity level of A_0 .

4 Empirical Experiments

We evaluate CAN using three classification tasks: ultrafine entity typing (Choi et al., 2018), and dialogue-based relation extraction (DialogRE; Yu et al., 2020). We compare off-the-shelf classifiers (BASELINE) and our method.⁷ We select CAN’s hyperparameters (α, d, τ) using the official validation sets (val) and evaluate on the official test sets (). The challenging subsets of the val and test are identified as the source predictions using top- k entropy. The rest of val, which has low ambiguity, is used as the reference set A_0 .⁸

Results Table 1 summarizes the experimental results. CAN offers consistent performance gains for different classifiers and datasets by re-adjusting the uncertain examples only, especially when the task has many classes. Table 1 suggests more effective classifiers benefit less from CAN, but still see improvements. For example, in the ultrafine entity typing experiment, we observe a larger improvement for the multitask version (Choi et al., 2018) compared with the denoised one (Onoe and Durrett, 2019), which uses a cleaned up version of the training data.

Tuning CAN Figure 3 shows the how the values of the hyperparameters (α, d, τ) affect CAN on the val set using ultrafine entity typing multitask BASELINE. We observe that the effect of d and τ diminish gradually as α grows, because α controls how

⁷Appendix B.1 provides details for BASELINE methods.

⁸Appendix A and B provide implementation and reproducibility details.




Figure 3: Effect of α , recursion depth d , and τ on the macro F1 scores of ultrafine entity typing `val` set with the BASELINE multitask classifier.

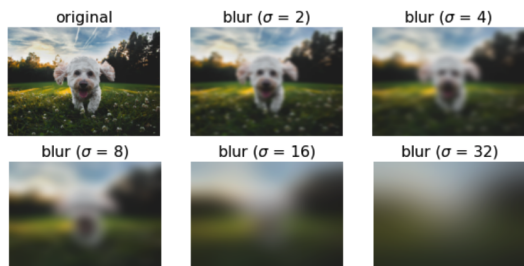


Figure 4: Image blurring examples.

quickly CAN transforms the source predictions to high-confidence ones. For a fixed recursion depth d , the performance does not always improve using a larger α . This suggests that larger α can deteriorate the performance through over-calibration. We also see that a small value of recursion depth d yields the best results in general.

5 Analysis

We hypothesize that classifiers with better performance require less re-adjustment. We test this hypothesis by controlling the number of “hard” examples, therefore controlling the performance of the BASELINE classifier. We use ImageNet (Deng et al., 2009) because of the existence of established ways for image perturbation without modifying the image semantics. We use a ResNet-50 (He et al., 2016) model from the default pretrained models in torchvision (Paszke et al., 2017) as BASELINE.

We systematically make the task harder by convolving the images with a Gaussian function with zero-mean and various values of standard deviation σ . Higher values of σ results in more blurred image, which emulate examples with higher uncertainty and lower BASELINE classifier performance. Figure 4 shows example images with different values of $\sigma \in \{2, 4, 8, 16, 32\}$.

We assess the effect of CAN by comparing two variants of CAN with BASELINE using this setup: (a) CAN-NAIVE: $\alpha = 1.0$, recursion depth is 1; and (b) CAN-BEST: an upperbound-version of CAN with α and depth d optimized on the test data.

Figure 5 presents the absolute top-1 accuracy




Figure 5: Effect on accuracy of making examples “harder” by blurring.

(%) gain on the hard subset of `val`. Optimizing α and d (CAN-BEST) indeed offers larger performance gains across six different image settings comparing with CAN-NAIVE. We observe a general positive correlation between the number of hard examples and the relative gain using CAN, confirming that our method benefits classifiers that receive harder data. However, this advantage diminishes when a classifier significantly under-performs, as the accuracy gain dropped for Blur ($\sigma = 8$) onward and its top-1 accuracy is only 3.25% on hard subset.

6 Related Work

Rational Speech Act (RSA) RSA is a framework for pragmatic reasoning, where a speaker and a listener generate and understand utterances by reasoning about the understanding and intentions of their interaction partner (Frank and Goodman, 2012; Goodman and Frank, 2016). Both agents are probabilistic, and RSA uses alternating normalization in a recursive process. Our technique is motivated by the type of reasoning provided by RSA, whereas interpretation and generation of messages are considered within the context of other interpretations and messages to resolve ambiguities.

We adapt the RSA technique to post-process the output of probabilistic classifiers. The classifier takes a similar role to the RSA listener in the alternating normalization process. The matrices S_d and L_d from Equations 4 and 6 are aligned with the stochastic matrices for the derived speaker and listener in RSA. The row norm (Equation 4) mirrors the special case of a speaker that considers the cost of generating a message as zero. This assumption

is reasonable for classification. It is also common practice when working with a finite set of intents in RSA (Monroe et al., 2017; Zarrieß and Schlangen, 2019). The column norm (Equation 6) describes a mathematical formulation aligned with how a pragmatic listener infers speaker expectations.

Previous work has applied RSA to systems that generate and understand language (Andreas and Klein, 2016; Mao et al., 2016; Vedantam et al., 2017; Cohn-Gordon et al., 2018; Zarrieß and Schlangen, 2019) in both referential games (Frank and Goodman, 2012; Goodman and Frank, 2016; Monroe et al., 2017) and sequential decision-making systems (Fried et al., 2018a,b). Our method departs from these applications by focusing on the ambiguity avoidance property of the listener agent as applied to generic classification tasks.

Confidence Calibration Similar to confidence calibration techniques (Platt et al., 1999; Zadrozny and Elkan, 2002; Guo et al., 2017; Kumar et al., 2019), our method rescales the posterior distribution produced by the classifier at test time. However, the aim of calibration is to make the output probabilities more representative of the correctness likelihood, whereas our’s is to resolve ambiguity.

7 Conclusions and Future Work

We propose Classification with Alternating Normalization (CAN) as a simple and light-weight post-processing step. Our method adjusts the predicted class distribution of “edge cases” for a generic classifier during test time. Via experiments on both simulated and real-world NLP tasks, we show that CAN helps improve performance using a fixed reference set with low ambiguity and increases the performance of standard classifiers. Future work may further investigate the properties of CAN. For example, Appendix 5 describes an initial study on vision examples showing the benefit of CAN increases as examples become noisier, even with the same classifier. Advancing this study and generalizing it to language is an important direction for future work. Finally, one could study improving CAN, for example by selecting A_0 strategically, or applying it during training.

Acknowledgements

We would like to thank the members of the Cornell NLP group, and anonymous reviewers for their helpful feedback. Photo of Figure 4 by Joe Caione on [Unsplash](#).

References

- Jacob Andreas and Dan Klein. 2016. Reasoning about pragmatics with neural listeners and speakers. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1173–1182, Austin, Texas. Association for Computational Linguistics.
- Eunsol Choi, Omer Levy, Yejin Choi, and Luke Zettlemoyer. 2018. Ultra-fine entity typing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 87–96, Melbourne, Australia. Association for Computational Linguistics.
- Reuben Cohn-Gordon, Noah Goodman, and Christopher Potts. 2018. Pragmatically informative image captioning with character-level inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 439–443, New Orleans, Louisiana. Association for Computational Linguistics.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Michael C Frank and Noah D Goodman. 2012. Predicting pragmatic reasoning in language games. *Science*, 336(6084):998–998.
- Daniel Fried, Jacob Andreas, and Dan Klein. 2018a. Unified pragmatic models for generating and following instructions. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1951–1963, New Orleans, Louisiana. Association for Computational Linguistics.
- Daniel Fried, Ronghang Hu, Volkan Cirik, Anna Rohrbach, Jacob Andreas, Louis-Philippe Morency, Taylor Berg-Kirkpatrick, Kate Saenko, Dan Klein, and Trevor Darrell. 2018b. Speaker-follower models for vision-and-language navigation. In *Advances in Neural Information Processing Systems*, volume 31, pages 3314–3325. Curran Associates, Inc.
- Noah D Goodman and Michael C Frank. 2016. Pragmatic language interpretation as probabilistic inference. *Trends in cognitive sciences*, 20(11):818–829.

- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. 2017. On calibration of modern neural networks. In *International Conference on Machine Learning*, pages 1321–1330. PMLR.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. [Deep residual learning for image recognition](#). In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778.
- Ananya Kumar, Percy S Liang, and Tengyu Ma. 2019. [Verified uncertainty calibration](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Xiao Ling and Daniel S. Weld. 2012. Fine-grained entity recognition. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, AAAI’12, page 94–100. AAAI Press.
- Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan Yuille, and Kevin Murphy. 2016. [Generation and comprehension of unambiguous object descriptions](#). In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11–20.
- Will Monroe, Robert X.D. Hawkins, Noah D. Goodman, and Christopher Potts. 2017. [Colors in context: A pragmatic neural model for grounded language understanding](#). *Transactions of the Association for Computational Linguistics*, 5:325–338.
- Yasumasa Onoe and Greg Durrett. 2019. [Learning to denoise distantly-labeled data for entity typing](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2407–2417, Minneapolis, Minnesota. Association for Computational Linguistics.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in pytorch.
- John Platt et al. 1999. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Advances in large margin classifiers*, 10(3):61–74.
- Sonse Shimaoka, Pontus Stenetorp, Kentaro Inui, and Sebastian Riedel. 2016. [An attentive neural architecture for fine-grained entity type classification](#). In *Proceedings of the 5th Workshop on Automated Knowledge Base Construction*, pages 69–74, San Diego, CA. Association for Computational Linguistics.
- Richard Sinkhorn and Paul Knopp. 1967. Concerning nonnegative matrices and doubly stochastic matrices. *Pacific Journal of Mathematics*, 21(2):343–348.
- Ramakrishna Vedantam, Samy Bengio, Kevin Murphy, Devi Parikh, and Gal Chechik. 2017. [Context-aware captions from context-agnostic supervision](#). In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1070–1079.
- Dian Yu, Kai Sun, Claire Cardie, and Dong Yu. 2020. [Dialogue-based relation extraction](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4927–4940, Online. Association for Computational Linguistics.
- Arianna Yuan, Will Monroe, Yu Bai, and Nate Kushman. 2018. [Understanding the rational speech act model](#). In *Annual Meeting of the Cognitive Science Society (CogSci)*.
- Bianca Zadrozny and Charles Elkan. 2002. [Transforming classifier scores into accurate multiclass probability estimates](#). In *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD ’02, page 694–699, New York, NY, USA. Association for Computing Machinery.
- Sina Zarrieß and David Schlangen. 2019. [Know what you don’t know: Modeling a pragmatic speaker that refers to objects of unknown categories](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 654–659, Florence, Italy. Association for Computational Linguistics.

A Implementation Details

A.1 Extension to Multilabel Problems

The problem we describe in Section 2 is supervised multi-class classification. For multi-label classification task, given n input examples and m classes, a classifier usually produces an array of size $n \times m$ where the value in each position represents the predicted probability of one input-class pair. We transform this array into shape $nm \times 2$, where each row is the binary probability distribution for each input-class pair. This is similar to the common binary cross-entropy loss for this task.

A.2 Simulation Metrics Definitions

Under the framework described in Section 2, we define two metrics to evaluate the expected classifier performance. *Expected performance gain* measures the performance of \mathbf{b}_1 w.r.t. \mathbf{b}_0 . Let \mathcal{Y} , $\hat{\mathcal{Y}}$ be two random variables. \mathcal{Y} represents the true label of a given instance, and $\hat{\mathcal{Y}}$ is the predicted class label out of a set of classes C . The *expected accuracy* of \mathbf{b}_0 would be:

$$\mathbb{E}_0[Acc] = \frac{1}{|C|} \sum_{c \in C} \mathbb{P}(\mathcal{Y} = c) \mathbb{P}(\hat{\mathcal{Y}} = c) \quad (8)$$

$$= \frac{1}{|C|} \mathbf{q}^T \mathbf{b}_0, \quad (9)$$

where each entry at position c of $\mathbf{q} \in \mathbb{R}^{|C|}$ is the probability mass function (PMF) of \mathcal{Y} when $\mathcal{Y} = c$. \mathbf{q} is the main diagonals of the randomly generated Λ_q . Each entry at position c of \mathbf{b}_0 is the PMF of $\hat{\mathcal{Y}}$ when $\hat{\mathcal{Y}} = c$. The relative expected performance gain is then defined as:

Definition 1 (Expected Performance Gain). The relative *expected performance gain* of \mathbf{b}_1 w.r.t. \mathbf{b}_0 is:

$$\delta(\mathbf{b}_0, \mathbf{b}_1) = \frac{\mathbb{E}_1[Acc] - \mathbb{E}_0[Acc]}{\mathbb{E}_0[Acc]}. \quad (10)$$

δ quantifies the performance on individual example. Next we introduce a second metric, *accuracy gain*, to measure a set of input examples.

Definition 2 (Accuracy Gain). Given $B_1 \in \mathbb{R}^{n \times m}$, where the i^{th} row represent a re-adjusted predicted distribution, denoted as \mathbf{b}_1^i , the overall performance of B_1 w.r.t. B_0 is evaluated as:

$$\Delta(B_0, B_1) = \frac{1}{n} \sum_i \mathbb{1}\{\delta_i > 0, \arg \max_{\mathbf{b}_0^i} \neq \arg \max_{\mathbf{b}_1^i}\} \quad (11)$$

where $\delta_i = \delta(\mathbf{b}_0^i, \mathbf{b}_1^i)$. We consider CAN is successful when the predicted classes change from \mathbf{b}_1

to \mathbf{b}_0 and $\delta(\mathbf{b}_0, \mathbf{b}_1) > 0$. Δ is used for the primary metric in our simulation study.

B Reproducibility Details

The experiments in this work do not require training or GPU. We either download the publicly available pre-trained model checkpoints, or obtain the model output on `val` and `test` sets from the researchers who propose the BASELINES.

B.1 Datasets and BASELINES

The statistics of the evaluated tasks and the associated datasets are shown in Table 2. All evaluation and optimization protocols are based on the practices of the corresponding BASELINES for each task.

Ultrafine Entity Typing This task is to predict a set of semantic types of a given entity mention within a sentence. The dataset (Choi et al., 2018) contains 10,331 entity types, including coarse, fine and ultra-fine grained classes. Since each entity can have more than one types, this task is a multi-label classification problem. Following previous practices (Ling and Weld, 2012; Choi et al., 2018), we adopt the loose Macro and loose Micro F1 score as metrics. We optimize the loose Macro F1 score on the `val` set to select hyperparameters. Two existing classifiers are evaluated: (1) Multitask (Choi et al., 2018): a model proposed along with the dataset and utilizes a LSTM-based AttentiveNER model (Shimaoka et al., 2016); (2) Denoised (Onoe and Durrett, 2019): uses denoised distant training data.

DialogRE The task of Dialogue-based relation extraction (DialogRE) (Yu et al., 2020) dataset is to predict one or more types of relations between two entities mentioned in dialogues. We applied CAN to the two baselines proposed in the paper: BERT, and BERTs. Each experiment has five runs following the practices of Yu et al. (2020), so we report the average Macro F1 and Micro F1. The Micro F1 is used for the parameter optimization following Yu et al. (2020).

B.2 Hyperparameters

A hold-out validation set is required by our method. In practice, this can be the same set for fine tuning other hyperparameters during training neural classifiers. Our assumption is that the different splits of a dataset (`train`, `val`, `test`) are sampled from the same distribution. The training

Task	Datasets	BASELINE	Cls. Type	# Classes	train/val/test
Entity Typing	Ultrafine Entity Typing (Choi et al., 2018)	Multitask (Choi et al., 2018) Denoised (Onoe and Durrett, 2019)	Multi-label	10331	1998 / 1998 / 1998
Relation Extraction	DialogRE (Yu et al., 2020)	BERT (Yu et al., 2020) BERTs (Yu et al., 2020)	Multi-label	36	5997 / 1914 / 1862
Image Classification	ImageNet (Deng et al., 2009)	ResNet-50 (He et al., 2016)	Single-label	1000	1,281,167 / 50,000 / -

Table 2: Evaluation tasks and off-the-shelf BASELINE methods used. Multiple BASELINE architectures types are evaluated in our experiments.

Task	Datasets	BASELINE	α	d	τ
Entity Typing	Ultrafine Entity Typing	Multitask	16	1	0.25
		Denoised	22	2	0.75
Relation Extraction	DialogRE	BERT	5, 0.7, 8, 5, 0.6	2, 3, 1, 3, 2	0.75, 0.25, 0.25, 0.75, 0.5
		BERTs	0.8, 7, 4, 7, 4	5, 1, 1, 1, 1	0.25, 0.5, 0.5, 0.75, 0.25
Image Classification	ImageNet-Original	ResNet50	0.6	1	0.75
	ImageNet-Blur ($\sigma = 2$)		0.7	2	
	ImageNet-Blur ($\sigma = 4$)		0.9	5	
	ImageNet-Blur ($\sigma = 8$)		0.9	3	
	ImageNet-Blur ($\sigma = 16$)		0.5	1	
	ImageNet-Blur ($\sigma = 32$)		0.6	1	

Table 3: Hyper-parameters used for all experiment. We report details of all the output from five runs of DialogRE dataset.

Task	Datasets	BASELINE	val	test
Entity Typing	Ultrafine Entity Typing	Multitask	14269 0.07%	19007 0.09%
		Denoised	4454 0.02%	4601 0.02%
Relation Extraction	DialogRE	BERT	1478 2.15%	1430 2.13%
			4598 6.67%	4512 6.73%
			4126 5.99%	3993 5.96%
			1210 1.76%	1148 1.71%
			2610 3.79%	1148 1.71%
		BERTs	3725 5.41%	3841 5.73%
			1539 2.23%	1513 2.26%
			1579 2.29%	1547 2.31%
			885 1.28%	859 1.28%
			3131 4.54%	3153 4.70%
Image Classification	ImageNet-Original	ResNet50	13484 26.97%	-
	ImageNet-Blur ($\sigma = 2$)		24788 49.58%	-
	ImageNet-Blur ($\sigma = 4$)		34421 68.84%	-
	ImageNet-Blur ($\sigma = 8$)		41230 82.46%	-
	ImageNet-Blur ($\sigma = 16$)		45721 91.44%	-
	ImageNet-Blur ($\sigma = 32$)		48889 97.78%	-

Table 4: Number of ambiguous examples in val and test sets (absolute | relative (%)).

class distribution is utilized as prior Λ_q . We use the original val set of the datasets to optimize for $\alpha \in \{0.1, \dots, 0.9, 1, \dots, 35\}$, iteration number $d \in \{1, \dots, 5\}$, and the ambiguous threshold $\tau \in \{0.25, 0.5, 0.75\}$. Table 3 summarizes all the hyperparameters used in this paper. Table 4 presents resulting number of ambiguous distribu-

tions to re-adjust. We find that the optimized values of d are mostly ≤ 2 (77.3 % of all experiments), which is in line with the discussions previously in the main text.