Understanding Out-of-distribution: A Perspective of Data Dynamics

Anonymous Author(s)
Affiliation
Address
email

Abstract

Despite machine learning models' success in Natural Language Processing (NLP) tasks, predictions from these models frequently fail on out-of-distribution (OOD) samples. Prior works have focused on developing state-of-the-art methods for detecting OOD. The fundamental question of how OOD samples differ from in-distribution samples remains unanswered. This paper explores how data dynamics in training models can be used to understand the fundamental differences between OOD and in-distribution samples in extensive detail. We found that syntactic characteristics of the data samples that the model consistently predicts incorrectly in both OOD and in-distribution cases directly contradict each other. In addition, we observed preliminary evidence supporting the hypothesis that models are more likely to latch on trivial syntactic heuristics (e.g., overlap of words between two sentences) when making predictions on OOD samples. We hope our preliminary study accelerates the data-centric analysis on various machine learning phenomena.

1 Introduction

Detecting out-of-distribution (OOD) has become one of the key bottlenecks in building reliable open-world systems [2], leading to new SOTA approaches meant to mitigate this problem [4, 12, 14, 16, 19]. Previous works have explored the discrepancy between OOD and in-distribution test performance from the model and algorithm perspective. Le Lan & Dinh (2020) [13] observed the limitation of density estimation for anomaly detection. Zhang et al. (2021) [28] suggested that the cause of OOD failure is the combination of the model architecture and maximum likelihood objective. Choi et al. (2018) [5], Just & Ghosal (2019) [11], Fetaya et al. (2020) [8], Kirichenko et al. (2020) [1], Zhang et al. (2020) [27], and Wang et al. (2020) [25] concluded that significant mismatch between the real and estimated distribution gives rise to the performance discrepancy. Nevertheless, what fundamental differences in OOD and in-distribution samples cause them to be separated in their data distributions?

In this work, we propose a new perspective on this discrepancy and examine the differences between OOD and in-distribution data samples by considering the *training dynamics of the model*. We study this problem within the Natural Language Inference (NLI): the task of determining whether a premise sentence entails (i.e., implies the truth of) a hypothesis sentence [6, 7, 3].

We perform two types of data-dynamic analyses: We first examine the difference between OOD and in-distribution samples distributions in the data cartography space [23] at each training epoch. Cartography space allows us to distinguish which samples are harder to learn (i.e., the model often misclassifies them) based on training dynamics measures. Second, we add syntactic heuristic [18] as another component in this analysis, which enables us to identify what data characteristics tend to be harder to learn. For instance, we hypothesize whether more word overlap between a premise and a hypothesis text influences model's ability to infer their label correctly.

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

Our analyses suggested that the syntactic heuristic that the model deems hard-to-learn during training directly contradicts the characteristic of OOD samples are hard-to-learn during inference. We also found preliminary evidence suggesting the model's tendency to make inferences based on trivial syntactic heuristic is higher in the OOD case. We hope that this study will drive more effort to better understand the difference between in-distribution and OOD samples and, consequently, develop more informed and data-centric OOD detection and generalization methods.

2 Our Hypotheses

This section presents our two initial hypotheses to characterize OOD data samples. Motivated by previous works [9, 10], we perform a comparative analysis between OOD and in-distribution samples over *dynamic information of model training* (e.g., confidence, prediction variance across epoch) as training progresses. Moreover, to interpret the quantitative difference of these measures, we mark data samples based on their *syntactic heuristics*, such as word overlap between two texts. We hypothesize that the difference between OOD and in-distribution samples will be more clearly visible in the combination of multiple training dynamics (§2.1) and sample heuristics (§2.2).

2.1 H1: Are OOD and in-distribution samples different with their training dynamics?

Data cartography [23] is a tool to group data samples into three regions: easy-to-learn, hard-to-learn, and ambiguous (see Figure 1), enabled by following two training dynamics as axes: (i) *Confidence* is defined as the mean model probability of the true label (y_i^*) across epochs (equation 1). The term $p_\theta^{(e)}(y_i^*|x_i)$ denotes the model's probability with parameters θ at the end of eth epoch. (ii) *Variability* measures the spread of $p_\theta^{(e)}(y_i^*|x_i)$ across epochs, and is defined by equation 2.

$$\hat{\mu}_i = \frac{1}{E}\sum_{e=1}^{E} p_\theta^{(e)}(y_i^*|x_i) \qquad (1) \qquad\qquad \hat{\sigma}_i = \sqrt{\frac{\sum_{e=1}^{E}(p_\theta^{(e)}(y_i^*|x_i) - \hat{\mu}_i)^2}{E}} \qquad (2)$$

Intuitively, higher variability implies a high range of probability outputted by the model for the same sample. As seen in Figure 1, the rightmost region (i.e., high variability) is the ambiguous region. The easy-to-learn region is characterized by high confidence and low variability (i.e., correct prediction with high assigned probability across epochs), and hard-to-learn region samples have low confidence and low variability (i.e., incorrect prediction across epochs).

As models' performance in inferring in-distribution samples far exceeds that of the OOD case, we hypothesize that most OOD samples will stay in hard-to-learn regions. In contrast, most in-distribution samples will eventually move to the easy-to-learn region, results in a *noticeable difference between the two distributions in the cartography space.* We hope to interpret this difference using sytactic heuristic (§2.2).

2.2 H2: Do models tend to adopt syntactic attributes of OOD samples more readily?

Past works have shown that machine learning models often adopt 'shortcut' data characteristics instead of the generalization that humans would learn and use to perform the same task. For example, a model trained on NLI might assign a label of contradiction to any input containing the word "not" [20, 22], as this heuristic often applies in standard NLI training sets. As OOD samples come from a different distribution, they will have less common generalizable characteristics with the training set than the in-distribution samples. Thus, we hypothesize that *models will rely even more upon the shortcut characteristics (i.e., heuristics) when making inferences on OOD samples.*

We focus on lexical overlap heuristic [18], as it is one of the simplest and most common across NLI datasets. This heuristic assumes that a premise *entails* all hypotheses constructed from words in the premise. Below are examples:

Premise	Hypothesis	Label	Type
The judge was paid by the actor.	The actor paid the judge.	Entailed	Support
The actor was paid by the judge.	The actor paid the judge.	Not entailed	Contradict

A sample can either: (i) supporting, (ii) contradicting (i.e., possessing the property of heuristic and

labels *entailed* and *not entailed* respectively), or (iii) having no heuristic. We mark each sample's heuristics, as illustrated by the different circle colors of Figure 1.

To quantify this, we use the following measure: $m2 = \frac{|s1 \bigcap s2|}{|s2|}$. Essentially, $m2$ measures how much percentage of words in the premise is found in the hypothesis. This aligns with the lexical overlap definition: all words in the hypothesis (i.e., $s2$) is contained in the premise (i.e. $s1$).

To measure the model's tendency to adapt the heuristic, we calculate the correlation between $m2$ and confidence $\hat{\mu}_i$ and observe the trend across epochs. We calculate this correlation for all samples and also for samples from each class independently. We then observe how these trends vary as training progresses. Intuitively, a high heuristic adoption is indicated by high absolute correlation values in the individual class samples (e.g., the model learns that high words overlap means the two sentences are entailed, vice versa).

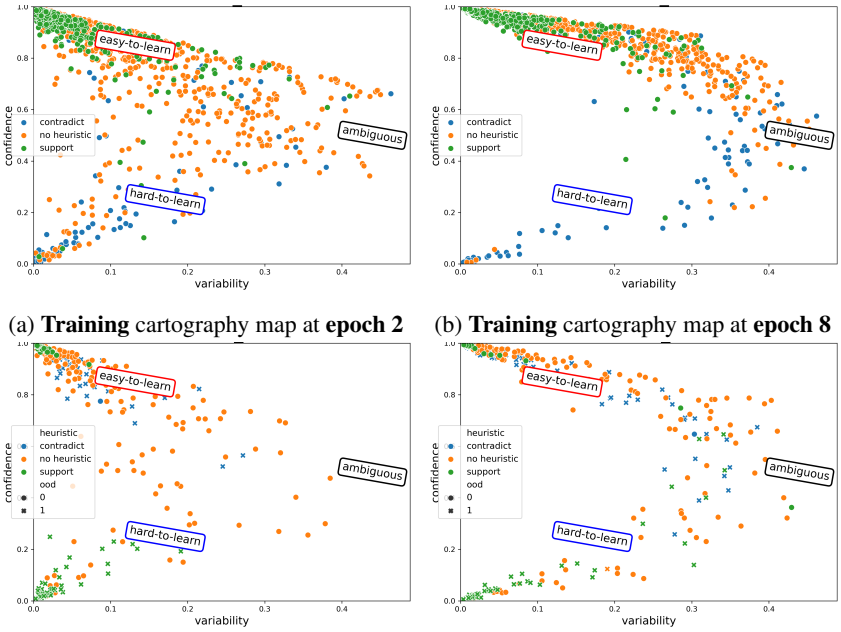

(a) **Training** cartography map at **epoch 2** (b) **Training** cartography map at **epoch 8**

(c) **Evaluation** cartography map at **epoch 2** (d) **Evaluation** cartography map at **epoch 8**

Figure 1: Training and evaluation cartography maps (train and in-dist: MNLI, OOD: WNLI). The different colors illustrate different sample heuristics and the different shapes in figures (c) and (d) illustrate OOD sample ('o' for in-distribution and 'x' for OOD).

3 Experiments

Setup: In all experiments, we trained Roberta [17] with batch size 20 and learning rate $1.1e^{-5}$, and initialized with the same random seed. All experiments were carried out using PyTorch [21] with one NVidia Tesla K20X GPU. Experiments were carried out using three datasets (MNLI [26], RTE [24], WNLI [15]) which combinations are shown in table 1. Our extensive experiments results can be found in Appendix A.3, Appendix A.1, and Appendix A.2.

Table 1: Datasets combination in our experiments.

Training	Eval (In-distribution)	Eval (OOD)
MNLI (train)	MNLI (dev matched)	WNLI (train)
MNLI (train)	MNLI (dev matched)	RTE (dev)
RTE (train)	RTE (dev)	WNLI (train)

3.1 OOD vs in-distribution on training dynamics information

Result for this section are the cartography maps in Figure 1. For visualization clarity, the points showed are a sampled fraction from the original set. We observe the following phenomena: A contradicting pattern between the trajectories of training and OOD samples. By comparing figures (a) and (b) (beginning and end of the training, respectively), we observed samples that **contradict the heuristic** (blue circles) either stays in the hard-to-learn region or move towards the ambiguous region. On contrary, the evaluation maps (figures (c) and (d)) shows that OOD samples that **support the heuristic** (green 'x's) stay in the hard-to-learn region. This implies that samples characteristics

(i.e., lexical overlap heuristic) that the model deems hard/easy to learn during training *completely flips in the OOD case* (especially during initial epochs).

At the end of epoch 8, we observe relatively more color mix at the ambiguous region for the OOD case (blue and green 'x's at the ambiguous region of map (d)). While in the training epoch 8, this mix is more apparent towards the easy-to learn region (blue and green circles). We conjecture that a *more generalizable knowledge is learned as more training epoch is performed* (indicated by the samples mix). Although, it is slower to be adopted by the model in inferring OOD samples.

3.2 OOD vs in-distribution on syntactic characteristics

For the second hypothesis, we find the following observations: Correlation values are relatively low in the plot for all samples. This applies to all cases (train, in-distribution, and OOD test). However, an obvious divergence is observed in the correlation between OOD samples confidence score and the heuristics measure for the entailment samples. This observation confirms our hypothesis in 2.2 that models are *more prone to adopt syntactic heuristics when making inference on OOD samples*.

All samples' trends adhere to the per-label trends in the train and in-distribution test case but contradict the OOD case. In the OOD case, we can see that the negative and non-zero slopes are only found in all samples trend, but the correlation keeps increasing (with a huge difference in value) in the entailment samples trend. This might indicate that knowledge used by the model to infer class labels is *consistent for all classes in train and in-distribution but conflicts in the OOD case*. More supplementary plots of this analysis (more experiments on different datasets and the trend plots for **non-entailment** samples) can be found in appendix A.4.

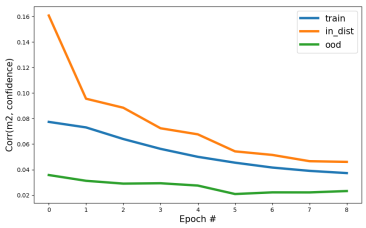

(a) Correlation between **m2** and **all** samples $\hat{\mu}_i$

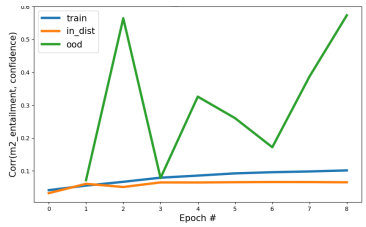

(b) Correlation between **m2** and **entailment** samples $\hat{\mu}_i$

Figure 2: Hypothesis 2.2: Training and in-distr. test samples are MNLI. OOD samples are WNLI.

4 Limitations and Future Directions

This data-centric approach to understanding the difference between OOD and in-distribution samples offers an intuition of failure mode in OOD. However, the results and approach presented in this paper have some limitations.

Although all NLI sets in our experiments came from different topic domains and had distinct characteristics, a more careful selection is needed to ensure maximum separation of in-distribution vs. OOD data distribution. Zhang et al. (2021) [28] highlights the problem with experiments setup where OOD regions can lie in support of the data distribution: high logit values assigned to certain OOD samples due to model estimation error. By ensuring this criterion is met, we can achieve a more accurate confidence score, which in turn will improve our results in section 3.1.

More controlled experiment is needed for results in section 3.2 to ensure that the correlation between confidence and m2 is not spurious (i.e., there exists no unobserved confounding variable). One way to remedy this is by running the same measurements on random subsets and see if the correlation pattern still holds (i.e., randomized controlled experiments).

Future work may include remedies for the two mentioned limitations. Furthermore, investigating the root cause of contradicting syntactic heuristics of OOD and in-distribution samples that the model fails to predict correctly (section 3.1) is also an exciting direction. Adding more syntactic heuristics for analysis (e.g., coreference, negation) can further corroborate the analysis and conjecture presented in this paper. Finally, we hope that the insights provided by this data-centric view can motivate more informed development of OOD detection and generalization methods.

References

[1] Why normalizing flows fail to detect out-of-distribution data. *Advances in Neural Information Processing Systems*, 2020-December, 2020.

[2] Abhijit Bendale and Terrance Boult. Towards open world recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1893–1902, 2015.

[3] Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 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, September 2015. Association for Computational Linguistics.

[4] Jiefeng Chen, Yixuan Li, Xi Wu, Yingyu Liang, and Somesh Jha. Atom: Robustifying out-of-distribution detection using outlier mining.

[5] Hyun-Jae Choi, Eric Jang, and Alexander Amir Alemi. Waic, but why? generative ensembles for robust anomaly detection. *arXiv: Machine Learning*, 2018.

[6] Cleo Condoravdi, Dick Crouch, Valeria de Paiva, Reinhard Stolle, and Daniel G. Bobrow. Entailment, intensionality and text understanding. In *Proceedings of the HLT-NAACL 2003 Workshop on Text Meaning*, pages 38–45, 2003.

[7] Ido Dagan, Oren Glickman, and Bernardo Magnini. The pascal recognising textual entailment challenge. In Joaquin Quiñonero-Candela, Ido Dagan, Bernardo Magnini, and Florence d'Alché Buc, editors, *Machine Learning Challenges. Evaluating Predictive Uncertainty, Visual Object Classification, and Recognising Tectual Entailment*, pages 177–190, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.

[8] Ethan Fetaya, Jörn-Henrik Jacobsen, Will Grathwohl, and Richard Zemel. Understanding the limitations of conditional generative models. *arXiv preprint arXiv:1906.01171*, 2019.

[9] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *Proceedings of International Conference on Learning Representations*, 2017.

[10] Dan Hendrycks, Xiaoyuan Liu, Eric Wallace, Adam Dziedzic, Rishabh Krishnan, and Dawn Song. Pretrained transformers improve out-of-distribution robustness. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2744–2751, Online, July 2020. Association for Computational Linguistics.

[11] John Just and Sambuddha Ghosal. Deep generative models strike back! improving understanding and evaluation in light of unmet expectations for ood data. *arXiv preprint arXiv:1911.04699*, 2019.

[12] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, page 6405–6416, Red Hook, NY, USA, 2017. Curran Associates Inc.

[13] Charline Le Lan and Laurent Dinh. Perfect density models cannot guarantee anomaly detection. *arXiv preprint arXiv:2012.03808*, 2020.

[14] Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. *Advances in neural information processing systems*, 31, 2018.

[15] Hector Levesque, Ernest Davis, and Leora Morgenstern. The winograd schema challenge. In *Thirteenth International Conference on the Principles of Knowledge Representation and Reasoning*, 2012.

[16] Weitang Liu, Xiaoyun Wang, John D Owens, and Yixuan Li. Energy-based out-of-distribution detection. *arXiv preprint arXiv:2010.03759*, 2020.

[17] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.

[18] Tom McCoy, Ellie Pavlick, and Tal Linzen. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy, July 2019. Association for Computational Linguistics.

[19] Sina Mohseni, Mandar Pitale, JBS Yadawa, and Zhangyang Wang. Self-supervised learning for generalizable out-of-distribution detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 5216–5223, 2020.

[20] Aakanksha Naik, Abhilasha Ravichander, Norman Sadeh, Carolyn Rose, and Graham Neubig. Stress test evaluation for natural language inference. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2340–2353, Santa Fe, New Mexico, USA, August 2018. Association for Computational Linguistics.

[21] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.

[22] Ivan Sanchez, Jeff Mitchell, and Sebastian Riedel. Behavior analysis of NLI models: Uncovering the influence of three factors on robustness. 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 1975–1985, New Orleans, Louisiana, June 2018. Association for Computational Linguistics.

[23] Swabha Swayamdipta, Roy Schwartz, Nicholas Lourie, Yizhong Wang, Hannaneh Hajishirzi, Noah A. Smith, and Yejin Choi. Dataset cartography: Mapping and diagnosing datasets with training dynamics. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9275–9293, Online, November 2020. Association for Computational Linguistics.

[24] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. GLUE: A multi-task benchmark and analysis platform for natural language understanding. 2019. In the Proceedings of ICLR.

[25] Ziyu Wang, Bin Dai, David Wipf, and Jun Zhu. Further analysis of outlier detection with deep generative models, 2020.

[26] Adina Williams, Nikita Nangia, and Samuel Bowman. 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, New Orleans, Louisiana, June 2018. Association for Computational Linguistics.

[27] Hongjie Zhang, Ang Li, Jie Guo, and Yanwen Guo. Hybrid models for open set recognition. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 102–117, Cham, 2020. Springer International Publishing.

[28] Lily Zhang, Mark Goldstein, and Rajesh Ranganath. Understanding failures in out-of-distribution detection with deep generative models. In *International Conference on Machine Learning*, pages 12427–12436. PMLR, 2021.

A Appendix

A.1 Supplementary plot: OOD vs in-distribution on training dynamics information (Training and in-dis: RTE; OOD: WNLI)

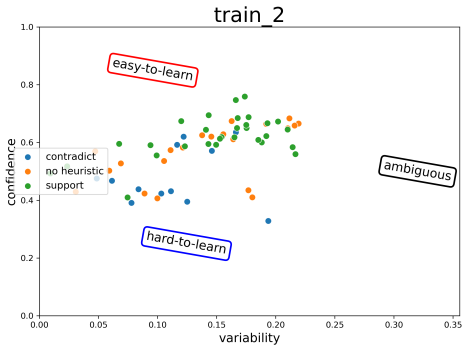

(a) **Training** cartography map at **epoch 2**

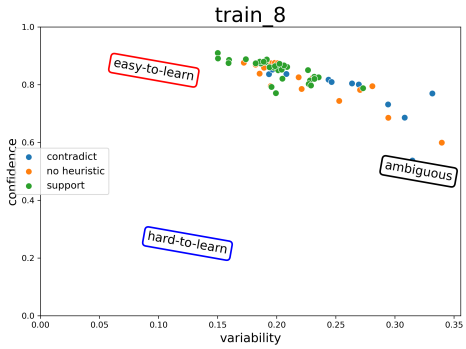

(b) **Training** cartography map at **epoch 8**

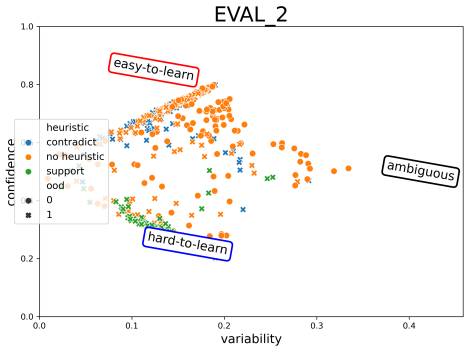

(c) **Evaluation** cartography map at **epoch 2**

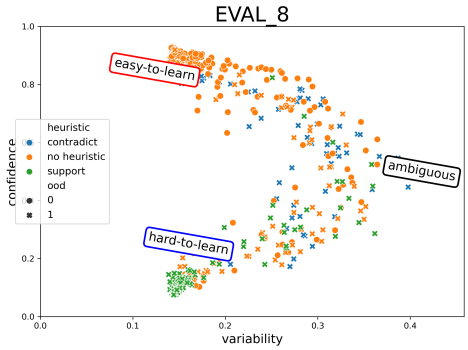

(d) **Evaluation** cartography map at **epoch 8**

Figure 3: Training cartography maps (training set: RTE). The number of heuristics related samples in RTE is small.

A.2 Supplementary plot: OOD vs in-distribution on syntactic characteristics (entailment)

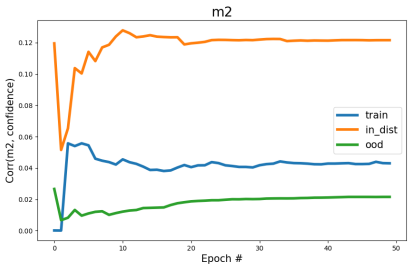

(a) Correlation between **m2** and **all** samples $\hat{\mu}_i$

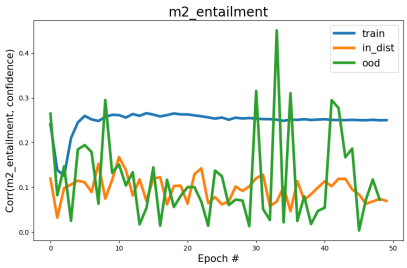

(b) Correlation between **m2** and **entailment** samples $\hat{\mu}_i$

Figure 4: Results for hypothesis 2.2. Training and in-distribution test samples are RTE, and OOD samples are WNLI.

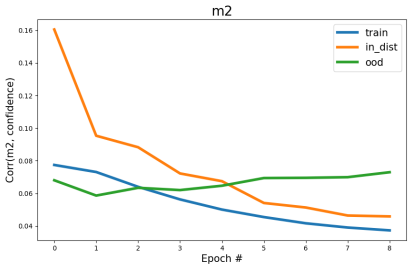

(a) Correlation between **m2** and **all** samples $\hat{\mu}_i$

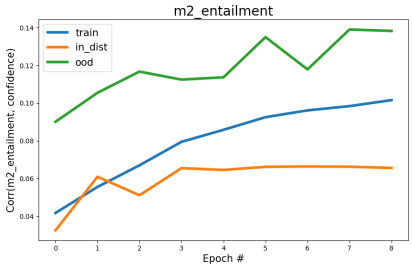

(b) Correlation between **m2** and **entailment** samples $\hat{\mu}_i$

Figure 5: Results for hypothesis 2.2. Training and in-distribution test samples are MNLI, and OOD samples are RTE.

Results presented are at the end of epoch 8 for MNLI training and the end of epoch 50 for RTE training. This is based on the epoch in which the training error has converged (around 0.02).

A.3 Supplementary plot: OOD vs in-distribution on training dynamics information (Training and in-dis: MNLI; OOD: RTE)

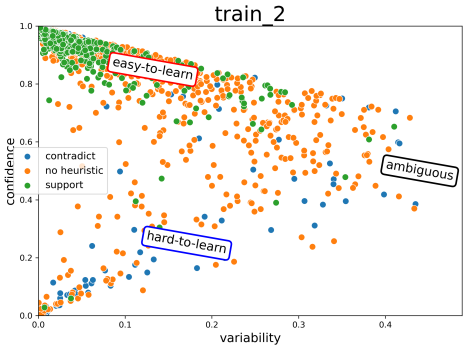

(a) **Training** cartography map at **epoch 2**

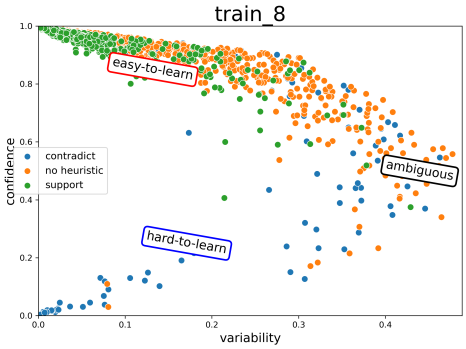

(b) **Training** cartography map at **epoch 8**

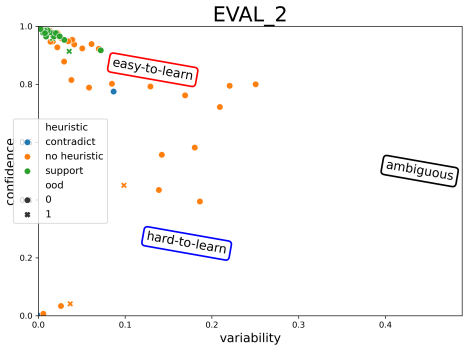

(c) **Evaluation** cartography map at **epoch 2**

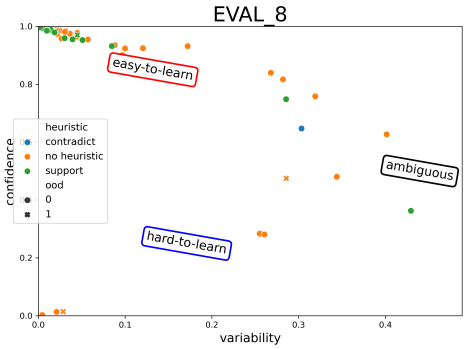

(d) **Evaluation** cartography map at **epoch 8**

Figure 6: Training and evaluation cartography maps (train: MNLI, evaluation: RTE). The number of heuristics related samples in RTE is small.

A.4 Supplementary plot: OOD vs in-distribution on syntactic characteristics (non-entailment)

This section shows plots for correlation between confidence scores ($\hat{\mu}_i$) of **non entailment** samples and m2

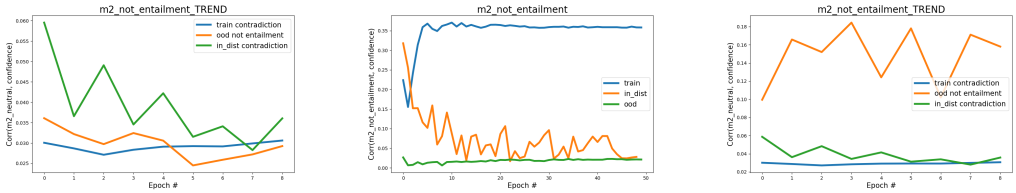

(a) Train & in-distribution: MNLI, OOD: WNLI

(b) Train & in-distribution: RTE, OOD: WNLI

(c) Train & in-distribution: MNLI, OOD: RTE

Figure 7: Supplementary results for 3.2. Correlation between $\hat{\mu}_i$ of **non-entailment** samples and m2

A.5 Supplementary material: Extra lexical overlap measure

We also added another measure to quantify tendency to adopt lexical overlap heuristic. We calculated $m1 = \frac{|s1 \bigcap s2|}{|s1|}$. Essentially, this measures how much percentage of words found in the premise ($s1$) can also be found in the hypothesis ($s2$).

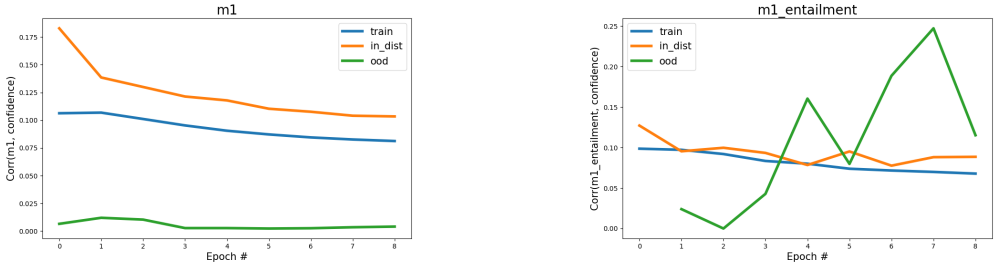

(a) Correlation between **m2** and **all** samples $\hat{\mu}_i$

(b) Correlation between **m2** and **entailment** samples $\hat{\mu}_i$

Figure 8: Results for hypothesis 2.2. Training and in-distribution test samples are MNLI, and OOD samples are WNLI.

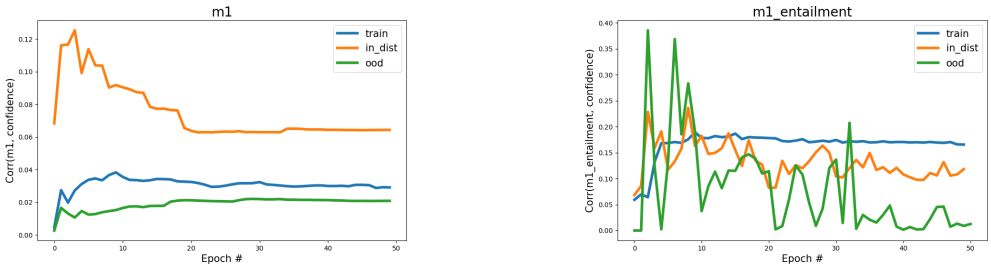

(a) Correlation between **m2** and **all** samples $\hat{\mu}_i$

(b) Correlation between **m2** and **entailment** samples $\hat{\mu}_i$

Figure 9: Results for hypothesis 2.2. Training and in-distribution test samples are RTE, and OOD samples are WNLI.

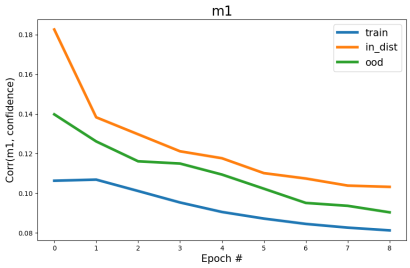

(a) Correlation between **m2** and **all** samples $\hat{\mu}_i$

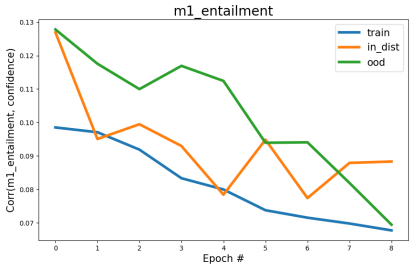

(b) Correlation between **m2** and **entailment** samples $\hat{\mu}_i$

Figure 10: Results for hypothesis 2.2. Training and in-distribution test samples are MNLI, and OOD samples are RTE.