

Stop Measuring Calibration When Humans Disagree

Joris Baan¹, Wilker Aziz¹, Barbara Plank^{2,3,4}, Raquel Fernández¹

¹University of Amsterdam, ²IT University of Copenhagen, ³MCML Munich, ⁴LMU Munich
{j.s.baan,w.aziz,raquel.fernandez}@uva.nl, b.plank@lmu.de

Abstract

Calibration is a popular framework to evaluate whether a classifier *knows when it does not know*—*i.e.*, its predictive probabilities are a good indication of how likely a prediction is to be correct. Correctness is commonly estimated against the human majority class. Recently, calibration to human majority has been measured on tasks where humans inherently disagree about which class applies. We show that measuring calibration to human majority given inherent disagreements is theoretically problematic, demonstrate this empirically on the ChaosNLI dataset, and derive several instance-level measures of calibration that capture key statistical properties of human judgements—class frequency, ranking and entropy.¹

1 Introduction

Neural text classifiers are becoming more powerful but increasingly difficult to interpret (Rogers et al., 2020). In response, the demand for transparency and trust in their predictions is growing (Yin et al., 2019; Bansal et al., 2019; Bianchi and Hovy, 2021). One step towards understanding when to trust predictions is to evaluate whether models *know when they do not know*—*i.e.*, whether predictive probabilities are a good indication of how likely a prediction is to be correct—known as *calibration*. This is crucial in user-facing and high-stake applications.

Calibration, and particularly the Expected Calibration Error (ECE; Naeini et al., 2015; Guo et al., 2017), is widely studied in Machine Learning and Computer Vision (Mena et al., 2021), and is gaining increased attention in Natural Language Processing (NLP; Desai and Durrett, 2020; Kong et al., 2020; Jiang et al., 2021; Dan and Roth, 2021).

An important implicit *assumption* in the widely used definition of perfect calibration proposed by Guo et al. (2017) is that predictions are *either right*

or wrong—in other words, that the true class distribution, *i.e.*, human judgement distribution, is deterministic (one-hot). However, for many problems, while categories exist, their boundaries are fluid: there exists *inherent disagreement* about labels. This means that gold labels are at best an idealization—as irreconcilable disagreement is abundant (Plank et al., 2014; Aroyo and Welty, 2015; Jamison and Gurevych, 2015; Palomaki et al., 2018; Pavlick and Kwiatkowski, 2019). Evidence for this can be found in various tasks, including those which involve linguistic and subjective judgements (Akhtar et al., 2020; Basile et al., 2021). Surprisingly, however, while limitations of calibration are studied (§2), this fundamental assumption is ignored.

In this work, we show that popular calibration metrics—such as ECE—are not applicable to data with inherent human disagreement (§3). We propose an alternative, instance-level notion of calibration based on human uncertainty, and operationalize it with several measures that capture key statistics of the human judgement distribution other than matching the majority vote (§4). Finally, we verify our theoretical claims with a case study on the ChaosNLI dataset, and investigate temperature scaling—a popular post-hoc calibration method—through the lens of human uncertainty (§5).

2 Background

Data We have data $\mathcal{D} = \{(x_n, y_n)\}_{n=1}^N$ where x_n is an instance (*i.e.*, text or texts) and $y_n \in [C]$ is a category. For any instance $X = x$,² we assume that human annotators draw their labels independently from the same Categorical distribution with class probabilities $\pi(x) \in \Delta_{C-1}$. That

²We use capital letters for random variables (*e.g.*, X) and lowercase letters for their assignments (*e.g.*, $X = x$), $[C]$ is short for $\{1, \dots, C\}$, $[a = b]$ is the Iverson bracket (*i.e.*, it evaluates to 1 if the predicate $a = b$ is True, 0 otherwise), and $\Delta_{C-1} \subset \mathbb{R}^C$ is the simplex (*i.e.*, set of C -dimensional vectors whose coordinates are positive and sum to 1).

¹Code available at <https://github.com/jsbaan/calibration-on-disagreement-data>.

is, the probability $\Pr(Y = c | X = x)$ that a human should label x an instance of $c \in [C]$ is $\pi_c(x)$. For observed x , an estimate of $\pi(x)$ can be obtained via maximum likelihood estimation (MLE). This estimate $\bar{\pi}(x)$ is the vector whose coordinate $\bar{\pi}_c(x) = \frac{\sum_{n=1}^N \mathbb{1}_{[x_n=x][y_n=c]}}{\sum_{n=1}^N \mathbb{1}_{[x_n=x]}}$ is the relative frequency with which x is labeled as $c \in [C]$. Oftentimes, $\pi(x)$ is assumed to be one-hot (*i.e.*, the task is unambiguous), in such cases, a single human judgement per instance is sufficient for an exact estimate.

Classification A probabilistic classifier approximates $\pi(x)$ with a trained parametric function (*e.g.*, BERT; Devlin et al., 2019) that maps an input x to a vector $\mathbf{f}(x)$ of class probabilities. After training, and given an instance x , we typically map the model’s output $\mathbf{f}(x)$ to a single decision $\hat{y} \in [C]$. More often than not, this is the mode of the model distribution: $\hat{y} = \arg \max_c f_c(x)$. The correctness of \hat{y} is assessed against the observed human ‘gold standard’ decision $y^* = \arg \max_c \bar{\pi}_c(x)$.

Calibration A classifier is *multi-class calibrated* (Vaicenavicius et al., 2019; Kull et al., 2019) if, for all instances mapped to the same vector \mathbf{q} , the relative frequency with which c is correct (assessed against the gold standard) is q_c for every c :

$$\Pr(Y^* = c \mid \mathbf{f}(X) = \mathbf{q}) = q_c \quad \forall c \in [C] \quad (1)$$

Consider a problem with three classes. A model is multi-class calibrated if, for all instances mapped to the same vector, *e.g.*, $(0.90, 0.07, 0.03)^\top$, predicting the first class would result in a correct decision for 90% of these instances, the second class for 7%, and the third class for 3%. Estimation of the left-hand side (LHS) of Eq(1) by counting is difficult as it requires observing multiple instances mapped to the same probability vector.

A weaker notion of calibration (popular in NLP; Desai and Durrett, 2020; Jiang et al., 2021) is *confidence calibration* (Guo et al., 2017):³

$$\Pr(Y^* = \hat{Y} \mid \max(\mathbf{f}(X)) = p) = p \quad (2)$$

A model is confidence calibrated if, for all instances mapped to a maximum probability value p (*e.g.*, 0.9), the most probable class under the model is correct for 90% of these instances.

³The maximum class probability assigned to a random text X is itself a random variable, we denote it by $\max(\mathbf{f}(X)) := \max_{c \in [C]} f_c(X)$. As for any other random variable, $\max(\mathbf{f}(X)) = p$ identifies the set of texts for which the maximum probability predicted by the model is exactly p .

Expected Calibration Error is most often used to measure (confidence) calibration in practice. Naeini et al. (2015) originally proposed ECE for binary classification and Guo et al. (2017) later adapted it to a multi-class setting:

$$\text{ECE} = \sum_{m=1}^M \frac{|B_m|}{N} |\text{acc}(B_m) - \text{conf}(B_m)| \quad (3)$$

ECE estimates the confidence calibration error—absolute difference between the LHS and the RHS of Eq(2)—in expectation by discretizing the probability of the model decision into a fixed number M of intervals (or bins). Each prediction vector $\mathbf{f}(x)$ is assigned to a bin B_m based on its highest probability $\max_c f_c(x)$. The ECE is the weighted average of the difference between the average confidence and accuracy per bin. To obtain zero calibration error, if 90 out of 100 instances that received a highest probability between 0.8 and 1.0 are correctly classified, the average confidence on those 100 instances must be 0.9.

Several recent studies identify and address problems with ECE—mostly with its binning scheme and implicit decision rule (*e.g.*, Kumar et al., 2018; Nixon et al., 2019; Widmann et al., 2019; Gupta et al., 2021; Si et al., 2022). Instead, in this work we identify a fundamental problem in the definition of perfect calibration when applying it to setups where there exists no real gold label.

3 Calibration & Disagreement Pathology

It is common practice to handle human disagreement with majority voting or other aggregation methods (Dawid and Skene, 1979; Artstein and Poesio, 2008; Paun et al., 2022). Aggregate (gold) labels are then used to evaluate a classifier’s accuracy. We now illustrate the problem this poses when measuring calibration.

Desideratum: Any classifier \mathbf{g} that, given an instance x , predicts the human judgement distribution $\mathbf{g}(x) = \pi(x)$ should be perfectly calibrated.

Consider the oracle classifier that has access to the MLE $\bar{\pi}(x)$ of $\pi(x)$ for any instance x in a validation set. For each x , the oracle is able to predict the human labeling uncertainty $\pi(x)$. This estimate is unbiased and becomes more precise the more judgements we have access to. By definition, when human majority voting is used, this classifier achieves perfect accuracy—its highest confidence prediction always matches the gold standard.

However, according to ECE, the oracle classifier is miscalibrated (this is true for other definitions of calibration to accuracy, including multi-class and classwise). Recall that the calibration error is the absolute difference between accuracy and average confidence per bin. The accuracy of the oracle classifier is always 1. On data where humans disagree, the average confidence will be lower than 1.

This mismatch results in high calibration error (as demonstrated in §5) and exposes a problem with using ECE to measure calibration on disagreement data. An important takeaway is that, even if we can train a classifier that perfectly models the human judgement distribution, this classifier would still be severely miscalibrated. To achieve perfect calibration, its probabilities must drift towards an unfaithful representation of human confidence. Therefore, we argue that human majority accuracy is a bad estimate of correctness to calibrate against.

4 Calibration to Human Uncertainty

To get a faithful probabilistic classifier, we expect it to predict the uncertainty the human population exhibits on any given x . Notions of calibration to accuracy (e.g., multi-class, classwise, confidence) are defined *marginally* (i.e., for instances grouped by a property of model predictions such as their probability). Instead, we argue for a direct assessment of calibration at the instance level. Given x , perfect calibration to human uncertainty requires:

$$\Pr(Y = c \mid X = x) = f_c(x) \quad \forall c \in [C] \quad (4)$$

This is our desideratum of §3 re-expressed for a practical classifier $\mathbf{f}(\cdot)$. In words, a model is calibrated for x if it predicts probability $f_c(x)$ equal to the probability $\Pr(Y = c \mid X = x)$ with which humans label x as c . With multiple human judgements (whether or not they disagree), the LHS can be estimated by $\bar{\pi}(x)$ —the relative frequency of the observed labels. Assessing the degree to which Eq(4) holds in expectation across instances gives us a tool to criticize classifiers in terms of their overall calibration to human uncertainty in a given task. This is appealing because we can assess the trustworthiness both globally (overall calibration) as well as on individual predictions.

Distance Measures To operationalize our notion of human calibration, we propose three distance measures—each capturing a different key statistic. First, Human Entropy Calibration Error:

$$\text{EntCE}(x) = H(\mathbf{f}(x)) - H(\bar{\pi}(x)) \quad (5)$$

This captures the alignment between disagreement among humans and a model’s *indecisiveness*. It is sensitive to average confusion, but not to class ranking. Second, Human Ranking Calibration Score:

$$\text{RankCS} = \frac{1}{N} \sum_{n=1}^N [\text{argsort}(\mathbf{f}(x_n)) - \text{argsort}(\bar{\pi}(x_n))] \quad (6)$$

RankCS is a global measure that can be viewed as a stricter alternative to majority vote accuracy. It is sensitive to class ranking but not to magnitude of probability, complementing entropy calibration.⁴ Third, Human Distribution Calibration Error—the strictest and most informative measure of the three:

$$\text{DistCE}(x) = \text{TVD}(\mathbf{f}(x), \bar{\pi}(x)) \quad (7)$$

We opt for the popular total variation distance (TVD) between the predictive distribution and the human judgement distribution.⁵ One could compare other statistics of interest, and we encourage the community to do so. For example, a more fine-grained classwise analysis (see Appendix B).

Advantages First, unlike ECE, human calibration naturally handles disagreement data—in fact, it requires multiple annotations to reliably estimate the human judgement distribution. Second, DistCE and EntCE measure the calibration of *individual* predictions, which is a powerful tool to aid decision making. Crucially, this avoids the need for a binning scheme, often criticized in ECE (Nixon et al., 2019; Gupta et al., 2021). Third, DistCE ensures full multi-class calibration: there is no implicit decision rule and a classifier’s underlying statistical model is directly evaluated on its ability to match the entire human judgement distribution. Fourth, unlike ECE and its variants, DistCE is a proper scoring rule, which comes with a range of desirable properties (Gneiting and Raftery, 2007).

Related Work Several recent studies on *soft evaluation* evaluate or optimize for a quantity similar to our DistCE. However, we are the first to propose a general notion of calibration in disagreement

⁴We leave aside specifying a method for handling cases where two or more classes obtain an equal number of votes, which makes multiple rankings correct.

⁵ $\text{TVD}(\mathbf{q}, \mathbf{p}) = 1/2 \|\mathbf{q} - \mathbf{p}\|_1$ is convenient for various reasons: it is a metric (hence, symmetric) and defined for all discrete distributions (including sparse distributions); being expressed directly in (absolute difference in) probability, it is dimensionless (or unitless) and bounded between 0 and 1; it quantifies the maximum discrepancy in probability across the event space. For a more technical discussion, see Appendix E.

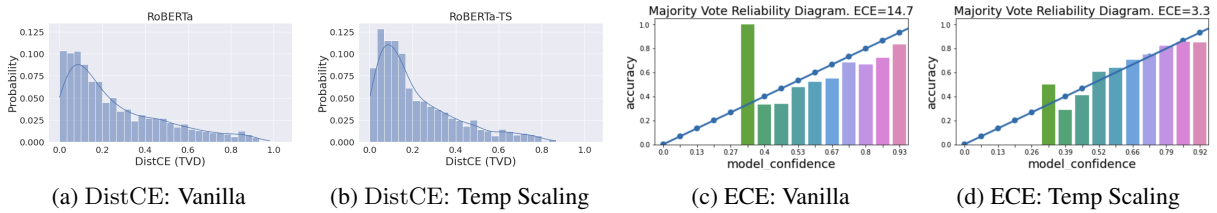


Figure 1: Left: Distribution over instance-level calibration errors (DistCE). While temperature scaling (TS) causes fewer severely miscalibrated instances—illustrated by the right tail of the distribution retracting from (a) to (b)—there are *also* fewer instances that are perfectly calibrated, see drop of first bar in (b). Right: Reliability Diagrams indicate that TS improves ECE enormously, because the bars in (d) move towards the diagonal.

settings. Nie et al. (2020) introduce ChaosNLI (see §5) and report the Kullback–Leibler (KL) and Jensen–Shannon (JS) divergences from estimates of the human judgement distributions. Follow-up work explores decreasing this divergence. For example, by fine-tuning on them directly (Meissner et al., 2021; Zhang et al., 2021); using Bayesian-inspired methods (Zhou et al., 2022); methods popularized to bring down ECE, such as temperature scaling or label smoothing (Zhang et al., 2021; Wang et al., 2022); or framing the task as regression (Chen et al., 2020).

We show that the human majority class is not a meaningful statistic to calibrate against in disagreement settings. In §5, we empirically demonstrate that *human calibration* is more faithful, and a useful tool to gain insights into calibration errors.

5 Case Study

5.1 Experimental Setup

Dataset We use the ChaosNLI dataset (Nie et al., 2020) as case study. It contains English natural language inference instances selected from the development sets of SNLI (Bowman et al., 2015), MNLI (Williams et al., 2018) and AbductiveNLI (Bhagavatula et al., 2020) for having a borderline annotator agreement, *i.e.*, at most 3 out of 5 human votes for the same class. ChaosNLI collects an additional 100 independent annotations for each of the roughly 1,500 instances per dataset, resulting in $T = 100$ human votes distributed over $C = 3$ classes per premise-hypothesis pair for a total of $N = 4,645$ instances. The dataset was collected very carefully and with strict annotation guidelines. This ensures that disagreement cannot easily be discarded as noise (Pavlick and Kwiatkowski, 2019; Nie et al., 2020). The task description and examples can be found in Appendix A.

Method We fine-tune RoBERTa (Liu et al., 2019) on SNLI following the standard procedure described by Desai and Durrett (2020). We evaluate on the ChaosNLI-SNLI split. To investigate the value of human calibration, we inspect Wang et al. (2022)’s claim that ECE is a good alternative to measuring divergence to the human judgement distribution—and that temperature scaling (TS; Guo et al., 2017) is a suitable calibration method to do so. We discuss temperature scaling and how we choose a temperature in Appendix C.

5.2 Results

Table 1 shows accuracy, ECE,⁶ RankCS, and summary statistics of the instance-level EntCE and DistCE metrics for RoBERTa, temperature scaled RoBERTa-TS and the oracle classifier.

Oracle is miscalibrated Indeed, the oracle classifier is severely miscalibrated according to ECE—even more so than RoBERTa (0.25 vs 0.14), demonstrating the problem we highlight in §3. Instead, on all our human calibration metrics, the oracle is perfectly calibrated.

Inspecting Error Distributions Applying TS to RoBERTa results in a sharp decrease in ECE (from 0.14 to 0.03). The reliability diagrams⁷ in Figures 1c and 1d confirm this, suggesting that TS successfully calibrates probability values.

However, TS only causes a very small change in mean DistCE (from 0.26 to 0.22). Though the practical significance of this shift might not be immediately obvious (that is also true for other metrics, such as ECE), human calibration allows us to inspect how errors are distributed *across instances*.

⁶We observe a similar trend for classwise-ECE, a popular improvement to standard ECE, in Appendix D.

⁷A reliability diagram (DeGroot and Fienberg, 1983; Naeni et al., 2015) visualizes ECE by plotting accuracy against confidence (intervals). As bars approach the diagonal, ECE converges to 0.

	RoBERTa	RoBERTa-TS	Oracle
Acc \uparrow	0.74 \pm 0.01	0.74 \pm 0.01	1.00
ECE \downarrow	0.14 \pm 0.01	0.03 \pm 0.01	0.25
RankCS \uparrow	0.62 \pm 0.01	0.62 \pm 0.01	1.00
$\mathbb{E}[\text{EntCE}] \downarrow$	0.30 \pm 0.02	0.21 \pm 0.02	0.00
$\mathbb{E}[\text{DistCE}] \downarrow$	0.26 \pm 0.00	0.22 \pm 0.00	0.00

Table 1: Accuracy & ECE versus our RankCS and summary (mean) of instance-level $|\text{EntCE}|$ & DistCE on ChaosNLI-SNLI. Results shown over 3 random seeds.

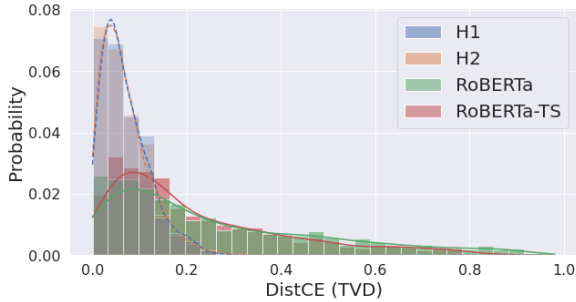


Figure 2: DistCE error distributions for two “*as-good-as-it-realistically-gets*”-classifiers and two RoBERTas (vanilla and TS). There are big differences between but not within groups.

This is an important tool to gain more insight into the effects of a method such as TS on calibration.

The global DistCE error distributions in Figure 1a and 1b reveal that perfectly-calibrated instances are sacrificed to reduce poorly-calibrated instances. This corroborates our intuition that TS artificially compresses the predicted probability range, which, arguably is not desirable. For more extensive and fine-grained analyses, including out-of-distribution evaluation on the ChaosNLI-MNLI set, see Appendix B.

How Good is Good? A naturally arising question is how *good* the shape of a DistCE error distribution is. To answer this, we need a target: what does the error distribution of a “*as-good-as-it-realistically-gets*”-classifier look like?

To approximate such a classifier, for each premise-hypothesis pair, we sub-sample 20 votes and use them to construct a higher-variance MLE of the underlying human judgement distribution $\pi(x)$. We construct two such classifiers (H1 and H2) and plot their DistCE distribution alongside RoBERTa and RoBERTa-TS in Figure 2. We evaluate the classifiers against the 100 available annotations.

As expected, we barely observe any difference between the two sub-sampled human classifiers.

However, there is a massive difference between the RoBERTa-based and human-based classifiers—with RoBERTas stretching to much higher instance level calibration errors (x-axis) than humans—thereby providing a sense of *scale*. To quantify these differences, we compute KL-divergences between DistCE error distributions. We opt for KL-divergence because it is asymmetric; weighting differences in bins that are not probable under the human model less than those that are. We also report TVD.

Table 2 shows that KL divergences from RoBERTa or RoBERTa-TS to an ideal classifier’s error distribution are 150-170x bigger (0.611 and 0.688) compared to the control group (one ideal classifier to another; 0.004). Even though RoBERTa-TS shows slightly reduced KL-divergence compared to the vanilla model, it is nowhere near an ideal classifier, and it is unclear whether the observed reduction in KL translates to a meaningful or practical difference, *i.e.*, something a practitioner would care about.

DistCE error distribution	$\text{KL}(\text{H1}, \cdot)$	$\text{TVD}(\text{H1}, \cdot)$
H2	0.004	0.022
RoBERTa	0.688	0.500
RoBERTa-TS	0.611	0.454

Table 2: Quantifying divergence between DistCE error distributions for human sub-populations (H1 and H2) and neural classifiers (RoBERTa and RoBERTa-TS).

6 Conclusion

We demonstrate a fundamental problem with measuring calibration to the human majority vote in settings with inherent disagreement in human labels. We propose an alternative, instance-level notion based on the full human judgement distribution and operationalize this notion with three metrics. We study temperature scaling RoBERTa on the ChaosNLI dataset using these metrics and conclude that they—and crucially, the ability to inspect them in distribution—provide a more robust and faithful lens to analyze classifier calibration in disagreement settings.

Human uncertainty can be used to evaluate many other calibration techniques—we only performed a preliminary analysis for temperature scaling—and we encourage the community to look into those, in addition to exploring other datasets with inherent disagreements.

Limitations

A reliable estimate of the human judgement distribution is an important requirement for human calibration. For the ChaosNLI dataset, the reliability is endorsed by the large number of annotations per instance and strict quality control (Nie et al., 2020). Most datasets do not provide this. We believe, however, that the advantages of collecting additional annotation outweigh its cost, since, without it, datasets are likely to under-represent human disagreement. We therefore advocate future datasets to include multiple annotations per instance (at least for a small test set), as recently also advocated by, e.g., Prabhakaran et al. (2021), and a better understanding of *how many* annotations are required for good estimates of human uncertainty. An important challenge is to distinguish inherent disagreement from noise, for example due to spammers, which negatively affects data quality (Raykar and Yu, 2012; Aroyo et al., 2019; Klie et al., 2022).

Another limitation is that uncertainty estimates from one population cannot be said to be universally correct. The notion that, given an instance, one unique human distribution governs all annotators is a simplification. Even a single collection of votes from one experiment might contain sub-populations, in which case the marginal distribution is not representative of individual components.

Acknowledgements

We thank the anonymous reviewers and meta-reviewer for their invaluable feedback. This project has received funding by the ELLIS Amsterdam Unit. BP is supported by the Independent Research Fund Denmark (DFF) Sapere Aude grant 9063-00077B and ERC Consolidator Grant 101043235. RF is supported by ERC Consolidator Grant 819455.

References

- Sohail Akhtar, Valerio Basile, and Viviana Patti. 2020. [Modeling annotator perspective and polarized opinions to improve hate speech detection](#). *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, 8(1):151–154.
- Lora Aroyo, Lucas Dixon, Nithum Thain, Olivia Redfield, and Rachel Rosen. 2019. [Crowdsourcing subjective tasks: The case study of understanding toxicity in online discussions](#). In *Companion Proceedings of The 2019 World Wide Web Conference, WWW '19*, page 1100–1105, New York, NY, USA. Association for Computing Machinery.
- Lora Aroyo and Chris Welty. 2015. Truth is a lie: Crowd truth and the seven myths of human annotation. *AI Magazine*, 36(1):15–24.
- Ron Artstein and Massimo Poesio. 2008. Inter-coder agreement for computational linguistics. *Computational linguistics*, 34(4):555–596.
- Gagan Bansal, Besmira Nushi, Ece Kamar, Walter S Lasecki, Daniel S Weld, and Eric Horvitz. 2019. Beyond accuracy: The role of mental models in human-ai team performance. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 7, pages 2–11.
- Valerio Basile, Michael Fell, Tommaso Fornaciari, Dirk Hovy, Silviu Paun, Barbara Plank, Massimo Poesio, Alexandra Uma, et al. 2021. We need to consider disagreement in evaluation. In *1st Workshop on Benchmarking: Past, Present and Future*, pages 15–21. Association for Computational Linguistics.
- Chandra Bhagavatula, Ronan Le Bras, Chaitanya Malaviya, Keisuke Sakaguchi, Ari Holtzman, Hannah Rashkin, Doug Downey, Wen tau Yih, and Yejin Choi. 2020. [Abductive commonsense reasoning](#). In *International Conference on Learning Representations*.
- Federico Bianchi and Dirk Hovy. 2021. [On the gap between adoption and understanding in NLP](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3895–3901, Online. Association for Computational Linguistics.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Tongfei Chen, Zhengping Jiang, Adam Poliak, Keisuke Sakaguchi, and Benjamin Van Durme. 2020. [Uncertain natural language inference](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8772–8779, Online. Association for Computational Linguistics.
- Soham Dan and Dan Roth. 2021. [On the effects of transformer size on in- and out-of-domain calibration](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2096–2101, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Alexander Philip Dawid and Allan M Skene. 1979. Maximum likelihood estimation of observer error-rates using the em algorithm. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 28(1):20–28.

- Morris H. DeGroot and Stephen E. Fienberg. 1983. [The comparison and evaluation of forecasters](#). *Journal of the Royal Statistical Society. Series D (The Statistician)*, 32(1/2):12–22.
- Shrey Desai and Greg Durrett. 2020. [Calibration of pre-trained transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 295–302, Online. Association for Computational Linguistics.
- 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.
- Luc Devroye and Gábor Lugosi. 2001. *Combinatorial methods in density estimation*. Springer Science & Business Media.
- Tilmann Gneiting and Adrian E Raftery. 2007. Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, 102(477):359–378.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. 2017. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML’17*, page 1321–1330. JMLR.org.
- Kartik Gupta, Amir Rahimi, Thalaisyasingam Ajanthan, Thomas Mensink, Cristian Sminchisescu, and Richard Hartley. 2021. [Calibration of neural networks using splines](#). In *International Conference on Learning Representations*.
- Emily Jamison and Iryna Gurevych. 2015. [Noise or additional information? Leveraging crowdsourced annotation item agreement for natural language tasks](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 291–297, Lisbon, Portugal. Association for Computational Linguistics.
- Zhengbao Jiang, Jun Araki, Haibo Ding, and Graham Neubig. 2021. [How Can We Know When Language Models Know? On the Calibration of Language Models for Question Answering](#). *Transactions of the Association for Computational Linguistics*, 9:962–977.
- Jan-Christoph Klie, Bonnie Webber, and Iryna Gurevych. 2022. [Annotation Error Detection: Analyzing the Past and Present for a More Coherent Future](#). *Computational Linguistics*, pages 1–42.
- Lingkai Kong, Haoming Jiang, Yuchen Zhuang, Jie Lyu, Tuo Zhao, and Chao Zhang. 2020. [Calibrated language model fine-tuning for in- and out-of-distribution data](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1326–1340, Online. Association for Computational Linguistics.
- Meelis Kull, Miquel Perelló-Nieto, Markus Kängsepp, Telmo de Menezes e Silva Filho, Hao Song, and Peter A. Flach. 2019. [Beyond temperature scaling: Obtaining well-calibrated multi-class probabilities with dirichlet calibration](#). In *NeurIPS*, pages 12295–12305.
- Aviral Kumar, Sunita Sarawagi, and Ujjwal Jain. 2018. [Trainable calibration measures for neural networks from kernel mean embeddings](#). In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 2805–2814. PMLR.
- 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. *arXiv preprint arXiv:1907.11692*.
- Johannes Mario Meissner, Napat Thumwanit, Saku Sugawara, and Akiko Aizawa. 2021. [Embracing ambiguity: Shifting the training target of NLI models](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 862–869, Online. Association for Computational Linguistics.
- José Mena, Oriol Pujol, and Jordi Vitrià. 2021. [A Survey on Uncertainty Estimation in Deep Learning Classification Systems from a Bayesian Perspective](#). *ACM Comput. Surv.*, 54(9).
- Mahdi Pakdaman Naeni, Gregory F. Cooper, and Milos Hauskrecht. 2015. Obtaining Well Calibrated Probabilities Using Bayesian Binning. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI’15, page 2901–2907. AAAI Press.
- Yixin Nie, Xiang Zhou, and Mohit Bansal. 2020. [What can we learn from collective human opinions on natural language inference data?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9131–9143, Online. Association for Computational Linguistics.
- Jeremy Nixon, Michael W. Dusenberry, Linchuan Zhang, Ghassen Jerfel, and Dustin Tran. 2019. Measuring calibration in deep learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
- Jennimaria Palomaki, Olivia Rhinehart, and Michael Tseng. 2018. A case for a range of acceptable annotations. In *SAD/CrowdBias@ HCOMP*, pages 19–31.
- Silviu Paun, Ron Artstein, and Massimo Poesio. 2022. Statistical methods for annotation analysis. *Synthesis Lectures on Human Language Technologies*, 15(1):1–217.

- Ellie Pavlick and Tom Kwiatkowski. 2019. [Inherent disagreements in human textual inferences](#). *Transactions of the Association for Computational Linguistics*, 7:677–694.
- Barbara Plank, Dirk Hovy, and Anders Søgaard. 2014. Linguistically debatable or just plain wrong? In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 507–511.
- Vinodkumar Prabhakaran, Aida Mostafazadeh Davani, and Mark Diaz. 2021. [On releasing annotator-level labels and information in datasets](#). In *Proceedings of The Joint 15th Linguistic Annotation Workshop (LAW) and 3rd Designing Meaning Representations (DMR) Workshop*, pages 133–138, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Vikas C. Raykar and Shipeng Yu. 2012. Eliminating spammers and ranking annotators for crowdsourced labeling tasks. *J. Mach. Learn. Res.*, 13:491–518.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020. A primer in BERTology: What we know about how BERT works. *Transactions of the Association for Computational Linguistics*, 8:842–866.
- Chenglei Si, Chen Zhao, Sewon Min, and Jordan Boyd-Graber. 2022. Revisiting calibration for question answering. *arXiv preprint arXiv:2205.12507*.
- Juozas Vaicenavicius, David Widmann, Carl Andersson, Fredrik Lindsten, Jacob Roll, and Thomas Schön. 2019. Evaluating model calibration in classification. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 3459–3467. PMLR.
- Yuxia Wang, Minghan Wang, Yimeng Chen, Shimin Tao, Jiabin Guo, Chang Su, Min Zhang, and Hao Yang. 2022. [Capture human disagreement distributions by calibrated networks for natural language inference](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1524–1535, Dublin, Ireland. Association for Computational Linguistics.
- David Widmann, Fredrik Lindsten, and Dave Zachariah. 2019. Calibration tests in multi-class classification: A unifying framework. *Advances in Neural Information Processing Systems*, 32.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). 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 1112–1122. Association for Computational Linguistics.
- Ming Yin, Jennifer Wortman Vaughan, and Hanna Wallach. 2019. Understanding the effect of accuracy on trust in machine learning models. In *Proceedings of the 2019 CHI conference on human factors in computing systems*, pages 1–12.
- Shujian Zhang, Chengyue Gong, and Eunsol Choi. 2021. [Learning with different amounts of annotation: From zero to many labels](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7620–7632, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Xiang Zhou, Yixin Nie, and Mohit Bansal. 2022. [Distributed NLI: Learning to predict human opinion distributions for language reasoning](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 972–987, Dublin, Ireland. Association for Computational Linguistics.

Appendix

A ChaosNLI

Table 3 and Table 4 show examples with very low or very high agreement. The task description provided by Nie et al. (2020) is in Figure 3.⁸

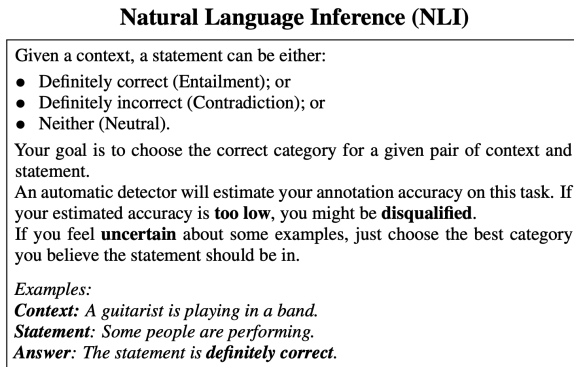


Figure 3: The task description for natural language inference by Nie et al. (2020) used to collect ChaosNLI-SNLI and ChaosNLI-MNLI.

B Additional Analyses

This section provides additional analyses and figures. Each figure shows a comparison between model predictions and human predictions using the measures we propose in §4. We use three different views to compare statistics of human judgements to statistics of model predictions.⁹

The first view shows two *marginal* or dataset-level histograms: one of a human statistic and one of a model statistic, *e.g.*, entropy. This view is useful to compare the global distribution over an instance-level statistic between human judgements and model predictions (Figure 4 and 5).

The second view shows one histogram of the *conditional* instance-level error between a model and human statistic, *e.g.*, entropy (Figure 8 and 9). This is interesting for diagnosing a classifier’s under or over confidence. Instances centered around zero have zero error, instances in the positive range exhibit over-confidence, and instances in the negative range under-confidence.

The third view is similar to the previous, (it is also *conditional*, *e.g.*, it compares an instance-level statistic between humans and a model) but shows

⁸ChaosNLI can be downloaded at https://www.dropbox.com/s/h4j7dqszmpt2679/chaosNLI_v1.0.zip.

⁹Note that these views do not apply to the DistCE (row 3 column 1) and reliability diagrams (row 4), which will not change across views.

the *absolute* errors (Figure 6 and 7). This is useful to spot general miscalibration, regardless of the direction (*i.e.*, under-confidence vs over-confidence).

In each figure, the top row shows a histogram of the predicted probability for class 0 (entailment), 1 (neutral) and 2 (contradiction). The second row shows predicted probability for the k th highest predicted probability, *i.e.*, the first, second and third guess from either the model or human distribution (note that the corresponding classes are not necessarily the same for the model and humans—this row is informative to compare the magnitude of the probability for the same rank). The third row shows the histogram of DistCE (TVD) from §5 (left) and the histogram for EntCE (right). The fourth row shows conventional reliability diagrams that visualize ECE, and the number of instances per bin (note that this is not normally shown—though we find it very insightful).

B.1 Beyond DistCE

Figure 4 shows a vanilla RoBERTa on ChaosNLI-SNLI. We can see that all histograms are very different between humans and model. It is clear that the model predictions are drawn from a different distribution than the human predictions. This signals bad calibration. Figure 5 shows a temperature scaled RoBERTa. According to all plots, the predictions still appear drawn from a different distribution, even though the ECE drops significantly. TS seems to transform the distributions, but they are not clearly closer to the human distribution. The entropy figure seems to indicate that TS *overshoots* to the other end of the spectrum (*i.e.*, from over-confidence to under-confidence). The TS model is unable to match the extreme left and right end of the human certainty spectrum for class 0, 1 and 2, which corroborates our intuition that the probability range is compressed. Another observation is that the human distribution rarely put much mass on the third guess (*i.e.*, the predicted class with the lowest probability). However, after TS models actually do so—which is undesirable.

We next compare the histogram of (non absolute) instance-level errors in Figure 8 and 9. TS brings entropy error median and mean from -.26 to .12. The TS model therefore *overshoots* also on the instance-level (recall this plot is the instance-level error, unlike the previous marginal figure we discussed) and becomes more much more uncertain than humans are. TS causes less class 0 predictions

Premise	Hypothesis	Label Distribution
A man running a marathon talks to his friend	There is a man running.	Entailment: 100, Neutral: 0, Contradiction: 0
A young girl plays with a neon-colored Slinky in a crowd of people on a street lined with flags.	The girl is five years old.	Entailment: 0, Neutral: 100, Contradiction: 0
A well built black man stands in the subway, listening to headphones.	A man with headphones on is standing in a subway.	Entailment: 100, Neutral: 0, Contradiction: 0

Table 3: High agreement examples from the ChaosNLI dataset.

Premise	Hypothesis	Label Distribution
The important thing is to realize that it’s way past time to move it.	It cannot be moved, now or ever.	Entailment: 34, Neutral: 32, Contradiction: 34
An elderly woman crafts a design on a loom.	The woman is sewing.	Entailment: 35, Neutral: 31, Contradiction: 34
Number 13 kicks a soccer ball towards the goal during children’s soccer game.	A player passing the ball in a soccer game.	Entailment: 36, Neutral: 33, Contradiction: 31

Table 4: Low agreement examples from the ChaosNLI dataset.

to have 0 error (which is bad). It also narrows the spread slightly for class 1 and increases errors for class 3.

We next compare the histogram of absolute errors in Figure 6 and 7. We see that TS reduces error tail of class 0, but also reduces instances with 0 error (similar trend as discussed in §B about DistCE). Similarly for class 2 and 3, and entropy: it seemingly reduces the mean and median by cutting off the tail. Finally, the human rank calibration error (RCE) shown in the title of each figure, shows that models are not good at matching the human rank at all—as expected, much worse than matching the majority vote.

In general, it appears that the median and mean error on most metrics go down with TS, mainly by removing instances from the tail (those are extremely miscalibrated examples). However, they seem to sacrifice predictions that were well/perfectly aligned with human judgement probabilities. This illustrated by the mode of the error distributions moving towards the right (meaning more predictions with a higher error). Arguably, this is not desirable—and our metrics provide tools to expose such behaviors.

B.2 Out of Distribution Evaluation

The OOD setting is interesting to evaluate uncertainty estimates, because it is especially important to have reliable uncertainty estimates for examples that are especially difficult (*e.g.*, because the classifier has not seen them during training, and might not reflect the learned distribution). In such cases, it is desirable that a classifier is more uncertain on such examples. In fact, OOD detection is often used to evaluate uncertainty estimates, next to or instead of calibration. Models are often found to be more badly miscalibrated for OOD datasets (Desai and Durrett, 2020), which is an observation we confirm.

We observe similar trends as in the in-distribution analysis. The marginal distributions from Figure 10 to Figure 11 seem to match the human marginal slightly better than on the in-distribution dataset. However, inspecting the error distributions in Figure 14, 15, 12 and 13, we see that the distributions are transformed somewhat, but we do not believe that to be evidence that TS is a good method to improve calibration. The instances are still obviously drawn from a different distribution.

C Temperature Scaling

Temperature scaling is a simple method that uses a single temperature parameter t to scale the output logits of a classifier (Guo et al., 2017). The standard way to choose a temperature is to perform a search on a range of possible values for t on a development set. However, sometimes, the temperature is tuned directly on the test set. This is commonly referred to as the *oracle temperature*. Indeed, we use this method to obtain our temperature, because we consider it an (unrealistic) upper bound on what TS can do. For OOD evaluation, we use the temperature tuned on the ID evaluation. In our experiments in §5, we found a temperature of 2.0 to result in the lowest ECE.

D Classwise-ECE

Classwise-ECE (Nixon et al., 2019) is based on the notion of classwise-calibration (Vaicenavicius et al., 2019; Kull et al., 2019). The main difference with ECE—based on the notion of confidence calibration (§2)—is that it removes the dependency on a decision rule on top of the classifier, and computes the calibration separately for each class. Classwise-ECE, then, is the average calibration error over classes:

$$\frac{1}{K} \sum_{k=1}^K \sum_{m=1}^M \frac{|B_{mk}|}{N} |\text{acc}(m, k) - \text{conf}(m, k)| \quad (8)$$

Table D shows a similar trend for classwise-ECE as for ECE: the oracle classifier is severely miscalibrated. This confirms our claims that the general notion of calibration is not suited for data on which humans inherently disagree about a class—and is not restricted to Guo et al. (2017) notion of confidence calibration.

	RoBERTa	RoBERTa-TS	Oracle
Classwise-ECE	10	5	16

Table 5: Classwise-ECE on ChaosNLI-SNLI in %.

E Total Variation Distance

The total variation distance $\text{TVD}(\mathbf{q}, \mathbf{p})$ between two Categorical distributions with parameters $\mathbf{q}, \mathbf{p} \in \Delta_{C-1}$ is defined as:

$$\max_{\mathcal{A} \in \mathcal{P}(\{C\})} |\mathbb{Q}(\mathcal{A}) - \mathbb{P}(\mathcal{A})| \quad (9a)$$

which can also be expressed as

$$= \frac{1}{2} \underbrace{\sum_{c=1}^C |q_c - p_c|}_{=\|\mathbf{q}-\mathbf{p}\|_1} . \quad (9b)$$

where $[C]$ is the sample space, $\mathcal{P}([C])$ is the event space (for generality, we use the powerset of $[C]$, the set of all subsets of outcomes in the sample space), $\mathcal{A} \in \mathcal{P}([C])$ is any event in the event space, \mathbb{Q} and \mathbb{P} are the probability measures prescribed by each of the Categorical distributions (*i.e.*, $\mathbb{Q}(\mathcal{A}) = \sum_{c \in \mathcal{A}} q_c$ and $\mathbb{P}(\mathcal{A}) = \sum_{c \in \mathcal{A}} p_c$). For a complete technical results with definitions, proofs, and various properties see [Devroye and Lugosi \(2001, Chapter 5\)](#).

Properties. TVD is defined for any two probability vectors, whether dense or sparse. It is a metric (hence symmetric and minimised only for identical distributions) and bounded:

$$\text{TVD}(\mathbf{q}, \mathbf{p}) = \text{TVD}(\mathbf{p}, \mathbf{q}) \quad (10a)$$

$$\text{TVD}(\mathbf{q}, \mathbf{p}) = 0 \quad \text{iff } \mathbf{q} = \mathbf{p} \quad (10b)$$

$$\text{TVD}(\mathbf{q}, \mathbf{p}) \in [0, 1] . \quad (10c)$$

Interpretations. The identity in Eq(9b), which expresses TVD in terms of the L1 norm, gives us a rather practical (linear-time) algorithm to compute it by summing half the absolute difference in probability for the outcomes in the sample space $[C]$ of the random variable. This means that TVD is expressed in units of absolute difference in probability. That definition, Eq(9a), also helps interpretation, it shows that TVD quantifies the maximum discrepancy in probability between the two measures over their entire event spaces.

ChaosNLI SNLI roberta-base 0 dev test temp=1.0 ECE=14.70, HCE=0.26, Acc=0.74, HRE=0.64

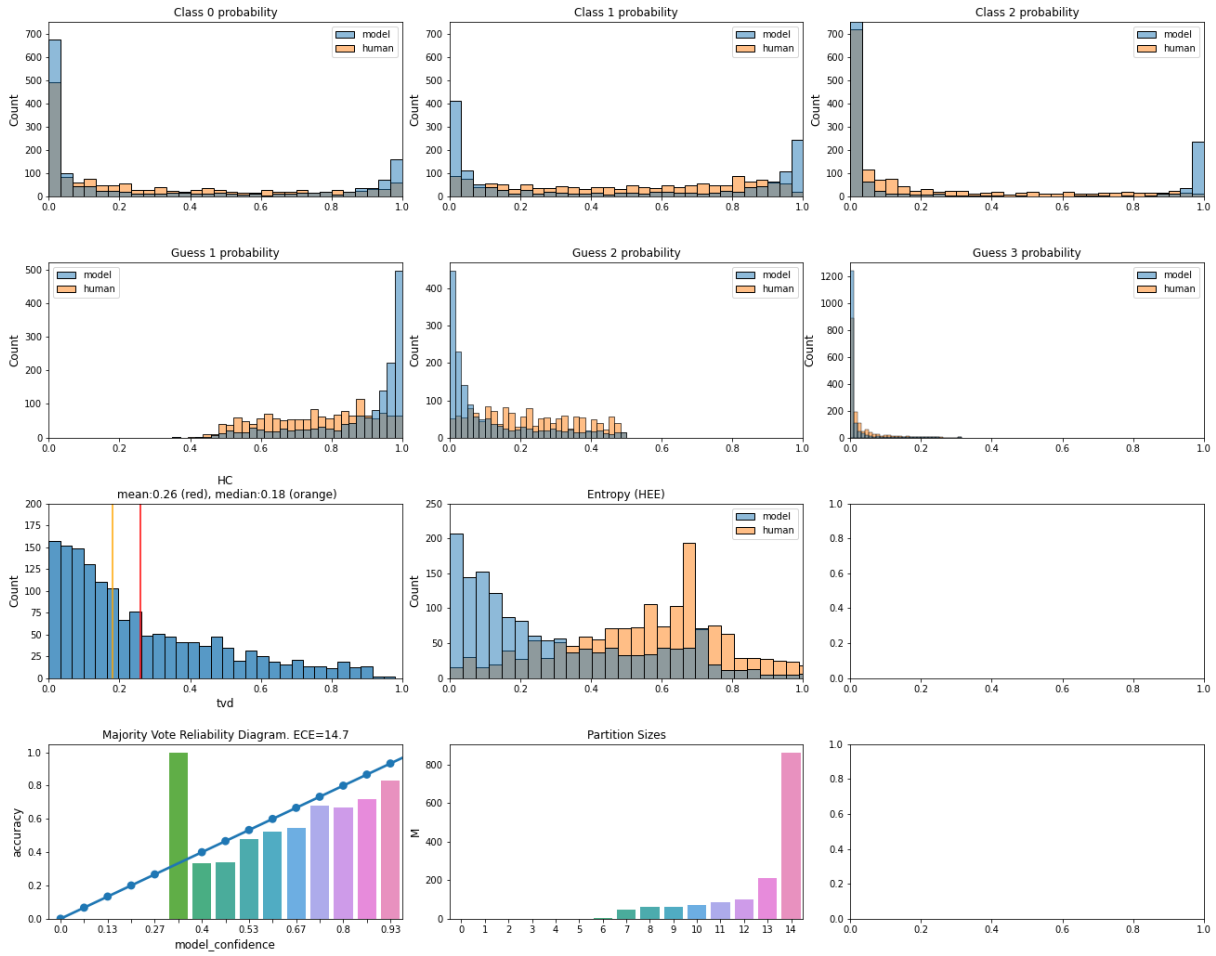


Figure 4: RoBERTa-0 on the ChaosNLI-SNLI dev+test set. Several figures comparing human uncertainty to model uncertainty using TVD, confidence, entropy, and reliability diagrams. This figure shows the distribution over instance-based absolute errors between probabilities for each class (top row) or the model vs human k th guess (i.e., the highest model probability versus the highest human probability on each instance). See Appendix B for more information.

ChaosNLI SNLI roberta-base 0 dev test temp=2.0 ECE=3.33, HCE=0.22, Acc=0.74, HRE=0.64

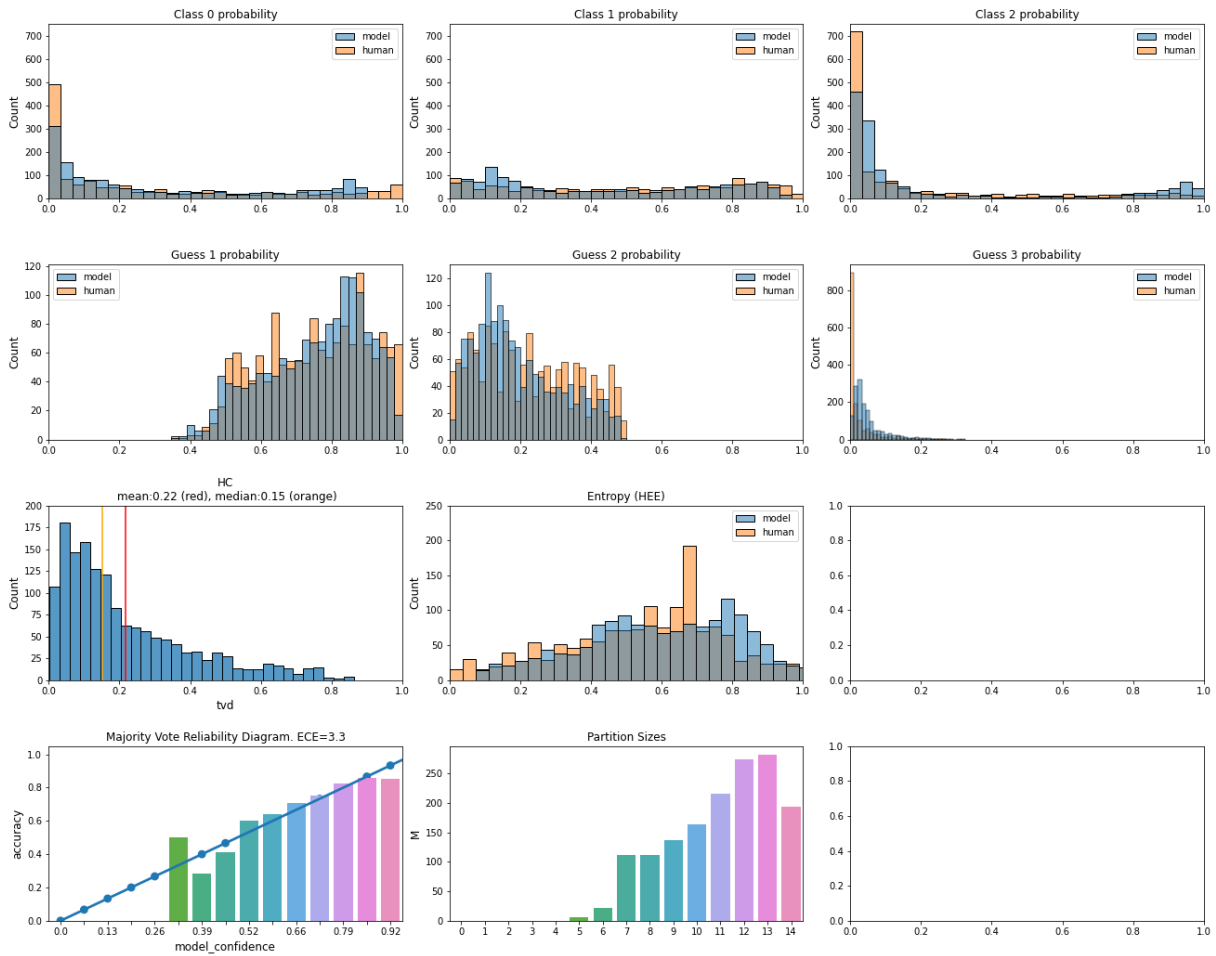


Figure 5: RoBERTa-0 with oracle temperature scaling on the ChaosNLI-SNLI dev+test set. Several figures comparing human uncertainty to model uncertainty using TVD, confidence, entropy, and reliability diagrams. This figure shows the human and model distribution over the probability, entropy or TVD range. Top row shows distribution over probability magnitudes for class 0, 1 and 2, while the second row shows the distribution for first, second and third guess for the model vs human k th guess (i.e., the highest model probability versus the highest human probability on each instance). See Appendix B for more information.

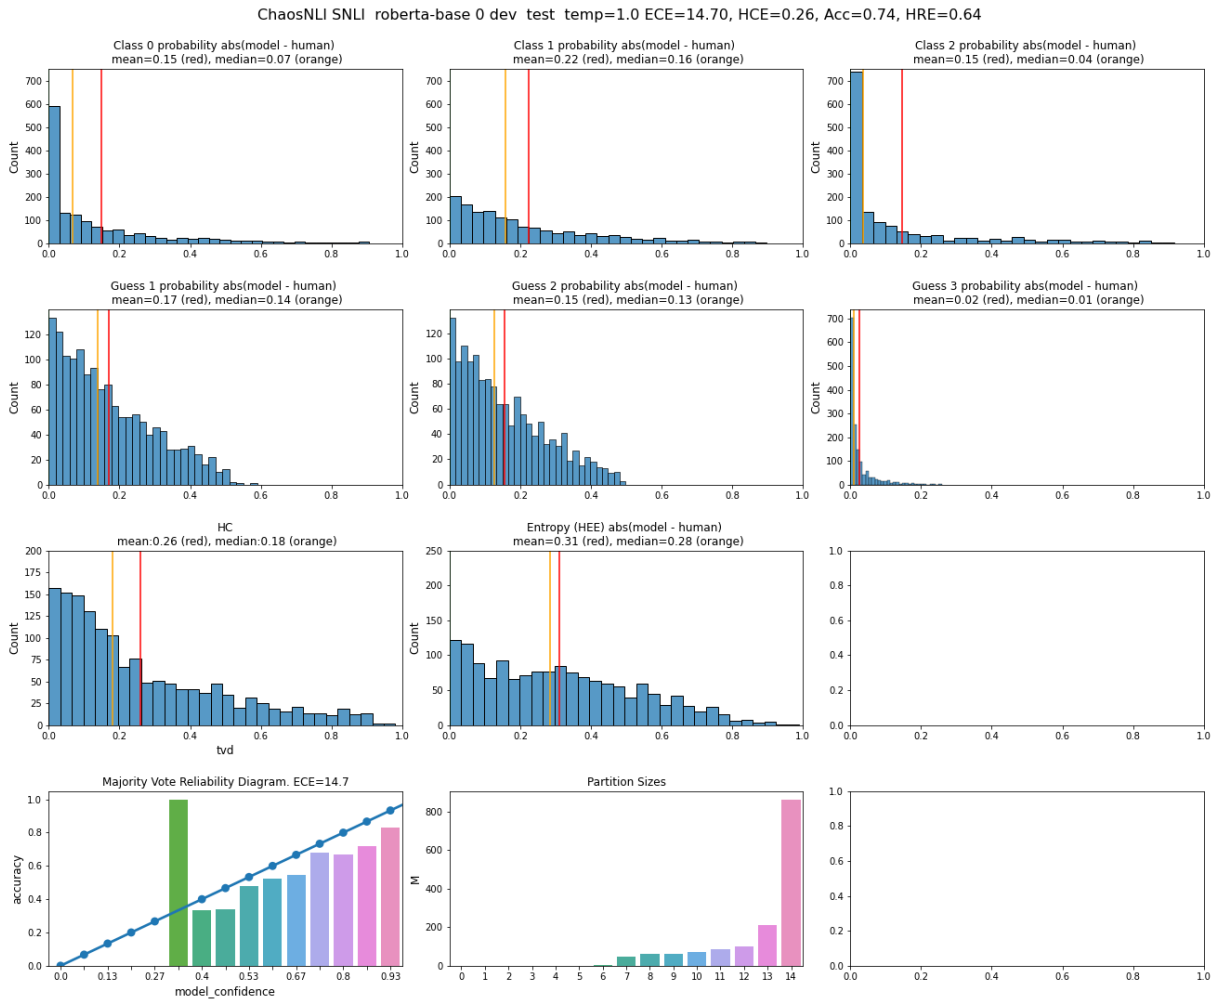


Figure 6: RoBERTa-0 on the ChaosNLI-SNLI dev+test set. Several figures comparing human uncertainty to model uncertainty using TVD, confidence, entropy, and reliability diagrams. This figure shows the distribution over instance-based absolute errors between probabilities for each class (top row) or the model vs human k th guess (i.e., the highest model probability versus the highest human probability on each instance). See Appendix B for more information.

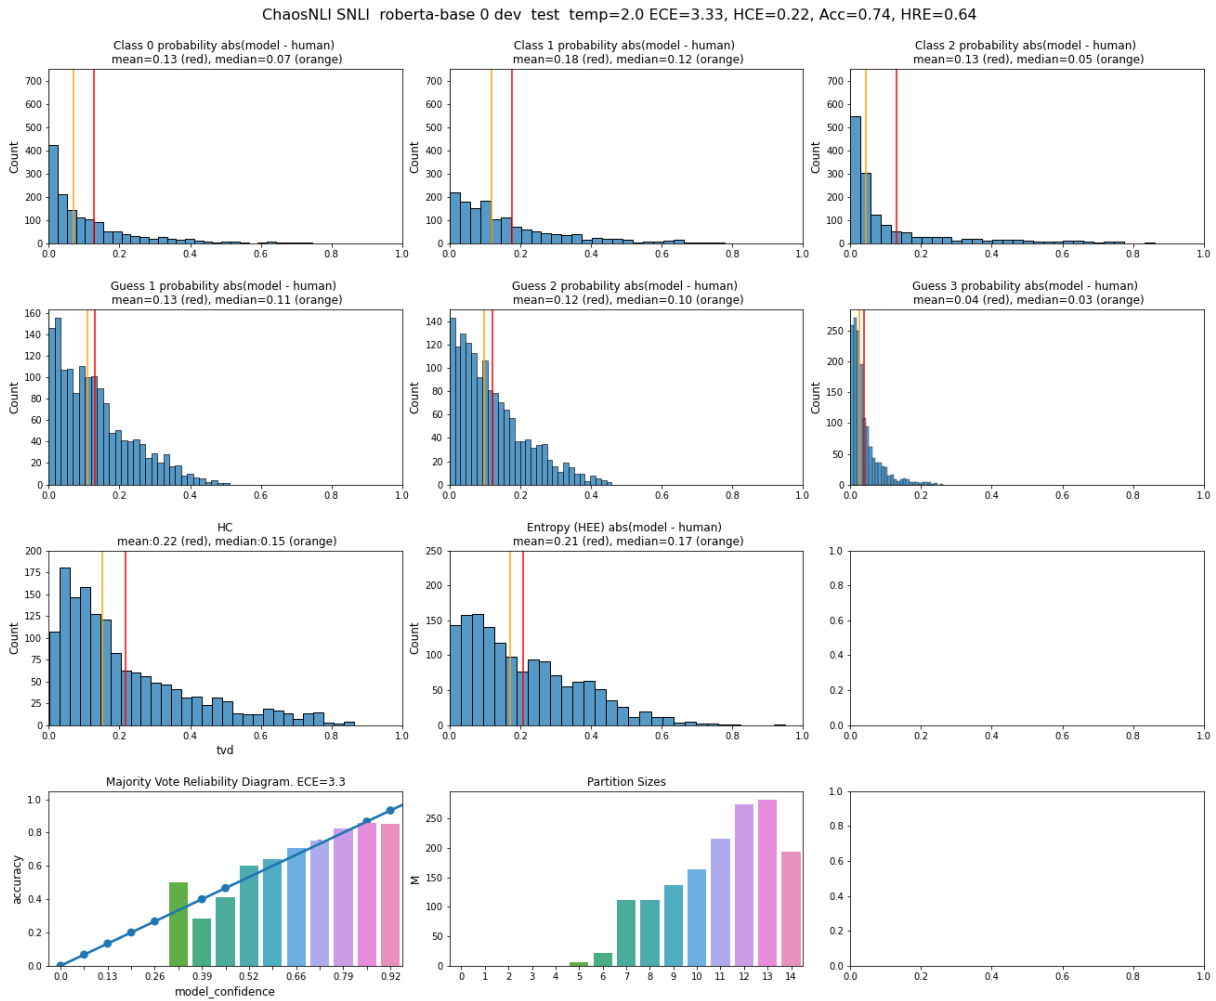


Figure 7: RoBERTa-0 with oracle temperature scaling on the ChaosNLI-SNLI dev+test set. Several figures comparing human uncertainty to model uncertainty using TVD, confidence, entropy, and reliability diagrams. This figure shows the distribution over instance-based absolute errors between probabilities for each class (top row) or the model vs human k th guess (i.e., the highest model probability versus the highest human probability on each instance). See Appendix B for more information.

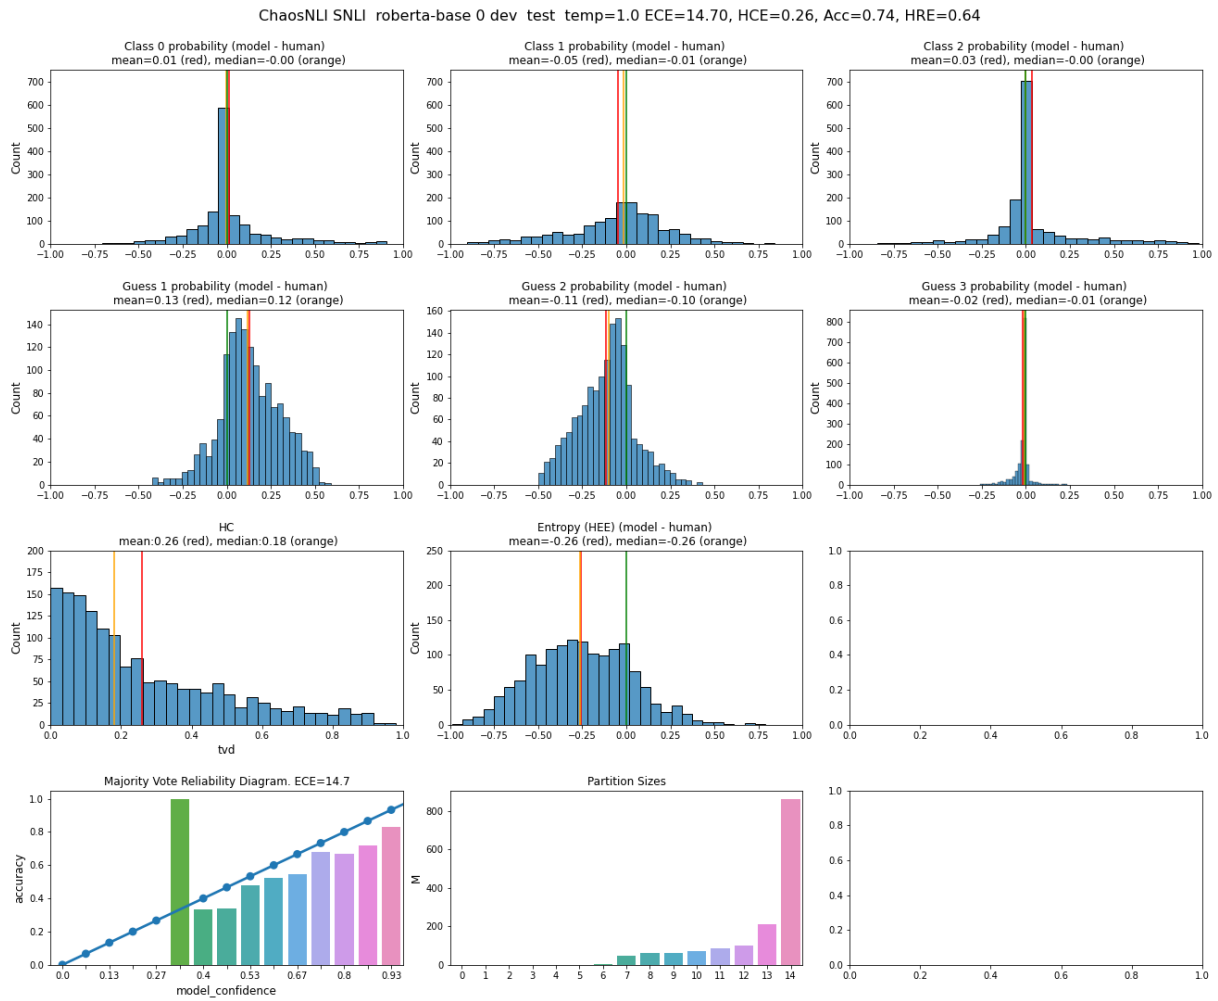


Figure 8: RoBERTa-0 on the ChaosNLI-SNLI dev+test set. Several figures comparing human uncertainty to model uncertainty using TVD, confidence, entropy, and reliability diagrams. This figure shows the distribution over instance-based absolute errors between probabilities for each class (top row) or the model vs human k th guess (i.e., the highest model probability versus the highest human probability on each instance). See Appendix B for more information.

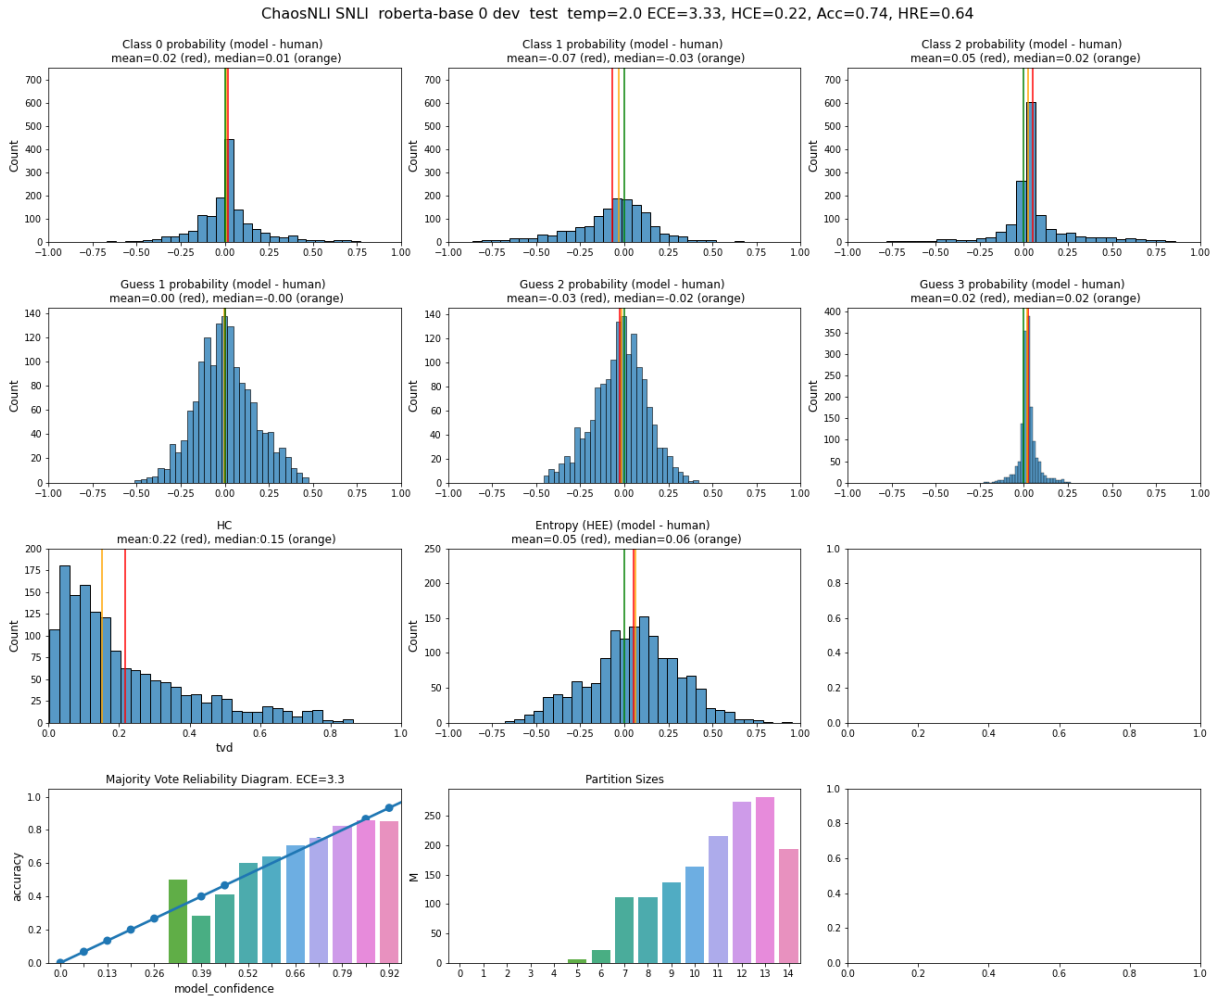


Figure 9: RoBERTa-0 with oracle temperature scaling on the ChaosNLI-SNLI dev+test set. Several figures comparing human uncertainty to model uncertainty using TVD, confidence, entropy, and reliability diagrams. This figure shows the distribution over instance-based absolute errors between probabilities for each class (top row) or the model vs human k th guess (i.e., the highest model probability versus the highest human probability on each instance). See Appendix B for more information.

ChaosNLI MNLi roberta-base 0 dev test temp=1.0 ECE=24.57, HCE=0.35, Acc=0.54, HRE=0.44

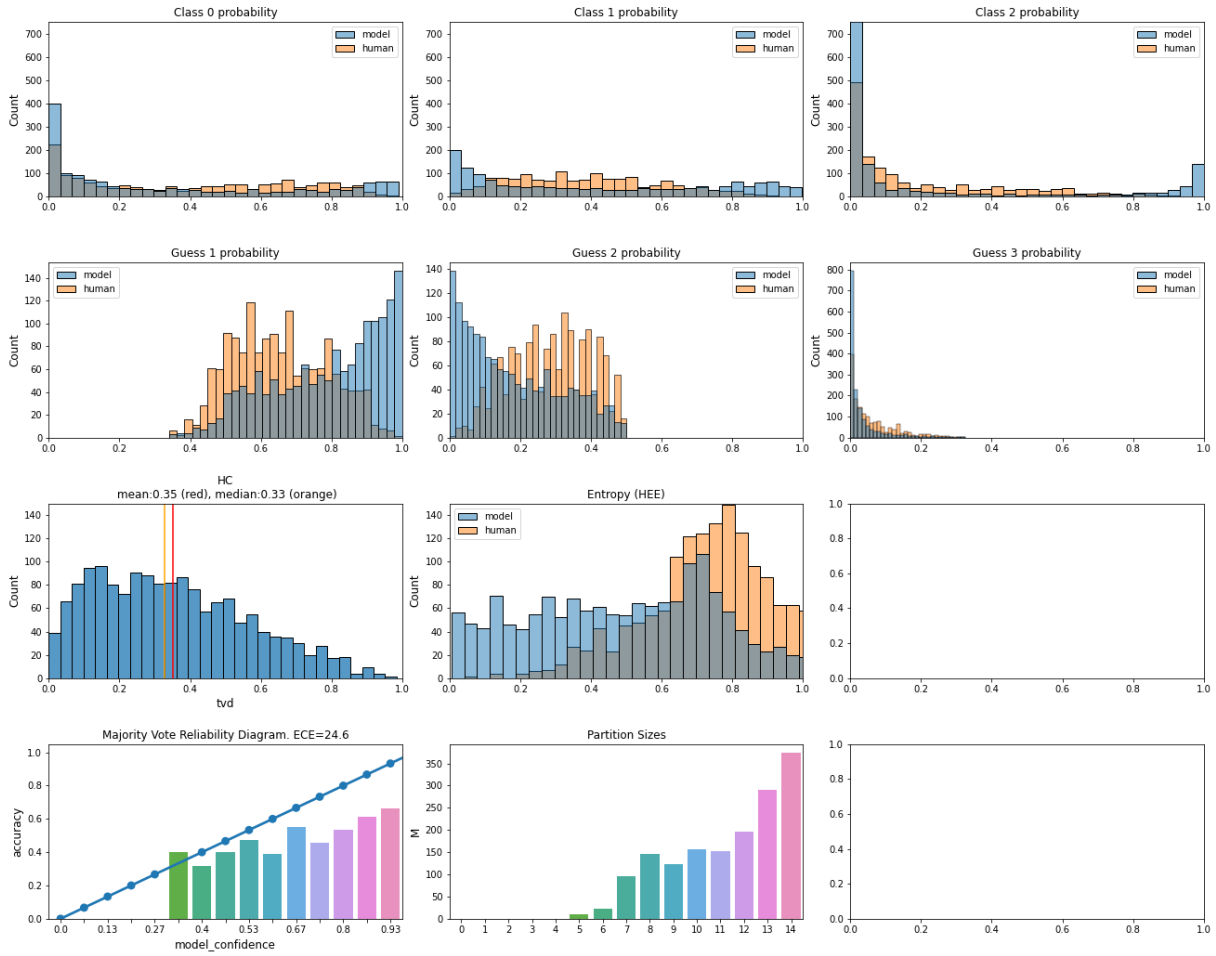


Figure 10: RoBERTa-0 with (ID) temperature scaling on the OOD ChaosNLI-MNLi dev+test set. Several figures comparing human uncertainty to model uncertainty using TVD, confidence, entropy, and reliability diagrams. This figure shows the distribution over instance-based absolute errors between probabilities for each class (top row) or the model vs human k th guess (i.e., the highest model probability versus the highest human probability on each instance). See Appendix B for more information.

ChaosNLI MNLi roberta-base 0 dev test temp=2.0 ECE=10.12, HCE=0.28, Acc=0.54, HRE=0.44

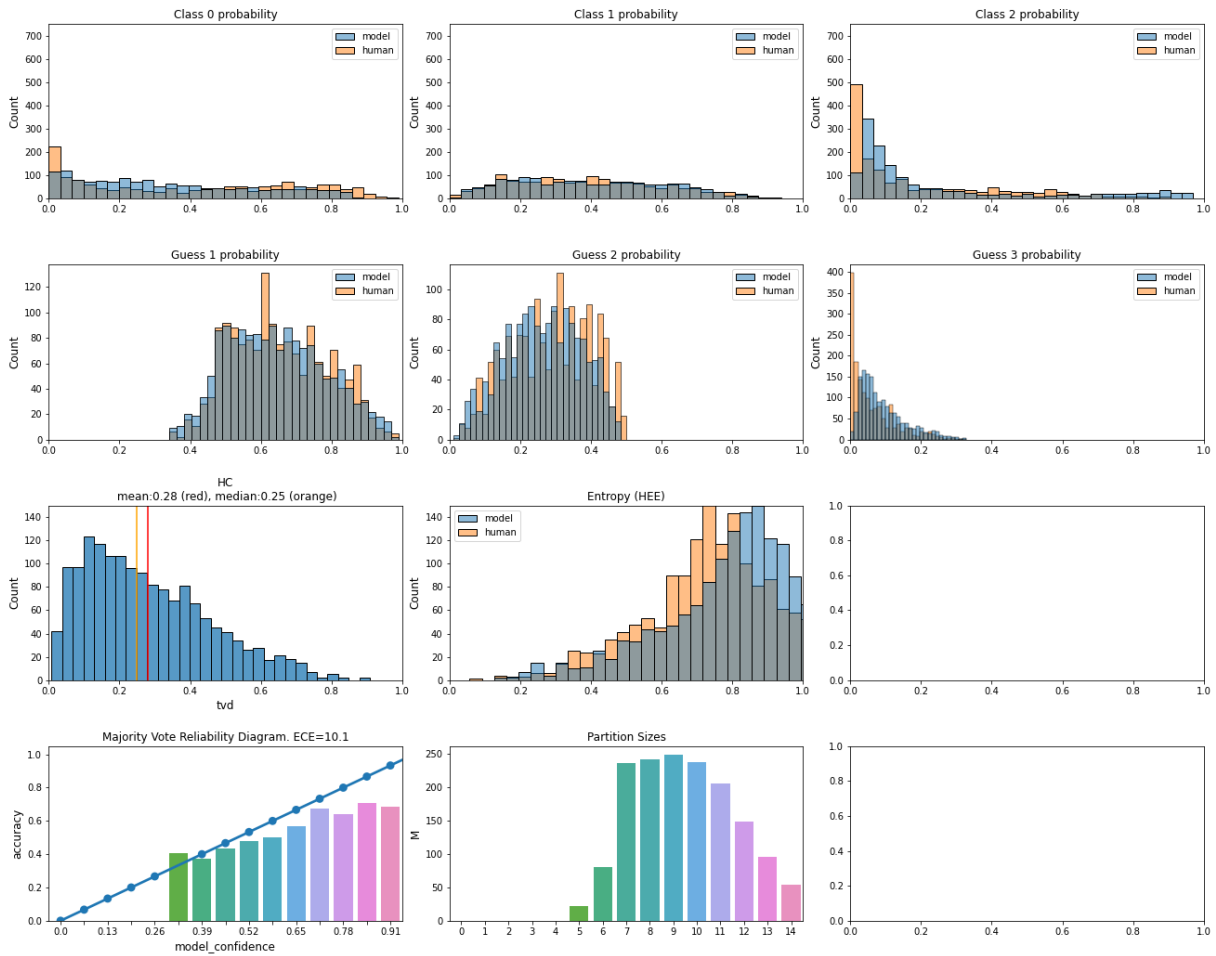


Figure 11: RoBERTa-0 with (ID) temperature scaling on the OOD ChaosNLI-MNLi dev+test set. Several figures comparing human uncertainty to model uncertainty using TVD, confidence, entropy, and reliability diagrams. This figure shows the human and model distribution over the probability, entropy or TVD range. Top row shows distribution over probability magnitudes for class 0, 1 and 2, while the second row shows the distribution for first, second and third guess for the model vs human k th guess (i.e., the highest model probability versus the highest human probability on each instance). See Appendix B for more information.

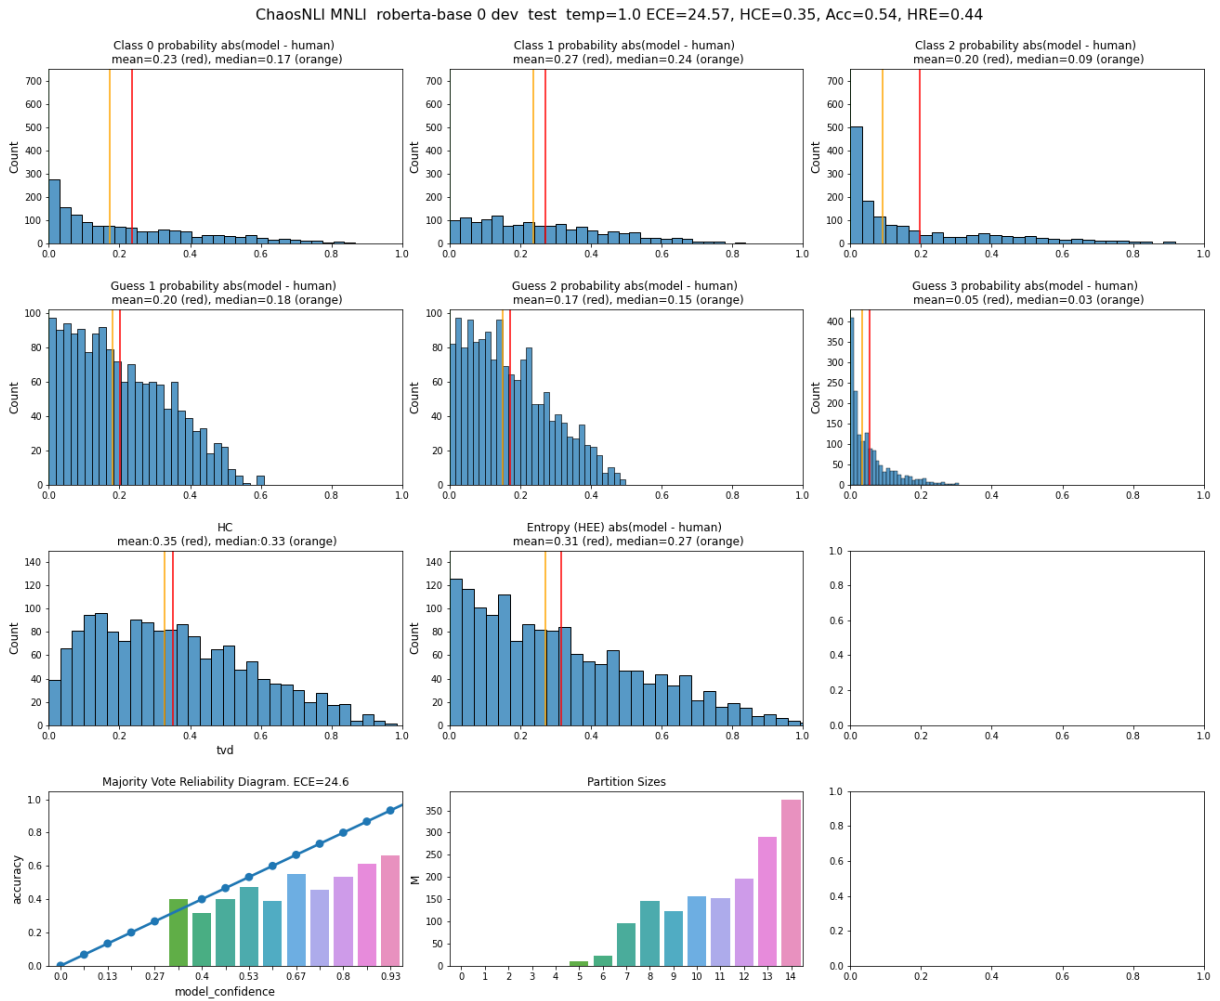


Figure 12: RoBERTa-0 with (ID) temperature scaling on the OOD ChaosNLI-MNLI dev+test set. Several figures comparing human uncertainty to model uncertainty using TVD, confidence, entropy, and reliability diagrams. This figure shows the distribution over instance-based absolute errors between probabilities for each class (top row) or the model vs human k th guess (i.e., the highest model probability versus the highest human probability on each instance). See Appendix B for more information.

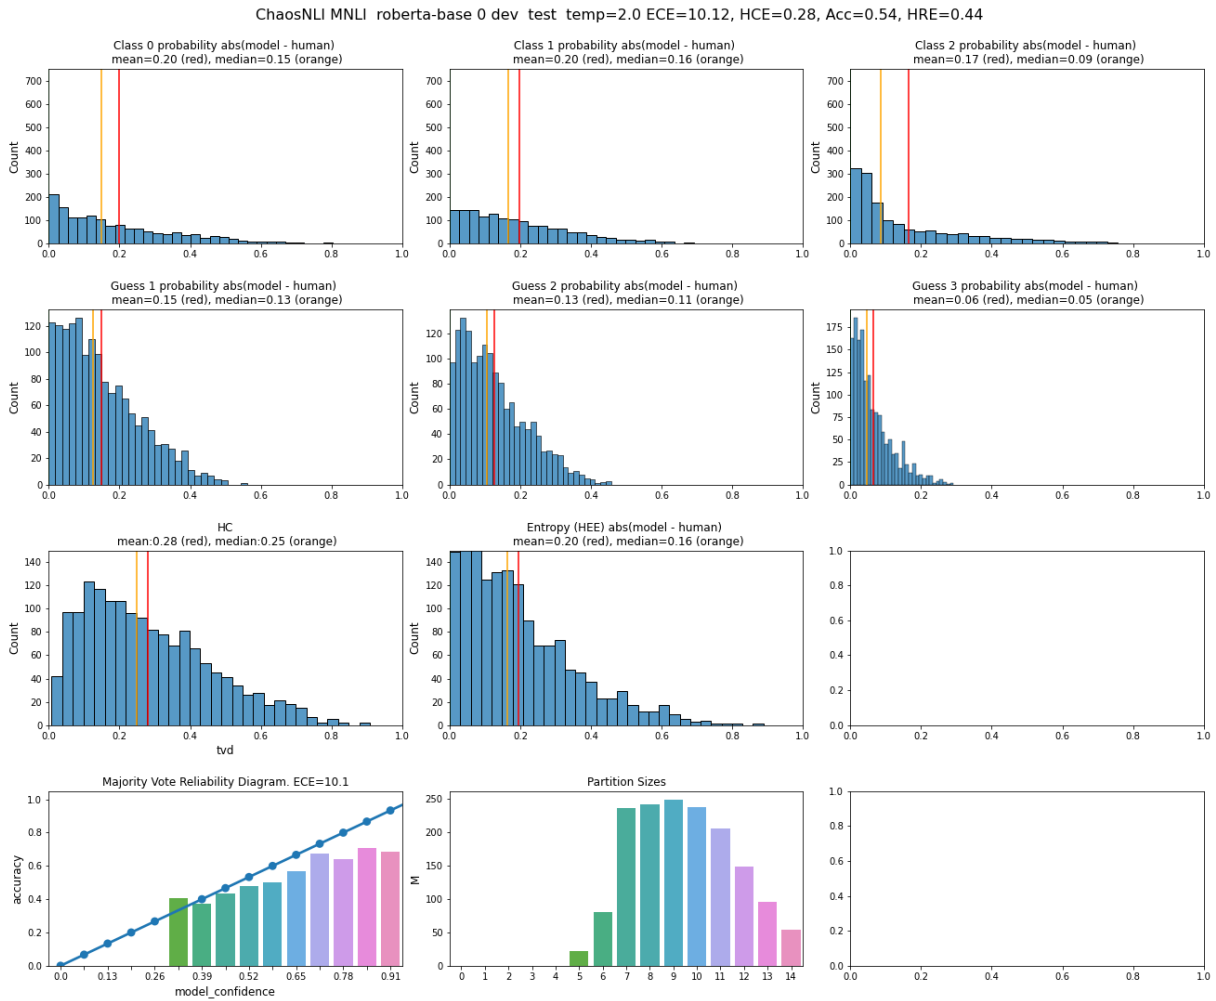


Figure 13: RoBERTa-0 with oracle temperature scaling on the ChaosNLI-SNLI dev+test set. Several figures comparing human uncertainty to model uncertainty using TVD, confidence, entropy, and reliability diagrams. This figure shows the distribution over instance-based absolute errors between probabilities for each class (top row) or the model vs human k th guess (i.e., the highest model probability versus the highest human probability on each instance). See Appendix B for more information.

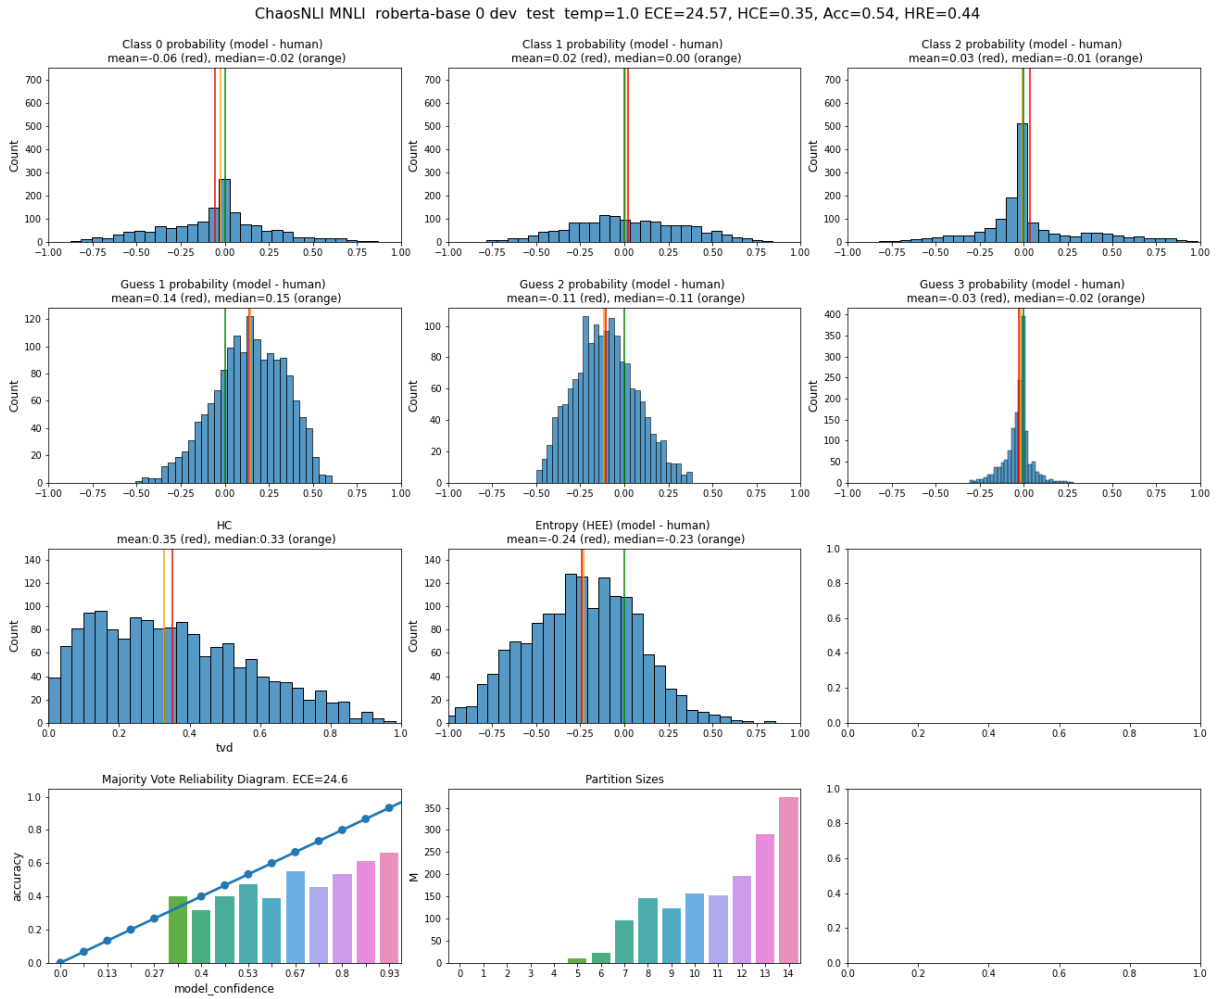


Figure 14: RoBERTa-0 with (ID) temperature scaling on the OOD ChaosNLI-MNLI dev+test set. Several figures comparing human uncertainty to model uncertainty using TVD, confidence, entropy, and reliability diagrams. This figure shows the distribution over instance-based absolute errors between probabilities for each class (top row) or the model vs human k th guess (i.e., the highest model probability versus the highest human probability on each instance). See Appendix B for more information.

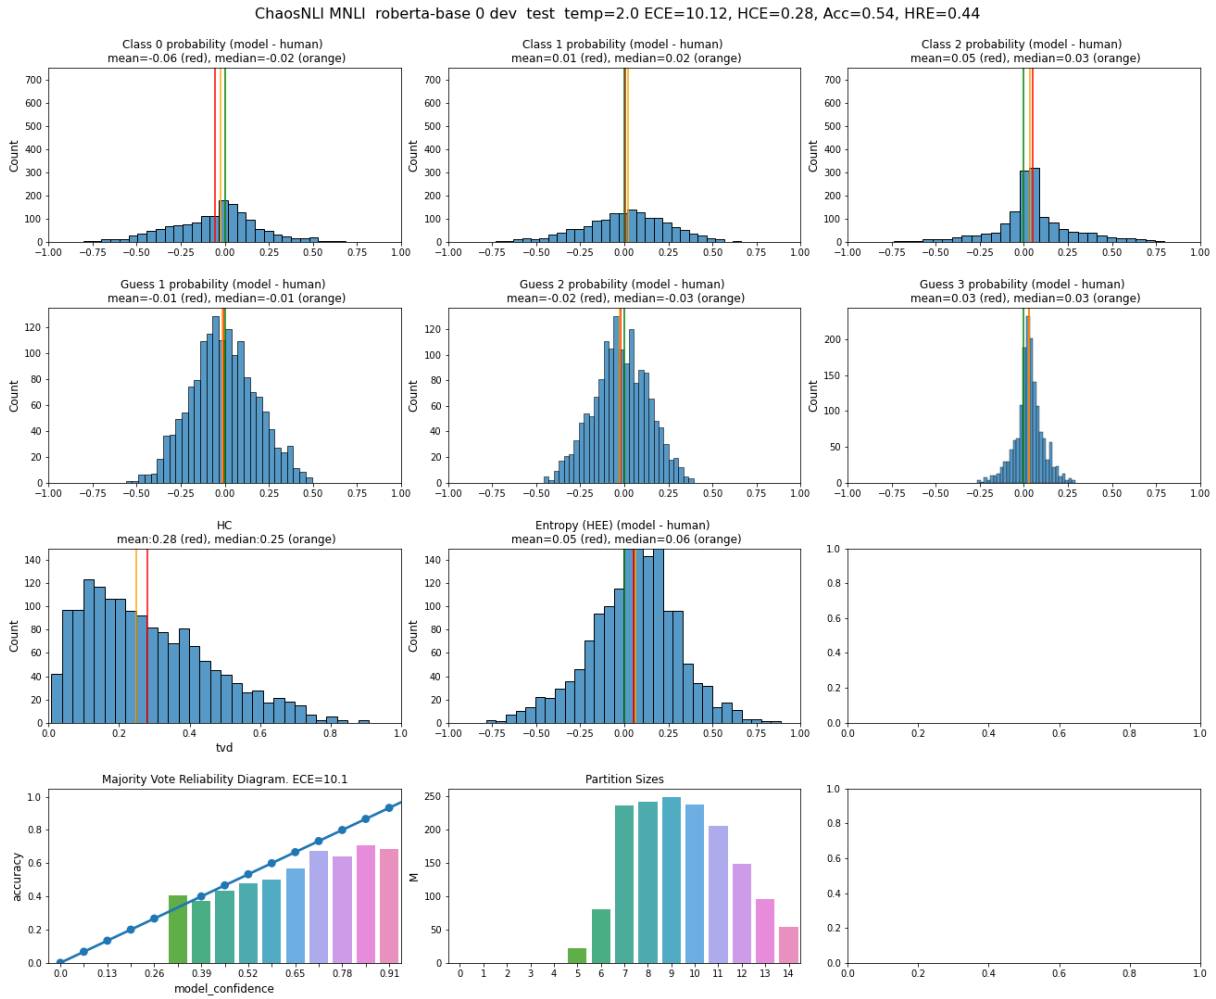


Figure 15: RoBERTa-0 with (ID) temperature scaling on the OOD ChaosNLI-MNLI dev+test set. Several figures comparing human uncertainty to model uncertainty using TVD, confidence, entropy, and reliability diagrams. This figure shows the distribution over instance-based absolute errors between probabilities for each class (top row) or the model vs human k th guess (i.e., the highest model probability versus the highest human probability on each instance). See Appendix B for more information.