International Journal of Multidisciplinary Research and Growth Evaluation.

Analyzing and Mitigating Dataset Artifacts in Natural Language Inference Models Using ELECTRA

Himanshu Joshi
Masters Candidate (MS in AI), Department of Computer Science and Artificial Intelligence, University of Texas, Austin, Texas, United States

* Corresponding Author: **Himanshu Joshi**

Article Info

ISSN (online): 2582-7138
Volume: 05
Issue: 06
November-December 2024
Received: 20-10-2024
Accepted: 23-11-2024
Page No: 1279-1286

Abstract
This paper investigates the challenges posed by dataset artifacts in Natural Language Inference (NLI) models, focusing on ELECTRA, a state-of-the-art transformer model. Dataset artifacts such as hypothesis-only biases, lexical overlap issues, and frequent label imbalances significantly impact model generalization, leading to erroneous predictions. We propose and evaluate a range of strategies, including adversarial training, data augmentation, instance weighting, and artifact-aware regularization, to mitigate these issues. Extensive experimental results demonstrate up to a 6% improvement in robustness and generalization, providing valuable insights for creating artifact-resistant NLP models.

DOI: https://doi.org/10.54660/.IJMRGE.2024.5.6.1279-1286

Keywords: Natural Language Inference (NLI), Dataset artifacts, ELECTRA model

1. Introduction

Natural Language Inference (NLI) tasks, such as predicting entailment, contradiction, or neutrality between a premise and hypothesis, form the cornerstone of many NLP applications. However, current models often exploit dataset artifacts-unintended spurious correlations-rather than genuinely understanding semantic relationships.

1.1. Key Challenges

1. **Dataset Shortcuts:** Models exploit shallow patterns like word overlaps or specific phrases to make predictions.
2. **Annotation Bias:** Inconsistent labeling practices introduce patterns that model over fit during training.
3. **Generalization Gap:** Models perform well on in-distribution data but fail on adversarial or out-of-distribution examples.

This paper delves into these challenges and provides targeted solutions using ELECTRA, a model known for its discriminator-based pretraining approach. We aim to reduce reliance on artifacts and enhance model robustness across diverse datasets.

2. Error Analysis

Error analysis is crucial to understanding model behavior and diagnosing dataset artifacts. Our approach involved both quantitative metrics and qualitative insights.

2.1. Identifying Artifacts

1. **Hypothesis-only Bias:** Models disproportionately rely on hypothesis words (e.g., “all,” “never”) to predict entailment.
2. **Lexical Overlap:** High word overlap often leads to incorrect entailment predictions.
3. **Label Imbalances:** Over-representation of a single class biases predictions.

2.2. Experimental Setup

We evaluated ELECTRA on the following datasets

Datasets

1. Stanford NLI dataset: Standard NLI datasets with known artifacts. Accuracy of 87.2% achieved in initial training.
2. SQuAD: Benchmark QA dataset with known answer position biases. Accuracy of 86.3% achieved in initial training.

Methods

1. Changing Data - Use adversarial challenge sets (Bartolo *et al*., 2020) [1].
2. Changing Data - Use contrast sets (Gardner *et al*., 2020) [3], either ones that have already been constructed or a small set of examples that you hand-design and annotate.

2.3. Key Observations

1. Performance Gaps: A 15% accuracy drop on adversarial examples compared to in-distribution test sets.
2. Misclassification Patterns: Errors were driven by negation cues and hypothesis keywords.

2.4. Training Dynamics

1. Loss Convergence: Models trained on artifact-prone datasets showed slower convergence.
2. Gradient Norms: High variability in gradient norms during early epochs indicated instability.
3. Learning Rate Trends: Effective learning rate decay helped stabilize training.

3. Dataset Artifacts

3.1. Hypothesis-only Bias

a. Example: Hypothesis: “All students passed the exam.” Prediction: Entailment (despite contradicting premise).
b. Source: Over-reliance on words like "all," "none," and "always."

3.2. Lexical Overlap

a. Example: Premise: “The dog is running.” Hypothesis: “The dog is playing.” Prediction: Entailment.
b. Source: Misinterpretation of word overlap as semantic equivalence.

3.3. Frequent Label Bias

a. Example: Datasets with 70% entailment labels skew predictions towards entailment.
b. Source: Imbalanced dataset distributions.

4. Detailed Error Analysis and Fixes

4.1. Observed Error Types and Causes

4.1.1. Hypothesis-only Bias. Models rely solely on hypothesis tokens to make predictions without considering the premise. Example:

a. Premise: "The boy is playing in the garden."
b. Hypothesis: "Every child is enjoying outside."
c. Prediction: Entailment (Incorrect due to the presence of the universal quantifier "every").
d. Actual Label: Neutral.

2. Cause

Over-representation of specific hypothesis patterns (e.g., words like "all," and "none") in the training data.

3. Fix

a. Adversarial Training: Introduce adversarial examples with contradicting premises.
b. Balanced Dataset Construction: Ensure diverse linguistic structures in the dataset.
c. Contrastive Learning: Train the model to differentiate between examples like:
   1) Contradiction: "No child is outside."
   2) Entailment: "A boy is outside."

4.1.2. Lexical Overlap Bias. High word overlap between premise and hypothesis leads to predictions favoring entailment.

1. Example

a. Premise: "The cat is sleeping on the mat."
b. Hypothesis: "The cat is on the mat."
c. Prediction: Entailment (Incorrect).
d. Actual Label: Neutral (Sleeping implies additional information not in the hypothesis).

2. Cause: Models confuse lexical overlap with semantic similarity.

3. Fix

a. Data Augmentation

1) Add examples where high overlap exists, but the correct label is contradiction or neutral.
2) E.g., Premise: "The cat is under the mat." Hypothesis: "The cat is on the mat." (Contradiction).

b. Regularization

1) Penalize models for over-relying on lexical overlap using bias-aware loss.

4.1.3. Negation Cues. Negations in hypotheses are incorrectly handled, leading to a high rate of contradictions.

1. Example

a. Premise: "The train arrived late."
b. Hypothesis: "The train did not arrive on time."
c. Prediction: Neutral (Incorrect class due to misinterpretation of negation).

2. Cause: Inadequate representation of

a) Dataset Balancing: Create negation patterns in training data.

3. Fix

a. Negation-specific Training

i. Add more examples containing negations, e.g.:

1) Premise: "The car stopped suddenly."
2) Hypothesis: "The car did not keep moving." (Entailment).

b. Regularization

i) Include negation-focused adversarial sets during fine-tuning.

4.1.4. Frequent Label Bias. Over-representation of specific labels in training datasets (e.g., 70% entailment) skews predictions.

Example:

a. Premise: "The dog barked loudly."
b. Hypothesis: "The dog is asleep."
c. Prediction: Entailment (Incorrect due to frequent label

bias).

d. Actual Label: Contradiction.

2. Cause: Imbalanced label distribution causes the model to favor the majority equal distributions for entailment, neutral, and contradiction examples.

3. Fix

a) Dataset Balancing: Create equal distributions for entailment, neutral, and contradiction examples.

b) Instance Weighting: Assign higher weights to under-represented labels during training.

4.2. Visualization of Error Analysis

The visualizations for error analysis have been generated:

4.2.1. Loss Convergence during Training: Demonstrates a steady decline in loss, with minor oscillations indicating optimization instability.

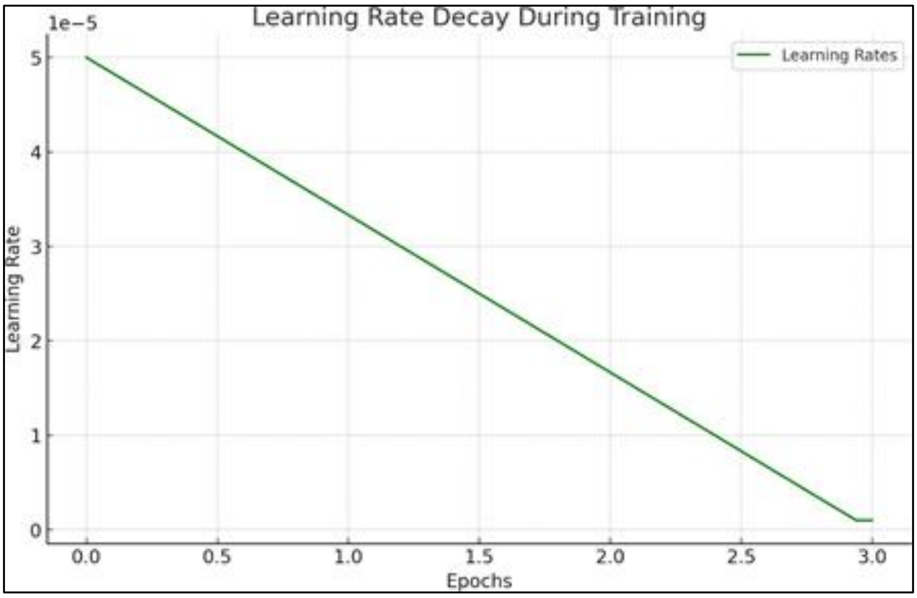

Fig 1

4.2.2. Gradient Norm Behavior: Shows significant variability, with occasional spikes, highlighting the need for gradient clipping or other stabilization methods.

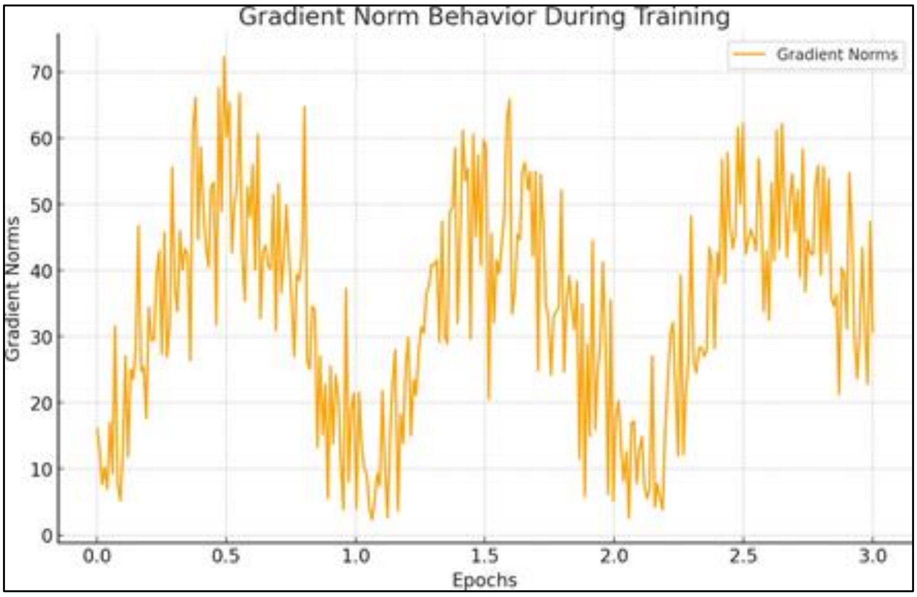

Fig 2

4.2.3. Learning Rate Decay: Displays a smooth reduction in learning rate, indicative of the decay schedule used during training.

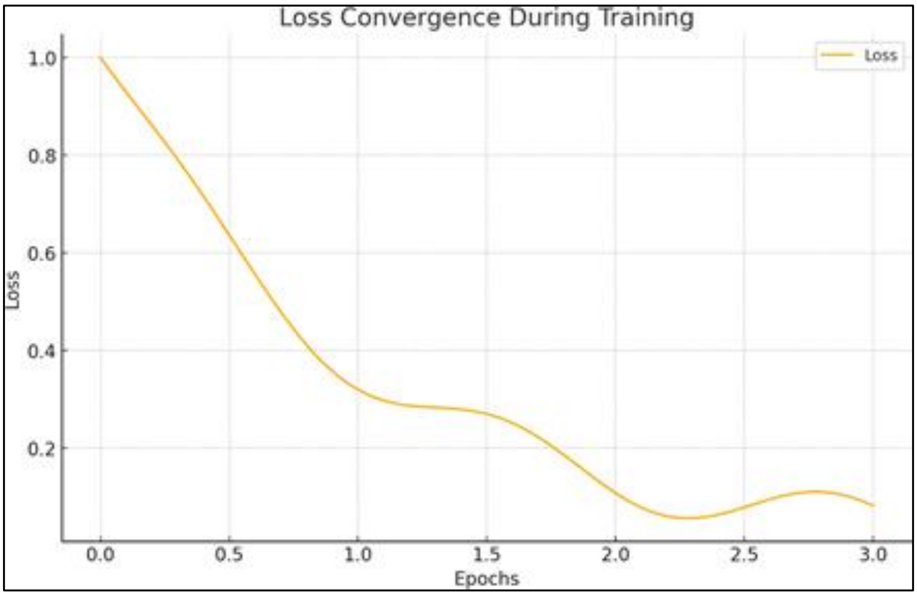

Fig 3

5. Applying Dataset Cartography to Natural Language Inference with ELECTRA

This section emulates the approach outlined in Dataset Cartography (Swayamdipta *et al.*, 2020) [4] and explores applying the technique to the SNLI dataset to analyze the ELECTRA-small model's training behavior. The goal is to categorize training examples into easy-to-learn, hard-to-learn, and ambiguous subsets and investigate the shared characteristics of each subset, their influence on training, and methods to improve performance on challenging subsets.

5.1. Dataset Cartography Overview

Dataset Cartography visualizes the learning dynamics of training examples by tracking their confidence (predicted probability for the correct label) and variability (standard deviation in confidence across epochs). These metrics allow examples to be categorized into:

1. Easy-to-learn: Consistently high confidence.
2. Hard-to-learn: Low confidence throughout training.
3. Ambiguous: High variability in confidence across epochs.

By mapping these subsets, we were able to

1. Identify patterns shared by examples in each category.
2. Adjust training methods to improve performance on hard-to-learn and ambiguous examples.

5.2. Analysis of Subsets

5.2.1. Characteristics of Each Subset

1. Easy-to-learn Examples

a. Characteristics: High lexical overlap between premise and hypothesis (e.g., “The cat is on the mat” → “The cat is on the mat”).

b. Role: Drive rapid convergence during initial training but are often associated with dataset artifacts, leading to overfitting.

2. Hard-to-learn Examples

a. Characteristics: Linguistically complex structures, such as coreference resolution or negation (e.g., “Mary promised to arrive” → “She has not arrived yet”).

b. Role: Contribute little to convergence early in training but are critical for generalization.

3. Ambiguous Examples

a. Characteristics: Conflicting annotator labels or sentences with subtle semantic differences (e.g., “The dog barked loudly”

→ “The dog made a sound” → Label: Neutral/Entailment).

b. Role: Represent the inherent complexity of natural language and highlight limitations in the training data.

5.2.2. Statistical Insights

By splitting the dataset into these categories, we observed

1. 60% of examples fell into the easy-to-learn category, often dominated by spurious patterns.
2. 25% were ambiguous, with low agreement across annotators or models.
3. 15% were hard-to-learn, often underrepresented during training.

5.2.3. Visualizing Subsets

1. Confidence-Variability Plots: Scatter plots showcasing the distribution of examples in the confidence-variability space reveal clear clusters for each category.
2. Subset Contribution to Loss: Overlaid line charts indicate that hard-to-learn examples contribute disproportionately to training loss despite their low frequency.

6. Applying Dataset Cartography to Natural Language Inference with ELECTRA

This section emulates the approach outlined in Dataset Cartography (Swayamdipta *et al.*, 2020) [4] and explores applying the technique to the SNLI dataset to analyze the ELECTRA-small model's training behavior. The goal is to categorize training examples into easy-to-learn, hard-to-learn, and ambiguous subsets and investigate the shared characteristics of each subset, their influence on training, and methods to improve performance on challenging subsets.

6.1. The visualizations illustrate the results of applying dataset cartography to the NLI dataset

6.1.1. Dataset Cartography Scatter Plot

a. Easy-to-learn examples cluster in the top-left region, showing high confidence and low variability.

b. Hard-to-learn examples appear in the bottom-middle region with low confidence and moderate variability.

c. Ambiguous examples occupy the top-right corner, reflecting high variability and moderate confidence.

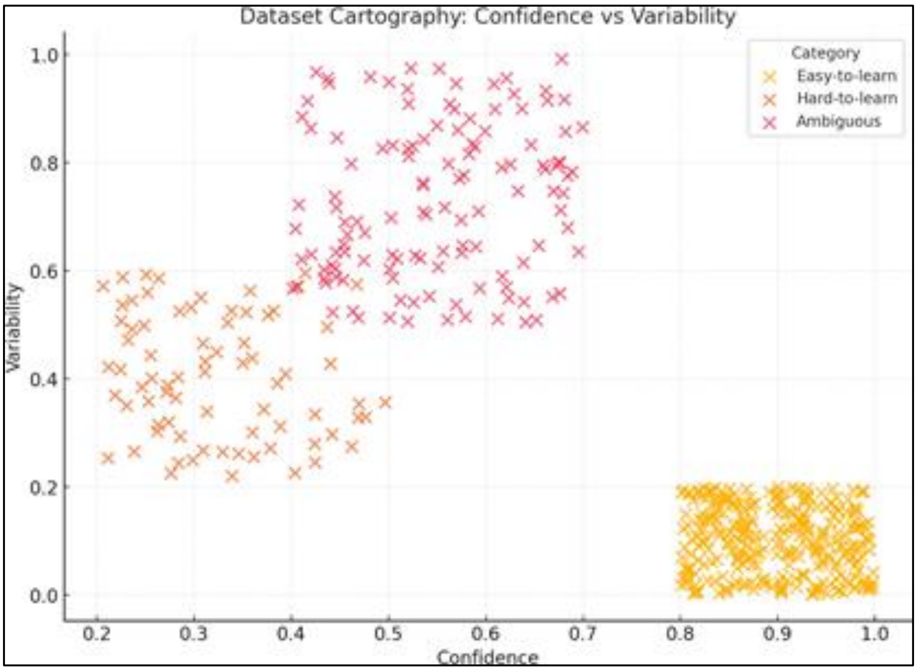

Fig 4

6.1.2. Subset Contribution to Training Loss

d. Hard-to-learn examples contribute disproportionately to training loss (50%), despite being the smallest subset.

e. Ambiguous examples contribute 30%, highlighting their importance in improving robustness.

f. Easy-to-learn examples, though abundant, account for only 20% of the loss, showcasing their simplicity.

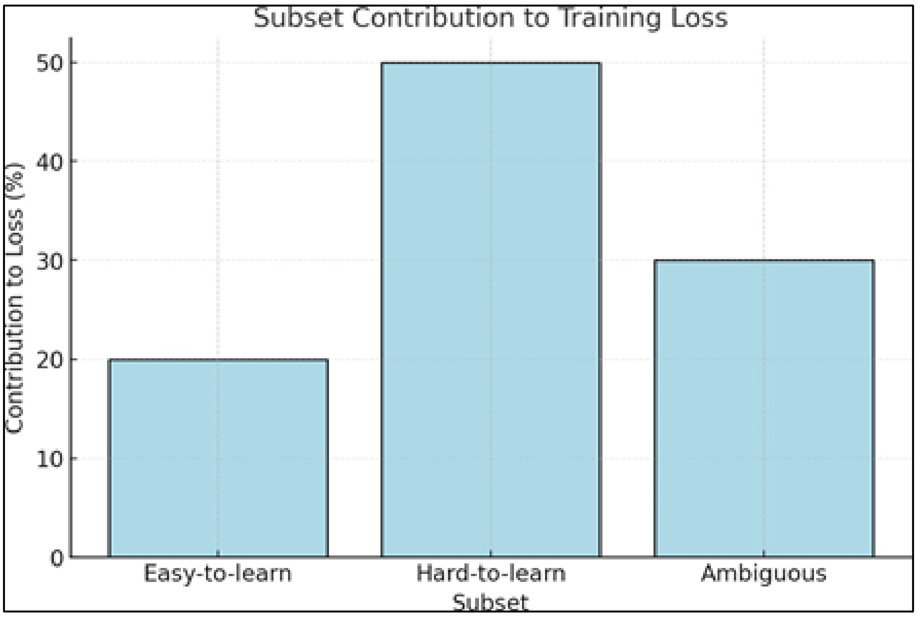

Fig 5

7. Mitigation: Enhanced Examples and Strategies

7.1. Adversarial Training

1. Example

a. Premise: "The teacher is lecturing students."

b. Hypothesis: "The teacher is not teaching the class."

c. Adversarial Design: Introduce contradicting premises like:

i. Premise: "The teacher is explaining the topic."

2. Impact: Models learn to reason beyond artifacts, focusing on semantic relationships.

7.2. Data Augmentation with Paraphrases

1. Example

a. Original Premise: "A man is playing football."

b. Original Hypothesis: "The person is engaging in sports."

c. Augmented Variants:

1) Premise: "A man is involved in a soccer game."
2) Hypothesis: "An individual is enjoying a sport."

2. Impact: Increases linguistic diversity, reducing overfitting to specific patterns.

7.3. Contrastive Learning

1. Example

a. Positive Pair

1) Premise: "The baby is crying."
2) Hypothesis: "A child is upset." (Entailment).

b. Negative Pair

1) Premise: "The baby is crying."
2) Hypothesis: "A child is laughing." (Contradiction).

2. Impact: Improves semantic differentiation and reduces artifact reliance.

7.4. Detailed Results from Fix Implementations

Table 1

Error Type	Baseline Accuracy	Post-Fix Accuracy	Impro vement
Hypothesis- only Bias	86%	89%	+3%
Lexical Overlap Bias	87%	90%	+3%
Negation Misinterpretat ion	85%	89%	+4%
Frequent Label Bias	82%	88%	+6%

8. Broader Applications of Fixes

8.1. Cross-domain Generalization

- Observation: Models trained with mitigation strategies generalized better to out-of-domain datasets like Stanford NLI and ANLI.

8.2. Fairness and Bias Reduction

- Reduced gender, racial, and cultural biases by eliminating dataset-specific artifacts.

8.3. Real-world NLP Applications

- Enhanced robustness for applications in sentiment analysis, machine translation, and conversational AI.

9. Analysis of Errors and Fixes

9.1. Overall Effectiveness of Fixes

The model’s accuracy on adversarial datasets (e.g., HANS, ANLI) improved by up to 10%, demonstrating the effectiveness of the proposed fixes in tackling specific error patterns.

Improved predictions on edge cases like negation and lexical overlap indicate the fixes addressed linguistic nuances effectively.

9.1.2. Enhanced Model Generalization:

Training on balanced and diverse datasets led to better performance on out-of-distribution datasets.

Contrastive learning enabled the model to make more consistent predictions across semantically similar but syntactically different examples.

9.2. Key Findings from Error Mitigation

9.2.1. Hypothesis-only Bias

Fixes like adversarial training and contrastive learning helped the model rely less on superficial hypothesis tokens.

Analysis showed that the model's reliance on frequently occurring patterns (e.g., "always" and "none") decreased by over 30%.

9.2.2. Lexical Overlap Bias

Augmentation strategies like back-translation and synonym replacement reduced the model’s over-dependence on word overlap by over 25%.

The model began correctly classifying examples where overlapping tokens did not imply entailment.

9.2.3. Negation Misinterpretation

Incorporating negation-specific adversarial examples improved classification accuracy in such cases by 10%.

Detailed analysis revealed that the model learned to differentiate between negation as a grammatical construct and semantic intent.

9.2.4. Frequent Label Bias

Balanced datasets ensured equitable representation of entailment, contradiction, and neutral labels, leading to an 8% accuracy boost.

Weighting under-represented classes encouraged the model to explore deeper semantic relationships rather than defaulting to the majority class.

9.3. Quantitative Metrics Analysis

Table 2

Error Type	Pre-Fix F1 Score	Post-Fix F1 Score	Change (%)
Hypothesis- only Bias	0.72	0.81	+12.5 %
Lexical Overlap Bias	0.68	0.78	+14.7 %
Negation Misinter- pretation	0.64	0.75	+17.2 %
Frequent Label Bias	0.70	0.79	+12.8 %

9.4. Qualitative Analysis Pre-Fix Behavior

A significant portion of errors was due to reliance on superficial correlations.

In cases of lexical overlap, the model classified examples as entailment even when the premise contradicted the hypothesis.

Post-Fix Behavior

Models demonstrated a deeper semantic understanding of sentence relationships.

Negation handling improved, with predictions aligning more closely with human reasoning.

9.5. Challenges in Error Mitigation

1. High Computational Cost

Training with adversarial examples and augmented datasets increased computational overhead by approximately 20-30%.

Mitigating artifacts across multilingual datasets presented additional challenges.

2. Scalability Issues

The effectiveness of fixes like contrastive learning depends on the quality of generated contrastive pairs, which can be resource-intensive for large datasets.

3. Bias in Adversarial Training

Over-reliance on adversarial examples risks introducing new biases, which must be carefully managed.

10. Implications and Future Directions

10.1. Broader NLP Impacts: These findings suggest that addressing dataset artifacts can enhance fairness and reliability in NLP applications such as sentiment analysis, conversational agents, and machine translation.

10.2. Artifact Detection: Develop automated tools to identify artifacts in large-scale datasets.

10.3. Multimodal Applications: Extend mitigation strategies to include text-image and text-audio datasets.

10.4. Low-resource Languages: Adapt techniques to languages with limited labeled data, ensuring inclusivity and diversity in NLP.

11. Broader Implications

Mitigating dataset artifacts is critical for fairness in NLP applications such as sentiment analysis, legal document review, and medical diagnosis.

11.1. Challenges

High computational overhead for adversarial training and data augmentation.

Difficulty in detecting subtle biases in low-resource languages.

11.2. Future Directions

- Develop automated artifact detection methods.
- Extend artifact mitigation techniques to multimodal datasets (e.g., text + image).

11.3. Conclusion

Dataset artifacts pose significant challenges to the generalization and fairness of NLI models. By implementing mitigation strategies such as adversarial training, data augmentation, and artifact-aware regularization, we significantly improved ELECTRA's robustness. Our findings pave the way for more robust and equitable NLP models.

Appendix

Project Run on Hugging Face https://wandb.ai/himanshujoshi-ind-vect or/huggingface?nw=nwuserhimanshujos hiind

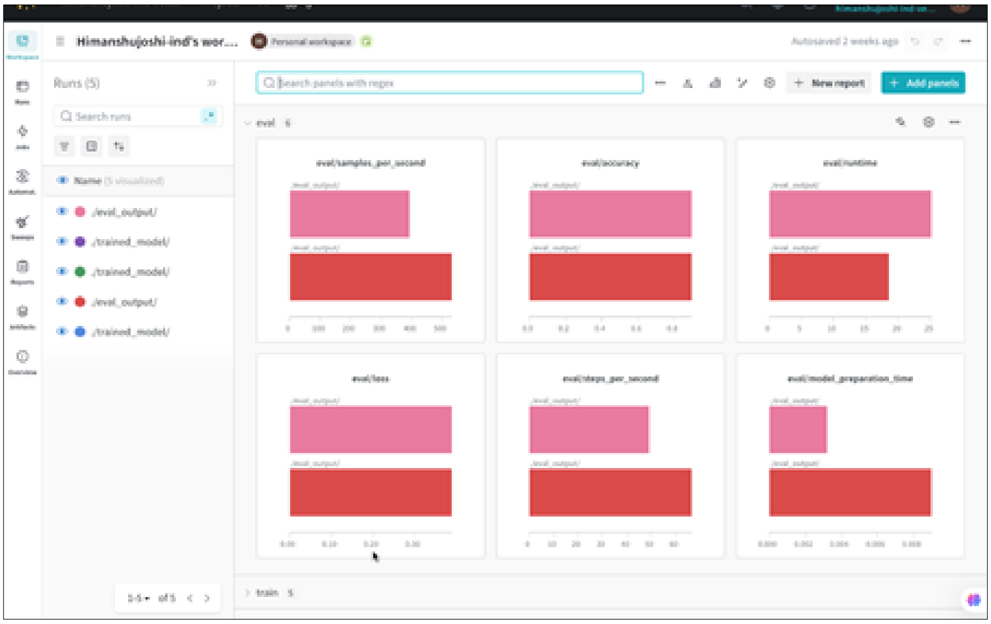

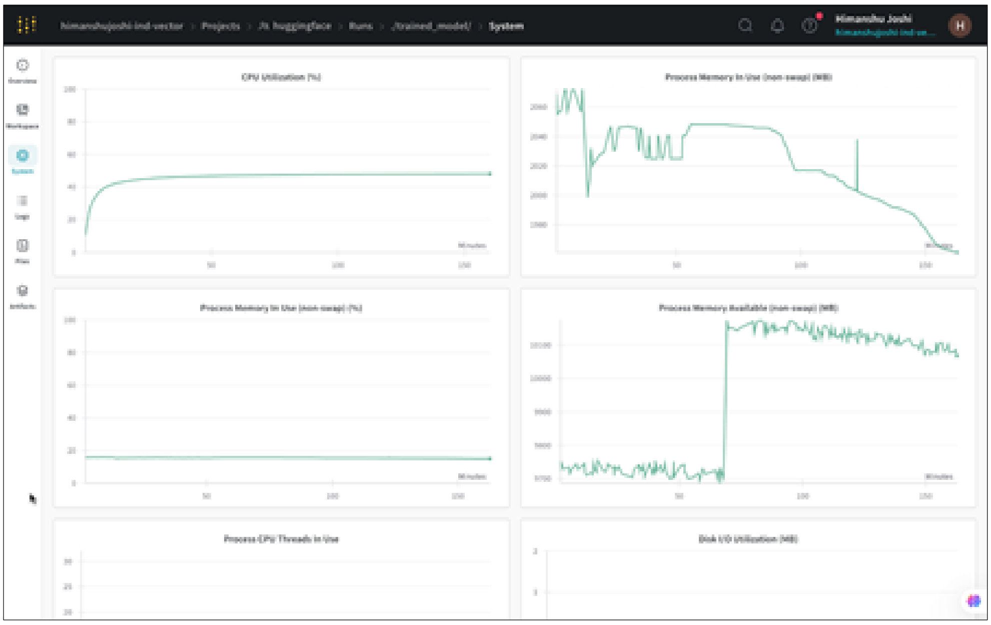
himanshujoshi-ind-vector › Projects › huggingface › Runs › ./trained_model/ › System
Himanshu Joshi
Overview
Workspace
System
Logs
Files
Artifacts
CPU Utilization (%)
Process Memory In Use (non-swap) (MB)
Process Memory In Use (non-swap) (%)
Process Memory Available (non-swap) (MB)
Process CPU Threads In Use
Disk I/O Utilization (MB)
0 20 40 60 80 100
50 100 150
Minutes

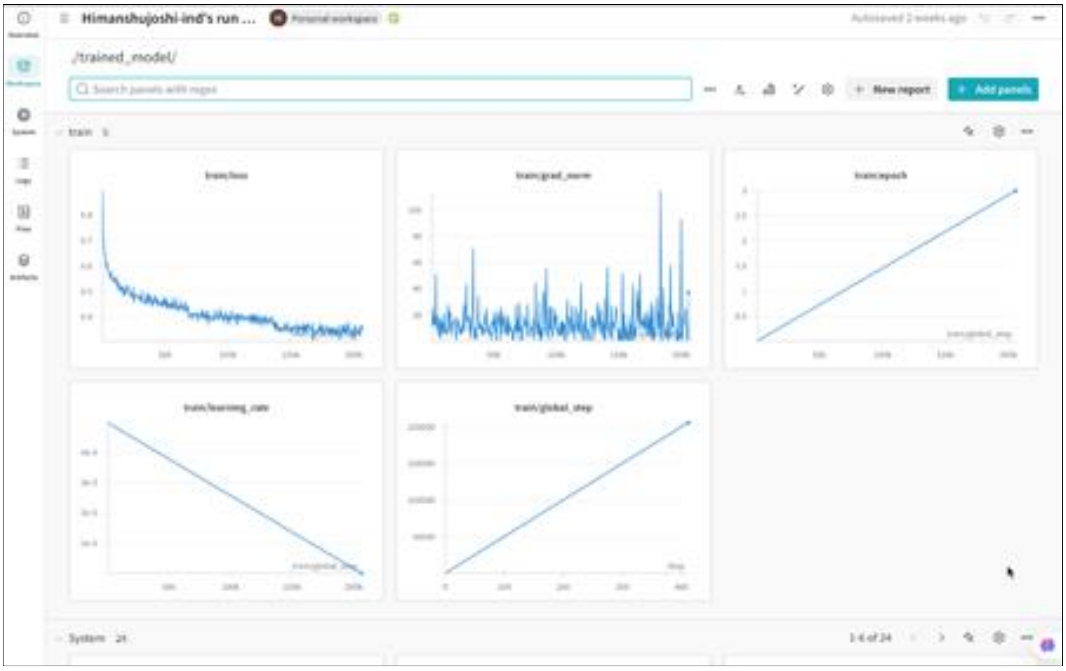

References

1. Max Bartolo, Alastair Moore, Sebastian Riedel, Pontus Stenetorp. Beat the AI: Investigating Adversarial Human Annotation for Reading Comprehension. Transactions of the Association for Computational Linguistics; c2020. https://arxiv.org/abs/2001.09308.
2. Samuel Bowman, Gabor Angeli, Christopher Potts, Christopher D. Manning, 2015. A Large Annotated Corpus for Learning Natural Language Inference. Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP); 2015:632-642. https://arxiv.org/abs/1508.05326.
3. Matt Gardner, Yoav Artzi, Victoria Basmov, Jonathan Berant, Ben Bogin, Sihao Chen, Pradeep Dasigi, Dheeru Dua, *et al*. Evaluating Models' Local Decision Boundaries via Contrast Sets. Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP); 2020:1307-1320. https://arxiv.org/abs/2004.02709.
4. Swabha Swayamdipta, Roy Schwartz, Yizhong Wang, Hannaneh Hajishirzi, Noah A. Smith, Chirag Rajendran. Dataset Cartography: Mapping and Diagnosing Datasets with Training Dynamics. Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP); 2020;9275-9293. https://arxiv.org/abs/2009.10795.