



Master's Thesis

Enhancing Paraphrase Type Generation: The Impact of DPO and RLHF Evaluated with Human-Ranked Data

Master's Thesis for a degree in Applied Data Science (M.Sc.)
Chair for Scientific Information Analytics
University of Göttingen

Submission date: December 24st, 2024

By: Christopher Lee Lübbers

From: Göttingen

Matriculation number: 20676955

First supervisor: Prof. Dr. Bela Gipp

Second supervisor: Dr. Terry Lima Ruas

Supervisor: Dominik Meier

University of Göttingen
Chair for Scientific Information Analytics

CONTENTS

1	Introduction	1
2	Related Work	3
3	Methodology	4
3.1	Research Questions and Problem Definition	4
3.2	Paraphrase Type Generation	4
3.3	Paraphrase Type Detection	6
4	Results	7
4.1	Effect of DPO on Paraphrase Type Generation Accuracy	7
4.2	Human Preferences for DPO-Generated Paraphrases	9
4.3	Automated Metrics vs. Human Preferences	9
4.4	Paraphrase Type Detection with a Novel Model	9
4.5	Impact of DPO on Broader NLP Tasks	11
5	Conclusion	12
	Acknowledgments	13
	References	13
A	Technical Information	15
B	Datasets	15
C	Models	16
D	Evaluation	17

I hereby declare that I have written this thesis independently without any help from others and without the use of documents or aids other than those stated. I have mentioned all used sources and cited them correctly according to established academic citation rules.

Göttingen, December 24st, 2024

Enhancing Paraphrase Type Generation: The Impact of DPO and RLHF Evaluated with Human-Ranked Data

CHRISTOPHER LEE LÜBBERS, University of Göttingen, Germany

ABSTRACT

Paraphrasing re-expresses meaning to improve text simplification, machine translation, and question-answering. Specific paraphrase types enable accurate semantic analysis and robust language models. However, current paraphrase-type generation methods do not fully align with human preferences because they rely on automated metrics and limited human-annotated training data. These shortcomings obscure crucial aspects of semantic fidelity and linguistic transformations defined by paraphrase types.

We address this gap by leveraging a human-ranked paraphrase-type dataset and integrating Direct Preference Optimization (DPO) to align model outputs directly with human judgments. DPO-based training increases paraphrase-type generation accuracy by 3 percentage points over a supervised baseline and raises human preference ratings by 7 percentage points. Our newly created human-annotated dataset supports more rigorous future evaluations. We also provide a paraphrase-type detection model that achieves F1 scores of 0.91 for addition/deletion, 0.78 for same polarity substitution and 0.70 for punctuation changes.

These findings demonstrate that preference data and DPO training produce more reliable, semantically accurate paraphrases, enabling downstream applications such as improved summarization and more robust question-answering. The PTD model surpasses automated metrics and provides a more reliable framework for evaluating paraphrase quality. This approach advances paraphrase-type research toward richer, user-aligned language generation and establishes a stronger foundation for future evaluations grounded in human-centric criteria.

1 INTRODUCTION

Paraphrasing transforms a sentence’s form while preserving its core meaning [2], enabling more effective machine translation, summarization, and question-answering. For example, ‘The scientist explained the results clearly’ can be paraphrased as ‘The scientist clarified the results.’ By generating faithful yet varied reformulations, models become more robust, better able to adapt to shifting domains, and more attuned to user expectations in interactive systems [13, 18].

Within paraphrasing research, atomic paraphrase types (APT) define fine-grained linguistic transformations [36]. Addition (‘The scientist *thoroughly* explained the results clearly’), or same polarity substitution, where a synonym preserves meaning (‘The scientist *described* the results clearly’) exemplify transformations that specify precise linguistic modifications rather than broad similarity. Building on APT, Paraphrase Type Generation (PTG), and Paraphrase Type Detection (PTD) [37] drive the field forward by focusing model outputs and analyses on specific transformations. Despite the importance of PTG and PTD, models rarely produce or detect paraphrases that consistently convey intended meanings while following subtle linguistic shifts.

This deficiency arises from a fundamental gap: a lack of high-quality, human-ranked data that specifies which paraphrases users find preferable and why. Widely used metrics like BLEU and ROUGE [19, 25] focus on lexical overlap rather than semantic fidelity and cannot detect if paraphrasing was applied correctly [31]. Creating human-ranked paraphrase datasets requires extensive effort [3], and the scarcity of such data restricts progress. Without sufficient human annotation that ranks paraphrases on subtle qualities, models struggle to internalize these preferences, especially for complex transformations [22, 23].

This research addresses that gap by leveraging human-ranked data to guide paraphrase generation and detection. We¹ employ Reinforcement Learning from Human Feedback (RLHF), Direct Preference Optimization

¹‘We’ is used to align with academic conventions, reflecting collective research efforts, though this thesis is the work of a sole author.

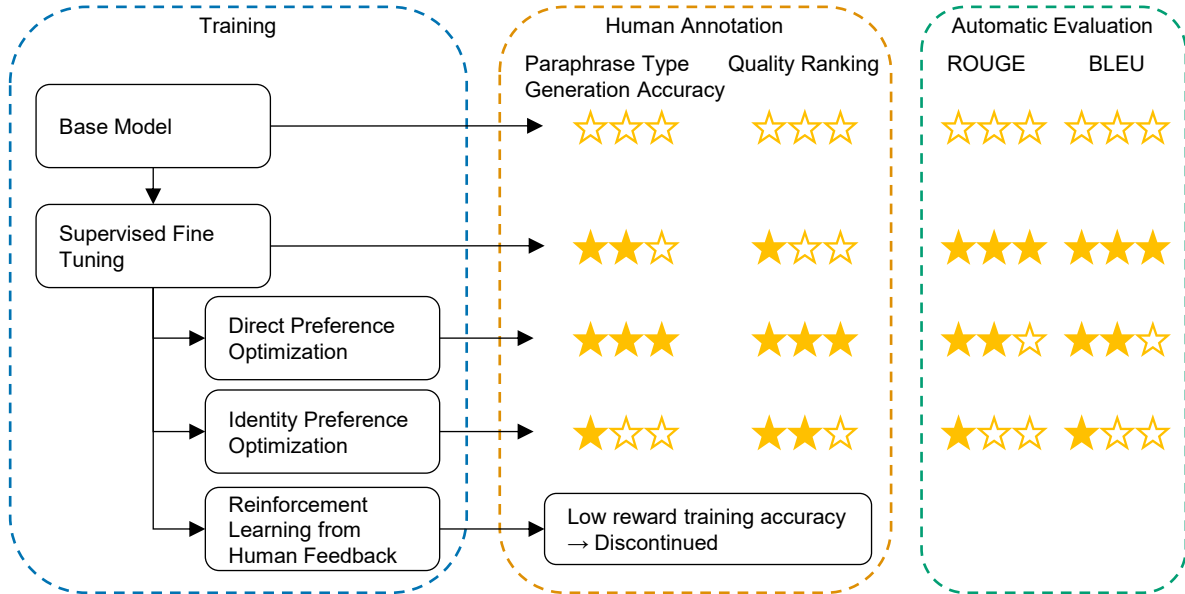


Fig. 1. Comparison of direct preference optimization (DPO), identity preference optimization (IPO), and reinforcement learning from human feedback (RLHF) against supervised fine-tuning and a baseline. We used Annotated Paraphrase Types (APTY) and Extended Typology Paraphrase Corpus (ETPC) datasets. The diagram outlines paraphrase-type generation steps, evaluation metrics, and the incorporation of human feedback. Stars indicate increasing alignment quality, and the discontinued pathway marks where RLHF performed suboptimally

(DPO), and Identity Preference Optimization (IPO) [1] to generate paraphrases that incorporate target APT-defined transformations more faithfully. Working with the Annotated Paraphrase Types (APTY) dataset [22], we guide models to produce paraphrases that exhibit the desired APT-defined transformations and reflect authentic human preferences. Figure 1 illustrates the approach, comparing optimization techniques, showing data flow, and incorporating human feedback. Our evaluations confirm that DPO-based models yield more accurate and user-preferred paraphrases, reveal the shortcomings of traditional evaluation metrics, and highlight the value of human-centric training for improved language understanding.

Key contributions include:

- (1) Enhanced paraphrase-type generation accuracy (section 4.1): DPO training on APTY increases human-annotated accuracy by 3 % over a supervised baseline, aligning outputs with nuanced linguistic transformations.
- (2) Improved user-aligned quality (section 4.2): Human evaluators favor these improved outputs 7 % more than baseline paraphrases, underscoring enhanced semantic fidelity and stylistic appropriateness.
- (3) A new human-ranked dataset: The dataset we produce enables a more rigorous, fine-grained evaluation of paraphrase quality and paves the way for future research.
- (4) Exposing metric limitations (section 4.3): Weak correlations (Spearman’s $r < 0.3$) between automated metrics and human rankings motivate the development of richer evaluation frameworks.
- (5) Improved paraphrase-type detection (section 4.4): Our PTD model achieves F1 scores of 0.91 on addition/deletion, 0.78 on same polarity substitution, and 0.70 for punctuation changes, enabling more granular assessments.

- (6) Improved reasoning (section 4.5): PTG boosts multistep soft reasoning (MuSR) task performance by 38 %, demonstrating broader benefits for language generation and reasoning tasks.

These advances provide a framework that better aligns PTG and PTD with human needs, benefiting translation, summarization, and conversational AI applications.

2 RELATED WORK

Paraphrase generation aims to produce meaning-preserving reformulations of text, yet state-of-the-art methods struggle to reflect human-defined preferences for nuanced linguistic transformations. Multiple studies introduced typologies and evaluation methods, yet most approaches still fail to incorporate human preferences at a granular level.

Vila et al. [36] established the APT taxonomy, identifying transformations like addition, deletion, and substitution. The APT typology provides foundational insight into linguistic variation in paraphrases, surpassing generic similarity-based approaches. Building on APT, Wahle et al. [37] defined PTG and PTD tasks. They demonstrated that models could generate and identify these APT-defined transformations but found that outputs did not align with human judgments.

This mismatch underscored the need to incorporate user-centric criteria into model training. To address this shortfall, Meier et al. [22] introduced the APTY-ranked dataset, which included human rankings of paraphrases. They discovered that large models like ChatGPT performed well on simpler paraphrase types but struggled with complex transformations. Their findings revealed the critical role of direct human feedback in guiding models toward more nuanced, preference-aligned outputs. This reveals a critical limitation: models generate typologically diverse paraphrases that fail to align with human preferences. These alignment gaps motivate our research, aiming to incorporate human rankings into model training effectively.

Another major challenge involves evaluating paraphrase quality. Traditional measures such as BLEU [25] and ROUGE [19] rely on lexical overlap, failing to capture semantic fidelity or detect actual paraphrasing. Shen et al. [31] found that reference-free metrics surpass reference-based ones. They propose ParaScore, a metric that explicitly models lexical divergence but weakly correlates with human judgments. Similarly, Zhou et al. [39] show that under-representing paraphrase types and using single references reduce evaluation reliability. Oh et al. [24] showed that multiple pseudo-references improved approximations of human preferences but still lacked an entirely human-centered framework.

Pre-trained models such as BERT [4] and GPT [27] improved paraphrase fluency, as Natsir et al. [23] found, yet they optimized surface-level similarities rather than deep semantic fidelity. Huang et al. [11] introduced ParaAMR to enhance paraphrase syntactic diversity but did not incorporate human preference signals. Although reinforcement learning emerged to encode human evaluations, its application to fine-grained paraphrase types remains limited. Prior efforts focused on fluency and syntax [18], not human-judged transformations. Rafailov et al. [28] proposed DPO to align model outputs with human rankings without complex reward models. [1] modified the loss function of DPO to counter overfitting. DPO and IPO allow optimizing for user preferences at finer levels, addressing PTG gaps.

These gaps, including typology without human alignment, unreliable metrics, and the absence of fine-grained optimization, underscore the need for a new framework. Unlike previous works, our approach combines DPO with APTs to generate paraphrases aligned with nuanced human preferences. This framework aligns with human preferences and supplies robust evaluation and fine-grained optimization techniques for practical applications.

3 METHODOLOGY

This methodology addresses the stated research questions by proposing a formal problem setting, detailing datasets, presenting step-by-step procedures, and describing evaluation protocols. We integrate human-ranked data, DPO, and a PTD model to enhance the generation and evaluation of paraphrase types.

3.1 Research Questions and Problem Definition

We consider three key research questions: (RQ1) Can integrating human-ranked data with DPO improve PTG accuracy and user preference compared to baseline and reward-based approaches? (RQ2) Can a PTD model verify the presence and correctness of fine-grained transformations, offering a more nuanced evaluation than traditional metrics? (RQ3) Do improvements in PTG and PTD, driven by human-centric optimization, enhance performance on broader NLP tasks that require semantic fidelity?

We use PTG and PTD definitions from Wahle et al. [37]. We define PTG as follows: Given an original sentence x and a target transformation $l_i \in L$, where L is a set of APT, produce a paraphrase \tilde{x} that preserves the meaning of x while applying transformation l_i .

For PTD, given a pair (x, \tilde{x}) , the task is to identify which APT categories l_i were applied to x . PTD enables reference-free evaluation of transformations and provides deeper insight into paraphrase quality.

3.2 Paraphrase Type Generation

The PTG pipeline involved three steps, as illustrated in fig. 2. First, we fine-tuned the base model with supervised fine-tuning (SFT) on the Extended Typology Paraphrase Corpus (ETPC) [14]. Second, we further trained the fine-tuned model on APTY-ranked using reward training, DPO, and IPO. Finally, we evaluated the resulting models with ROUGE, BLEU, and human annotations on ETPC examples.

3.2.1 Datasets. We use datasets in the English language. ETPC contains 5,800 paraphrase examples categorized by APT. With extensive coverage and prior validation [37], ETPC ensures a solid typological foundation during SFT and prepares the model for integrating human preferences. We employed a 70-30 training-test split and used the same prompt as Wahle et al. [37]:

Given the following sentence, generate a paraphrase with the following type.
Sentence: ['original']
Paraphrase Types: ['APT'].
Answer:

The preprocessed APTY-ranked dataset, consisting of 333 human-ranked paraphrase examples, provides fine-grained human rankings, addressing ETPC’s limitations. APTY’s ranked preferences introduce a human-centric dimension, representing a key innovation over prior work that relied on static reference-based comparisons. As the only dataset with human-ranked paraphrase types, APTY-ranked offers a novel perspective on the quality and diversity of paraphrase generation. We preprocessed APTY-ranked by fixing encoding issues, stripping whitespace, normalizing punctuation, and applying an 80-20 split stratified by paraphrase types. We used the same prompt as in ETPC (section 3.2.1). APTY-ranked complements ETPC by incorporating human judgments into training.

3.2.2 Models and Training. We chose LLaMA-3.1-8B [6] and LLaMA-2-7B [35] models for open-source availability. These models scale for large-scale NLP tasks, offer parameter efficiency, and enable resource-efficient training for diverse NLP applications. We employed Low-Rank Adaptation (LoRA) [10] layers on Llama models to optimize resource usage. This approach reduces computation and allows focused adaptation, supporting reproducibility

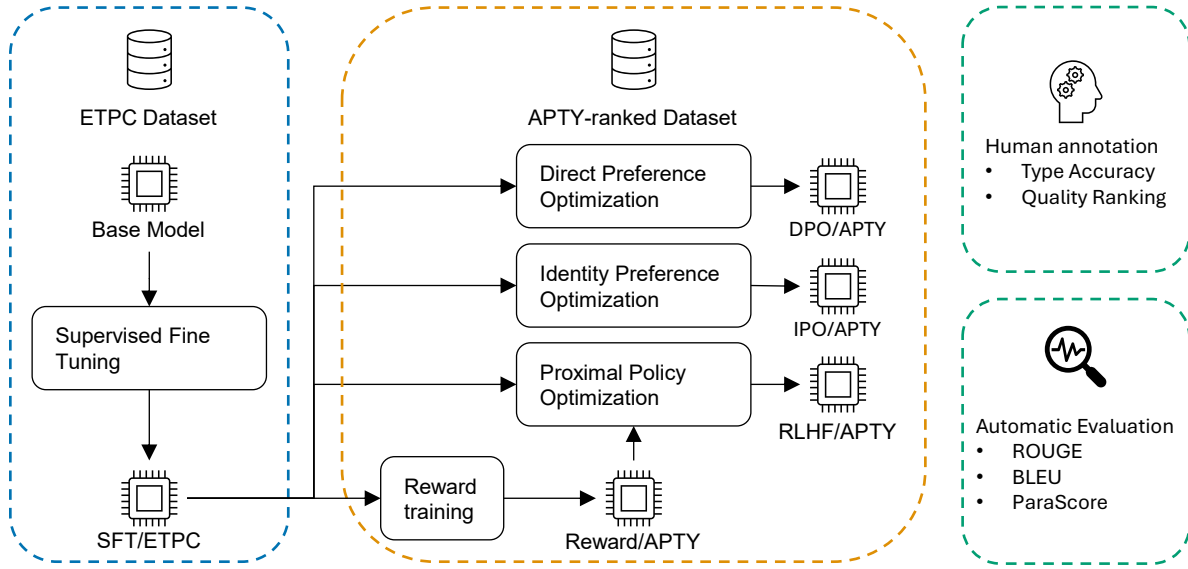


Fig. 2. Paraphrase Type Generation (PTG) workflow. The model is first fine-tuned on the Extended Typology Paraphrase Corpus (ETPC) to learn atomic paraphrase transformations, then further refined with the Annotated Paraphrase Type (APTY) dataset using reward training, direct preference optimization (DPO), and identity preference optimization (IPO). Final evaluation employs ROUGE, BLEU, and human annotations to confirm accurate, preference-aligned paraphrases.

and broader adoption. However, LoRA fine-tuning may limit generalization for intricate reasoning [32]. We also included BART [17] for ParaScore comparisons since ParaScore supports limited models.

The training involved multiple stages.

- (1) SFT on ETPC: We trained BART on ETPC to establish a baseline understanding of paraphrase types. Wahle et al. [37] provided a Llama-3.1-8B model fine-tuned on ETPC.
- (2) After SFT, we refined the models using APTY-ranked. We explored three optimization techniques:
 - **Reward Training:** Uses a reward model to guide PTG toward higher-quality outputs, depending on a predefined reward structure. This model would serve in Proximal Policy Optimization [30], a standard approach for transformer-based RLHF.
 - **DPO aligns outputs with human rankings, bypassing a separate reward model.** fig. 3 shows an overview of the method. We manually tuned hyperparameters to maximize reward margins, representing the mean difference between chosen and rejected rewards. Besides maximizing reward margins, we balanced accuracy and stable loss stability. A trial table is available in the repository [20].
 - **IPO counters overfitting by adjusting the loss function of DPO.** For the IPO model, we also tuned hyperparameters to maximize reward margins and maintain accuracy and stable loss [20].

3.2.3 Evaluation. We conducted a thorough evaluation combining automatic metrics and human judgments. This dual evaluation bridges gaps between automatic metrics and human preferences. We computed ROUGE-1, ROUGE-2, ROUGE-L, and BLEU from ETPC references. Due to tool constraints, we applied ParaScore only to BART outputs. While these metrics provide baselines, their known misalignment with human judgment demands more human-centric approaches.




Fig. 3. Direct Preference Optimization (DPO) framework, adapted from Rafailov et al. [28]. DPO aligns model outputs with human rankings by directly optimizing for preferred responses, eliminating the need for a separate reward model or reinforcement learning. This streamlined approach results in paraphrases more closely matching user-defined quality criteria.

Two human reviewers evaluated outputs for semantic fidelity and type adherence, providing insights into user alignment and complementing metrics. The reviewers were one master’s student in Applied Data Science and one doctoral student. Multiple annotators mitigated subjectivity and bias. This step validates whether metric improvements correspond to human preferences. The reviewers received a sentence with an APT and model-generated paraphrases. Annotators decided if the specified APT was correctly and exclusively applied. They ranked valid paraphrases on a scale of 1 (best) to 4 (worst). If the generated sentence was not paraphrased or the specified type was absent, annotators assigned ranked 4. All valid paraphrases were ranked from 1 downward. We used Pearson correlation [26] to evaluate the correlation between rank annotations and metrics. Our annotations schema assigns ‘4’ to all incorrect paraphrases, creating a rigid valid-invalid boundary. This discontinuity mismatches the discrete human annotation scale and the continuous ROUGE/BLEU scores, reducing correlation. We applied a logistic transformation to smooth ranks and align them with continuous metrics. We used

$$TransformedScore = \frac{1}{1 + \exp(OriginalScore - 2.5)}$$

with 2.5 as the annotation midpoint. This transformation introduces score continuity.

We compared broader performance on Open LLM Leaderboard v2 [7] tasks: Instruction-Following Evaluation (IFEval) [40], Big Bench Hard (BBH) [34], MATH Lvl 5 [9], Graduate-Level Google-Proof Q&A Benchmark (GPQA) [29], Multistep Soft Reasoning (MuSR) [33], Massive Multitask Language Understanding - Professional (MMLU-Pro) [38].

3.3 Paraphrase Type Detection

Evaluations often rely on references, which may be limited or unavailable. PTD verifies which transformations a given paraphrase applies, enabling reference-free evaluation and complementing PTG. If a model excels at generating certain paraphrase types, can a detection model confirm these transformations? Figure 4 shows the PTD pipeline.

3.3.1 Datasets. For PTD, we first trained a model on Quora Question Pairs (QQP) [12] to learn paraphrase vs. non-paraphrase distinctions. QQP provides a widely used benchmark, offering a solid foundation before fine-tuning on more nuanced data. We used the train split, which contains 364k rows, for training, and the validation split, which contains 40.4k rows, for testing.

Next, we fine-tuned on a filtered ETPC set, focusing on the top 10 paraphrase types: addition/deletion, change of order, derivational changes, inflectional changes, punctuation changes, same polarity substitution (contextual), semantic-based, spelling changes, subordination and nesting changes, synthetic/analytic substitution (table 4). This prevents performance loss for rare types. It aligns with APTY’s focus on the top 10 types. We tracked




Fig. 4. Paraphrase Type Detection pipeline. Initial training on the Quora Question Pairs (QQP) dataset teaches paraphrase recognition, followed by fine-tuning on a filtered ETPC subset to identify specific atomic transformations. Stratified sampling and weighted losses address class imbalances, ensuring robust detection of diverse paraphrase types.

each example’s paraphrase types, enabling type-based organization and count maintenance. A balanced split prioritized rare paraphrase types, ensuring representation in training and testing. We used an 80-20 training-testing split, preserving the distribution of paraphrase types. Handling duplicates proved critical, as examples could belong to multiple paraphrase types. We used sets to avoid repeated multiple subset assignments. A weighted BCEWithLogitsLoss addresses imbalances, giving rare types receive proper attention.

3.3.2 Model and Training. We selected DeBERTa [8] for its advanced architecture and strong classification performance. Our two-stage fine-tuning, binary classification on QQP, and multilabel classification on ETPC are innovative approaches to bridging basic paraphrase recognition with fine-grained type detection. We optimized class weights, learning rate, weight decay, and batch size with Optuna to guarantee robust results. Meticulous tuning ensures accurate type detection, supporting targeted model evaluations.

3.3.3 Evaluation. We used Macro-F1 scores on ETPC to measure balanced performance across all paraphrase types. Macro-F1 prevents the dominance of common types. Weighted loss and type-stratified sampling address type imbalances, providing equitable detection performance. A validation set and systematic hyperparameter tuning ensured reliable improvements.

4 RESULTS

We tested the hypothesis that incorporating human-ranked preferences via DPO improves PTG accuracy for specific transformations, increases human satisfaction with the outputs, reveals limitations of current automated metrics, and enhances performance in complex reasoning tasks. We conducted a series of experiments to validate each component of this hypothesis, using multiple models, human annotations, and both established and novel evaluation approaches. The code is on GitHub [20] and trained models on Huggingface [21].

4.1 Effect of DPO on Paraphrase Type Generation Accuracy

We analyzed the accuracy of trained models in generating specific APT. Preliminary Llama-2-7B experiments showed a significant accuracy increase. While Llama-2-7B and SFT models achieved 12 % accuracy, DPO reached 35 % (fig. 10). We continued to evaluate Llama-3.1-8B (baseline), SFT/ETPC, DPO/APTY, and IPO/APTY. Because the reward model accuracy was only 49 %, we discontinued RLHF, highlighting its limitations. Each model generated single-type paraphrases on 260 base sentences. This produced 1040 human-annotated paraphrases covering the top 10 APTs. The base sentences are available on Github². We defined accuracy as the proportion of generated paraphrases correctly exhibiting the target transformation without semantic loss.

²https://github.com/worta/generate_apt_paraphrases




Fig. 5. Accuracy of paraphrase-type generation across methods. The DPO/APTY model achieves 57 % accuracy, surpassing both the baseline Llama-3.1-8B (8 %) and a supervised fine-tuning (SFT/ETPC) approach. Accuracy reflects the proportion of paraphrases that apply the targeted transformation without semantic loss, validated by human annotations. Certain transformations, such as punctuation and word-order changes, remain challenging for all models. The base model Llama-3.1-8B could not generate a correct paraphrase for some paraphrase types, resulting in missing bars.

DPO/APTY model achieved 57 % mean accuracy (SD=9), outperforming Llama-3.1-8B (8 %, SD=14), SFT/ETPC (54 %, SD=8) and IPO/APTY (52 %, SD=10). A one-way ANOVA confirmed significant differences ($F(3, 1036) = 49.4, p < 10^{-12}$). Inter-annotator agreement was substantial (Cohen's $\kappa = 0.69$ [16]), ensuring reliable human assessments.

DPO's human-aligned approach improves handling complex types (fig. 5). Despite many examples, challenging types (semantic transformations, inflectional changes, word-order alterations, punctuation changes) remained difficult. DPO/APTY's gains show that human-ranked data yields better paraphrase type alignment than baselines and alternatives. While accuracy shows DPO's technical skill, human preference analysis reveals practical effectiveness.

4.2 Human Preferences for DPO-Generated Paraphrases

We next asked whether these accuracy gains translate into outputs that humans favor. Human annotators ranked 260 paraphrase sets, each containing 1 paraphrase by each model, from best (1) to worst (4), assigning the worst rating to invalid or incorrectly transformed paraphrases.




Fig. 6. Human preference rankings of model-generated paraphrases. The DPO/APTY model produces 49 % top-ranked paraphrases (rank 1) compared to 33 % for the baseline SFT/ETPC. Rankings (1–4) consider both correctness and adherence to the specified paraphrase type. The DPO/APTY model’s consistently superior rankings highlight its closer alignment with human judgments.

DPO/APTY achieved 40 % top-rank vs. baseline’s 6 % and SFT/ETPC’s 33 % (fig. 6), $\chi^2(9) = 231.9, p < 10^{-44}$. table 5 details the rankings. Krippendorff’s [15] $\alpha = 0.63$ indicated moderate inter-annotator agreement.

DPO’s advantage persisted across categories, including complex types (e.g., punctuation, word-order changes). DPO/APTY aligns outputs with human preferences, improving quality across types. These findings validate our approach and support critiques of traditional metrics [31]. By prioritizing semantic accuracy and nuanced transformations, DPO aligns with Vila et al. [36] and Wahle et al. [37], advancing human-aligned paraphrase generation. Although humans favor DPO, we must assess automated metrics for alignment.

4.3 Automated Metrics vs. Human Preferences

We correlated human rankings with ROUGE and BLEU scores to assess whether standard metrics recognize these improvements. Annotators ranked 30 random paraphrase sets with base sentences and references taken from ETPC. As shown in Figure 7, Pearson correlations fell below 0.3 ($p < 10^{-10}$).

Despite human preference for DPO/APTY, even advanced metrics like ParaScore showed weak alignment (table 1).

These results show standard metrics fail as proxies for human perception [24, 31]. Our findings confirm that DPO/APTY models produce human-preferred paraphrases but not necessarily higher lexical scores. Automatic metrics fail to detect whether the generated sentence qualifies as a paraphrase. Improved metrics aligned with human judgment are essential. We developed a PTD model for a more nuanced evaluation.

4.4 Paraphrase Type Detection with a Novel Model

To move beyond traditional metrics, we developed a PTD model that classifies paraphrase transformations without relying on reference sentences.




Fig. 7. Correlation between automated metrics and human quality judgments. For 120 paraphrases from four models, Pearson correlations of ROUGE and BLEU with human rankings remain below 0.3, indicating weak alignment. Human rankings were transformed to a logistic scale for compatibility. These results underscore the need for more human-centered evaluation methods.

Table 1. Automated metric scores (ROUGE, BLEU, ParaScore) for an illustrative paraphrase from SFT/ETPC and DPO/APTY fine-tuned BART models. Despite the SFT/ETPC output not constituting a valid paraphrase, automated metrics remain high due to lexical overlap with the reference. This example demonstrates the limitations of standard metrics in capturing genuine paraphrase quality.

Metric / Base Sentence	SFT/ETPC	DPO/APTY
Those reports were denied by the interior minister, Prince Nayef.	Those reports were denied by Prince Nayef, the interim minister of education.	Those reports, however, were denied by the minister, Prince Nayef
ROUGE-1	0.47	0.42
ROUGE-2	0.09	0.11
ROUGE-L	0.28	0.31
BLEU	0.04	0.05
ParaScore Base	0.87	0.84
ParaScore Free	0.86	0.83

Trained on the top 10 APTs, this model achieved a weighted F1 of 0.71, performing best on simpler transformations (e.g., addition/deletion: F1=0.91, 95% CI [0.90, 0.93]) and less well on complex semantic shifts (fig. 8). Details are noted in table 6. We evaluated the model on ETPC, where each example can have multiple APT. We also evaluated how this PTD model agrees with human annotation on the 260 single APT transformation sets we used for accuracy and ranking (fig. 11).




Fig. 8. F1 scores for paraphrase type detection (PTD) across ten atomic transformations. The PTD model excels at simpler types (Addition/Deletion: F1=0.91; Same Polarity Substitution: F1=0.78) but struggles with complex semantic shifts. Scores are derived from 1,040 human-annotated examples, highlighting variability in detection difficulty.

Although not a complete solution, the PTD model provides more granular insights, enabling fine-grained evaluation that traditional metrics fail to offer. This supports the central claim by demonstrating the value of tools aligned with the conceptual framework of APTs for evaluating paraphrase quality.

4.5 Impact of DPO on Broader NLP Tasks

Finally, we evaluated whether improvements driven by human-ranked paraphrase-type data extend to broader tasks. We tested all models on Open LLM Leaderboard v2 benchmarks. While DPO/APTY showed marginal declines on some tasks (fig. 9), it improved MuSR team allocation performance by up to 38% (table 2). The PTG models ranked 7th among all 85 Llama-3.1-8B models on the Open LLM Leaderboard MuSR ranking (2024-12-20).

Table 2. Detailed Multistep Soft Reasoning (MuSR) task results. Scores indicate the model’s ability to handle multi-step reasoning processes and integrate contextual information, reflecting performance on complex inference tasks. Optimized models excel on team allocation tasks.

Model	Murder mysteries	Object placements	Team allocation
Llama-3.1-8B	0.54	0.35	0.25
SFT/ETPC	0.51	0.35	0.43
DPO/APTY	0.52	0.34	0.42
IPO/APTY	0.51	0.35	0.43




Fig. 9. Model performance on various NLP benchmarks, including Multistep Soft Reasoning (MuSR) and Big Bench Hard (BBH). The DPO/APTY model notably improves MuSR performance, which requires complex reasoning, though it shows minor declines on other tasks. The bars represent model scores, illustrating how human preference alignment influences performance across multiple domains.

This result indicates that human-centered PTG training benefits paraphrasing and certain aspects of reasoning, bolstering the overarching hypothesis that human preference alignment enhances language models across multiple dimensions.

5 CONCLUSION

This study addressed a critical gap in PTG by applying DPO and using human-ranked data to produce paraphrases that align with nuanced human preferences. While earlier work often relied on automated metrics and limited human feedback, we now see that integrating direct human judgments improves transformation-specific accuracy by 3 % and increases user preference by 16 %. These results answer the central research question: DPO-trained models yield more reliable, context-sensitive paraphrases than those relying solely on baseline methods.

Our PTD model provides a rigorous evaluation approach. This model verified the presence of specific transformations and surpassed conventional metrics such as ROUGE and BLEU. Although the detector excelled at identifying more straightforward changes, it struggled with subtler linguistic shifts, reflecting the complexity of capturing nuanced semantics. This outcome suggests that future methods should refine detection approaches and incorporate richer semantic cues.

However, certain limitations remain. We discontinued RLHF due to low reward model accuracy, which highlights the limitations of this application. A reward model is trained on a classification task. The differences in human preferences in APTY might be too subtle for the model. The APTY dataset’s limited size and subjective annotations reduced generalizability, and subtle semantic transformations and complex word-order alterations remained difficult. Future work should expand training corpora, incorporate multiple languages and modalities, and develop evaluation metrics that more effectively capture semantic, contextual, and stylistic fidelity. While

LoRA improved computational efficiency, it may have reduced complex paraphrasing performance. Addressing this challenges requires systematically evaluating DPO with different architectures or fine-tuning strategies. Standardized frameworks blending multiple metrics weighted by human correlation could better reflect linguistic aspects.

This study provides a foundation for producing semantically rich, human-preferred paraphrases. Integrating human judgments in training and evaluation brings NLP closer to user preferences. These advances offer a blueprint for bridging the gap between technical optimization and real-world language demands.

ACKNOWLEDGMENTS

Dominik Meier and Dr. Terry Lima Ruas from the Chair for Scientific Information Analytics (Prof. Dr. Bela Gipp) at the University of Göttingen provided the idea for this project. Jan Philip Wahle implemented supervised fine-tuning of Llama models on the ETPC dataset as part of the EMNLP'23 paper "Paraphrase Types for Generation and Detection"[37]. The doctoral student Dominik Meier supervised the project at the University of Göttingen.

I am deeply grateful to my daughters and my wife for their endless encouragement, love, and patience. Their tireless support and belief in me carried me through late nights and tough moments, making this thesis possible.

REFERENCES

- [1] Mohammad Gheshlaghi Azar, Mark Rowland, Bilal Piot, Daniel Guo, Daniele Calandriello, Michal Valko, and Rémi Munos. 2023. A General Theoretical Paradigm to Understand Learning from Human Preferences. arXiv:2310.12036 [cs.AI] <https://arxiv.org/abs/2310.12036>
- [2] Rahul Bhagat and Eduard Hovy. 2013. What Is a Paraphrase? *Computational Linguistics* 39, 3 (09 2013), 463–472. https://doi.org/10.1162/COLI_a_00166 arXiv:https://direct.mit.edu/coli/article-pdf/39/3/463/1801912/coli_a_00166.pdf
- [3] David Chen and William Dolan. 2011. Collecting Highly Parallel Data for Paraphrase Evaluation. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, Dekang Lin, Yuji Matsumoto, and Rada Mihalcea (Eds.). Association for Computational Linguistics, Portland, Oregon, USA, 190–200. <https://aclanthology.org/P11-1020>
- [4] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. arXiv:1810.04805 [cs.CL] <https://arxiv.org/abs/1810.04805>
- [5] William B. Dolan and Chris Brockett. 2005. Automatically Constructing a Corpus of Sentential Paraphrases. In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*. <https://aclanthology.org/I05-5002>
- [6] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, and others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783* (2024).
- [7] Clémentine Fourrier, Nathan Habib, Alina Lozovskaya, Konrad Szafer, and Thomas Wolf. 2024. Open LLM Leaderboard v2. https://huggingface.co/spaces/open-llm-leaderboard/open_llm_leaderboard.
- [8] Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. DeBERTa: Decoding-enhanced BERT with Disentangled Attention. arXiv:2006.03654 [cs.CL] <https://arxiv.org/abs/2006.03654>
- [9] Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring Mathematical Problem Solving With the MATH Dataset. arXiv:2103.03874 [cs.LG] <https://arxiv.org/abs/2103.03874>
- [10] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. LoRA: Low-Rank Adaptation of Large Language Models. arXiv:2106.09685 [cs.CL] <https://arxiv.org/abs/2106.09685>
- [11] Kuan-Hao Huang, Varun Iyer, I-Hung Hsu, Anoop Kumar, Kai-Wei Chang, and Aram Galstyan. 2023. ParaAMR: A Large-Scale Syntactically Diverse Paraphrase Dataset by AMR Back-Translation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 8047–8061. <https://doi.org/10.18653/v1/2023.acl-long.447>
- [12] Shankar Iyer, Nikhil Dandekar, , and Kornél Csernai. 2017. First Quora Dataset Release: Question Pairs. Accessed: 2024-12-24. <https://quoradata.quora.com/First-Quora-Dataset-Release-Question-Pairs>
- [13] Mohit Iyyer, John Wieting, Kevin Gimpel, and Luke Zettlemoyer. 2018. Adversarial Example Generation with Syntactically Controlled Paraphrase Networks. arXiv:1804.06059 [cs.CL] <https://arxiv.org/abs/1804.06059>
- [14] Venelin Kovatchev, M. Antònia Martí, and Maria Salamó. 2018. ETPC - A Paraphrase Identification Corpus Annotated with Extended Paraphrase Typology and Negation. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Nicoletta Calzolari, Khalid Choukri, Christopher Cieri, Thierry Declerck, Sara Goggi, Koiti Hasida, Hitoshi Isahara, Bente Maegaard, Joseph Mariani, Hélène Mazo, Asuncion Moreno, Jan Odijk, Stelios Piperidis, and Takenobu Tokunaga (Eds.). European Language Resources Association (ELRA), Miyazaki, Japan. <https://aclanthology.org/L18-1221>

- [15] Klaus Krippendorff. 2024. *Content Analysis: An Introduction to Its Methodology* (fourth edition ed.). SAGE Publications, Inc., Thousand Oaks, California. <https://doi.org/10.4135/9781071878781>
- [16] J Richard Landis and Gary G. Koch. 1977. The measurement of observer agreement for categorical data. *Biometrics* 33 1 (1977), 159–74. <https://api.semanticscholar.org/CorpusID:11077516>
- [17] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. arXiv:1910.13461 [cs.CL] <https://arxiv.org/abs/1910.13461>
- [18] Zichao Li, Xin Jiang, Lifeng Shang, and Hang Li. 2018. Paraphrase Generation with Deep Reinforcement Learning. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Brussels, Belgium, 3865–3878. <https://doi.org/10.18653/v1/D18-1421>
- [19] Chin-Yew Lin. 2004. ROUGE: A Package for Automatic Evaluation of Summaries. In *Text Summarization Branches Out*. Association for Computational Linguistics, Barcelona, Spain, 74–81. <https://www.aclweb.org/anthology/W04-1013>
- [20] Christopher L. Luebbers. 2024. dpo-rlhf-paraphrase-types. Accessed: 2024-12-24. <https://github.com/cluebbbers/dpo-rlhf-paraphrase-types>
- [21] Christopher L. Luebbers. 2024. Enhancing Paraphrase Type Generation Huggingface Collection. Accessed: 2024-12-24. <https://huggingface.co/collections/cluebbbers/enhancing-paraphrase-type-generation-673ca8d75dfe2ce962a48ac0>
- [22] Dominik Meier, Jan Philip Wahle, Terry Ruas, and Bela Gipp. 2024. Towards Human Understanding of Paraphrase Types in ChatGPT. <http://arxiv.org/abs/2407.02302>
- [23] Affan Hilmy Natsir, Indriana Hidayah, and Teguh Bharata Adji. 2023. Deep Learning in Paraphrase Generation: A Systematic Literature Review. In *2023 IEEE 7th International Conference on Information Technology, Information Systems and Electrical Engineering (ICITISEE)*. Institute of Electrical and Electronics Engineers, 118–123. <https://doi.org/10.1109/ICITISEE58992.2023.10405123>
- [24] Shinhyeok Oh, Hyojun Go, Hyeongdon Moon, Yunsung Lee, Myeongho Jeong, Hyun Seung Lee, and Seungtaek Choi. 2023. Evaluation of Question Generation Needs More References. In *Findings of the Association for Computational Linguistics: ACL 2023*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 6358–6367. <https://doi.org/10.18653/v1/2023.findings-acl.396>
- [25] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, Pierre Isabelle, Eugene Charniak, and Dekang Lin (Eds.). Association for Computational Linguistics, Philadelphia, Pennsylvania, USA, 311–318. <https://doi.org/10.3115/1073083.1073135>
- [26] Karl Pearson. 1895. Note on Regression and Inheritance in the Case of Two Parents. *Proceedings of the Royal Society of London Series I* 58 (Jan. 1895), 240–242.
- [27] Alec Radford and Karthik Narasimhan. 2018. Improving Language Understanding by Generative Pre-Training. <https://api.semanticscholar.org/CorpusID:49313245>
- [28] Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. 2023. Direct Preference Optimization: Your Language Model is Secretly a Reward Model. <http://arxiv.org/abs/2305.18290>
- [29] David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R. Bowman. 2023. GPQA: A Graduate-Level Google-Proof Q&A Benchmark. arXiv:2311.12022 [cs.AI] <https://arxiv.org/abs/2311.12022>
- [30] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal Policy Optimization Algorithms. arXiv:1707.06347 [cs.LG] <https://arxiv.org/abs/1707.06347>
- [31] Lingfeng Shen, Lema Liu, Haiyun Jiang, and Shuming Shi. 2022. On the Evaluation Metrics for Paraphrase Generation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (Eds.). Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 3178–3190. <https://doi.org/10.18653/v1/2022.emnlp-main.208>
- [32] Reece Shuttleworth, Jacob Andreas, Antonio Torralba, and Pratyusha Sharma. 2024. LoRA vs Full Fine-tuning: An Illusion of Equivalence. arXiv:2410.21228 [cs.LG] <https://arxiv.org/abs/2410.21228>
- [33] Zayne Sprague, Xi Ye, Kaj Bostrom, Swarat Chaudhuri, and Greg Durrett. 2024. MuSR: Testing the Limits of Chain-of-thought with Multistep Soft Reasoning. arXiv:2310.16049 [cs.CL] <https://arxiv.org/abs/2310.16049>
- [34] Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V. Le, Ed H. Chi, Denny Zhou, and Jason Wei. 2022. Challenging BIG-Bench Tasks and Whether Chain-of-Thought Can Solve Them. arXiv:2210.09261 [cs.CL] <https://arxiv.org/abs/2210.09261>
- [35] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open Foundation and Fine-Tuned Chat Models. arXiv:2307.09288 [cs.CL] <https://arxiv.org/abs/2307.09288>
- [36] Marta Vila, M. Antònia Martí, and Horacio Rodríguez. 2014. Is This a Paraphrase? What Kind? Paraphrase Boundaries and Typology. *Open Journal of Modern Linguistics* 04, 01 (2014), 205–218. <https://doi.org/10.4236/ojml.2014.41016>

- [37] Jan Philip Wahle, Bela Gipp, and Terry Ruas. 2023. Paraphrase Types for Generation and Detection. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 12148–12164. <https://doi.org/10.18653/v1/2023.emnlp-main.746>
- [38] Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhramil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, et al. 2024. MMLU-Pro: A More Robust and Challenging Multi-Task Language Understanding Benchmark. arXiv:2406.01574 [cs.CL] <https://arxiv.org/abs/2406.01574>
- [39] Chao Zhou, Cheng Qiu, and Daniel E. Acuna. 2022. Paraphrase Identification with Deep Learning: A Review of Datasets and Methods. <http://arxiv.org/abs/2212.06933>
- [40] Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. 2023. Instruction-Following Evaluation for Large Language Models. arXiv:2311.07911 [cs.CL] <https://arxiv.org/abs/2311.07911>

A TECHNICAL INFORMATION

A requirements file is available on Github [20]. The table (table 3) and requirements file ensure reproducibility and transparency for exact replication.

Table 3. Python libraries and corresponding versions used in the experimental environment. Specifying exact versions ensures reproducibility and facilitates independent validation of the reported results.

Library	Version
python	3.11.10
NumPy	2.1.3
pandas	2.2.3
PyTorch	2.5.1
torchvision	0.19.1
torchaudio	2.4.1
cuda	12.1
scikit-learn	1.5.2
spacy	3.7.5
scipy	1.14.1
seaborn	0.13.2
matplotlib	3.9.2
transformers	4.46.2
datasets	3.1.0
trl	0.12.0
accelerate	1.1.1
bitsandbytes	0.44.1
peft	0.13.2
rouge	1.0.1
nltk	3.9.1

B DATASETS

ETPC refines Microsoft Research Paraphrase Corpus [5] by annotating it with the Extended Paraphrase Typology (EPT). This annotation introduces granular distinctions between paraphrase and non-paraphrase, including contextual, habitual, and negation relations. ETPC enhances MRPC utility by enabling detailed evaluation and error analysis and supporting tasks like semantic similarity and entailment. We used the Oct 2, 2023 version. table 4 shows APT frequencies.

The APTY dataset [22] extends ETPC with human-ranked preferences for APT-based paraphrases. APTY includes APTY-base (correctness and errors) and APTY-ranked (human preferences). The dataset columns include 'original', 'chosen', and 'rejected'. We used the July 8, 2024 version. table 4 shows APT frequencies after preprocessing.

Table 4. Frequency counts of paraphrase types in the Extended Typology Paraphrase Corpus (ETPC) and the Annotated Paraphrase Type (APTY) dataset. In ETPC, multiple instances of a paraphrase type can occur within a single sentence, explaining total vs. unique counts. In APTY, multiple chosen and rejected examples may relate to the same original sentence.

	ETPC total	ETPC unique	APTY total	APTY unique
Addition/Deletion	5722	2988	40	9
Change of format	240	207		
Change of order	860	766	60	10
Converse substitution	43	42		
Coordination changes	48	47		
Derivational Changes	187	181	22	6
Diathesis alternation	162	161		
Direct/indirect style alternations	66	66		
Ellipsis	66	64		
Entailment	81	81		
Identity	3878	3870		
Inflectional Changes	613	544	12	7
Modal Verb Changes	184	180		
Negation switching	20	20		
Non-paraphrase	842	605		
Opposite polarity substitution (contextual)	15	12		
Opposite polarity substitution (habitual)	4	4		
Punctuation changes	833	748	28	9
Same Polarity Substitution (contextual)	4173	2511	56	10
Same Polarity Substitution (habitual)	840	681		
Same Polarity Substitution (named ent.)	536	448		
Semantic-based	337	328	59	10
Spelling changes	636	534	31	9
Subordination and nesting changes	473	448	18	5
Syntax/discourse structure changes	308	305		
Synthetic/Analytic Substitution	897	806	7	5

C MODELS

We trained models with parameter configurations enhancing paraphrase generation. Key parameters include:

- LoRA target modules for Llama-3-8B: up_proj, down_proj, k_proj, o_proj, v_proj, gate_proj, and q_proj
- DPO parameters: learning rate = 1e-6, weight decay = 4e-1, $\beta = 0.2$, max_grad_norm = 200, lr_scheduler = cosine
- IPO parameters; warmup ratio = 0.2, weight decay = 0.02, learning rate = 5e-6, $\beta = 0.2$, lr_scheduler = 'reduce learning rate on plateau'

D EVALUATION

We annotated generated paraphrases with Label Studio 1.10.0 available at <https://labelstud.io/>. Our labeling instructions:

Only evaluate the first paraphrase.

All incorrect paraphrases are ranked "wrong".

All correct paraphrases are ranked from "best" downwards.

We conducted preliminary DPO optimization with Llama-2-7B. Mean accuracy rose from 12 % (base, SFT/ETPC) to 35 % (DPO/APTY). Accuracy by APT is shown in fig. 10.




Fig. 10. Accuracy of paraphrase-type generation by Llama-2-7B and a DPO/APTY model, based on human annotations. While DPO/APTY outperforms the baseline overall, challenging transformations (e.g., 'change of order', 'derivational changes') remain difficult for both models.

Table 5. Distribution of human rankings for paraphrase quality: The table shows each model’s human rankings (1 = best, 4 = worst). DPO/APTY achieves the highest top-rank proportion; Llama-3.1-8B mostly ranks 4. A Chi-Square test (92.34, $p=5.5e-16$) confirms DPO/APTY’s superiority.

	1	2	3	4
Llama-3.1-8B	6	2	0	91
SFT/ETPC	33	13	9	45
DPO/APTY	40	10	8	43
IPO/APTY	34	11	8	47

Table 6. F1 scores for the paraphrase type detection model across multiple transformation categories. The model excels at Addition/Deletion (F1=0.91), Same Polarity Substitution (0.77), and Punctuation (0.70), but achieves lower accuracy for more complex transformations, indicating varied detection difficulty.

Class	F1 Score	CI Lower	CI Upper	Support
Addition/Deletion	0.91	0.90	0.93	1327
Change of order	0.46	0.41	0.52	303
Derivational Changes	0.12	0.01	0.23	33
Inflectional Changes	0.48	0.39	0.57	130
Punctuation Changes	0.70	0.64	0.76	223
Same Polarity Substitution (contextual)	0.78	0.75	0.80	1113
Semantic-Based	0.16	0.07	0.25	63
Spelling Changes	0.33	0.25	0.40	144
Subordination and Nesting Changes	0.34	0.25	0.44	94
Synthetic/Analytic substitution	0.29	0.24	0.35	248




Fig. 11. F1 scores for agreement of paraphrase type detection (PTD) with human annotators. The PTD model scores better, because PTG was performed on single transformation.

AI Usage Card for *Enhancing Paraphrase Type Generation: The Impact of DPO and RLHF*

Evaluated with Human-Ranked Data



CORRESPONDENCE(S)

Christopher L. Lübbers

CONTACT(S)

c.luebbbers@stud.uni-goettingen.de

AFFILIATION(S)

University of Göttingen, Chair for Scientific Information Analytics

PROJECT NAME

Enhancing Paraphrase Type Generation: The Impact of DPO and RLHF Evaluated with Human-Ranked Data

KEY APPLICATION(S)

Natural Language Processing, Paraphrase Generation

MODEL(S)

ChatGPT
ChatGPT
ChatGPT
Grammarly

DATE(S) USED

2024/06/24 - 2024/12/20
2024/09/12 - 2024/12/04
2024/12/05 - 2024/12/20
2024/06/24 - 2024/12/20

VERSION(S)

4o
o1-preview
o1

IDEATION

ChatGPT o1-preview, o1

IMPROVING EXISTING IDEAS

Reviewed and refined the research proposal, identifying areas for improvement.

LITERATURE REVIEW

ChatGPT 4o

FINDING LITERATURE

Expanded the list of known literature by identifying additional relevant sources.

METHODOLOGY

ChatGPT o1-preview, o1

FINDING ITERATIVE OPTIMIZATIONS

Enhanced the proposed methodology by addressing potential gaps and increasing rigor.

WRITING

ChatGPT 4o, Grammarly

GENERATING NEW TEXT BASED ON INSTRUCTIONS

Transformed bullet points into complete, coherent sentences using ChatGPT.

ASSISTING IN IMPROVING OWN CONTENT

Grammarly refined written content for correctness, clarity, and engagement.

PARAPHRASING RELATED WORK

Summarized related works partially with ChatGPT's assistance.

PRESENTATION

ChatGPT 4o

GENERATING NEW ARTIFACTS

Created initial versions of figures and tables, which were later manually improved.

CODING

ChatGPT 4o

GENERATING NEW CODE BASED ON DESCRIPTIONS OR EXISTING CODE

Collaboratively developed code in a pair programming setup.

REFACTORING AND OPTIMIZING EXISTING CODE

Debugged and optimized code for functionality and efficiency.

ETHICS

WHAT ARE THE IMPLICATIONS OF USING AI FOR THIS PROJECT?

AI accelerated analysis and improved clarity for writing.

WHAT STEPS ARE WE TAKING TO MITIGATE ERRORS OF AI FOR THIS PROJECT?

Rigorous manual review, cross-validation, and expert verification ensure accuracy and reliability.

WHAT STEPS ARE WE TAKING TO MINIMIZE THE CHANCE OF HARM OR INAPPROPRIATE USE OF AI FOR THIS PROJECT?

Sentences for paraphrase generation were carefully chosen manually and annotated by experienced researchers.

THE CORRESPONDING AUTHORS VERIFY AND AGREE WITH THE MODIFICATIONS OR GENERATIONS OF THEIR USED AI-GENERATED CONTENT

All authors reviewed and approved all AI-generated or AI-modified content.