

Generalization in NLI: Ways (Not) To Go Beyond Simple Heuristics

Prajjwal Bhargava¹, Aleksandr Drozd^{2,3}, Anna Rogers^{3,4}

¹ The University of Texas at Dallas

² RIKEN Center for Computational Science

³ Tokyo Institute of Technology

⁴ Center for Social Data Science, University of Copenhagen

¹prajjwalin@protonmail.com, ²alex@blackbird.pw, ³arogers@sodas.ku.dk

Abstract

Much of recent progress in NLU was shown to be due to models’ learning dataset-specific heuristics. We conduct a case study of generalization in NLI (from MNLI to the adversarially constructed HANS dataset) in a range of BERT-based architectures (adapters, Siamese Transformers, HEX debiasing), as well as with subsampling the data and increasing the model size. We report 2 successful and 3 unsuccessful strategies, all providing insights into how Transformer-based models learn to generalize.

1 Introduction

Many popular NLP datasets contain spurious patterns, which get learned instead of the actual task (Gururangan et al., 2018; Belinkov et al., 2019; Rogers et al., 2020a; Gardner et al., 2021). This raises the issue of learning methods that would avoid that problem. We present a case study of generalization to adversarial data in Natural Language Inference (NLI), reporting both positive and negative results for a range of BERT-based approaches.

2 Methodology

Data. NLI is a 3-class classification task: does the premise entail, contradicts, or is neutral with respect to the hypothesis? MNLI (Williams et al., 2018) is one of the most popular resources for this task, but it has been shown to suffer from both annotation artifacts (Gururangan et al., 2018; Poliak et al., 2018) and annotator bias (Geva et al., 2019). A cartography (Swayamdipta et al., 2020) map of MNLI (fig. 1) suggests that most of its examples are easy to learn, which explains why vanilla fine-tuning with modern models is sufficient to achieve high accuracy on MNLI benchmark.

We measure generalization of models fine-tuned on MNLI with HANS (McCoy et al., 2019b), a synthetic dataset targeting *lexical overlap*, *subsequence* and *constituent* heuristics. According to




Figure 1: MNLI data map (with RoBERTa-large)

McCoy et al. (2019b), a model trained on MNLI is likely to learn these heuristics and thus predict the “entailment” label for most HANS examples. E.g. it would incorrectly predict that “*The doctor was paid by the actor*” entails “*The doctor paid the actor*”, simply because these sentences contain the same words. See Appendix A for more examples.

Methodology. We experiment with variants¹ of BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), and ALBERT (Lan et al., 2020). Our implementation² is based on Transformers (Wolf et al., 2019) and Pytorch (Paszke et al., 2019), and for two experiments we also report results with a custom Pytorch-Lightning trainer³. HANS has 30K examples used only for testing, where we report the accuracy⁴. MNLI test set is not public, and we report accuracy on the “matched” dev set (20K examples, 393K for training).

¹Most models we used were provided by HuggingFace Transformers library. In scope of this project we ported the smaller BERT versions by Turc et al. (2019) for that library.

²https://github.com/prajjwalin/generalize_lm_nli

³<https://github.com/vecto-ai/langmo>

⁴Since HANS contains only two labels (entailment, non-entailment), and a model trained on MNLI would have three (entailment, contradiction, neutral), a completely random model would be biased towards the “non-entailment”. For direct comparison with MNLI we report the average accuracy for the two HANS labels, unless specified otherwise.

3 Experiments

There are two main directions for solving the generalization challenge: modifying the training distribution and the model itself. For the former we experimented with subsampling (§3.1), and for the latter – with bottlenecking with Siamese Transformers (§3.2) and adapters (§3.3), explicit debiasing (§3.4), and increasing model size (§3.5). This section presents the motivation and setup for all experiments, and all the results are shown in §4.

§3.1 Subsampling the training data with cartography. Data cartography (Swayamdipta et al., 2020) characterizes training data points via the model’s confidence in the true class, and the variability of this confidence across epochs. Figure 1 shows that MNLI examples form a spectrum: some are consistently “easy” (high-confidence) and “hard” (low-confidence) across epochs. “Ambiguous” samples have midrange confidence and high variability. If much of MNLI is “easy”, presumably these samples are less informative.

Experiments. We partition MNLI based on the training dynamics of RoBERTa-large and BERT-base, and train the respective models on varying amounts of “hard” and “ambiguous” examples (preceded by 25% of “easy” samples for 2 epochs). See appendix B for more details.

§3.2 Siamese Networks. In this architecture predictions are based on a pair of inputs (Chopra et al., 2005; Koch et al., 2015). It was successful on NLI (Chen et al., 2017) and in conjunction with BERT encoders (Reimers and Gurevych, 2019). One of their properties is forcing the model to consider the relation between two sequences in a more holistic way than with cross-attention between concatenated premise+hypothesis (as in standard BERT fine-tuning). Intuitively, encoding premise and hypothesis separately could bottleneck⁵ their interaction and encourage learning more abstract patterns, which is what we need here: ideally, an NLI model would learn logical rules rather than numerous lexical or syntactic patterns.

Experiments. Our Siamese Transformer consists of a MLP and two BERT encoders which receive hypotheses and premises in a segregated manner. We used mean-pooled outputs of last transformer

⁵The information bottleneck idea (Tishby et al., 2000; Alemi et al., 2016) has recently been successfully adapted for BERT fine-tuning to avoid overfitting in a low-resource setting by Mahabadi et al. (2020), who propose a regularization term suppressing the learning of irrelevant information.

layer (CLS embedding yielded similar results) combined as $[U, V, U - V, U * V]$ as inputs to MLP classifier. We experiment with base and large BERTs, with both frozen and trainable encoders.

§3.3 Adapter Networks. Intuitively, standard fine-tuning of BERT changes the amount of task-independent linguistic knowledge that the model can store, and may corrupt it (if the supervised task has significant artifacts). Therefore, by adding separate task-specific components rather than changing the language model weights, we could expect to increase the amount of non-task-specific knowledge in the model. This could be done with adapters (Houlsby et al., 2019; Pfeiffer et al., 2020): trainable MLPs inserted within each sub-layer of encoder. Concretely, in a transformer layer l , additional adapter parameters ϕ_l are introduced to learn task specific parameters while keeping pre-trained parameters intact. Having smaller trainable components should also bottleneck the model and encourage it to learn higher-level patterns.

Experiments. We add adapters in each sub-layer as proposed in Houlsby et al. (2019) to BERT and RoBERTa with the configuration defined in Pfeiffer et al. (2020). The adapter consists of two linear layers (up and down) with a bottleneck of reduction factor of 16 and the ReLU non-linearity.

§3.4 Explicit De-biasing. If MNLI ‘teaches’ to rely on superficial features, we could try to avoid them. Following Zhou and Bansal (2020), we use HEX projection (Wang et al., 2019). The system includes the main Transformer encoder and a ‘naive’ model learning superficial features. HEX orthogonally projects the Transformer representation into the affine space the most different from the ‘naive’ representation, hopefully removing the bias.

Experiments. We extract pooled representations from our main model (BERT-base). The ‘naive’ model is a CBOW model with a self-attention layer (Vaswani et al., 2017) to capture co-occurrence information from the sequence with input and token embeddings. See Wang et al. (2019) for more details on the method, and Appendix C for implementation and hyperparameter details. During inference, we use logits from BERT only.

§3.5 Increasing Model Size. Scaling language models to massive amounts of data has been a reliable source of success on NLP leaderboards, and yielded some interesting emergent properties (Brown et al., 2020; Raffel et al., 2020). If pre-

Architecture	Encoder	HF trainer			Custom Trainer		
		MNLI/std	HANS/std	runs	MNLI/std	HANS/std	runs
Siamese networks / frozen encoder	BERT-base	51.43	50.74	1	57.2 / 0.2	51.3 / 0.1	3
	BERT-large	51.72	51.12	1	61.4 / 0.1	51.6 / 0.1	5
Siamese networks / trainable encoder	BERT-base	58.9	52.79	1	76.5 / 0.03	51.3 / 0.03	3
	BERT-large	59.9	51.21	1	78.7	52.5	1
Adapter networks	BERT-base	82.6	50.97	1			
	BERT-large	84.75	57.17	1			
	RoBERTa-base	86.33	57.21	1			
	RoBERTa-large	90.4	75.93	1			
HEX debiasing	BERT-base	56.25	50.58	1			
Vanilla finetuning: increased model size	BERT-tiny (4.4M)	64.48/0.24	50/0	3	67.4 / 0.2	50 / 0.02	5
	BERT-mini (11.3M)	72.3/0.29	50.97/0.04	3	76.3 / 1	52.3 / 0.3	10
	BERT-small (29.1M)	76.48/0.12	50.39/0.14	3	78.4 / 0.5	51.1 / 0.3	5
	BERT-medium (41.7M)	79.64/0.14	51.02/0.26	3	80 / 0.3	52 / 0.4	5
	BERT-base (110M)	83.74/0.04	53.98/0.78	3	84 / 0.2	69 / 4	16
	BERT-large (340M)	85.9/0.02	72.04/1.97	3	86.5 / 0.1	77.8 / 2.4	3
	RoBERTa-base (125M)	87.46/0.1	73.11/1.13	3	87.5 / 0.3	77.7 / 1.7	10
	RoBERTa-large (355M)	90.3/0.07	79.95/0.56	3	90 / 0.4	82.05 / 1	3
	ALBERT-base-v2 (11M)	83.06/0.13	66.6/0.78	3	84.2 / 0.6	69.2 / 2.2	4
	ALBERT-large-v2 (17M)	85.08/0.3	70.64/2.91	3	85.5 / 0.9	70.5 / 1.6	4

Table 1: Generalization from MNLI to HANS in selected approaches.

training “teaches” transferable linguistic knowledge, the models absorbing more data could be expected to generalize better.

Experiments. We perform standard fine-tuning on MNLI with variants of BERT: tiny, mini, small, medium by [Turc et al. \(2019\)](#), base and large by [Devlin et al. \(2019\)](#), as well as RoBERTa ([Liu et al., 2019](#)) and ALBERT ([Lan et al., 2020](#)). In this and the Siamese network experiment we report not only the results obtained with the HuggingFace Trainer, but also with our custom implementation based on Pytorch Lightning (also with the AdamW optimizer and with similar learning rates).

4 Results and Discussion

4.1 Negative Results

Table 1 shows that Siamese networks and HEX debiasing perform at chance level on HANS. Adapters work better, but do not match vanilla fine-tuning of their base models. While it is impossible to prove the negative, our experience suggests that, given a reasonable amount of effort, these approaches are not the most promising for the generalization problem we considered. The paper is accompanied by code for our implementations.

Our Siamese model would be expected to fail if high performance of vanilla BERT was largely due to cross-attention across [premise + hypothesis], enabling it to learn many specific patterns (such as negation in the hypothesis). Our bottleneck MLP would not have the capacity to do that, and it

clearly also fails to find a more abstract pattern in the representations it receives. Further experiments are needed to verify this hypothesis. Whether or not overall we would like our NLI models to be able to operate with independent representations of premise and hypothesis rather than cross-attention within one representation, is an open question.

For HEX, [Zhou and Bansal \(2020\)](#) suggest that the problem might be that it has access only to the final output of BERT, which could contain more information about the predicted NLI labels than the input sequence as such. Then there would be little to debias. Our results support this hypothesis, but more qualitative research is needed to verify it.

The RoBERTa-large MNLI results of our adapter implementation is on par with the recent state-of-the-art Compacter adapters on T5 ([Mahabadi et al., 2021](#)), but generalization in both BERT and RoBERTa is overall worse than with vanilla fine-tuning. Following on the recent report of adapter efficacy in low-resource setting ([He et al., 2021](#)), we conducted an additional experiment with adapters and RoBERTa-large, where the model had to learn from a small, more informative subsample. At 1024 training examples adapters performed better when the MNLI subsample was diverse (selected with K-means-based clustering, see [appendix D](#)) rather than randomly selected: 80.7% vs 85%. But the generalization to HANS was still not very impressive: 67.8% vs 57.5%, respectively. This strategy does seem to select more informative examples for MNLI distribution, but not for HANS.




Figure 2: Generalization to HANS at varying stages of training on “hard” MNLI samples. Graph labels indicate % of MNLI training data used.

4.2 (Cautiously) Positive Results

Figure 2 shows that when trained on the “hard” samples, for RoBERTa-large there does exist a MNLI subsample (at about 25% training data) yielding good performance on both HANS and MNLI. Further 13% addition of biased MNLI data makes the model lose its performance on HANS. But we could not find such a sample for BERT-base, although the cartography samples were model-specific. This also does not happen for either model when training on the “ambiguous” subsamples: RoBERTa initially “learns” HANS at 5% of training data, but “loses” it before reaching even 60% accuracy on MNLI (see fig. 4 in the Appendix).

The most encouraging results come from the increased model size with our custom trainer, as shown in fig. 3. For BERT, RoBERTa and ALBERT, the “large” versions generalize consistently better than the “base” versions. Concurrent work (Anonymous, 2021) focusing specifically on the effect of model size on the learning of lexical overlap heuristic came to a similar conclusion.

However, “larger is better” is not the whole story. The improvement occurs only past a certain threshold: going from BERT-tiny to BERT-medium does not help generalization. At the same time, both ALBERTs have fewer parameters than BERT-small, but they do generalize, which suggests that their parameter sharing is truly effective. Also RoBERTa-base learns to generalize more consistently than BERT-large, which may be either due to some inherent superiority of RoBERTa, or because this instance of RoBERTa happens to be better than this instance of BERT. One point that is clear is that **better generalization also requires longer fine-tuning**, which interestingly barely affects the core MNLI performance on the larger models, but makes a lot of difference for the smaller BERTs.




Figure 3: Generalization from MNLI to HANS: model size effect (custom Pytorch Lightning implementation). The dots on the line for each model indicate performance after training epochs 1-4.

5 Related work

Several studies have reported successful generalization from MNLI to HANS. Among data-based strategies, it has been achieved via augmenting MNLI data with predicate-argument structures (Moosavi et al., 2020) and syntactic transformations (Min et al., 2020). Although there are many reports of syntactic knowledge in pre-trained BERT (Rogers et al., 2020b), Min et al. (2020) suggest that pre-training does not yield a strong inductive bias to *use* syntax in downstream tasks, and augmentation “nudges” the model towards that.

Theoretically, subsampling that reduces the saliency of spurious patterns should have a similar effect, but our cartography-based subsampling did not work consistently, possibly because MNLI has little counter-evidence to spurious patterns, and the right subsample is hard to find reliably. We have additional negative results for subsampling with K-means clustering (see Appendix D for details).

The idea of using shallow models to identify biases *before* training and “teach” the model to treat them differently has been successfully explored by Utama et al. (2020), Clark et al. (2020), and Sanh et al. (2021). Our negative results with HEX debiasing *after* training complements these reports.

Our results corroborate that generalization is improved by larger models (Anonymous, 2021) and longer fine-tuning (Tu et al., 2020). The latter work shows that this happens thanks to the few HANS-like samples in MNLI: they take longer to learn, and without them longer fine-tuning does not help.

A general challenge for DL-based NLP research is variability due to extraneous factors. Generalization from MNLI to HANS may be much improved simply with a lucky fine-tuning initialization (Mc-

Coy et al., 2019a). For QA Crane (2018) show that there are many other external factors (down to linear algebra library version) that also play a role, and for Transformers overall implementations make a big difference (Narang et al., 2021). Our work provides an illustration of that phenomenon in NLI. The reported HANS performance of vanilla fine-tuned BERT-base varies in the published studies from 50.0% to 62.5%. Our Pytorch-Lightning implementation at 4 epochs achieves 69% (avg. of 16 runs), not due to any architectural differences. Overall it also has higher variability between runs, possibly due to batch size differences.

6 Analysis: Bias Under Low Confidence

Our overall ratio of positive to negative results illustrates the difficulty of the spurious patterns problem. Once the model learns that some pattern is a strong signal for a label, it will over-rely on it. But how much heuristic-matching evidence does it need?

In this experiment we fine-tune the base versions of BERT and RoBERTa for 4 epochs. We use the dataset cartography to identify the “hard” training samples for both models. As shown in Figure 1, the classifier has overall low confidence for the “hard” samples. We corrupt these “hard” samples by inserting extra characters randomly in 30% content words in the sequence. For example:

$$\begin{array}{ccccccc} & & \text{ythink} & \text{about} & & & \\ & & \underbrace{\hspace{1.5cm}} & \underbrace{\hspace{1.5cm}} & & & \\ \text{Premise:} & \text{do it now,} & \text{think} & \text{bout} & \text{it later} & & \\ & \underbrace{\hspace{1.5cm}} & & \underbrace{\hspace{1.5cm}} & & & \\ \text{Hypothesis:} & \text{think} & \text{about it now,} & \text{do it} & \text{later} & & \end{array}$$

The corrupted sequences remain relatively readable for humans, but this reduces the signal from direct lexical matches seen by the model (even with BERT tokenization). Note that the model has already seen these samples 4 times before corruption. We repeated this experiment with substituting, deleting and swapping characters.

Since the classifier confidence for the “hard” examples is low, and the perturbations are random, they could be expected to just flip the predictions in random directions, equally across MNLi labels. Instead, with all corruption strategies and for both models we see the pattern shown in Table 2: the accuracy drops significantly for contradiction (by 10-20 points), and improves significantly for entailment (by 10-30 points). For the neutral class the change is not as large (mostly gaining 5-13 points).

These results suggest that **in low-confidence context even decreased lexical overlap still**

Corruption	Labels	BERT	RoBERTa
Character insert	Entailment	+18.2	+11.9
	Neutral	+13.78	+0.8
	Contradiction	-28.89	-8.4
Character substitute	Entailment	+35.5	+20.4
	Neutral	+1.6	+5.9
	Contradiction	-23.9	-17.6
Character swap	Entailment	+23.8	+18.1
	Neutral	-1.6	+3.3
	Contradiction	-15.5	-13.9
Character delete	Entailment	+31.73	+18.4
	Neutral	-11.2	+5.8
	Contradiction	-10.3	-16.3

Table 2: “Hard” samples: changes in prediction accuracy for MNLi classes by BERT-base and RoBERTa-base after random character corruption.

nudges the model towards entailment rather than contradiction. This runs contrary to the overall strategy shown by HANS, and it is not due to the majority class bias (as MNLi train set is balanced between entailment and contradiction). One possible explanation is that if it is non-entailment that the generalizing models slowly learn from the little supporting evidence in MNLi (Tu et al., 2020), then corruption would make that already-difficult job even harder for the model, decreasing the accuracy on non-entailment. On the other hand, even corrupted MNLi examples still have some lexical overlap, and so the model, unable to recognize any subtler patterns, might default to that.

This finding has implications for high-cost-of-error applications where false positives are preferable to false negatives. If the data has spurious patterns, the model may score well on a generalization benchmark, but be still biased towards a certain label when its confidence is low. Consider e.g. most of COVID detection models are “at high risk of bias” due to noisy data (Wynants et al., 2020).

7 Conclusion

Most supervised datasets are biased in one way or the other, and task-independent techniques to improve NLP model generalization are crucial for further advances. We experimented with 5 strategies to improve generalization of BERT-class models for NLI task: explicit debiasing, bottlenecking the model, adapters, data subsampling, and increasing model size. We find the latter strategy the most promising, but we also report all the negative results, which contribute to the overall knowledge about generalization in BERT-based models.

Acknowledgments

This work was partially supported by JST CREST Grant Number JPMJCR19F5, Japan. We thank the anonymous reviewers for their insightful comments, T. McCoy and T. Linzen for the help with HANS data, and E. Vatai for the help with langmo implementation.

References

- Alexander A. Alemi, Ian Fischer, Joshua V. Dillon, and Kevin Murphy. 2016. [Deep Variational Information Bottleneck](#).
- Anonymous. 2021. Largely, right for better reasons:lexical generalization improves with model size. (*Under review*).
- Yonatan Belinkov, Adam Poliak, Stuart Shieber, Benjamin Van Durme, and Alexander Rush. 2019. [Don't take the premise for granted: Mitigating artifacts in natural language inference](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 877–891, Florence, Italy. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. 33:1877–1901.
- Qian Chen, Xiaodan Zhu, Zhen-Hua Ling, Si Wei, Hui Jiang, and Diana Inkpen. 2017. [Enhanced LSTM for natural language inference](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1657–1668, Vancouver, Canada. Association for Computational Linguistics.
- Sumit Chopra, Raia Hadsell, and Yann LeCun. 2005. [Learning a similarity metric discriminatively, with application to face verification](#). In *Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05) - Volume 1 - Volume 01*, CVPR '05, page 539–546, USA. IEEE Computer Society.
- Christopher Clark, Mark Yatskar, and Luke Zettlemoyer. 2020. [Learning to Model and Ignore Dataset Bias with Mixed Capacity Ensembles](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3031–3045, Online. Association for Computational Linguistics.
- Matt Crane. 2018. [Questionable Answers in Question Answering Research: Reproducibility and Variability of Published Results](#). *Transactions of the Association for Computational Linguistics*, 6:241–252.
- 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.
- Matt Gardner, William Merrill, Jesse Dodge, Matthew E. Peters, Alexis Ross, Sameer Singh, and Noah Smith. 2021. [Competency Problems: On Finding and Removing Artifacts in Language Data](#). *arXiv:2104.08646 [cs]*.
- Mor Geva, Yoav Goldberg, and Jonathan Berant. 2019. [Are We Modeling the Task or the Annotator? An Investigation of Annotator Bias in Natural Language Understanding Datasets](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1161–1166, Hong Kong, China. Association for Computational Linguistics.
- 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.
- Ruidan He, Linlin Liu, Hai Ye, Qingyu Tan, Bosheng Ding, Liying Cheng, Jiawei Low, Lidong Bing, and Luo Si. 2021. [On the Effectiveness of Adapter-based Tuning for Pretrained Language Model Adaptation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2208–2222, Online. Association for Computational Linguistics.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. [Parameter-efficient transfer learning for NLP](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2790–2799, Long Beach, California, USA. PMLR.
- Angelos Katharopoulos and François Fleuret. 2018. [Not all samples are created equal: Deep learning with importance sampling](#). In *Proceedings of the 35th International Conference on Machine Learning*,

- ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 2530–2539. PMLR.
- Gregory Koch, Richard Zemel, and Ruslan Salakhutdinov. 2015. Siamese neural networks for one-shot image recognition. In *ICML deep learning workshop*, volume 2. Lille.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. [Albert: A lite bert for self-supervised learning of language representations](#). In *International Conference on Learning Representations*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized BERT pretraining approach](#). *CoRR*, abs/1907.11692.
- Rabeeh Karimi Mahabadi, Yonatan Belinkov, and James Henderson. 2020. [Variational Information Bottleneck for Effective Low-Resource Fine-Tuning](#). In *International Conference on Learning Representations*.
- Rabeeh Karimi Mahabadi, James Henderson, and Sebastian Ruder. 2021. [Compacter: Efficient low-rank hypercomplex adapter layers](#).
- R. Thomas McCoy, Junghyun Min, and Tal Linzen. 2019a. [BERTs of a feather do not generalize together: Large variability in generalization across models with similar test set performance](#). *arXiv:1911.02969 [cs]*.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019b. [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.
- Junghyun Min, R. Thomas McCoy, Dipanjan Das, Emily Pitler, and Tal Linzen. 2020. [Syntactic data augmentation increases robustness to inference heuristics](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2339–2352, Online. Association for Computational Linguistics.
- Nafise Sadat Moosavi, Marcel de Boer, Prasetya Ajie Utama, and Iryna Gurevych. 2020. [Improving Robustness by Augmenting Training Sentences with Predicate-Argument Structures](#). *arXiv:2010.12510 [cs]*.
- Sharan Narang, Hyung Won Chung, Yi Tay, William Fedus, Thibault Fevry, Michael Matena, Karishma Malkan, Noah Fiedel, Noam Shazeer, Zhenzhong Lan, Yanqi Zhou, Wei Li, Nan Ding, Jake Marcus, Adam Roberts, and Colin Raffel. 2021. [Do Transformer Modifications Transfer Across Implementations and Applications?](#) *arXiv:2102.11972 [cs]*.
- 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. 2019. [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.
- Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. 2020. [Adapterfusion: Non-destructive task composition for transfer learning](#).
- 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.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Nils Reimers and Iryna Gurevych. 2019. [SentenceBERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Anna Rogers, Olga Kovaleva, Matthew Downey, and Anna Rumshisky. 2020a. [Getting Closer to AI Complete Question Answering: A Set of Prerequisite Real Tasks](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 8722–8731.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020b. [A Primer in BERTology: What We Know About How BERT Works](#). *Transactions of the Association for Computational Linguistics*, 8:842–866.
- Victor Sanh, Thomas Wolf, Yonatan Belinkov, and Alexander M. Rush. 2021. [Learning from others’ mistakes: Avoiding dataset biases without modeling them](#). In *International Conference on Learning Representations*.
- Swabha Swayamdipta, Roy Schwartz, Nicholas Lourie, Yizhong Wang, Hannaneh Hajishirzi, Noah A. Smith, and Yejin Choi. 2020. [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. Association for Computational Linguistics.
- Naftali Tishby, Fernando C. Pereira, and William Bialek. 2000. [The information bottleneck method](#). *arXiv:physics/0004057*.
- Lifu Tu, Garima Lalwani, Spandana Gella, and He He. 2020. [An Empirical Study on Robustness to Spurious Correlations using Pre-trained Language Models](#). *Transactions of the Association for Computational Linguistics*, 8:621–633.
- Iulia Turc, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Well-read students learn better: On the importance of pre-training compact models](#).
- Prasetya Ajie Utama, Nafise Sadat Moosavi, and Iryna Gurevych. 2020. [Towards debiasing NLU models from unknown biases](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7597–7610, Online. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Haohan Wang, Zexue He, Zachary L. Lipton, and Eric P. Xing. 2019. [Learning robust representations by projecting superficial statistics out](#). In *International Conference on Learning Representations*.
- 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.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R’emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. [Huggingface’s transformers: State-of-the-art natural language processing](#). *ArXiv*, abs/1910.03771.
- Laure Wynants, Ben Van Calster, Gary S. Collins, Richard D. Riley, Georg Heinze, Ewoud Schuit, Marc M. J. Bonten, Darren L. Dahly, Johanna A. Damen, Thomas P. A. Debray, Valentijn M. T. de Jong, Maarten De Vos, Paula Dhiman, Maria C. Haller, Michael O. Harhay, Liesbet Henckaerts, Pauline Heus, Michael Kammer, Nina Kreuzberger, Anna Lohmann, Kim Luijken, Jie Ma, Glen P. Martin, David J. McLernon, Constanza L. Andaur Navarro, Johannes B. Reitsma, Jamie C. Sergeant, Chunhu Shi, Nicole Skoetz, Luc J. M. Smits, Kym I. E. Snell, Matthew Sperrin, René Spijker, Ewout W. Steyerberg, Toshihiko Takada, Ioanna Tzoulaki, Sander M. J. van Kuijk, Bas C. T. van Bussel, Iwan C. C. van der Horst, Florian S. van Royen, Jan Y. Verbakel, Christine Wallisch, Jack Wilkinson, Robert Wolff, Lotty Hooft, Karel G. M. Moons, and Maarten van Smeden. 2020. [Prediction models for diagnosis and prognosis of covid-19: Systematic review and critical appraisal](#). 369:m1328.
- Xiang Zhou and Mohit Bansal. 2020. [Towards robustifying NLI models against lexical dataset biases](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8759–8771, Online. Association for Computational Linguistics.

A Additional details and samples from the used datasets

We use MNLI to perform training of LMs and evaluate their generalization capabilities on HANS. See Table 3 for some sample sentences from MNLI and HANS. MNLI has three classes (“entailment”, “contradiction” and “neutral”), while HANS only has “entailment” and “non-entailment”. HANS targets the three heuristics (“lexical overlap”, “subsequence” and “constituent”) which are usually adopted by pre-trained LMs such as BERT.

MNLI contains 393K and 20K examples in the train and dev sets respectively (the test set is not publicly available). HANS contains 30K examples split across 10K across each heuristic, which were used entirely for testing.

Label	Premise	Hypothesis
Entailment	A member of my team will - execute your orders with immense precision. This information belongs to them.	One of our number will carry out - your instructions minutely. How do you know? - All this is their information again.
Neutral	Product and geography are - what make cream skimming work The speaker doesn’t know who it is.	Conceptually cream skimming has two - basic dimensions - product and geography. Who could there be ?
Contradiction	I ignored Ben. He only muttered something about - splitting the sky.	Hello, Ben. He distinctly said - you were to repair the sky.

(a) Examples of sentences from MNLI (Williams et al., 2018)

Heuristic	Premise	Hypothesis	Label
Lexical overlap	The banker near the judge saw the actor.	The banker saw the actor.	E
	The judge by the actor stopped the banker.	The banker stopped the actor.	N
Subsequence	The artist and the student called the judge.	The student called the judge.	E
	The senator near the lawyer danced.	The lawyer danced.	N
Constituent	Before the actor slept, the senator ran.	The actor slept.	E
	If the actor slept, the judge saw the artist.	The actor slept.	N

(b) Examples of sentences from HANS (McCoy et al., 2019b).

The *label* column shows the correct label for the sentence pair; *E* stands for *entailment* and *N* stands for *non-entailment*. A model relying on the heuristics would label all examples as *entailment* (incorrectly for those marked as N)

Table 3: The NLI data used in this study

B Cartography

Cartography is a recent method for characterizing the difficulty of training samples. It describes data points as “easy”, “ambiguous” and “hard” through analyzing predictions of the model during training on those samples (training dynamics).

Sampling. In our experiments, we sample the datasets in the ranked order. For example, the easiest example is at the top, so it is first in the sampled batch of dataset. Our implementation of obtaining ranked ordering is based on the original implementation by the authors of the method⁶.

Role of “easy” examples. During vanilla fine-tuning on randomly subsampled data, the model encounters all three kinds of examples some of which may provide meaningful signals to learn and some may aid in out-of-domain generalization. In our experiments, we find that training solely on “ambiguous” and “hard” examples do not aid the network in improving the performance. This finding is consistent with the observation from (Swayamdipta et al., 2020) wherein they showed that “easy” examples aid in optimization of the network during initial stages and are crucial for training. Therefore, in our experiments the models were trained first on 25% “easy” examples, and then on subsets of “hard” examples containing the top 1%, 5%, 10%, 17%, 25%, 33%, 50% and 75% of the “hard” data (for consistency with the experiment by Swayamdipta et al. (2020), see sec. 4). The results of this experiment are shown in fig. 2. The same experiment was repeated with the “ambiguous” samples, shown in fig. 4.




Figure 4: Evaluating generalization at varying stages of training on “ambiguous” samples. Percentages on marker represents percentage of MNL train data used as training progresses.

⁶<https://github.com/allenai/cartography>

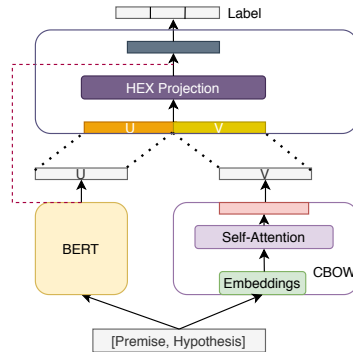


Figure 5: Orthogonal debiasing with HEX projection

C Stabilizing HEX

Here we provide more details about how HEX is being used. The self-attention output from BERT is pooled and passed through two MLP layers to get an individual representation of each input sequence, as shown in fig. 5. We feed the pooled representation of BERT and the intermediate representation of CBOW into two MLPs to obtain vectors U and V . We use f to represent classification layer parameterized by ξ . The output vectors $F_A = f([U, V]; \xi)$ and $F_G = f([0, V]; \xi)$ are concatenated along the non-batch dimension.

$$\begin{aligned} F_A &= f([U, V]; \xi), \\ F_G &= f([0, V]; \xi) \end{aligned} \quad (1)$$

where F_A, F_G denotes both concatenated representations and zero matrix prepended with network B’s representation $[,]$ denotes concatenation operation along the non-batch dimension.

Following Vaswani et al. (2017), we project F_A :

$$F_L = (I - F_G(F_G^T F_G + \lambda I)^{-1} F_G^T) F_A \quad (2)$$

Table 4 shows hyperparameter search for λ . During inference, we use logits obtained through BERT only.

We follow a slight variation of Equation 2 to smoothen the process of optimization. The addition of λ hyper-parameter has been done to ensure that inverse is carried out on a non-singular matrix. The value of λ plays a significant role in determining how these representations are being learned. In our experiments, $1e - 4$ worked well and was used for initializing it after which it was set as a model’s parameter. We observe that pseudo-inverse of $F_G^T F_G$ is unstable and can make optimization

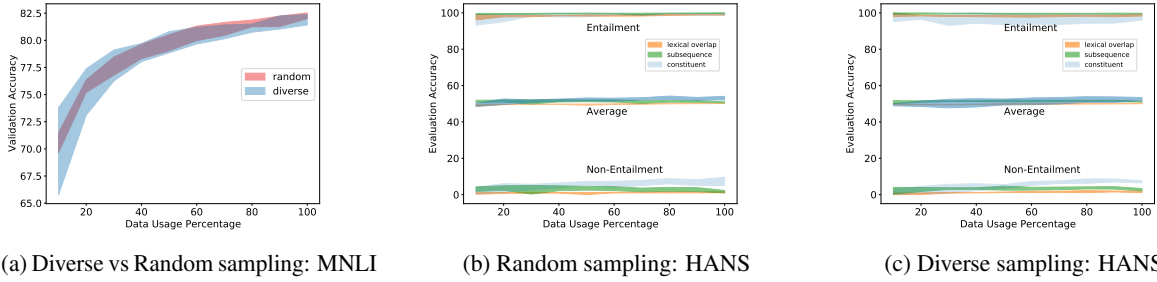


Figure 6: Fine-tuning BERT-base on varied amounts of MNLI data: in-domain and generalization performance

λ	MNLI	HANS (Average)		
		L	S	C
1e-4	55.2	49.50 / 52.62 / 52.52		
2e-4	56.54	49.81 / 50.90 / 50.83		
3e-4	57.02	50.03 / 50.11 / 49.93		
4e-4	57.72	49.49 / 52.39 / 51.42		
5e-4	57.09	49.93 / 50.16 / 50.03		
6e-4	55.26	49.66 / 51.59 / 52.68		
7e-4	57.25	49.60 / 51.09 / 49.63		
8e-4	57.78	49.60 / 51.20 / 51.37		
9e-4	48.58	50.00 / 50.00 / 49.98		
1e-5	53.45	49.80 / 50.69 / 50.68		
2e-5	53.77	50.00 / 50.00 / 50.00		
3e-5	56.46	49.19 / 52.84 / 54.41		
4e-5	47.30	49.61 / 50.90 / 49.28		
5e-5	50.80	49.69 / 51.07 / 50.41		

Table 4: Performance on MNLI and HANS with HEX (BERT-base) with different values of λ . L, S, C denote lexical overlap, subsequence and constituent heuristic

process hard, so we make U and V square matrices to obtain inverse instead. Additionally, during inference time, we directly feed outputs from the main network to the MLP layer to obtain logit vectors instead of using F_L . It has been reported (Wang et al., 2019) that this doesn’t have any profound impact on the logit vector and makes inference faster. We also applied L1 and L2 normalization on U and V to account for differences in scale but did not see any noticeable improvement. We found that values of λ greater than 0.0001 do not aid the network in learning.

D Subsampling the Training Data with K-means clustering

Motivation. Fundamentally, the problem is the mismatch between MNLI and HANS distributions. For a biased dataset, one solution could be to find such a subsample that would enable the model to perform well on both distributions.

Experiments. We encode MNLI examples as BERT [CLS] embeddings and cluster them in 512

clusters using K-means. We then fine-tune BERT-base on varying amounts of MNLI data, progressively increasing the amount of training examples by 10%. The data in the sub-sample is selected (a) randomly (as a control), (b) so as to maximize the diversity of examples within the sample (Katharopoulos and Fleuret, 2018). At the smallest subsample size we sample the data from all clusters. As data size increases, smaller clusters are exhausted while the larger ones remain, so the smallest subsamples are the most diverse, and the diversity decreases as the sample size increases. The experiment is repeated with 5 random seeds.

Results. Figure 6 shows that on MNLI, diverse sampling yields much more variation with small amounts of data than random sampling, but as the subsample approaches the full dataset the performance also becomes the same. Neither subsampling strategy improves generalization: the model still predicts “entailment” for most HANS examples. Thus overall the result for this strategy is negative.