

# Partial-input baselines show that NLI models can ignore context, but they don't.

Neha Srikanth

Department of Computer Science  
University of Maryland, College Park  
nehasrik@umd.edu

Rachel Rudinger

Department of Computer Science  
University of Maryland, College Park  
rudinger@umd.edu

## Abstract

When strong partial-input baselines reveal artifacts in crowdsourced NLI datasets, the performance of full-input models trained on such datasets is often dismissed as reliance on spurious correlations. We investigate whether state-of-the-art NLI models are capable of overriding default inferences made by a partial-input baseline. We introduce an evaluation set of 600 examples consisting of perturbed premises to examine a RoBERTa model's sensitivity to edited contexts. Our results indicate that NLI models are still capable of learning to condition on context—a necessary component of inferential reasoning—despite being trained on artifact-ridden datasets.

## 1 Introduction

Natural language inference (NLI) is integral to building systems that are broadly capable of language understanding (White et al., 2017). In a traditional NLI setup, models are provided with a *premise* as context and a corresponding *hypothesis*. They must then determine whether the premise entails, contradicts, or is neutral in relation to the hypothesis (Giampiccolo et al., 2007).

Researchers have shown that many NLI datasets contain statistical biases, or “annotation artifacts” (Gururangan et al., 2018; Herlihy and Rudinger, 2021) that systems leverage to correctly predict entailment. To diagnose such artifacts in datasets and provide a stronger alternative to majority-class baselines, Poliak et al. (2018) introduced partial-input baselines, a setting in which models are provided parts of NLI instances to predict an entailment relation. Poliak et al. (2018), Gururangan et al. (2018), and others posit that datasets containing artifacts may in turn produce models that are incapable of learning to perform true reasoning.

In this paper, we re-examine such dismissal of the reasoning capabilities of models trained on datasets containing artifacts. While a competitive partial-input baseline is sufficient to cast doubt

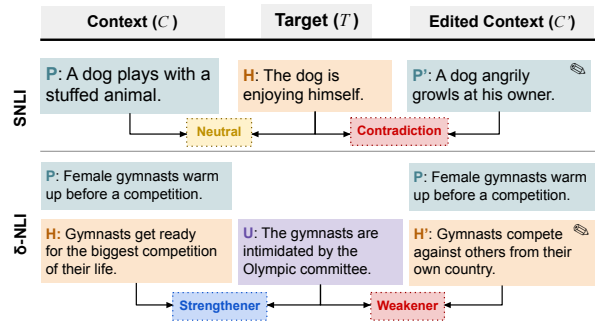


Figure 1: Given NLI and  $\delta$ -NLI instances that partial-input baselines correctly label, we edit their context to induce a different gold label. We use the edited instances to probe a full-input model's ability to leverage context to deviate from the predicted partial-input label.

on the inferential capabilities of full-input models trained on the same data, it is premature to conclude that full-input models are not capable of such reasoning at all. As a thought experiment, we imagine training a human to learn a task from a set of examples containing artifacts that allow them to cheat. If we force them to learn to perform the task by removing relevant context, they must rely on the artifacts to accurately perform the task. However, provided all the data, they may still learn to perform the task the “right way”.

Through two sets of experiments, we investigate whether NLI models are able to condition on the full input despite learning from artifact-ridden datasets. Section 3 investigates whether additional context strengthens a full-input model's confidence in the correct label, despite a partial-input model's correct prediction. In Section 4, we introduce an evaluation set to probe whether full-input models are sensitive to changes in context that flip the gold label in examples containing artifacts. We make our evaluation set and code publicly available.<sup>1</sup>

Our results indicate that full-input models are still successfully learning to utilize context, overriding strong signal in the partial input. With this

<sup>1</sup><https://github.com/nehasrikn/context-editing>

finding, we argue that while partial-input baselines are still a useful tool, they do not license the conclusion that models trained on datasets with artifacts do not learn to leverage context from the full input.

## 2 Background

Here we describe the two related tasks of natural language inference (NLI) and *defeasible* NLI; two corresponding datasets (SNLI and  $\delta$ -NLI, respectively); and annotation artifacts present in these datasets.

**Task Definitions.** Natural language inference (Giampiccolo et al., 2007; MacCartney, 2009; Bowman et al., 2015) is the task of determining whether a **premise** sentence ( $P$ ) entails a **hypothesis** sentence ( $H$ ). That is, given  $P$ , would a human conclude  $H$  is true (*entailment*), false (*contradiction*), or neither (*neutral*).

When  $H$  is neutral, the task of *defeasible* natural language inference (Rudinger et al., 2020) asks whether a third **update** sentence ( $U$ ) makes  $H$  more likely to be true ( $U$  is a *strengthenener*) or less likely to be true ( $U$  is a *weakenner*). In the following example,  $H$  is *neutral* given  $P$  (NLI task), and  $U$  is a *weakenner* given  $P$  and  $H$  (defeasible NLI task).

**Premise:** A man is sitting in a dim restaurant.  
**Hypothesis:** He is eating food.  
**Update:** He is browsing a menu.

In this work, we adopt the terms **context** ( $C$ ) and **target** ( $T$ ) for clarity when describing partial-input baselines for the two tasks. In NLI,  $C = P$  and  $T = H$ ; in defeasible NLI,  $C = (P, H)$ , and  $T = U$ . Thus, for either task, the partial-input baseline we use looks only at  $T$  and ignores  $C$ .

**SNLI.** SNLI (Bowman et al., 2015) is the first large-scale English NLI dataset, containing 570K labeled  $P$ - $H$  pairs. In SNLI, premises are derived from image captions (Young et al., 2014), and hypotheses for each label (*entailment*, *neutral*, *contradiction*) are elicited from crowdsource workers who are shown a premise.

**$\delta$ -NLI.** For the task of *defeasible* NLI, Rudinger et al. (2020) introduce the  $\delta$ -NLI dataset, which consists of extensions to three pre-existing English natural language reasoning datasets: SNLI (Bowman et al., 2015), ATOMIC (Sap et al., 2019), and SOCIAL-CHEM-101 (Forbes et al., 2020). To extend each dataset, instances (e.g.,  $P$ - $H$  pairs)

were presented to crowdworkers who then wrote an update sentence ( $U$ ) to strengthen or weaken the given hypothesis (as described in Task Definitions). The resulting binary classification task is to predict whether an update is a strengthener or weakener, given a  $(P, H, U)$  triple. In this work, we focus our evaluation on  $\delta$ -SNLI, the SNLI portion of the  $\delta$ -NLI dataset.

**Artifacts and Partial-Input Baselines.** Gururangan et al. (2018) and Poliak et al. (2018) observe that the crowdsourcing protocols adopted by Bowman et al. (2015) and others lead to the creation of data with annotation artifacts that enable partial-input baselines (e.g., hypothesis-only baselines) to perform well above a majority-class baseline. For SNLI, Poliak et al. (2018) report that an InferSent (Conneau et al., 2017) hypothesis-only baseline surpasses a majority baseline by 35 points. A similar effect is observed by Rudinger et al. (2020) in the  $\delta$ -NLI data, with an update-only RoBERTa (Liu et al., 2019) model achieving 15 points above a majority baseline.

## 3 Experiment 1: Context in NLI

An essential component of true reasoning is learning to leverage all parts of an example’s input to make a determination of entailment. Strong partial-input models demonstrate that full-input models do not necessarily need to utilize context to make correct predictions. However, we explore whether they do at all, and how access to such context shifts a full-input model’s confidence in the correct label. If, upon supplying  $C$ , a model strengthens its confidence in its prediction, we may conclude that it utilizes both  $C$  and  $T$  during inference as intended.

**Experimental Setup.** We finetune two sets of RoBERTa (Liu et al., 2019) models (a partial-input and a full-input model) on the train splits of SNLI and  $\delta$ -NLI (see Appendix A for dataset sizes). We utilize `roberta-base` from the Hugging Face library (Wolf et al., 2020), and finetune each model for two epochs. Appendix B further details our training setup. We then run inference on the test splits from SNLI and  $\delta$ -SNLI (the SNLI portion of  $\delta$ -NLI) using each pair of models. Table 1 reports accuracy of all models on the corresponding test splits. We calibrate RoBERTa models post-hoc using temperature scaling (Guo et al., 2017) as suggested by Desai and Durrett (2020), and examine confidence shifts in the correct label to understand

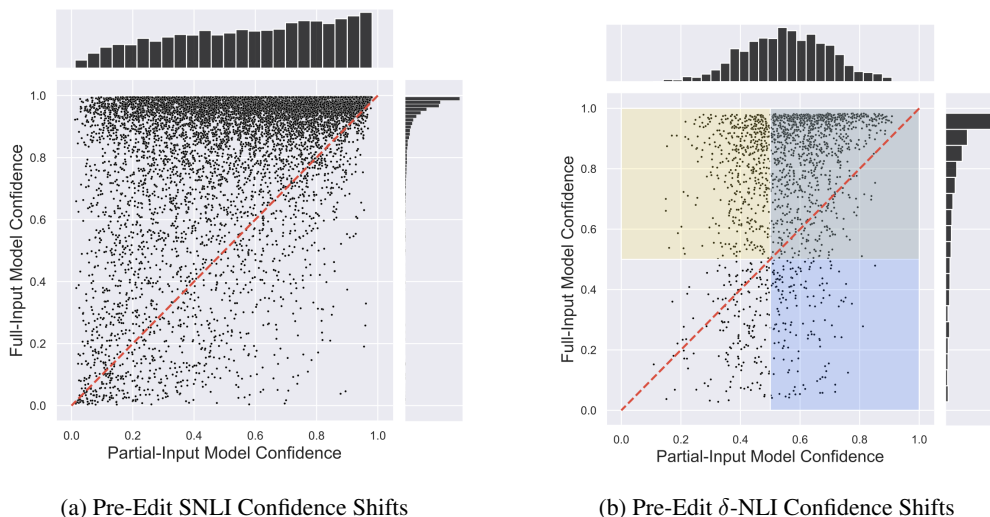


Figure 2: Visualizing the role of context as confidence shifts. 2a and 2b plot ordered pairs of confidence in the correct label  $l$  of partial-input vs. full-input RoBERTa models on examples from the SNLI and  $\delta$ -SNLI test splits. Blue and yellow regions mark correct predictions from the x and y-axis models respectively. Green regions represent correct predictions from both models. Diagonals indicate no shift in confidence.

	SNLI	$\delta$ -SNLI
Partial-Input ( $C$ )	0.70	0.65
Full-Input ( $C, T$ )	0.91	0.82

Table 1: Accuracy of both partial-input and full-input models on test splits of SNLI and  $\delta$ -SNLI.

whether full-input models utilize context at all.<sup>2</sup>

As shown in Fig. 2, we plot an ordered pair of each model’s confidence in the *correct label* for examples in the test splits, with the partial-input model’s confidence along the x-axis and the full-input model’s confidence along the y-axis. Density around the diagonal would indicate no change in confidence.

Evidenced by the density *above the diagonal* in Figures 2a and 2b, full-input models (i.e access to both  $C$  and  $T$ ) are more confident in the correct label than partial-input models.<sup>3</sup> While this behavior may seem unsurprising, partial-input baselines illustrate that models may show confidence in the correct label without needing to condition on context at all. Our results hint that full-input models may be successfully learning to leverage additional context instead of overgeneralizing on artifacts in the target. To probe this behavior directly, we intro-

<sup>2</sup>We observed minimal differences between pre and post-calibration results.

<sup>3</sup>The distribution difference between SNLI and  $\delta$ -SNLI may be attributed to the difference in the task difficulty, as well overall lower performance of RoBERTa models on  $\delta$ -NLI.

duce an evaluation set crafted by editing contexts in the following section.

## 4 Experiment 2: Context Editing

We investigate a model’s ability to leverage context despite the presence of artifacts by exploring how sensitive full-input models are to changes in *non-target* components of the input. We present an example modification scheme, illustrated in Figure 1, in which we edit context sentences from examples where a model correctly predicts the label  $l$  from the target  $T$  alone. Namely, while holding  $T$  constant, we introduce an edited context sentence  $C'$  that induces a different label  $l' \neq l$  on the new  $(C', T)$  pair. Using this scheme, we construct an evaluation set of 600 examples sourced from SNLI and  $\delta$ -NLI.

**Example Subselection.** We select SNLI and  $\delta$ -SNLI test examples to edit by running the partial-input RoBERTa models from Section 3 and full-input bag-of-words (BoW) models, implemented via `fasttext` with a maximum of 4-grams (Joulin et al., 2017). See Appendix B for training details on the bag-of-words models. We select examples to edit for which *either* the partial-input model or the BoW model predicted the correct label. This identifies the subset of examples likeliest to contain artifacts in  $T$ , lexical or otherwise.

**Editing SNLI Examples.** For a given SNLI example  $(P, H, l)$  and a new predefined target label

		$l'$		
		entailment	neutral	contradiction
$l$	entailment	–	0.76	0.76
	neutral	0.42	–	0.78
	contradiction	0.90	0.78	–

Table 2: A full-input RoBERTa model’s accuracy on the edited SNLI portion of our evaluation set.  $l$  and  $l'$  represent the original and target label respectively, before and after editing.

		$l'$	
		weakener	strengtheners
$l$	weakener	–	0.76
	strengtheners	0.75	–

Table 3: A full-input RoBERTa model’s accuracy on the edited  $\delta$ -SNLI portion of our evaluation set.  $l$  and  $l'$  again respectively represent the original and induced post-edit label.

$l'$ , we edit  $P$ , creating a modified SNLI example  $(P', H, l')$ . For each of the six directional pairs of labels (e.g., *entailment*  $\rightarrow$  *contradiction*), we randomly sample 50 examples from the subset to edit, resulting in 300 examples evenly distributed across label pairs.

**Editing  $\delta$ -SNLI Examples.** Given a  $\delta$ -SNLI instance  $(P, H, U, l)$ , we edit  $H$  while holding  $P$  and  $U$  constant to induce a new label  $l' \neq l$ , resulting in a modified example  $(P, H', U, l')$ . We edit 300 examples total, turning 150 *strengtheners* examples into *weakener* examples, and vice-versa.

Our final evaluation set consists of 600 examples containing edited context-target pairs split evenly across SNLI and  $\delta$ -SNLI. Figure 4 shows examples of editing contexts from both datasets. All examples were manually edited by one author and independently validated by another. During validation, we hide both the  $l$  and  $l'$ , and ask the annotator to label the text pair. Using Cohen’s Kappa (Cohen, 1960), we obtain an agreement measure of  $\kappa = 0.78$  and  $\kappa = 0.76$  for SNLI and  $\delta$ -SNLI examples in our test set respectively, indicating substantial agreement (Artstein and Poesio, 2008).

**Results.** Using the same full-input RoBERTa models trained in Section 3, we run inference on our edited evaluation set. Tables 2 and 3 show model performance stratified by original label  $l$  and target label  $l'$ . Our results show that full-input models are in fact sensitive to context modifications despite the presence of artifacts in  $T$ , consistently achieving above 70% accuracy on edited examples.

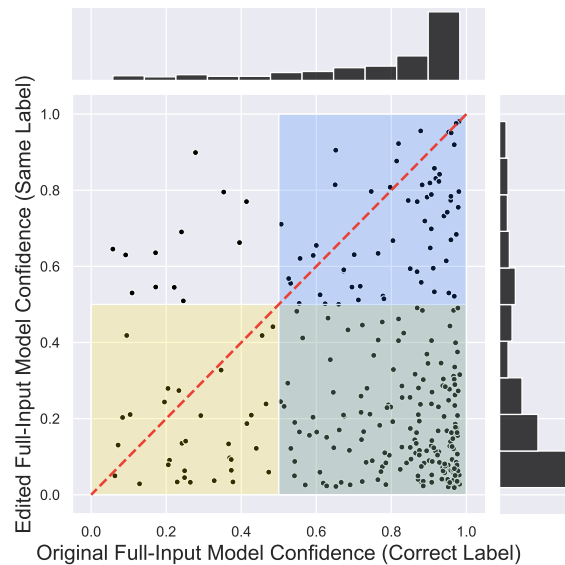


Figure 3: Full-input model’s confidence in  $l$  before vs. after editing  $C$  to induce a different label.

Thus, we conclude that these models are not over-generalizing on artifacts in the instance, learning to condition on context for prediction.

**Analyzing Post-Edit Model Confidence.** Similar to the analysis in Section 3, we inspect shifts in confidence upon editing contexts to shed more light on a full-input model’s utilization of  $C$ . For  $\delta$ -SNLI examples, we plot ordered pairs in Fig. 3 of a full-input model’s confidence in the correct label pre-edit, and its confidence in the *same label* post-edit (i.e the confidence in the now-incorrect label). The majority of mass is under the diagonal, indicating that our model is indeed sensitive to changes in context. The green bottom-left quadrant delineates ideal performance (correct before *and* after editing examples). We attribute the small cluster of examples in the blue quadrant (previously highly confident in the correct label and subsequently remained confident, but in the wrong label) to strong, non-lexical artifacts overriding additional signal from the context. Appendix C visualizes shifts in SNLI examples using simplex plots to accommodate the ternary label.

**Lexical Model Performance.** We do not retrain any of the models on edited examples prior to evaluating with our edited test set, precluding them from picking up on any newly-introduced non-lexical artifacts. However, to validate the absence of trivial lexical features that override the artifacts in  $T$  and to ensure sufficient difficulty, we run full-input BoW models on our edited evaluation set. A full-



	Context ( $C \rightarrow C'$ )	Target ( $I$ )	Label ( $l \rightarrow l'$ )
SNLI	<p><b>Premise:</b> A little girl in a pink hat is in a lush green field walking an ox.</p> <p><b>Edited Premise</b> 📄: A little girl in a pink hat sits on an ox carrying her through the middle of the Sahara.</p>	<p><b>Hypothesis:</b> A little girl is riding her ox in a desert.</p>	<p><b>Contradiction</b></p> <p>→</p> <p><b>Entailment</b></p>
	<p><b>Premise:</b> A girl in a yellow dress looks at a restaurant building.</p> <p><b>Edited Premise</b> 📄: A girl in a yellow dress enjoys some food at her favorite restaurant.</p>	<p><b>Hypothesis:</b> A girl debates eating at a certain restaurant.</p>	<p><b>Neutral</b></p> <p>→</p> <p><b>Contradiction</b></p>
$\delta$ -NLI	<p><b>Premise:</b> A boy in a red hooded top is smiling whilst looking away from his reflection.</p> <p><b>Hypothesis:</b> The boy doesn't want to see his reflection.</p> <p><b>Edited Hypothesis</b> 📄: The boy wants to see his reflection.</p>	<p><b>Update:</b> A woman's voice says to him 'Don't be afraid of it, it's harmless'</p>	<p><b>Strengthened</b></p> <p>→</p> <p><b>Weakened</b></p>
	<p><b>Premise:</b> A music concerto is happening on stage, with the lights on spot, while the audience listens in the dark.</p> <p><b>Hypothesis:</b> The concert is being held at night.</p> <p><b>Edited Hypothesis</b> 📄: The concert is being held during the day.</p>	<p><b>Update:</b> Blackout curtains cover the windows.</p>	<p><b>Weakened</b></p> <p>→</p> <p><b>Strengthened</b></p>

Figure 4: Editing SNLI and  $\delta$ -SNLI examples. We edit the premise (SNLI) and the hypothesis ( $\delta$ -NLI), while holding the target constant. The last column contains the original label  $l$  and the new induced label  $l'$ .

input `fasttext` model, often used for adversarial filtering (Zellers et al., 2018), achieves 16% and 24.3% accuracy on the SNLI and  $\delta$ -SNLI portions of our evaluation set respectively.

## 5 Related Work

In addition to work on partial-input baselines for NLI (Gururangan et al., 2018; Poliak et al., 2018; Tsuchiya, 2018), partial-input models have been studied for story completion (Cai et al., 2017) and reading comprehension (Kaushik and Lipton, 2018). Feng et al. (2019) observe that low-scoring partial-input baselines do not preclude other artifacts and heuristics, while Glockner et al. (2018) and McCoy et al. (2019) demonstrate lexical and syntactic examples of NLI heuristics. Adversarial editing has been explored for non-NLI tasks as well (Jia and Liang, 2017; Ribeiro et al., 2018). Finally, adversarial filtering has been proposed as a means of removing artifacts from datasets (Zellers et al., 2018; Le Bras et al., 2020).

## 6 Discussion and Conclusion

While partial-input models are useful tools for analysis, often leveling fair criticism of datasets, our results show it is hasty to conclude that models trained on such datasets are not capable of reasoning. Even though high-scoring partial-input baselines show that full-input models *could* ignore context, our experiments show that they can leverage this context quite effectively.

We argue that artifacts do not *necessarily* spell

disaster for a model’s reasoning capabilities. In particular, our context-editing experiments identify a set of instances that partial-input models fail (by design), but full-input models largely succeed at, displaying the capability of full-input models to leverage context to overcome SNLI and  $\delta$ -NLI artifacts in many, but not all, cases. Of course, we do not deny that artifacts *can* and *do* lead to models with exploitable heuristics, as demonstrated by Glockner et al. (2018) and McCoy et al. (2019).

While we do not attempt to define the *sufficient* conditions for a model to perform “true inference,” we demonstrate that these models can and do meet the *necessary* condition of leveraging the full input. Thus, we conclude that partial input baselines should be understood as agnostic warning signs: sufficient to conclude that full-input models *might* not be leveraging critical context, but insufficient to prove that they don’t.

This raises a number of interesting questions for follow-up work. If adversarial filtering (Le Bras et al., 2020) can identify instances containing artifacts, is it beneficial to remove these instances from the training set? Or could they be edited to flip the label and mitigate the spurious correlation? The edits we made in this work were done manually, but another interesting question is whether these edits could be made automatically or semi-automatically. Having a more efficient way of producing these examples would enable both rapid evaluation of models trained on datasets with artifacts (as in this work), as well as expansion of training sets to preemptively mitigate artifacts.

## References

- Ron Artstein and Massimo Poesio. 2008. Inter-coder agreement for computational linguistics. *Computational Linguistics*, 34(4):555–596.
- 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.
- Zheng Cai, Lifu Tu, and Kevin Gimpel. 2017. [Pay attention to the ending: strong neural baselines for the ROC story cloze task](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 616–622, Vancouver, Canada. Association for Computational Linguistics.
- Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1):37–46.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. [Supervised learning of universal sentence representations from natural language inference data](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 670–680, Copenhagen, Denmark. Association for Computational Linguistics.
- 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.
- Shi Feng, Eric Wallace, and Jordan Boyd-Graber. 2019. [Misleading failures of partial-input baselines](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5533–5538, Florence, Italy. Association for Computational Linguistics.
- Maxwell Forbes, Jena D. Hwang, Vered Shwartz, Maarten Sap, and Yejin Choi. 2020. [Social chemistry 101: Learning to reason about social and moral norms](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 653–670, Online. Association for Computational Linguistics.
- Danilo Giampiccolo, Bernardo Magnini, Ido Dagan, and Bill Dolan. 2007. [The third PASCAL recognizing textual entailment challenge](#). In *Proceedings of the ACL-PASCAL Workshop on Textual Entailment and Paraphrasing*, pages 1–9, Prague. Association for Computational Linguistics.
- Max Glockner, Vered Shwartz, and Yoav Goldberg. 2018. [Breaking NLI systems with sentences that require simple lexical inferences](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 650–655, Melbourne, Australia. Association for Computational Linguistics.
- 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.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. [Annotation artifacts in natural language inference data](#). 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 107–112, New Orleans, Louisiana. Association for Computational Linguistics.
- Marc Harper. 2015. [python-ternary: Ternary plots in python](#). *Zenodo 10.5281/zenodo.594435*.
- Christine Herlihy and Rachel Rudinger. 2021. [MedNLI is not immune: Natural language inference artifacts in the clinical domain](#). 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 1020–1027, Online. Association for Computational Linguistics.
- Robin Jia and Percy Liang. 2017. [Adversarial examples for evaluating reading comprehension systems](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2021–2031, Copenhagen, Denmark. Association for Computational Linguistics.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. [Bag of tricks for efficient text classification](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 427–431, Valencia, Spain. Association for Computational Linguistics.
- Divyansh Kaushik and Zachary C. Lipton. 2018. [How much reading does reading comprehension require? a critical investigation of popular benchmarks](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5010–5015, Brussels, Belgium. Association for Computational Linguistics.
- Ronan Le Bras, Swabha Swayamdipta, Chandra Bhagavatula, Rowan Zellers, Matthew Peters, Ashish Sabharwal, and Yejin Choi. 2020. [Adversarial filters of dataset biases](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 1078–1088. 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*.
- Bill MacCartney. 2009. *Natural language inference*. Stanford University.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. [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. Association for Computational Linguistics.
- Adam Poliak, Jason Naradowsky, Aparajita Haldar, Rachel Rudinger, and Benjamin Van Durme. 2018. [Hypothesis only baselines in natural language inference](#). In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pages 180–191, New Orleans, Louisiana. Association for Computational Linguistics.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2018. [Semantically equivalent adversarial rules for debugging NLP models](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 856–865, Melbourne, Australia. Association for Computational Linguistics.
- Alexey Romanov and Chaitanya Shivade. 2018. [Lessons from natural language inference in the clinical domain](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1586–1596, Brussels, Belgium. Association for Computational Linguistics.
- Rachel Rudinger, Vered Shwartz, Jena D Hwang, Chandra Bhagavatula, Maxwell Forbes, Ronan Le Bras, Noah A Smith, and Yejin Choi. 2020. Thinking like a skeptic: Defeasible inference in natural language. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 4661–4675.
- Maarten Sap, Ronan Le Bras, Emily Allaway, Chandra Bhagavatula, Nicholas Lourie, Hannah Rashkin, Brendan Roof, Noah A Smith, and Yejin Choi. 2019. Atomic: An atlas of machine commonsense for if-then reasoning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3027–3035.
- Masatoshi Tsuchiya. 2018. [Performance impact caused by hidden bias of training data for recognizing textual entailment](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Aaron Steven White, Pushpendre Rastogi, Kevin Duh, and Benjamin Van Durme. 2017. Inference is everything: Recasting semantic resources into a unified evaluation framework. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 996–1005.
- 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, New Orleans, Louisiana. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2:67–78.
- Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. 2018. [SWAG: A large-scale adversarial dataset for grounded commonsense inference](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 93–104, Brussels, Belgium. Association for Computational Linguistics.

Dataset	Split	Size
SNLI	train	550,152
	valid	10,000
	test	10,000
$\delta$ -NLI	train	200,694
	valid	14,968
	test	15,414
$\delta$ -SNLI	train	88,676
	valid	1,785
	test	1,837

Table 4: Split sizes across SNLI,  $\delta$ -NLI, and  $\delta$ -SNLI.

## A SNLI and $\delta$ -NLI Dataset Sizes

To train and tune our neural and lexical models, we utilize the train and validation splits from SNLI and  $\delta$ -NLI. We include train/validation/test split sizes in Table 4. Our  $\delta$ -NLI RoBERTa models were finetuned on examples from all 3 portions of the  $\delta$ -NLI dataset (ATOMIC, SNLI, and SOCIAL-CHEM-101), however for our evaluation set and analysis, we exclusively use the SNLI portion of the  $\delta$ -NLI dataset, abbreviated as  $\delta$ -SNLI. We include the split sizes for  $\delta$ -SNLI in Table 4 as well.

## B Model Training Setup

**Neural Models.** We use the Hugging Face library to train all of our RoBERTa models. We utilize `roberta-base`, which has 125M trainable parameters. All models were trained on a single NVIDIA 1080-TI GPU. After tuning on the validation set, all models were trained for two epochs with a learning rate of  $2e-5$  and a batch size of 32. Tables 2 and 3 report best accuracy across five runs.

**Lexical Models.** We use the `fasttext` library to implement our bag-of-words models. `fasttext` is an off-the-shelf text classification library. All lexical models were trained for 5 epochs with 4-grams as the maximum length of word ngrams. We use the default learning rate of 0.1.

## C Visualizing Distribution Shifts in SNLI Edited Examples

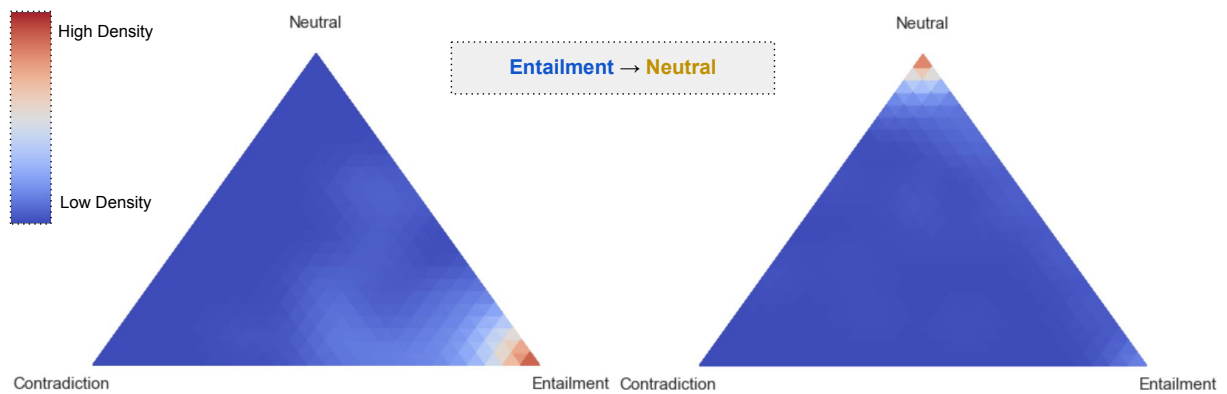
We are able to visualize confidence distribution shifts for  $\delta$ -NLI before and after editing using a 2D plane with ordered pairs, as in Figure 3, due to the label set containing only two update types—weakener and strengthener (so, a probability score  $> 0.5$  results in the predicted label). Since

the SNLI label set consists of *three* labels (*entailment*, *neutral*, *contradiction*), we choose to visualize shifts in the confidence distribution before and after editing contexts via ternary plots. For each directional label pair (*entailment*  $\rightarrow$  {*neutral*, *contradiction*}, *neutral*  $\rightarrow$  {*entailment*, *contradiction*}, *contradiction*  $\rightarrow$  {*entailment*, *neutral*}), we plot a heatmap of probabilities, or confidences, in each of the three classes on the simplex with Gaussian smoothing using `python-ternary`, a ternary plotting package (Harper, 2015). Figures 5, 6, and 7 show these visualizations. We include these plots mainly to help visualize information about the predicted labels of the incorrect examples (Table 2 only reports the accuracy on finer-grained buckets). We observe that for most classes of examples, the SNLI RoBERTa model utilizes the context, and correctly predicts the new induced gold label. However, the {*neutral*  $\rightarrow$  *entailment*} class of examples in particular proved difficult for the model, as evidenced by a large chunk of a mass remaining in the neutral corner of the simplex.

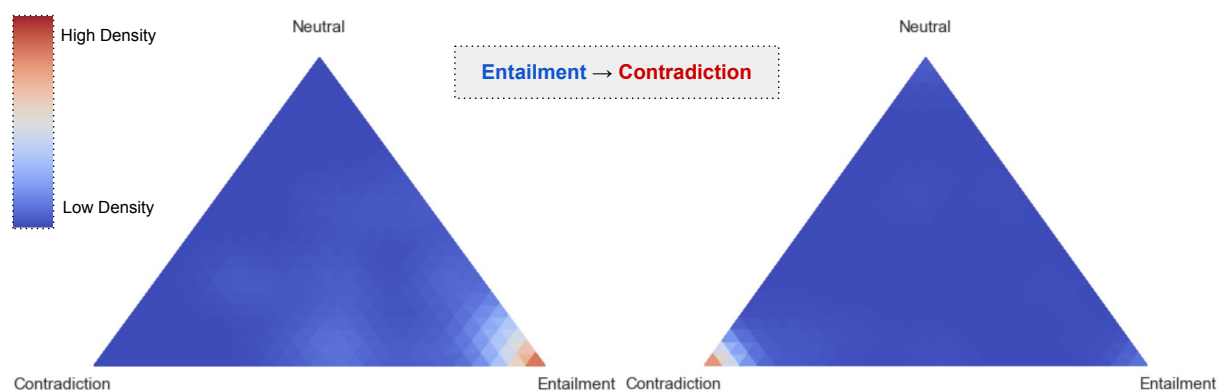
## D Dataset Limitations

In this work, we choose to explore the role of context with respect to the SNLI and  $\delta$ -NLI datasets. In particular, the proven presence of strong artifacts in SNLI made it an appealing dataset to explore a model’s behavior with respect to the utilization of context. We chose to include  $\delta$ -NLI in our analysis, since ultimately, we’d like reasoning systems to operate in complex and *dynamic* contexts. The ability to be sensitive to shifting contexts and understand when default inferences should be overridden by additional context (i.e more nuanced inference) is both central to our exploration and central to the task of defeasible reasoning itself. Our evaluation set does not include examples sourced from other NLI datasets such as MultiNLI (Williams et al., 2018). It also does not contain datasets across domains, such as MedNLI (Romanov and Shivade, 2018). However, we note that while the datasets may be different, others have shown artifacts present in such datasets (i.e (Herlihy and Rudinger, 2021; Gururangan et al., 2018)). Our goal was to utilize datasets containing high amounts of artifacts.



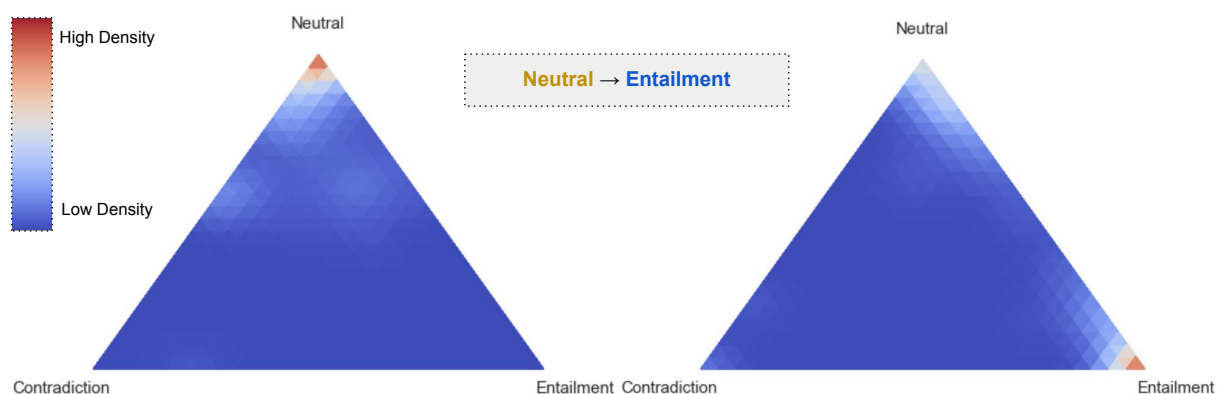


(a) Confidence distribution heatmap for examples with original relation as entailment, edited to induce a neutral relation.

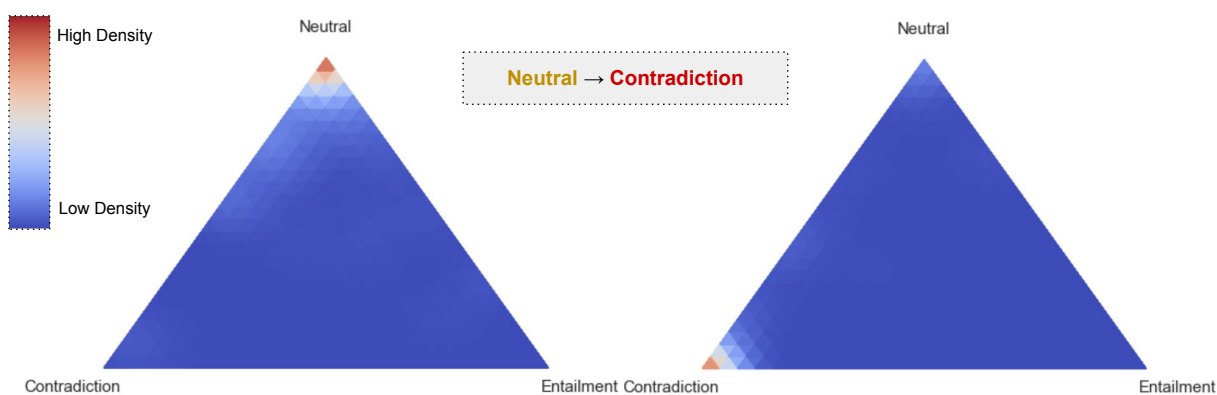


(b) Confidence distribution heatmap for examples with original relation as entailment, edited to induce a contradiction relation.

Figure 5: Confidence distribution heatmaps for SNLI examples before and after editing examples with original relation as entailment in our evaluation set. 5a shows entailment examples edited to induce a neutral relation, and 5b shows examples edited to induce a contradiction relation. Each  $l$  and  $l'$  is shown in the center box. We note that these distributions help visualize the accuracies presented in Table 2.

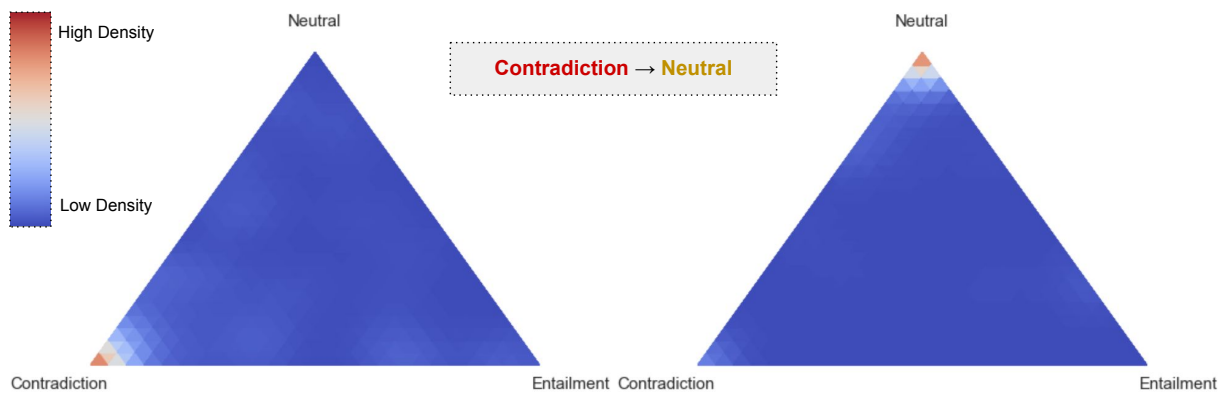


(a) Confidence distribution heatmap for examples with original relation as neutral, edited to induce an entailment relation. We note that this was the hardest edit category for RoBERTa models to flip, and draw attention to a substantial amount of mass still in the neutral corner of the simplex.

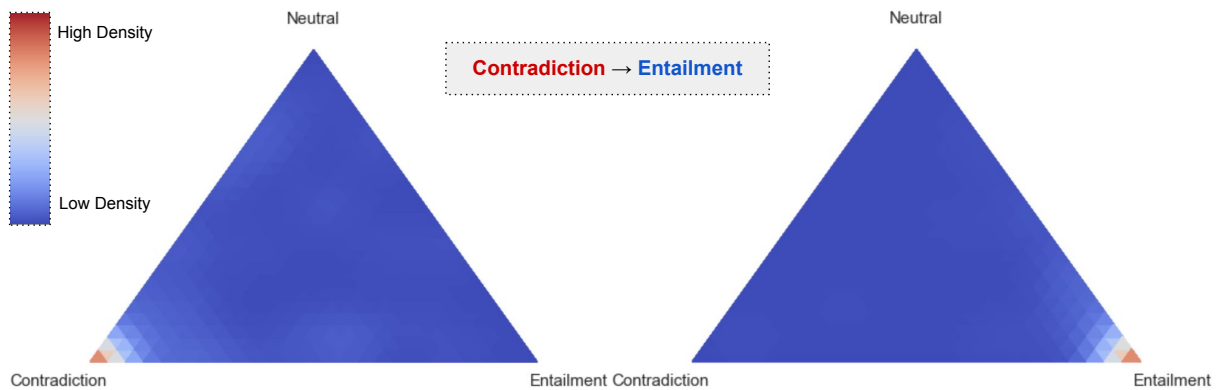


(b) Confidence distribution heatmap for examples with original relation as neutral, edited to induce a contradiction relation.

Figure 6: Confidence distribution heatmaps for SNLI examples before and after editing examples with original relation as neutral in our evaluation set. 6a shows examples edited to induce an entailment relation, and 6b shows examples edited to induce a contradiction relation. Each  $l$  and  $l'$  is shown in the center box. These distributions aid in visualization of the performance metrics reported in Table 2.



(a) Confidence distribution heatmap for examples with original relation as contradiction, edited to induce a neutral relation.



(b) Confidence distribution heatmap for examples with original relation as contradiction, edited to induce an entailment relation. We note that this particular class of edited examples achieved the highest accuracy, reflected in the low density of examples away from the entailment corner of the simplex.

Figure 7: Confidence distribution heatmaps for SNLI examples before and after editing examples with original relation as contradiction in our evaluation set. 7a shows examples edited to induce a neutral relation, and 7b shows examples edited to induce an entailment relation. Each  $l$  and  $l'$  is shown in the center box. These plots help visualize the performance metrics reported in Table 2.