EDITLENS: QUANTIFYING THE EXTENT OF AI EDIT-ING IN TEXT

Anonymous authors

000

001

002 003 004

010 011

012

013

014

016

017

018

021

025

026027028

029

031

033

034

037

040

041

042

043

044

045

046

047

048

051

052

Paper under double-blind review

ABSTRACT

A significant proportion of queries to large language models ask them to *edit* userprovided text, rather than generate new text from scratch. While previous work focuses on detecting fully AI-generated text, we demonstrate that AI-edited text is distinguishable from human-written and AI-generated text. First, we propose using lightweight similarity metrics to quantify the magnitude of AI editing present in a text given the original human-written text and validate these metrics with human annotators. Using these similarity metrics as intermediate supervision, we then train EDITLENS, a regression model that predicts the amount of AI editing present within a text. Our model achieves state-of-the-art performance on both binary (F1=94.7%) and ternary (F1=90.4%) classification tasks in distinguishing human, AI, and mixed writing. Not only do we show that AI-edited text can be detected, but also that the degree of change made by AI to human writing can be detected, which has implications for authorship attribution, education, and policy. Finally, as a case study, we use our model to analyze the effects of AI-edits applied by Grammarly, a popular writing assistance tool. To encourage further research, we commit to publicly releasing our models and dataset.

1 Introduction

Large language models (LLMs) generate text that is difficult to distinguish from human writing, enabling malicious applications such as academic plagiarism and fake review farms, thus motivating the need for accurate AI detection. While existing detectors frame the task as binary classification (fully human vs. fully AI-generated), mainstream LLM usage increasingly involves *co-writing*, where LLMs are used for editing and brainstorming via services like Grammarly, Sudowrite, or Google Docs' Gemini integration. In fact, a recent OpenAI study of over 1M ChatGPT conversations (Chatterji et al., 2025) shows that "about two-thirds of all Writing messages ask ChatGPT to modify user text (editing, critiquing, translating, etc.) rather than creating new text from scratch." Binary AI detection systems are not well-suited to detect such mixed-authorship texts: for example, Saha & Feizi (2025) find that binary detectors often flag AI-polished text as AI-generated, limiting their utility in situations where light AI editing is acceptable but fully AI-generated text is not.

In this paper, we develop EDITLENS, the first AI detector that estimates the extent of AI editing in a text as a continuous score. Previous work on detecting mixed AI and human text has treated the task as either a boundary detection problem (Kushnareva et al., 2024; Lei et al., 2025), a sentence-wise classification task (Wang et al., 2023), or a ternary classification problem between human, AI, and mixed text (Abassy et al., 2024; Wang et al., 2025). However, modern collaborative editing involves layered revisions, suggestions, and refinements that blur traditional notions of authorship, making it challenging to definitively attribute specific segments to either human or AI authors and rendering boundary detection and sentence-level tasks ill-posed. Although the ternary classification approach does not require assigning direct authorship to discrete segments, it is unable to quantify the degree or the magnitude of AI editing: Was the text lightly edited for spelling and grammar, or completely rewritten and restructured? Rather than classifying a text category, our model **directly regresses a score that indicates the degree of AI involvement in the production of the text as a whole.**

https://www.grammarly.com/

²https://sudowrite.com/

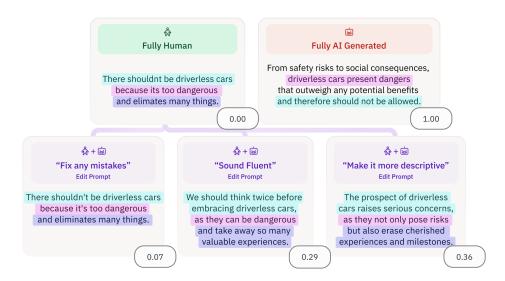


Figure 1: AI edits exist on a continuous spectrum from fully human written to fully AI generated.

Our contributions are the following:

- We introduce a comprehensive dataset spanning a full taxonomy of AI-edits to humanwritten texts.
- 2. We quantify the amount of AI editing applied to each text via lightweight similarity metrics, and validate that the similarity metrics correlate with the judgments of expert human annotators trained to detect AI writing styles.
- 3. We use these similarity metrics to finetune a regression head on an open-source large language model to detect the amount of AI-editing present given only the edited text.
- 4. When converted from a regression model to a binary or ternary classification model, we show that our model, EDITLENS, achieves state of the art performance, outperforming the best binary classifiers by 8%, and outperforming the best ternary classifiers by 16% (macro-F1).
- 5. We also show that unlike the discrete classifiers, the regression model is able to show nuance in progressively classifying more intense edits with higher scores, with case studies on APT-Eval, Beemo, and Grammarly.

Our findings have wide-ranging implications for AI text detection policy. By enabling measurement for the level of AI involvement, more flexible policies acceptable usage of generative AI models can be consistently enforced. Furthermore, our work can help mitigate false positives, a critical limitation of existing binary AI text classifiers. With the ability to control the amount of AI editing allowed, a much lower false positive rate can be achieved under the policy cap framework suggested by Jabarian & Imas (2025) for implementation in high-stakes settings such as academic integrity.

2 QUANTIFYING AI EDIT MAGNITUDE

2.1 Homogeneous vs. Heterogeneous Mixed Authorship

To better motivate our work, we first introduce the concepts of **heterogeneous** and **homogeneous** mixed authorship texts.

In the heterogeneous case, authorship of each segment of text can be directly attributed to a human or AI. An example of this is a situation where a human writes one paragraph and asks the AI to write the following paragraph. In cases like this, there exist one or more boundaries between human and AI segments. One can create token-level labels for heterogeneous mixed texts: every token was authored by either human or AI. Heterogeneous mixed text detection (also called fine-grained AI

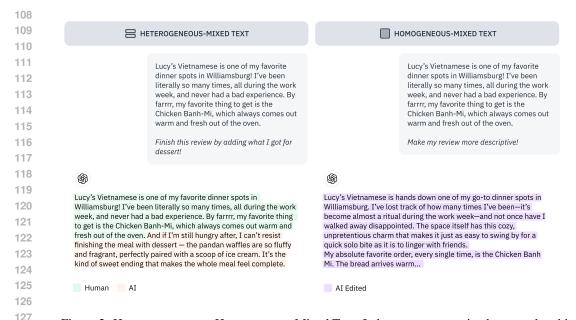


Figure 2: Heterogeneous vs. Homogeneous Mixed Text. In heterogeneous mixed text, authorship of each token is clearly attributable.

text detection) has been previously studied by Kushnareva et al. (2024), Wang et al. (2023), and Lei et al. (2025).

In the homogeneous case, authorship is entangled by the editing process. An example of this is a situation where a human writes a paragraph and asks an AI to paraphrase it. Even if AI replaces every word in the paragraph with a synonym, authorship is still mixed. As such, token-level binary labels are insufficient measures of authorship in this case, as both parties have provided input throughout the entire document. Despite its increasing prevalence, homogeneous mixed AI text is understudied, and we focus the rest of the paper on detecting this kind of mixed text.

2.2 TASK DEFINITION: HOMOGENEOUS MIXED TEXT

In many practical scenarios, a human-written document x is subsequently edited to yield a new document y, where multiple sequential edits may have been performed by one or more agents (human or AI) in an indistinguishable fashion to produce y. Unlike the heterogeneous mixed-text setting, where each segment is assumed to be authored wholly by either a human or an LLM, here authorship is latent and entangled within the editing process. Our objective is not to attribute authorship, but to predict the magnitude of change between x and y according to a similarity metric that agrees with expert judgments of the magnitude of AI writing style and semantics.

We model the edited text as the image of an editing operator \mathcal{E}_{λ} applied to x:

$$y = \mathcal{E}_{\lambda}(x; z), \qquad z \sim p(z), \quad \lambda \in \Lambda,$$

where z denotes a (latent) sequence of micro-edits (insertions, deletions, substitutions, reorderings) possibly performed by a mixture of editor types (humans or AIs) and λ summarizes an *edit intensity*. In the homogeneous setting, the editor identity within z is unobserved and not required at training or inference time. For simplicity, in this study, we focus on the case where a human text is edited in one pass by a single AI language model, but we also present results for multiple passes, and human-edited AI text as case studies in generalization.

Similarity-driven target. Let $sim : \mathcal{X} \times \mathcal{X} \to [0,1]$ be a fixed similarity functional. We define a change magnitude functional $\Delta : \mathcal{X} \times \mathcal{X} \to [0,1]$ by a monotone transformation of similarity (or distance):

$$\Delta(x,y) = g(\sin(x,y)), \quad \text{e.g.}, \quad g(s) = 1 - s$$

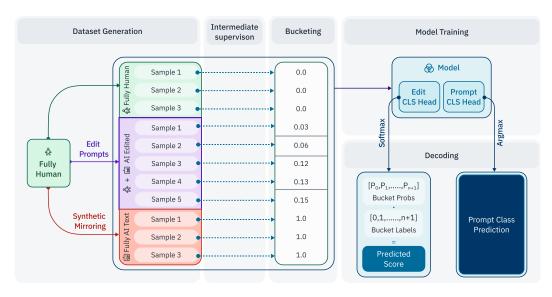


Figure 3: EDITLENS architecture. We use similarity metrics as intermediate supervision, and simultaneously train edit classification/regression and prompt classification tasks in multi-task learning.

where sim is a nonnegative distance. $\Delta(x,y)=0$ for identical texts (no edits) and increases as heavier editing is applied to form y. We motivate the particular choice of sim below via agreement with expert annotators' perception of the amount of AI pervasiveness within a text, and it is assumed known during training and evaluation.

Inference with edited text only. In most practical settings, only the edited document y is available at inference time. We therefore learn a single-input predictor that maps y directly to a change magnitude without reconstructing or retrieving a source x:

$$f_{\theta}^{\mathrm{ssi}}: \mathcal{X} \rightarrow [0,1], \qquad \hat{\Delta}(y) = f_{\theta}^{\mathrm{ssi}}(y).$$

Training remains supervised using pairs $\{(x^{(i)},y^{(i)})\}_{i=1}^N$ only to compute targets $\Delta^{(i)}=\Delta(x^{(i)},y^{(i)})$; the model never conditions on x at inference. Concretely, we optimize

$$\min_{\theta} \frac{1}{N} \sum_{i=1}^{N} \mathcal{L}\left(f_{\theta}^{\text{ssi}}(y^{(i)}), \Delta(x^{(i)}, y^{(i)})\right).$$

The Bayes-optimal predictor for this objective is the conditional expectation

$$f^{\star}(y) = \mathbb{E}[\Delta(X, y) \mid Y = y],$$

but crucially we *do not* estimate this expectation via reconstruction of x. Instead, $f_{\theta}^{\rm ssi}$ learns discriminatively from y alone, absorbing the necessary inductive biases (e.g., lexical volatility, style drift, fluency/consistency cues) to approximate f^* from labeled examples.

For additional discussion of the precise differences between homogeneous and heterogeneous mixed detection formulations, see the Appendix.

3 TRAINING A MODEL TO DETECT AI EDITS

3.1 Creating a Homogeneous Mixed Text Dataset

Because no dataset of homogeneous mixed AI-generated text exists *at scale*, we create a training set for this task.

We begin by collecting a source dataset of fully-human and fully-AI-generated texts. We select human-written texts from prior to the release of large language models in 2022 from 4 domains: reviews from Amazon (Zhang et al., 2015) and Google (Li et al., 2022), creative writing samples from

Reddit Writing Prompts (Fan et al., 2018), general educational web articles from FineWeb-EDU (Lozhkov et al., 2024), and news articles from XSum (Narayan et al., 2018) and CNN/DailyMail (See et al., 2017). As a holdout domain to measure out-of-distribution performance, we also include the Enron email dataset (Cohen, 2015).

Then, we generate an AI example corresponding to each human example following the synthetic mirroring procedure introduced in Emi & Spero (2024). We use GPT-4.1, Claude 4 Sonnet, and Gemini 2.5 Flash. We also include Llama-3.3-70B-Instruct-Turbo as a holdout LLM to measure performance on out-of-distribution LLMs. Our final train, test, and val splits contain 60k, 6k, and 2.4k examples respectively. We estimate the cost of creating this dataset to be roughly \$530. Additional dataset summary statistics can be found in Tables 10 to 14.

3.2 EDIT PROMPTS

We collected a set of editing prompts by first prompting ChatGPT 40, Claude Sonnet 4, and Gemini 2.5 Pro. When prompting the LLMs, we asked them to categorize the edit prompts, and we consolidated the categories across models into 9 prompt types. We also added a small number of human-written prompts written by the authors and self-categorized them. In total, we collected 303 editing prompts. The full list of prompts and summary statistics about the categories and contributors can be found in Tables 8 and 9. While this list of prompts is not exhaustive, it encompasses a significant coverage of the different ways that people use AI to edit texts. We split this list of prompts into train, test, and validation splits so that the model cannot overfit to a particular set of prompts.

3.3 Intermediate Supervision Metrics

We experiment with two methods for labeling the "difference" $\Delta(x,y)$ in a text before and after AI editing. The first is the cosine distance (1 - cosine similarity) between the Linq-Embed-Mistral (Choi et al., 2024) embeddings of the source text and the AI-edited version. We chose this embedding due to its strong all-around performance on the MTEB benchmark (Muennighoff et al., 2023).

The second is a precision-based method similar to the embedding-based ROUGE proposed by Ng & Abrecht (2015): given a minimum (a) and maximum (b) sequence length, we enumerate all phrases (including overlaps) of between a and b words in the source and edited texts. We compute the pairwise cosine similarity between phrases in the source and edited texts, then count the number of phrases in the edited text with a cosine similarity above a threshold τ for any phrase in the source text. This count is divided by the total number of phrases in the edited text, making it a precision-based metric. We refer to this metric as the $soft\ n$ -grams score throughout the paper. Soft n-grams reduces to n-gram overlap between source and target when $\tau=1$. We choose soft n-grams because it expresses similarity when the AI editor replaces a phrase or word with a semantically similar one rather than requiring exact matching. We note that this supervision metric is shortening-invariant, i.e., simply deleting text from the source still yields a soft n-grams score of 1.

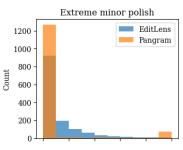
3.4 Human Agreement with Intermediate Supervision Metrics

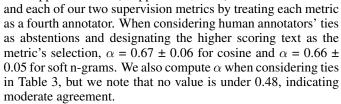
How well do these automatic metrics actually capture the extent of AI editing in a text? To support our choice of intermediate supervision metric, we conduct a study that asks human annotators to compare two AI-edited versions of the same text after findings by Russell et al. (2025) that humans are effective detectors of AI-authored text.

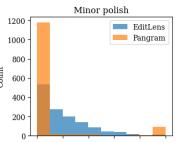
Task setup. Annotators are shown 3 texts side-by-side: a human written *source text* alongside 2 AI-edited versions of the source text. The labeling interface can be seen in Figure 7. Between the two AI texts, annotators are asked to select which text contains more AI edits. Annotators may also answer that there is a "Tie," i.e. both texts contain roughly the same amount of AI edits. Annotators have the option to leave freeform comments on each task, but were not required to do so. We recruited 3 annotators with extensive daily exposure to both human writing and AI-generated texts. Each annotator completed all 100 tasks in approximately 6 hours and was compensated at a rate of \$30 (USD) per hour.

Task generation procedure. We randomly sample 100 human-written texts of between 50 and 300 words from our test set. We then generate multiple AI-edited versions of each source text using randomly assigned prompts until we have two AI-edited versions that are within 15 words of the source text. We impose this length restriction on the edited texts to encourage annotators to consider the actual text, rather than simply the length when selecting the version with more AI edits.

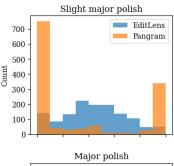
Agreement with metrics. We report Krippendorff's α (Krippendorff, 1980) for our 3 annotators



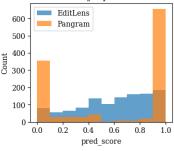




3.5 Modeling Details



Using QLoRA Dettmers et al. (2023), we finetune models of between 3 and 24B parameters from the Mistral and Llama families. We use QLoRA to sweep the widest possible range of sizes of base models to use as a backbone that fit in VRAM on a single GPU. We leave other finetuning and modeling architecture choices to future work. We experiment with directly training a regression head using MSE loss as well as training an n-way classification model, then decoding the output to a score, using weighted-average decoding rather than traditional argmax decoding. As an auxiliary task, we also train the model to classify the prompt category using an additional head trained with a cross-entropy loss. Additional details can be found in the Appendix.



4 RESULTS

EDITLENS demonstrates significantly more nuanced AI detection than existing classifiers through both quantitative metrics and qualitative analysis. We compare EDITLENS with several open- and closed-source AI detection baselines. On the AI Polish dataset (APT-Eval), EditLens achieves substantially stronger correlations with edit magnitude metrics compared to binary detectors (correlation 0.606), markedly outperforming the best binary baseline Pangram (correlation 0.491). This quantitative superiority is complemented by clear qualitative differences: while binary classifiers like Pangram predict scores clustered near 0 or 1, EditLens produces a nuanced distribution that appropriately tracks increasing levels of AI polish from minor to major edits. The model's regression-based approach enables it to achieve state-of-the-art performance across evaluation paradigms, delivering 94.0% accuracy in binary classification (human vs. any AI) and 90.2% accuracy in ternary classification (human vs. AI-edited vs. AI-generated), substantially

Figure 4: AI Polish distributions

outperforming existing binary and ternary detection methods. Additionally, EditLens generalizes effectively outside its training distribution: to unseen prompts, LLMs, and domains, to human-edited AI text in the BEEMO dataset, and to AI-edited AI text as well as multi-edited AI text.

4.1 AI POLISH DATASET

We first compare the performance on the AI Polish dataset (APT-Eval) of EDITLENS against the best-performing binary AI classifier, Pangram. APT-Eval contains both degree-based AI-edited text,

(a) Human vs. Any AI

Acc. (%)

69.1

68.6

80.7

93.8

94.0

F1

80.5

81.4

83.7

95.4

95.6

Table 1: Accuracy and F1-score on two binary classification tasks: (a) human vs. any AI generated or edited texts and (b) fully AI-generated texts vs. AI-edited and human texts. Thresholds were

calibrated using the val set. "SNG" denotes EDITLENS trained with soft n-grams supervised data.

with 4 discrete categories (extreme minor, minor, slight major, and major polish levels), as well as

percentage-based AI-edited text, where LLMs were asked to edit a certain percentage of the text,

Model (Threshold)

Binoculars 0.362

varying from 1-75%.

Pangram 0.001

FastDetectGPT_{0.009}

EDITLENS (Cosine) 0.039

EDITLENS (SNG) (0.041)

-7	-	44
_	_	
-2		

326 327 328 330

331 332

333

334

335 336 337 338



343 344 345



351 352 353

354 355

356 357 358

359 360 361

362 363 364

366 367

3	6	8
3	6	9
3	7	0

371 372 373

3	7	4
3	7	5
3	7	6

377

4.4 Out-of-Domain Performance

While there are no direct or exact labels, the score should generally monotonically increase as the amount of requested polish increases. In Figure 4, we qualitatively assess the distribution of the

model prediction scores on the degree-based edits. We can see a clear difference between the behavior of EDITLENS versus the behavior of Pangram. Pangram almost always predicts a score very close to 0 or 1, while EDITLENS is able to quantify the increasing levels of polish applied. We show the equivalent distributions for percentage-based polishing in the Appendix.

(b) Fully AI vs. AI-Edited + Human

Acc. (%)

90.6

31.4

92.3

94.3

96.4

F1

84.4

47.7

89.0

90.2

94.1

Model (Threshold)

Binoculars 0.601

Pangram 0.998

FastDetectGPT_{0.889}

EDITLENS (SNG) (0.998)

EDITLENS (Cosine) (0.960)

similarity metrics between source and target provided by APT-eval in Table 4. For EDITLENS and all binary classification baselines, we measure the Pearson correlation coefficient (r) between the prediction scores and the semantic similarity (-0.606), Levenshtein distance (0.799), and Jaccard distance (0.781) metrics between the pre-AI-polished and post-AI-polished documents. Stronger correlation values mean that the model is able to faithfully track edit magnitude across examples and assign higher scores as semantic similarity decreases (and Levenshtein/Jaccard distances increase), and lower scores when the edited text remains close to the source. EDITLENS exhibits a significant correlation between these similarity metrics and its scores, while the binary AI detectors correlate less strongly with these metrics.

Quantitatively, we also report the correlation value between the EDITLENS predicted score and the

4.2 PERFORMANCE AS A BINARY CLASSIFIER

In the binary classification setting, how does EDITLENS treat mixed text? Different use cases may have different standards for what they consider an acceptable amount of AI-generated text-a professor may allow the use of AI assistance for proofreading, but disallow fully AI-generated essays.

To measure the flexibility of our model and the baselines to be able to adjust to different sensitivity levels, we calibrate and compute the performance of each model on two settings: fully human-written vs. any AI-edited or AI-generated text, fully human-written and AI-edited text vs. AI-generated text. Model accuracy and F1-scores can be found in Table 1. Notably, EDITLENS outperforms our three binary baselines, FastDetectGPT, Binoculars, and Pangram, on our test set consisting of fully human-written, fully AI-generated, and AI-edited texts.

4.3 Performance as a Ternary Classifier

To compare with categorical mixed AI detection models, we evaluate each model on three classes: human, AI-generated, or AI-edited. To convert each binary classifier into a ternary classifier, we find two thresholds using the calibration procedure above on a held-out validation set, optimizing the F1 score between the human/mixed and mixed/AI classes. The decoding procedures for GPTZero and DetectAIve are detailed in the Appendix.

During dataset creation, we hold out both a model and a domain to test the ability of our model to generalize to out-of-distribution texts. We created an OOD model test set of 3k examples with Llama-3.3-70B-Instruct-Turbo generated and edited texts as well as an OOD domain test set using

Model	Туре	Accuracy	Macro-F1	Human F1	AI-Gen. F1	AI-Edited F1
Binoculars	Binary	48.5	50.1	54.7	58.4	37.4
FastDetectGPT	Binary	60.6	56.8	25.1	84.4	61.0
Pangram	Binary	73.0	69.5	76.4	89.0	43.2
DetectAIve	Ternary	57.8	52.5	79.9	16.1	61.5
GPTZero	Ternary	74.7	72.7	77.3	89.8	50.9
EDITLENS Soft N-Grams	Regression	89.7	89.9	89.6	94.1	86.1
EDITLENS Cosine	Regression	90.2	90.4	90.4	94.1	86.8

Table 2: Ternary classification performance across different model types.

the Enron email dataset (Cohen, 2015) as source texts, and measure the degradation in macro-F1 score of our best model, EDITLENS with cosine supervision.

On the OOD domain dataset, macro-F1 on the ternary classification task decreases from 0.904 to 0.866 (-0.038). On the OOD LLM dataset, macro-F1 on the ternary classification task decreases from 0.904 to 0.850 (-0.054).

4.5 Performance on multi-edited AI text

We also examine the case where multiple AI-edits have been applied to a single piece of text. We test our model on this case by applying a series of 5 edits to a piece of human-written text and measuring the EDITLENS score after each subsequent edit. In 5, we show that for each edit, the mean score increases.

4.6 GENERALIZATION TO AI-EDITED AI TEXT

To ensure EDITLENS estimates the extent of AI-editing rather than the presence of edits of any kind, we evaluate our detector's mean score difference on AI-edited, AI-generated text.

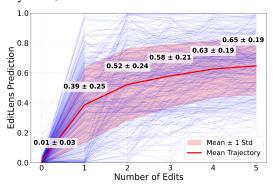


Figure 5: "Trajectory" of EDITLENS scores after

a number of subsequent AI edits to a single text.

We take synthetic mirrors of our original human dataset, considered to be 'AI-generated documents' and edit them using our held-out prompt set. On a dataset size of n=412, the mean score difference for a single edit pass on an originally human text is 0.38. The mean score difference for a single edit pass on an originally AI text is -0.05.

4.7 GENERALIZATION TO HUMAN-EDITED AI TEXT (BEEMO)

While the majority of our studies focus on AIedited human writing, we also evaluate the performance of EDITLENS on human-edited AI

text using the BEEMO (Artemova et al., 2024) dataset, which includes human expert-edited versions of AI model outputs. We find that the model adequately generalizes to human-edited AI text. The average decrease in score from the model output to the human-edited model output is 0.33 ± 0.30 , with the score decreasing after human-editing in 88.9% of the documents. More details are presented in the Appendix.

4.8 CASE STUDY: GRAMMARLY EDIT DATASET

Grammarly³ is a popular subscription-based AI writing assistant that allows users to edit text using both pre-filled and custom prompts within a native word processor. We manually collect a dataset of 1768 samples using 9 of the default prompts offered by Grammarly to simulate typical user queries

https://www.grammarly.com/

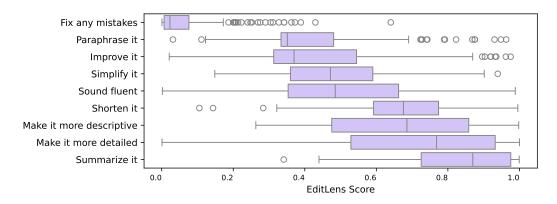


Figure 6: Distribution of EDITLENS scores on dataset from Grammarly by edit instruction.

for AI editing by sampling 197⁴ human-written source texts and applying each of the 9 edits to them in the Grammarly web interface. In Figure 6, we present the distributions of EDITLENS scores on examples from each editing instruction sorted by the median. Perhaps counterintuitively, EDITLENS considers "Fix any mistakes" the most minor of all edits, while "Summarize this" and "Make it more detailed" are the most invasive edits. In Figures 8 and 9 we show this is also true according to both the cosine and soft n-grams scores of the examples.

5 RELATED WORK

Binary AI-Generated Text Detectors. Several works have explored the binary setting of distinguishing fully-human from fully-AI-generated text. DetectGPT (Mitchell et al., 2023), FastDetect-GPT (Bao et al., 2024), DNA-GPT (Yang & Cheng, 2024), Binoculars (Hans et al., 2024), Ghostbusters (Verma et al., 2024), and numerous other open-works have been invented. Commercial closed-source solutions such as GPTZero (Tian & Cui, 2023) and Pangram (Emi & Spero, 2024) have also emerged as effective AI text classifiers. Heterogeneous Mixed Text Detection. As described above, previous work on mixed AI and human text detection focuses on the heterogeneous case: where distinct boundaries can be drawn between fully AI-generated and fully human-written segments. Examples of these works include AI Boundary Detection with RoFT Kushnareva et al. (2024), SeqXGPT Wang et al. (2023), HaCo-Det Su et al. (2025), and PALD Lei et al. (2025). Categorical Mixed Text Detection. Alternatively, some previous work has instead focused on mixed text as an additional category or categories in addition to human and AI. DetectAIve Abassy et al. (2024), HERO Wang et al. (2025), and GPTZero Tian & Cui (2023) are all examples where mixed categories have been added. We find the limitation of this approach is that the amount of editing cannot be quantified: all mixed text is treated as the same. Human-Edited AI Text. In addition to our problem setting of AI-edited human written text, there are also studies and datasets focusing on human-edited AI-generated text. Beemo Artemova et al. (2024) is a benchmark focusing on expert-edited AI text. LAMP Chakrabarty et al. (2025) is a corpus of LLM-generated paragraphs that have been improved by professional writers according to a defined taxonomy. Paraphrasers and Humanizers. Several previous works have studied the effects of automated paraphrasers (Krishna et al., 2023; Russell et al., 2025) and "humanizers" (Masrour et al., 2025) on how they degrade AI-generated text. We explore the effect of AI rewriting of AI outputs as it relates to our model in the results.

6 CONCLUSION

In this study, we introduce the task of continuous fine-grained AI edit prediction, and show ED-ITLENS, based on simple embedding-based supervision on a finetuned language model, significantly outperforms existing AI detection approaches. We release our dataset and models to encourage future research in this area.

⁴Occasionally, Grammarly would abstain, leaving us with fewer than 197 * 9 samples.

7 ETHICS STATEMENT

Our research involved using 3 human subjects to annotate the degree of AI-editing present in a text. We obtained informed consent from the subjects and fairly compensated them for their labor. We commit to maintaining their privacy.

Inaccurate AI detection software can cause harm as false accusations of AI misconduct can result in serious consequences, including emotional trauma, reputation damage, and undue punishments for academic misconduct. We acknowledge that our model has a non-zero error rate and its errors may result in such harms. We commit to continuing to engage with the academic community to educate others on appropriately contextualizing and communicating the results of AI detection software. We also commit to releasing the model for non-commercial use only and responsibly vetting access to researchers and educators.

We intend for our contribution to the research on AI detection to ultimately mitigate harm by providing a more nuanced picture of AI usage than binary AI detection classifiers. The ability to calibrate the sensitivity level of the regression model is also a step towards mitigating the false positive rate and lowering the overall number of false accusations of AI misconduct.

8 REPRODUCIBILITY STATEMENT

We will release the dataset, source code, and final model weights at a later date to be determined.

REFERENCES

Mervat Abassy, Kareem Elozeiri, Alexander Aziz, Minh Ngoc Ta, Raj Vardhan Tomar, Bimarsha Adhikari, Saad El Dine Ahmed, Yuxia Wang, Osama Mohammed Afzal, Zhuohan Xie, Jonibek Mansurov, Ekaterina Artemova, Vladislav Mikhailov, Rui Xing, Jiahui Geng, Hasan Iqbal, Zain Muhammad Mujahid, Tarek Mahmoud, Akim Tsvigun, Alham Fikri Aji, Artem Shelmanov, Nizar Habash, Iryna Gurevych, and Preslav Nakov. LLM-DetectAlve: a tool for fine-grained machine-generated text detection. In Delia Irazu Hernandez Farias, Tom Hope, and Manling Li (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 336–343, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-demo.35. URL https://aclanthology.org/2024.emnlp-demo.35.

- Anthropic. Introducing claude 4 (claude opus 4 and claude sonnet 4). Anthropic Blog Post, May 2025. URL https://www.anthropic.com/news/claude-4.
- Ekaterina Artemova, Jason Lucas, Saranya Venkatraman, Jooyoung Lee, Sergei Tilga, Adaku Uchendu, and Vladislav Mikhailov. Beemo: Benchmark of expert-edited machine-generated outputs. *arXiv preprint arXiv:2411.04032*, 2024.
- Guangsheng Bao, Yanbin Zhao, Zhiyang Teng, Linyi Yang, and Yue Zhang. Fast-detectgpt: Efficient zero-shot detection of machine-generated text via conditional probability curvature. In *The Twelfth International Conference on Learning Representations (ICLR)*, 2024. URL https://arxiv.org/abs/2310.05130.
- Tuhin Chakrabarty, Philippe Laban, and Chien-Sheng Wu. Can ai writing be salvaged? mitigating idiosyncrasies and improving human-ai alignment in the writing process through edits, 2025. URL https://arxiv.org/abs/2409.14509.
- Aaron Chatterji, Thomas Cunningham, David J. Deming, Zoe Hitzig, Christopher Ong, Carl Yan Shan, and Kevin Wadman. How people use chatgpt. Working Paper 34255, National Bureau of Economic Research, Cambridge, MA, September 2025. URL http://www.nber.org/papers/w34255. NBER Working Paper No. 34255.
- Chanyeol Choi, Junseong Kim, Seolhwa Lee, Jihoon Kwon, Sangmo Gu, Yejin Kim, Minkyung Cho, and Jy yong Sohn. Linq-Embed-Mistral Technical Report. *arXiv preprint arXiv:2412.03223*, December 2024. 15 pages.

- William W. Cohen. Enron email dataset. https://www.cs.cmu.edu/~enron/, 2015. Prepared by the CALO Project (Cognitive Assistant that Learns and Organizes). Originally made public by the Federal Energy Regulatory Commission.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient Finetuning of Quantized LLMs. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, NIPS '23, 2023.
 - Bradley Emi and Max Spero. Technical report on the pangram ai-generated text classifier, 2024. URL https://arxiv.org/abs/2402.14873.
 - Angela Fan, Mike Lewis, and Yann Dauphin. Hierarchical neural story generation. *arXiv preprint arXiv:1805.04833*, 2018.
 - Google. Gemini 2.5: Pushing the Frontier with Advanced Reasoning, Multimodality, Long Context, and Next Generation Agentic Capabilities. Technical report, Google DeepMind, 2025. URL https://storage.googleapis.com/deepmind-media/gemini/gemini_v2_5_report.pdf. Technical Report.
 - Abhimanyu Hans, Avi Schwarzschild, Valeriia Cherepanova, Hamid Kazemi, Aniruddha Saha, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Spotting llms with binoculars: Zero-shot detection of machine-generated text, 2024.
 - Brian Jabarian and Alex Imas. Artificial writing and automated detection. https://ssrn.com/abstract=5407424, September 2025. Chicago Booth Research Paper Forthcoming.
 - Klaus Krippendorff. *Content Analysis: An Introduction to Its Methodology*, chapter 12. Sage Publications, Beverly Hills, CA, 1980.
 - Kalpesh Krishna, Yixiao Song, Marzena Karpinska, John Wieting, and Mohit Iyyer. Paraphrasing evades detectors of ai-generated text, but retrieval is an effective defense. *arXiv preprint arXiv:2303.13408*, 2023.
 - Laida Kushnareva, Tatiana Gaintseva, German Magai, Serguei Barannikov, Dmitry Abulkhanov, Kristian Kuznetsov, Eduard Tulchinskii, Irina Piontkovskaya, and Sergey Nikolenko. Aigenerated text boundary detection with roft. In *1st Conference on Language Modeling (COLM)*, Philadelphia, United States, October 2024. URL https://openreview.net/forum?id=kzzwTrt04Z.
 - Eric Lei, Hsiang Hsu, and Chun-Fu Chen. Pald: Detection of text partially written by large language models. In *Proceedings of the Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=rWjZWHYPcz.
 - Jiacheng Li, Jingbo Shang, and Julian McAuley. Uctopic: Unsupervised contrastive learning for phrase representations and topic mining. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 51–62. Association for Computational Linguistics, 2022.
 - Anton Lozhkov, Loubna Ben Allal, Leandro von Werra, and Thomas Wolf. Fineweb-edu: the finest collection of educational content, 2024. URL https://huggingface.co/datasets/HuggingFaceFW/fineweb-edu.
 - Elyas Masrour, Bradley Emi, and Max Spero. Damage: Detecting adversarially modified ai generated text, 2025. URL https://arxiv.org/abs/2501.03437.
 - Meta. Llama 3.1 8b: Open foundation model with 128k context. Meta AI Announcement / Hugging Face Model Card, 2024a. URL https://huggingface.co/meta-llama/Llama-3.1-8B.
 - Meta. Llama 3.2: Small medium vision-enabled models (3b, 11b, 90b). Meta Connect 2024 Announcement, 2024b. URL https://ai.meta.com/blog/meta-llama-3-2.

- Meta. Llama 3.3 70b instruct (december 2024 release). Meta AI Announcement on X (Twitter) and Hugging Face Model Card, 2024c. URL https://huggingface.co/meta-llama/Llama-3.3-70B-Instruct.
 - Mistral. Mistral nemo: Our new best small model (12b, 128k context). Mistral AI Blog, July 2024. URL https://mistral.ai/news/mistral-nemo.
 - Mistral. Mistral small 3: A 24b latency-optimized model (apache 2.0). Mistral AI Blog, January 2025. URL https://mistral.ai/news/mistral-small-3.
 - Eric Mitchell, Yoonho Lee, Alexander Khazatsky, Christopher D. Manning, and Chelsea Finn. Detectgpt: Zero-shot machine-generated text detection using probability curvature. In *Proceedings of the 40th International Conference on Machine Learning*, ICML '23, 2023. URL https://arxiv.org/abs/2301.11305.
 - Niklas Muennighoff, Nouamane Tazi, Loïc Magne, and Nils Reimers. Mteb: Massive text embedding benchmark, 2023. URL https://arxiv.org/abs/2210.07316.
 - Shashi Narayan, Shay B. Cohen, and Mirella Lapata. Don't give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization. *ArXiv*, abs/1808.08745, 2018.
 - Jun-Ping Ng and Viktoria Abrecht. Better summarization evaluation with word embeddings for ROUGE. In Lluís Màrquez, Chris Callison-Burch, and Jian Su (eds.), *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pp. 1925–1930, Lisbon, Portugal, September 2015. Association for Computational Linguistics. doi: 10.18653/v1/D15-1222. URL https://aclanthology.org/D15-1222/.
 - OpenAI. Introducing gpt-4.1 in the api. OpenAI Product Blog, April 2025. URL https://openai.com/blog/introducing-gpt-4-1.
 - Nils Reimers, Iryna Gurevych, et al. all-minilm-l6-v2: Sentence-transformers model. Sentence-Transformers Model Card (Hugging Face Hub), 2021. URL https://huggingface.co/ sentence-transformers/all-MinilM-L6-v2.
 - Jenna Russell, Marzena Karpinska, and Mohit Iyyer. People who frequently use chatgpt for writing tasks are accurate and robust detectors of ai-generated text, 2025. URL https://arxiv.org/abs/2501.15654.
 - Shoumik Saha and Soheil Feizi. Almost ai, almost human: The challenge of detecting ai-polished writing. *arXiv preprint arXiv:2502.15666*, 2025.
 - Abigail See, Peter J. Liu, and Christopher D. Manning. Get to the point: Summarization with pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1073–1083, Vancouver, Canada, July 2017. Association for Computational Linguistics. doi: 10.18653/v1/P17-1099. URL https://www.aclweb.org/anthology/P17-1099.
 - Zhixiong Su, Yichen Wang, Herun Wan, Zhaohan Zhang, and Minnan Luo. HACo-det: A study towards fine-grained machine-generated text detection under human-AI coauthoring. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 22015–22036, Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 979-8-89176-251-0. doi: 10.18653/v1/2025.acl-long.1069. URL https://aclanthology.org/2025.acl-long.1069/.
 - Edward Tian and Alexander Cui. Gptzero: Towards detection of ai-generated text using zero-shot and supervised methods", 2023. URL https://gptzero.me.
 - Vivek Verma, Eve Fleisig, Nicholas Tomlin, and Dan Klein. Ghostbuster: Detecting text ghost-written by large language models. In Kevin Duh, Helena Gomez, and Steven Bethard (eds.), Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers), pp. 1702–1717, Mexico City, Mexico, June 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.naacl-long.95. URL https://aclanthology.org/2024.naacl-long.95/.

	Raw scores	4 buckets	5 buckets	6 buckets
Cosine Distance	0.67 ± 0.06	0.50 ± 0.05	0.52 ± 0.05	0.55 ± 0.05
Soft N-Grams	0.66 ± 0.05	0.48 ± 0.05	0.50 ± 0.05	0.52 ± 0.05

Table 3: Agreement (Krippendorff's α with bootstrap SE) between human annotators and proposed intermediate supervision metrics under different bucketing schemes for scores.

Pengyu Wang, Linyang Li, Ke Ren, Botian Jiang, Dong Zhang, and Xipeng Qiu. SeqXGPT: Sentence-level AI-generated text detection. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 1144–1156, Singapore, December 2023. Association for Computational Linguistics. URL https://aclanthology.org/2023.emnlp-main.73/.

Yitong Wang, Zhongping Zhang, Margherita Piana, Zheng Zhou, Peter Gerstoft, and Bryan A. Plummer. Real, fake, or manipulated? detecting machine-influenced text, 2025. URL https://arxiv.org/abs/2509.15350.

Xianjun Yang and Wei Cheng. Dna-gpt: Divergent n-gram analysis for training-free detection of gpt-generated text. In *Proceedings of the International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=Xlayxj2fWp.

Xiang Zhang, Junbo Zhao, and Yann LeCun. Character-level convolutional networks for text classification. In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015.

A DIFFERENCES WITH THE HETEROGENEOUS MIXED TEXT DETECTION TASK

The PALD (Lei et al., 2025) formulation considers a text x segmented as $x = x_1 \cdots x_n$, where each segment x_i is assumed to originate from either a human or an LLM, i.e., $x_i \sim P_{\text{human}}$ or $x_i \sim P_{\text{LLM}}$. The learning objective is to infer latent per-segment authorship labels $a_{1:n} \in \{\text{human}, \text{LLM}\}^n$ (and optionally segment boundaries), estimating $p_{\theta}(a_{1:n} \mid x)$ and predicting $\hat{a}_{1:n} = \arg\max p_{\theta}(a_{1:n} \mid x)$. In contrast, our homogeneous mixed text prediction task dispenses with provenance as supervision and regresses an authorship-agnostic edit magnitude aligned to a similarity metric. Given pre/post pairs only to derive targets, inference relies on a single-input predictor $f_{\theta}^{\text{ssi}}(y)$ that maps the edited text y directly to $\hat{\Delta}(y) \in [0,1]$, without segment labels, boundary inference, or reconstruction of the source. This reframing changes (i) the **assumptions** (binary authorship mixture vs. latent, entangled edits), (ii) the **outputs** (label sequence $a_{1:n}$ vs. scalar/regional magnitudes Δ), (iii) the **supervision** (segment-level authorship vs. metric-aligned change signals), and (iv) the **evaluation** (classification metrics such as accuracy/F1 vs. correlation and error against Δ , plus calibration).

B HUMAN AGREEMENT WITH INTERMEDIATE SUPERVISION METRICS

We compute the score for each pair of source and AI-edited texts, then assign each AI-edited text to one of n buckets according to the bucketing scheme described in Section C. All α values are reported in Table 3.

C MORE MODELING DETAILS

We use QLoRA to sweep both Llama and Mistral families of backbones between 3B and 24B parameters. We experiment with both a direct regression head and a N-way classification head with weighted-average decoding.

Determining thresholds for fully human and fully AI texts Some edits are too small to be detectable, such as adding a single comma, correcting a typo, etc. We choose a minimum threshold

of 0.03 for cosine distance threshold and 0.06 for soft n-grams in order to supervise it as AI-edited, chosen through manual inspection and validation of edits we would consider small enough such that the authorship is still entirely human.

Additionally, there are cases on the other end of the spectrum where AI was so pervasive in a text that it essentially rewrote the entire document and it became fully AI-generated. To measure the upper threshold where we would consider a text fully AI-generated, we analyzed the similarity metrics between the sources and their corresponding fully AI-generated synthetic mirrors. We selected thresholds that best separate fully AI-generated synthetic mirrors from the heaviest AI-edited text, which were 0.15 for cosine distance and 0.72 for soft n-grams.

C.1 REGRESSION FORMULATION

 Let s denote the raw similarity score and τ_{low} and τ_{high} be the low and high thresholds, respectively. We define the scaled similarity score as:

$$\tilde{s} = \begin{cases} 0.0 & \text{if } s \le \tau_{\text{low}} \\ 1.0 & \text{if } s \ge \tau_{\text{high}} \\ \frac{s - \tau_{\text{low}}}{\tau_{\text{high}} - \tau_{\text{low}}} & \text{otherwise} \end{cases}$$
 (1)

The regression model directly predicts the scaled similarity score \hat{s} using a mean squared error loss:

$$\mathcal{L}_{\text{MSE}} = \frac{1}{n} \sum_{i=1}^{n} (\tilde{s}_i - \hat{s}_i)^2 \tag{2}$$

where n is the number of training examples.

C.2 CLASSIFICATION FORMULATION

For the classification approach, we discretize the similarity scores into N buckets, where $N \in \{4,5,6\}$. Given minimum and maximum thresholds τ_{\min} and τ_{\max} , we define the bucket assignment function:

$$b(s) = \min\left(N - 1, \left| \frac{s - \tau_{\min}}{\tau_{\max} - \tau_{\min}} \cdot N \right| \right)$$
 (3)

The midpoint of bucket *j* is given by:

$$m_j = \tau_{\min} + \frac{(j+0.5) \cdot (\tau_{\max} - \tau_{\min})}{N} \tag{4}$$

We train the classification model using cross-entropy loss:

$$\mathcal{L}_{CE} = -\frac{1}{n} \sum_{i=1}^{n} \log p(b(s_i)|x_i)$$
 (5)

where $p(j|x_i)$ is the predicted probability for bucket j given input x_i .

During inference, we decode the final similarity score using a weighted average strategy:

$$\hat{s} = \sum_{j=0}^{N-1} p(j|x) \cdot m_j \tag{6}$$

where p(j|x) is the predicted probability for bucket j and m_j is the corresponding bucket midpoint.

C.3 ARCHITECTURE AND OPTIMIZATION

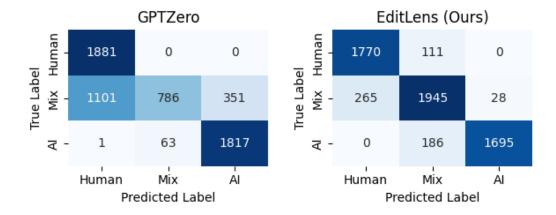
We train the model for 1 epoch with AdamW using a batch size of 24 and a constant learning rate of 3e-5. We initialize the model with pretrained weights from the base model and target *all* linear layers: self-attention QKV, output, and all linear layers in the MLP. We use a LayerNorm and single linear layer as the head for both prompt classification and edit heads and supervise both jointly in a multi-task learning routine. On 8 A100 GPUs, this takes approximately 8 hours for the largest model.

D TERNARY CLASSIFIER DECODING

GPTZero reports probabilities of three classes: "human", "AI", and "mixed," so we simply use argmax decoding to select the highest probability class. DetectAIve reports probabilities of four classes: "human", "AI", "AI Polished", and "AI humanized". We attempted to group "AI humanized" predictions with both the "AI" and the "AI Polished" categories for ternary classification, and found that grouping with "AI" produced a higher F1 score. Therefore, we group "AI Humanized" and "AI" into a single category for purposes of comparison.

E TERNARY CLASSIFICATION CONFUSION MATRICES

Analyzing the confusion matrix, we see that EDITLENS exhibits much stronger performance on the AI-edited text category than the strongest ternary classifier, GPTZero. While both EDITLENS and GPTZero are nearly perfect at distinguishing fully AI-generated text from fully human-written text, EDITLENS is the only model able to also consistently detect AI-edited text as a distinct category from fully human and fully AI.



F CORRELATION BETWEEN EDITLENS PREDICTIONS AND AI POLISH SIMILARITY METRICS

Model	r , semantic sim. \downarrow	r , Levenshtein \uparrow	r , Jaccard \uparrow
FastDetectGPT	-0.287	0.331	0.280
Binoculars	-0.291	0.326	0.274
Pangram	-0.491	0.615	0.556
EDITLENS (Ours)	-0.606	0.799	0.781

Table 4: Pearson correlation coefficients by model.

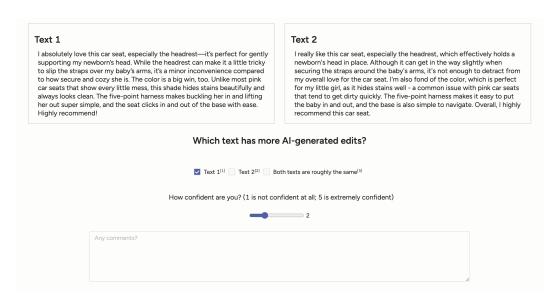


Figure 7: Data Annotator tasks were set up as above.

G DATA GENERATION MODELS

Checkpoint Name	API Provider	# Params	Open Source
claude-sonnet-4-20250514 Anthropic (2025)	Anthropic API	?	No
meta-llama/Llama-3.3-70B-Instruct-Turbo Meta (2024c)	Together.ai API	70B	Yes
gpt-4.1-2025-04-14 OpenAI (2025)	OpenAI API	?	No
gemini-2.5-flash Google (2025)	Gemini API	?	No

Table 5: Models used for dataset generation

H EMBEDDING MODELS

Checkpoint Name	# Params	Dim. Size
Ling-AI-Research/Ling-Embed-Mistral Choi et al. (2024)	7B	4096
sentence-transformers/all-MiniLM-L6-v2 Reimers et al. (2021)	22.7M	384

Table 6: Embedding models used for supervision

I BASE MODELS

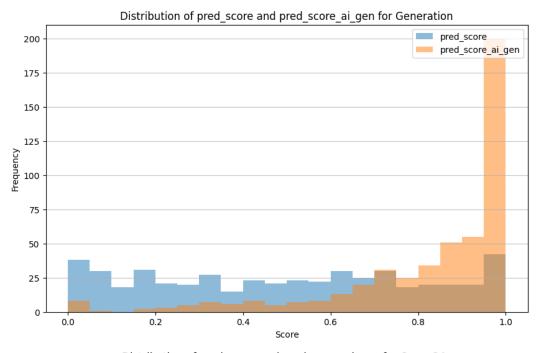
Checkpoint Name	# Params	Open Source
mistralai/Mistral-Nemo-Base-2407 Mistral (2024)	12B	Yes
mistralai/Mistral-Small-24B-Base-2501 Mistral (2025)	24B	Yes
meta-llama/Llama-3.1-8B Meta (2024a)	8B	Yes
meta-llama/Llama-3.2-3B Meta (2024b)	3B	Yes

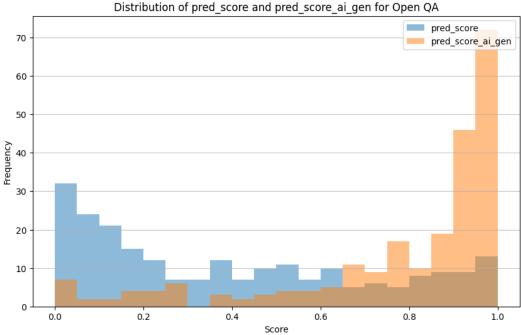
Table 7: Base models used for training

J More Results on Human-Edited AI Text

We focus on the human-edited AI versions of "Generation" and "OpenQA" categories of BEEMO, because the other categories, such as "Rewrite" and "Summarize", are already themselves AI-edited versions of human text, "Closed QA" the answers are so tightly constrained we would consider the answers to be human-written, and we would not consider the model outputs fully AI-generated. We also measure the correlation coefficient between our similarity metrics and the model scores. The intuition for this is that if the human edit is more invasive, we would expect the similarity metrics to increase, and the model score to decrease. As expected, we find a moderate negative correlation between our model's scores and the similarity, with -0.396 for cosine distance and -0.501 for soft n-grams.

Here we present the output distribution of EDITLENS for Beemo's Generation and OpenQA splits, on the fully AI-generated text (orange) and human-edited version (blue).





As is shown in the figures, the predicted score distribution moves significantly towards human-generated following editing, as expected.

K EDITING PROMPTS

Table 8: Full list of editing prompts by split with category and contributor.

TRAIN	
Editing Prompt	Contributor
Tone and Style Adjustments	
Write this in a way that a business person would get it	Human
Edit this to sound more polite	Human
Inject more personality and warmth into this text	Gemini 2.5 Pro
Adjust the tone to be more persuasive and convincing	Gemini 2.5 Pro
Make this sound more professional and authoritative	Gemini 2.5 Pro
Rewrite this to be more empathetic and understanding	Gemini 2.5 Pro
Make this sound more urgent and compelling	Gemini 2.5 Pro
Make this sound more objective and unbiased	Gemini 2.5 Pro
Adopt a more academic and scholarly tone	Gemini 2.5 Pro
Make this more direct and confrontational	Claude Sonnet
Make this more memorable and quotable	Claude Sonnet
Make this sound more diplomatic and tactful	Claude Sonnet
Make this more emotionally resonant	Claude Sonnet
Make this more formal	Claude Sonnet
Simplify for customers with no technical background	Claude Sonnet
Make this suitable for social media sharing	Claude Sonnet
Inject enthusiasm and energy into this writing	Claude Sonnet
Translate this for a teenage audience	Claude Sonnet
Soften the tone while maintaining the message	Claude Sonnet
Make this more relatable to the reader's experience	Claude Sonnet
Make this more inspiring and motivational	Claude Sonnet
Add gravitas and weight to this statement	Claude Sonnet
Adopt a more skeptical and questioning tone	Claude Sonnet
Make this more casual	Claude Sonnet
Adjust for a peer-reviewed academic journal	Claude Sonnet
Convert to a more analytical and logical approach	Claude Sonnet
Make this appropriate for C-suite executives	Claude Sonnet
Change this so it fits what a business person would want	ChatGPT 40
Make this sound more sure and strong	ChatGPT 40
Edit this for people who don't know the topic well	ChatGPT 40
Make this more direct and bold	ChatGPT 40
Make this sound fair and not take sides	ChatGPT 40
	ChatGPT 40
Make this sound more serious and important	ChatGPT 40
Make this sound more excited and energetic Make this easier for someone who doesn't know that much about it	
	ChatGPT 4o
Make this more accessible to a non-expert reader	ChatGPT 4o
Adapt this for readers with no prior background in the topic	ChatGPT 4o
Write this like you're talking to someone	ChatGPT 4o
Make this sound more doubtful and questioning	ChatGPT 4o
Adjust the voice to sound more academic.	ChatGPT 4o
Make this sound more relaxed and friendly	ChatGPT 4o
Write this in a way that top company leaders would like	ChatGPT 4o
Make this more formal and proper	ChatGPT 4o
Make this easier to remember and repeat	ChatGPT 4o
Tailor this message to suit a lay audience	ChatGPT 4o
Make this sound more exciting and well written	ChatGPT 4o
Make this sound more serious and proper	ChatGPT 4o
Use a smart and serious tone like in official stuff	ChatGPT 4o
Edit this to sound more urgent and important	ChatGPT 4o

Editing Prompt	Contributor
Make this more convincing and easier to understand	ChatGPT 4o
Change this so it's easy for a teen to read	ChatGPT 4o
Make this more logical and fact-based	ChatGPT 4o
Make the consequences feel more important	ChatGPT 4o
Make this sound uplifting and encouraging	ChatGPT 4o
Rewrite this to align with a formal tone	ChatGPT 4o
Make this more convincing and clear	ChatGPT 4o
Change this so a 5th grader can understand it	ChatGPT 4o
Take out hard words and explain them in a simple way	ChatGPT 4o
Fix this to make it more interesting	ChatGPT 4o
Change this to make it as strong as possible	ChatGPT 4o
Rewrite this to better suit a business audience	ChatGPT 4o
Use simpler words so anyone can understand this	ChatGPT 4o
Make this sound more like school writing	ChatGPT 4o
Make this sound nicer and more fun	ChatGPT 4o
Change this to sound more thoughtful	ChatGPT 4o
Make this funnier and more lighthearted	ChatGPT 4o
Use better words to make this sound smarter	ChatGPT 4o
Add some personality to this	ChatGPT 4o
Make this sound more like professional writing	ChatGPT 4o
-	
Adding Detail	
Make this more descriptive	Human
Make this more detailed	Human
Please add more details to make my argument better	Human
Add vivid imagery and sensory details to bring this to life	Claude Sonnet
Add backstory or context to enrich understanding	Claude Sonnet 4
Add depth and context to make this more comprehensive	Claude Sonnet 4
nclude specific measurements, colors, and physical characteristics	Claude Sonnet 4
Flesh out these ideas with supporting information	Claude Sonnet 4
Provide concrete examples to illustrate these points	Claude Sonnet 4
Add sensory details to make this more vivid	Claude Sonnet 4
Use more precise and colorful adjectives	Claude Sonnet 4
Add dialogue and quoted speech to make scenes more vivid	Claude Sonnet 4
Elaborate on the key points with concrete details	Claude Sonnet 4
Add storytelling elements to increase engagement	Claude Sonnet
Add descriptive metaphors and similes to enhance understanding	Claude Sonnet
Expand with real-world applications	Claude Sonnet
Use more evocative and powerful verbs	Claude Sonnet
Include expert opinions or research findings	Claude Sonnet 4
Incorporate specific brand names, locations, and proper nouns	Claude Sonnet
Paint a clearer picture with specific visual descriptions	Claude Sonnet
Use figurative language to make concepts more tangible	Claude Sonnet
Include personal anecdotes or case studies	Claude Sonnet 4
Give examples to help make this clearer	ChatGPT 4o
Edit this with clear examples to help explain this better	ChatGPT 4o
Add details that create a mood or feeling	ChatGPT 4o
Tell some of the story behind this to help understand it	ChatGPT 4o
Explain more about what this means and why it matters	ChatGPT 4o
Explain the main points using real examples	ChatGPT 4o
Add descriptions of sounds, smells, textures, and how things feel	ChatGPT 4o
Include exact sizes, colors, and what things look like	ChatGPT 4o
Add details about the setting and background	ChatGPT 4o
Add details that help readers see, hear, and feel what's happening	ChatGPT 4o
Add conversations and quotes to make scenes more real	ChatGPT 4o

Editing Prompt	Contributor
Develop this text further by explaining the implications	ChatGPT 4o
Help readers picture this more clearly with specific details	ChatGPT 4o
Describe how things look, sound, or feel more	ChatGPT 4o
Use more interesting and specific describing words	ChatGPT 4o
Add comparisons to help explain things better	ChatGPT 4o
Use real-life examples to show what you mean	ChatGPT 4o
Add background facts to back up these points	ChatGPT 4o
Add true stories or examples from real life	ChatGPT 4o
Use specific names of places, brands, and things	ChatGPT 4o
Add more ideas or facts that support what you're saying	ChatGPT 4o
Add details about the place and situation	ChatGPT 4o
Share what experts think or what research shows	ChatGPT 4o
Use real examples to explain this better	ChatGPT 4o
Use stronger, more exciting action words	ChatGPT 4o
Add illustrative examples to clarify these points	ChatGPT 4o
Fluency and Flow	
Rearrange this	Human
Can you make this sound fluent?	Human
Improve the transitions between the paragraphs	Gemini 2.5 Pro
Create a more effective and engaging opening	Gemini 2.5 Pro
Make this read like it was written by a native speaker	Claude Sonnet 4
Make this sound more conversational and engaging	Claude Sonnet 4
Improve the rhythm and readability of this writing	Claude Sonnet 4
Make the progression of ideas feel effortless	Claude Sonnet 4
Create smoother connections between these ideas	Claude Sonnet 4
Improve the natural rhythm of this text	Claude Sonnet 4
Smooth out the awkward phrasing in this passage	Claude Sonnet 4
Eliminate any choppy or awkward sentences	Claude Sonnet 4
Ensure the sentences transition smoothly from one idea to the next	ChatGPT 4o
Help the ideas move from one to the next easily	ChatGPT 40
Make the ideas connect better	ChatGPT 4o
Make the language flow more fluidly without sounding forced	ChatGPT 4o
Can you fix parts that sound choppy or off?	ChatGPT 4o
Make this writing smoother and better	ChatGPT 4o
Make the words fit together better	ChatGPT 4o
Make this easier and smoother to read	ChatGPT 4o
Help the ideas connect more smoothly	ChatGPT 4o
Make this sound like someone who speaks English well wrote it	ChatGPT 4o
Make the sentences flow better	ChatGPT 4o
Make the rhythm of the sentences better	ChatGPT 4o
Concision	
Rewrite to be more concise and powerful	Human
Clarify the main idea	Gemini 2.5 Pro
Remove any filler words	Gemini 2.5 Pro
Remove any jargon or technical terms and explain them in plain language	Gemini 2.5 Pro
Replace complex words with simpler alternatives	Gemini 2.5 Pro
Identify and eliminate any redundant phrases or words	Gemini 2.5 Pro
Make this more concrete and less abstract	Gemini 2.5 Pro
Simplify this text for a 5th-grade reading level	Gemini 2.5 Pro
Remove every unnecessary word and phrase	Claude Sonnet 4
Trim the fat without losing the muscle	Claude Sonnet 4
Make this more concise without losing important information	Claude Sonnet 4
Eliminate wordy expressions and redundancies	Claude Sonnet 4

Editing Prompt	Contributor
Tighten this writing by removing unnecessary words	Claude Sonnet
Remove all extra words and phrases	ChatGPT 4o
Trim this down while keeping the tone and meaning intact	ChatGPT 4o
Remove extra or repeated words	ChatGPT 4o
Use clearer words but keep the meaning	ChatGPT 4o
Get rid of long or confusing parts	ChatGPT 4o
Make this shorter and more direct	ChatGPT 4o
Take out anything extra but keep the good parts	ChatGPT 4o
Say this in fewer words but still make it strong	ChatGPT 4o
Take out words that aren't needed to make this better	ChatGPT 4o
Take out words that don't add anything	ChatGPT 4o
Cut this down but keep the same meaning and style	ChatGPT 4o
Structure and Organization	
Group related ideas together more effectively	Gemini 2.5 Pro
Ensure a clear introduction, body, and conclusion	Gemini 2.5 Pro
Arrange these points in order of importance	Claude Sonnet
Create better section breaks and headers	Claude Sonnet
Create a more compelling narrative arc	Claude Sonnet
Build toward a stronger climax or conclusion	Claude Sonnet
Reorganize this for better logical flow	Claude Sonnet
Use parallel structure to enhance readability	Claude Sonnet
Break this into clearer paragraphs with smooth transitions	Claude Sonnet
Rearrange this content for a clearer argument progression	ChatGPT 40
Put this in a better order	ChatGPT 4o
Put similar ideas together more clearly	ChatGPT 4o
Build up to a strong ending	ChatGPT 4o
Put this in an order that makes more sense	ChatGPT 40
Set this up in a way that's easier to read	ChatGPT 4o
Group ideas that go together and add connections	ChatGPT 4o
Move things around to make the main points stand out	ChatGPT 4o
Split this into paragraphs that connect better	ChatGPT 40
Make sentences match and sound good together	ChatGPT 4o
Put the most important stuff first	ChatGPT 4o
Organize this information in a more reader-friendly format	ChatGPT 4o
Tell the story in a more interesting way	ChatGPT 4o
General Improvement	2
Can you help my essay get a better grade?	Human
Rewrite this so it sounds good	Human
Make my essay better	Human
Make this essay look better	Human
Can you improve this?	Human
Make this an A paper	Human
Revise this to make it more engaging	Claude Sonnet
Enhance the overall effectiveness of this passage	Claude Sonnet
Refine this writing to make it more professional	Claude Sonnet
Enhance this text while maintaining the original meaning	Claude Sonnet
Optimize this text for maximum impact	Claude Sonnet
Strengthen this writing by improving word choice and structure	Claude Sonnet
Transform this into more compelling prose	Claude Sonnet
Polish this text for clarity and readability	Claude Sonnet
Upgrade the sophistication of this writing	Claude Sonnet
Make this work better overall	ChatGPT 4o
Make this writing more polished and effective	ChatGPT 4o

Editing Prompt	Contributor
Refine this text to improve its overall impact	ChatGPT 4o
Fix this but keep the same meaning	ChatGPT 4o
Fix this but keep the main point the same	ChatGPT 4o
Improve this passage while preserving its original intent	ChatGPT 4o
Paraphrasing	
Remake all of this in a different way	Human
Rewrite all of this in different words	Human
Recast these ideas in a different style	Claude Sonnet
Rephrase this text to avoid repetition	Claude Sonnet
Express these ideas using alternative phrasing	Claude Sonnet
Say the same thing but in a fresh way	Claude Sonnet
Present the same information from a fresh angle	Claude Sonnet
Reframe this argument using different terminology	Claude Sonnet
Restate this using different vocabulary and sentence structure	ChatGPT 40
Use easier words that mean the same thing	ChatGPT 40
Share this idea from a new point of view	ChatGPT 40
Say this in a new way	ChatGPT 40
Say this in a new way Use different words to say the same thing	ChatGPT 40 ChatGPT 40
Say this using different words and sentence types	ChatGPT 4o
Say this using different words and ideas	ChatGPT 4o
Say this in a new and interesting way	ChatGPT 4o
Paraphrase this to make it simpler and easier to understand	ChatGPT 4o
Clarity and Precision	
Can you fix the problems with my argument?	Human
Emphasize the key points	Human
Add precise measurements and timeframes	Claude Sonnet
Define any terms that might be unclear	Claude Sonnet
Eliminate any ambiguous or vague language	Claude Sonnet
Make the cause-and-effect relationships clearer	Claude Sonnet
Explain any hard words so people know what they mean	ChatGPT 4o
Fix parts that sound weird or hard to read	ChatGPT 4o
Say clearly who or what each word is talking about	ChatGPT 4o
Make this clear and easy to read	ChatGPT 4o
Make sure one idea leads to the next clearly	ChatGPT 4o
Make it obvious what causes what	ChatGPT 4o
Say this in a clear and simple way	ChatGPT 4o
Grammar and Mechanics	
Make my grammar sound better	Human
Fix any grammatical mistakes in this text	Gemini 2.5 Pro
Correct any grammar, punctuation, or spelling errors in this text	ChatGPT 4o
Use better words and fix how the sentences are written	ChatGPT 4o
VAL	
Tone and Style Adjustments	
Edit this into a blog post I can share online	Human
Rewrite this in a more conversational and approachable style	Gemini 2.5 Pro
Make this sound more confident and authoritative	Claude Sonnet
Adapt this for international readers	Claude Sonnet
Make this connect better with people's feelings	ChatGPT 40
Make this sound more like a friendly conversation	ChatGPT 4o
Make this easier for readers to relate to	ChatGPT 4o

Editing Prompt	Contributor
Make this sound more businesslike and serious	ChatGPT 4o
Paraphrasing	
Rewrite all of this	Human
Translate this into more accessible language	Claude Sonnet
Find alternative ways to express these concepts	Claude Sonnet
Rework this using varied sentence structures	Claude Sonnet
Change how the sentences are written	ChatGPT 4o
Say this in another way	ChatGPT 4o
Adding Detail	
Add atmospheric details to create mood and setting	Claude Sonnet
Add environmental and contextual descriptions	Claude Sonnet
Include contextual information to support these claims	ChatGPT 4o
Give more background so it's easier to understand	ChatGPT 4o
Use creative comparisons to make ideas clearer	ChatGPT 4o
Concision	
Rewrite this to be more direct and to the point	Gemini 2.5 Pro
Make this easier to understand	ChatGPT 4o
Use simpler language anyone can get	ChatGPT 4o
Structure and Organization	
Restructure this to emphasize the main points	Claude Sonnet
Add better breaks and section titles	ChatGPT 4o
Fluency and Flow	
Connect these thoughts more seamlessly	Claude Sonnet
Help the paragraphs connect better	ChatGPT 40
Grammar and Mechanics	
Proofread this for spelling and grammar errors	Human
General Improvement	
Elevate the quality of this writing	Claude Sonnet
Clarity and Precision	
Remove anything confusing or unclear	ChatGPT 4o
TEST	
Tone and Style Adjustments	
Lighten the tone and add a touch of humor	Gemini 2.5 Pro
Adjust the tone to be more friendly	Claude Sonnet
Increase the emotional stakes	Claude Sonnet
Make this sound more formal and school-like	ChatGPT 40
Make this nicer but keep the main point	ChatGPT 40
Change this so people from other countries can get it too	ChatGPT 40
Adding Detail	C.I.M.G. 1 10
Make this longer with more evidence	Human
Expand this text with more specific examples and details	Claude Sonnet
Show rather than tell by adding scene-setting details	Claude Sonnet
Include sounds, smells, textures, and other sensory elements	Claude Sonnet
Add a short story to make this more interesting	ChatGPT 40
Show how this works in real life	ChatGPT 40
	Charter 1 10
Concision Simplify this text	Human

Editing Prompt	Contributor
Make this more specific	Human
Reduce wordiness while amplifying impact	Claude Sonnet 4
Say this in fewer words without losing meaning	ChatGPT 4o
Edit this to be punchier and more direct	ChatGPT 4o
Fluency and Flow	
Make this text flow more naturally	Claude Sonnet 4
Improve how the sentences sound together	ChatGPT 4o
Refine the pacing and cadence of this paragraph	ChatGPT 4o
Write a better and more interesting beginning	ChatGPT 4o
Clarity and Precision	
Write this in a way that my teacher would get it	Human
Can you make my paper more persuasive?	Human
Make the main idea clearer	ChatGPT 4o
Improve this to make it stronger and clearer	ChatGPT 4o
Paraphrasing	
Paraphrase this	Human
Rewrite this in different words while keeping the same meaning	Claude Sonnet 4
Reword this to improve clarity while keeping the meaning	ChatGPT 4o
Structure and Organization	
Make sure there's a beginning, middle, and end	ChatGPT 4o
Group related ideas and use transitions to improve structure	ChatGPT 4o
Grammar and Mechanics	
Can you fix any spelling, grammar, or punctuation issues?	Human

	Train	Val	Test	All
Categories				
Tone and Style Adjustments	28.5% (69)	30.0% (9)	19.4% (6)	27.7% (84)
Adding Detail	19.8% (48)	16.7% (5)	19.4% (6)	19.5% (59)
Concision	9.9% (24)	10.0% (3)	16.1% (5)	10.6% (32)
Fluency and Flow	9.9% (24)	6.7% (2)	12.9% (4)	9.9% (30)
Paraphrasing	7.0% (17)	20.0% (6)	9.7% (3)	8.6% (26)
Structure and Organization	9.1% (22)	6.7% (2)	6.5% (2)	8.6% (26)
General Improvement	8.7% (21)	3.3% (1)	0.0% (0)	7.3% (22)
Clarity and Precision	5.4% (13)	3.3% (1)	12.9% (4)	5.9% (18)
Grammar and Mechanics	1.7% (4)	3.3% (1)	3.2% (1)	2.0% (6)
Contributors				
ChatGPT 4o	52.9% (128)	50.0% (15)	48.4% (15)	52.1% (158)
Claude Sonnet 4	31.4% (76)	33.3% (10)	25.8% (8)	31.0% (94)
Human	7.9% (19)	10.0% (3)	22.6% (7)	9.6% (29)
Gemini 2.5 Pro	7.9% (19)	6.7% (2)	3.2% (1)	7.3% (22)
TOTAL COUNTS	242	30	31	303

Table 9: Distribution of prompt categories and contributors across Train, Val, and Test splits shown as percentages with raw counts in parentheses.

	Unique source texts
Train	22690
Validation	1391
Test	3809

Table 10: Count of unique source texts across Train, Val, and Test datasets.

	Human	AI-edited	AI-generated	Total
Train	18810	22381	18809	60000
Validation	753	894	753	2400
Test	1881	2238	1881	6000

Table 11: Distribution of examples by label across Train, Val, and Test datasets.

		Claude Sonnet 4	GPT-4.1	Gemini 2.5 Flash
Train	AI-edited	7528	7521	7332
Train	AI-generated	6205	6243	6361
Validation	AI-edited	298	298	298
Validation	AI-generated	256	235	262
Test	AI-edited	739	795	704
Test	AI-generated	627	627	627

Table 12: Distribution of source LLM for AI-edited and AI-generated examples, across Train, Val, and Test datasets.

	Train	Validation	Test
Tone and Style Adjustments	6477	301	459
Adding Detail	4142	116	426
Fluency and Flow	2302	57	301
Structure and Organization	2110	71	176
Concision	2037	76	301
Paraphrasing	1679	176	229
Clarity and Precision	1229	33	290
Grammar and Mechanics	400	27	70
General Improvement	2005	37	0

Table 13: Composition of the AI-edited dataset, by split

		Mean Word Count	Min Word Count	Max Word Count
Train	Human	330.45	75	799
Train	AI-edited	328.11	75	799
Train	AI-generated	302.88	76	799
Validation	Human	331.02	75	798
Validation	AI-edited	324.29	75	799
Validation	AI-generated	299.81	79	797
Test	Human	241.47	75	799
Test	AI-edited	258.37	76	798
Test	AI-generated	218.05	76	792

Table 14: Word count statistics across splits for Human, AI-edited, and AI-generated texts in the dataset.

Table 15: Ternary Classification (Soft N-Grams)

Model	# buckets	# params	Acc.	Macro F1	Human F1	AI F1	AI-edited F1
Mistral-Small-24B-Base-2501	4	24B	0.897	0.899	0.896	0.941	0.861
Mistral-Small-24B-Base-2501	5	24B	0.873	0.876	0.913	0.875	0.840
Mistral-Small-24B-Base-2501	6	24B	0.882	0.886	0.886	0.918	0.853
Mistral-Small-24B-Base-2501	Regression	24B	0.861	0.865	0.880	0.887	0.828
Mistral-Nemo-Base-2407	4	12B	0.896	0.899	0.895	0.941	0.860
Llama-3.1-8B	4	8B	0.895	0.898	0.895	0.942	0.856
Llama-3.2-3B	4	3B	0.858	0.861	0.888	0.876	0.820

L LLM USAGE STATEMENT

Large Language Models (LLMs) were used in the experiments for the paper as described, to assist in writing the code to run the experiments, brainstorm the formalization of the task, assist in generating the figures for the paper, assist with LaTeX formatting, and review the paper to help the authors with constructive feedback. The authors did **not** use LLMs directly in the original writing of the manuscript, but did use LLMs to help with wording and phrasing in some sections. The authors take full responsibility for the factuality and originality of the content in this manuscript.

M GRAMMARLY SUPERVISION SCORES

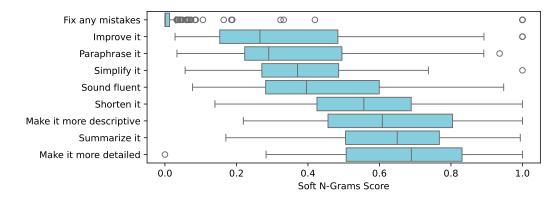


Figure 8: Soft N-Grams Score vs. Grammarly Prompt

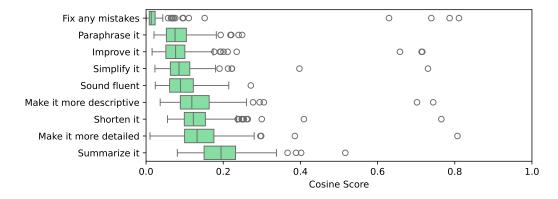


Figure 9: Cosine Similarity vs. Grammarly Prompt

Table 16: Ternary Classification (Cosine Similarity)

Model	# buckets	# params	Acc.	Macro F1	Human F1	AI F1	AI-edited F1
Mistral-Small-24B-Base-2501 (selected)	4	24B	0.902	0.904	0.904	0.941	0.868
Mistral-Small-24B-Base-2501	5	24B	0.874	0.878	0.880	0.912	0.841
Mistral-Small-24B-Base-2501	6	24B	0.855	0.859	0.872	0.882	0.822
Mistral-Small-24B-Base-2501	Regression	24B	0.857	0.861	0.883	0.879	0.820
Mistral-Nemo-Base-2407	4	12B	0.889	0.896	0.900	0.939	0.845
Llama-3.1-8B	4	8B	0.861	0.865	0.877	0.899	0.819
Llama-3.2-3B	4	3B	0.852	0.855	0.883	0.878	0.804

Table 17: Binary Classification (Soft N-Grams)

Model	# buckets	# params	Human vs Rest Acc.	Human vs Rest F1	AI vs Rest Acc.	AI vs Rest F1
Mistral-Small-24B-Base-2501	4	24B	93.967	95.542	94.333	90.151
Mistral-Small-24B-Base-2501	5	24B	94.400	95.857	92.983	87.533
Mistral-Small-24B-Base-2501	6	24B	93.033	94.969	95.233	91.833
Mistral-Small-24B-Base-2501	Regression	24B	92.683	94.664	93.433	88.419
Mistral-Nemo-Base-2407	4	12B	93.267	94.993	96.467	94.147
Llama-3.1-8B	4	8B	93.033	94.780	96.483	94.157
Llama-3.2-3B	4	3B	92.783	94.664	93.367	88.322

N PROMPT CLASSIFICATION RESULTS

We train EDITLENS with an additional classification head for the task of **edit prompt classification**, where the model is asked to predict the type of editing prompt that was used to generate the input text. On our test set, EDITLENS achieves an accuracy of 41.2% and a weighted F1 score of 41.9% on the 9-way edit prompt classification task. From the confusion matrix shown in Figure ??, we see that certain types of prompts are easier to detect than others—for instance, the model frequently classifies "Adding detail" prompts correctly, but struggles to classify "Grammar and mechanics" prompts, often classifying them as "Concision" prompts instead. We surmise that overlap between editing prompt categories (for example, the "Create a more compelling narrative arc" is labeled as "Structure and Organization" but could just as well be considered a "Fluency and Flow" prompt), the varying "General Improvement" prompts, and the small size of some editing prompt categories in the test set are hindering model training. We look to improve on the prompt classification task, in terms of formulation, data, and model performance in future work.

Table 18: Binary Classification (Cosine Similarity)

# buckets	# params	Human vs Rest Acc.	Human vs Rest F1	AI vs Rest Acc.	AI vs Rest F1
4	24B	93.783	95.377	96.433	94.062
5	24B	92.583	94.608	94.867	91.170
6	24B	92.167	94.332	93.317	88.202
Regression	24B	92.583	94.534	93.183	87.867
4	12B	93.475	94.568	96.448	93.847
4	8B	92.017	94.055	94.183	89.869
4	3B	92.133	94.046	93.083	87.776
	4 5 6 Regression 4 4	4 24B 5 24B 6 24B Regression 24B 4 12B 4 8B	4 24B 93.783 5 24B 92.583 6 24B 92.167 Regression 24B 92.583 4 12B 93.475 4 8B 92.017	4 24B 93.783 95.377 5 24B 92.583 94.608 6 24B 92.167 94.332 Regression 24B 92.583 94.534 4 12B 93.475 94.568 4 8B 92.017 94.055	4 24B 93.783 95.377 96.433 5 24B 92.583 94.608 94.867 6 24B 92.167 94.332 93.317 Regression 24B 92.583 94.534 93.183 4 12B 93.475 94.568 96.448 4 8B 92.017 94.055 94.183

