

ADPARAPHRASE: Paraphrase Dataset for Analyzing Linguistic Features toward Generating Attractive Ad Texts

Soichiro Murakami¹, Peinan Zhang¹, Hidetaka Kamigaito^{2,3},
Hiroya Takamura³, Manabu Okumura³

¹CyberAgent, Inc., ²Nara Institute of Science and Technology, ³Institute of Science Tokyo
{murakami_soichiro, zhang_peinan}@cyberagent.co.jp,
kamigaito.h@is.naist.jp, {takamura, oku}@pi.titech.ac.jp

Abstract

Effective linguistic choices that attract potential customers play crucial roles in advertising success. This study aims to explore the linguistic features of ad texts that influence human preferences. Although the creation of attractive ad texts is an active area of research, progress in understanding the specific linguistic features that affect attractiveness is hindered by several obstacles. First, human preferences are complex and influenced by multiple factors, including their content, such as brand names, and their linguistic styles, making analysis challenging. Second, publicly available ad text datasets that include human preferences are lacking, such as ad performance metrics and human feedback, which reflect people’s interests. To address these problems, we present ADPARAPHRASE, a paraphrase dataset that contains human preferences for pairs of ad texts that are semantically equivalent but differ in terms of wording and style. This dataset allows for preference analysis that focuses on the differences in linguistic features. Our analysis revealed that ad texts preferred by human judges have higher fluency, longer length, more nouns, and use of bracket symbols. Furthermore, we demonstrate that an ad text-generation model that considers these findings significantly improves the attractiveness of a given text. The dataset is publicly available at: <https://github.com/CyberAgentAILab/AdParaphrase>.¹

1 Introduction

Online advertising plays a significant role in digital marketing. Advertising aims to attract attention, spark interest, and encourage clicks and purchases for profit. To achieve this, writing an ad text relevant to user interests is essential; however, this alone is not sufficient. Simultaneously, the way an ad text is written, that is, the linguistic expression of the ad text to attract attention, is crucial

¹The dataset is provided under the CC BY-NC-SA 4.0 license.

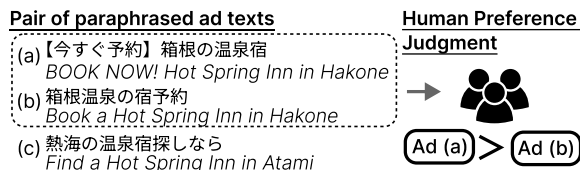


Figure 1: Overview of human preference judgments for a pair of paraphrase ad texts.

for the success of advertisements. Effective linguistic choices in ad texts can influence the ease of understanding and appeal to people, potentially leading to successful outcomes, such as improved click-through rates (CTR).

In this study, we aim to explore the linguistic features of an ad text that influences human preferences with the goal of maximizing the potential success of advertisements. Many studies have been conducted on the methods for generating attractive ad texts (Hughes et al., 2019; Kamigaito et al., 2021; Wei et al., 2022). However, progress in understanding the linguistic features that affect attractiveness, that is, human preferences, has been hindered by several obstacles. First, human preferences are complex and are influenced by various factors, including linguistic expressions and semantic content (Pryzant et al., 2018; Hu et al., 2023). For instance, when analyzing human preferences for two ad texts with different content and writing styles, such as ad texts (a) and (c) in Figure 1, it is difficult to determine the factors influencing preferences. They can be motivated by content such as place names (e.g., *Atami*) or writing styles such as uppercase text (e.g., *BOOK NOW!*). Second, publicly available ad text datasets containing human preference data are lacking (Murakami et al., 2023). Previous studies that analyzed attractive ad texts relied on log data, such as clicks and views, that reflect human preferences (Pryzant et al., 2018; Murakami et al., 2022). However, the datasets they used were proprietary to companies for confiden-

tiality reasons and were not publicly available. This restricts potential researchers to those with access to such data, hindering research on what makes ads attractive.

To explore the linguistic features that affect attractiveness by overcoming the aforementioned obstacles, we present ADPARAPHRASE, a novel paraphrase dataset that contains human preferences for pairs of ad texts that are semantically equivalent but differ in wording and style. We carefully constructed the dataset by collecting semantically similar ad texts, performing paraphrase identification with five judges per pair, and collecting human preference judgments with ten judges per pair, as shown in Figure 1. The dataset allows us to focus on differences in linguistic features between individual ad texts while minimizing the impact of differences in semantic content. Thus, we can analyze human preferences centered on these features.

Through a statistical analysis of the human preferences collected in our dataset, we found that an ad text that is more fluent, longer in length, contains more nouns, and uses bracket symbols tends to be preferred by most human judges. Based on the findings of this analysis, we explored various methods for generating more attractive ad texts in terms of their linguistic styles. In the experiments, we found that more attractive ad texts could be generated by considering our findings and preference judgments as few-shot examples; however, room for improvement still exists. We believe that our dataset will advance the understanding of linguistic features that influence human preferences in advertisements.

2 Construction of ADPARAPHRASE

We constructed ADPARAPHRASE, a paraphrase dataset of ad texts, to examine the influence of linguistic variations on human preferences. This dataset allowed us to collect and analyze human preferences by focusing on the linguistic differences between ad texts. Our two-step dataset construction process involves the collection of paraphrase candidates (§2.1) and manual annotation for paraphrase identification (§2.2).

2.1 Collecting Paraphrase Candidates

ADPARAPHRASE was constructed based on two publicly available Japanese ad text datasets, Ad Similarity (Zhang et al., 2024) and CAMERA

Source	Model	#Cand.	#Para.	#Non-para.
Ad Similarity	—	706	335	371
	Llama2	133	86	47
CAMERA	GPT-3.5	133	98	35
	GPT-4	133	81	52
	Human	133	125	8
Total		1,238	725	513

Table 1: Statistics of ADPARAPHRASE. #Cand. is the number of paraphrase candidates. #Para and #Non-Para denote the number of candidates identified as paraphrase and non-paraphrase, respectively.

(Mita et al., 2024).² We employed different strategies, as outlined below, to collect paraphrase candidates from the datasets based on their distinct formats. Table 1 summarizes the breakdown of the paraphrase candidates. We obtained a total of 1,238 candidates.

Ad Similarity Ad Similarity is a dataset that includes 6,332 pairs of ad texts rated on a five-point scale by three human evaluators. Higher scores indicate greater similarity between pairs, and vice versa. To collect the candidates, we set a threshold of four for the similarity score and extracted 1,088 pairs. However, we found that several extracted pairs contained different named entities such as date and price. To address this issue, we applied regular expression rules, removed 382 pairs, and finally obtained 706 candidate pairs.

CAMERA CAMERA is a benchmark dataset for ad text generation (ATG) tasks, in which an ad text is generated from user queries and source documents. Unlike with Ad Similarity, we created paraphrases from scratch using ad texts in the dataset as the source text. We asked two professional ad writers from an advertising agency to create a paraphrased text from a given source text. We provided two instructions: first, to rephrase its wording and style to enhance attractiveness without adding or deleting any information; and second, to limit the ad text to 15 full-width characters, which is the length limit for advertising delivery platforms such as Google Ads. Owing to time constraints, we requested approximately 100 paraphrases from human experts, resulting in 133 paraphrases. In addition, we used large language models (LLMs) that have shown remarkable paraphrasing capabilities (Cegin et al., 2023) to generate paraphrase

²Both are governed by the CC BY-NC-SA 4.0 license and we adhere to the intended use of both datasets.

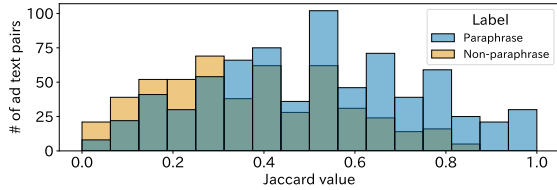


Figure 2: Distribution of Jaccard similarity for paraphrases and non-paraphrases in ADPARAPHRASE.

Paraphrase	「リピート確定」の声多数 (Many Voices: “Will Definitely Repeat”) これからも使い続けたいの声続出 (Many Voices Want to Keep Using It)
Non-paraphrase	最強スペックのAndroid端末 (The Most Powerful Android Phone) Google Pixel 史上最強スペック (The Most Powerful Google Pixel Ever)

Table 2: Examples of paraphrases and non-paraphrases.

candidates. To enable a direct comparison between LLMs and human experts in paraphrasing, we used the same source ad texts. We employed GPT-4, GPT-3.5, and Llama2 (OpenAI, 2024; Touvron et al., 2023) in a zero-shot setting and provided them with the instructions identical to those of human experts. This process yielded 399 pairs. A total of 532 pairs were collected from human experts and LLMs. We present the details of the paraphrase creation, including the instructions and model hyperparameters, in Appendices A and F.

2.2 Manual Annotation

For the paraphrase candidates, we conducted a manual annotation for paraphrase identification, which is a binary classification task, to determine whether each pair of ad texts is semantically equivalent (i.e., *paraphrase* or *non-paraphrase*). We defined the paraphrasing criteria at the sentence level rather than at the word or phrase level to assess whether the two sentences conveyed the same meaning. Annotation was performed by five workers with extensive experience in in-house advertising production. The gold label was determined by majority votes. For details of paraphrase identification, including quality control measures and the instructions presented to workers, please refer to Appendix B.

2.3 Data Analysis

The annotation results of the paraphrase identification are summarized in Table 1. Of the 1,238 candidates, 725 were paraphrases, and 513 were

non-paraphrases. Ad Similarity yielded 335 paraphrases, whereas human experts and LLMs produced 390 paraphrases. The average length of the ad texts is 7.1 words. The Fleiss’ kappa metric among the five workers was 0.462, indicating moderate agreement (Fleiss et al., 1971; Landis and Koch, 1977). Table 2 lists examples of paraphrases and non-paraphrases. Several cases determined to be non-paraphrases were caused by differences in named entities, such as product names.

To better understand ADPARAPHRASE, we calculated the Jaccard similarity to measure the lexical overlap between ad text pairs. Figure 2 illustrates the distributions of the metrics for the paraphrases and non-paraphrases. These results confirm that ADPARAPHRASE includes many paraphrases that are lexically different but semantically equivalent.

3 Collection of Human Preferences

We collected human preferences through human evaluations of the attractiveness of ad texts rather than relying on log data, such as CTR, for two reasons. First, human evaluations of attractiveness are widely used to measure the quality of ad texts in the ATG field (Murakami et al., 2023). Second, previous studies have shown that human attractiveness ratings and their predicted CTR (pCTR) are reasonably consistent, supporting the validity of using human evaluations as a proxy (Mita et al., 2024). Therefore, we expect this approach to serve as an alternative to not releasing such data publicly.

3.1 Evaluation Method

We used the 725 paraphrase pairs in Table 1 as the evaluation set and performed pairwise comparisons to determine the ad text that was more attractive in each pair. Each pair was evaluated by ten human judges, who are native Japanese speakers, recruited from a crowdsourcing platform. A skip option was provided for judges to use when they found the ad texts equally attractive. Appendix C provides details of the attractiveness evaluation, including the workers and interface of the evaluation tool.

3.2 Quality Control

Several strategies were introduced to ensure annotation quality. (1) To mitigate position bias in pairwise comparisons, we randomized the order of ad text pairs before presenting them to the judges. (2) Given the subjective nature of attractiveness, we provided clear evaluation criteria based on the

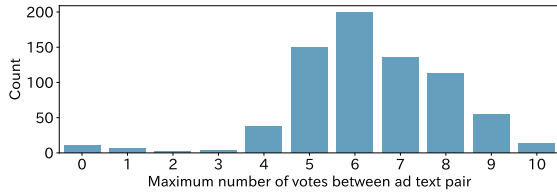


Figure 3: Distribution of maximum number of votes between ad text pair in attractiveness evaluation.

guidelines of Wang et al. (2021), offering examples of evaluation perspectives such as *catchy*, *memorable*, and *easy-to-read*. (3) To ensure the annotation quality across crowdsourced workers of varying skill levels, we incorporated dummy questions with clear answers. These dummy questions consisted of pairs of identical ad texts, expecting the judges to select a skip option. We rejected all responses from judges who failed to choose the skip option for the dummy questions.

3.3 Evaluation Results

Figure 3 shows a histogram of attractiveness evaluations from ten judges for ad text pairs. The x-axis indicates the maximum number of judges who selected the same ad text in a pair. For instance, a value of six indicates that six judges preferred the same ad text, whereas zero indicates that all judges skipped it, suggesting equal attractiveness of the ad text pair. The distribution shows many cases in which five to six judges selected the same ad text, suggesting inconsistent preferences for paraphrased ad texts. The inter-rater agreement measured by Fleiss’ kappa was 0.161, indicating a slight agreement.

Nevertheless, it should be noted that in 316 cases, seven or more judges agreed on their preference. In these cases, the inter-rater agreement was 0.307, indicating fair agreement.³ These findings suggest that linguistic differences between paraphrased ad texts influence the preferences of human judges to a certain extent.

4 Experiments

According to §3.3, the human preference judgments showed fair agreement in 316 evaluation cases. This raises the following question: *What linguistic differences between ad text pairs affect hu-*

³It is worth noting that the relatively low inter-annotator agreement for “attractiveness” can be attributed to the subjective nature of this evaluation criterion, as also reported in previous work (Mita et al., 2024).

Features		df	N	χ^2	p-value
Raw text features	<i>Text length</i>				
	character	1	248	13.32	< 0.01
	word	1	219	5.74	0.02
<hr/>					
<i>Content words</i>					
	noun	1	158	14.18	< 0.01
	verb	1	87	2.17	0.14
	adjective	1	11	0.74	0.39
	adjectival verb	1	40	0.95	0.33
	adverb	1	22	0.00	1.00
<i>Lexical choice</i>					
Lexical features	word frequency	1	301	0.54	0.46
	common noun	1	160	11.84	< 0.01
	proper noun	1	6	0.75	0.39
<i>Character type</i>					
	hiragana	1	215	0.01	0.92
	katakana	1	99	0.00	1.00
	kanji	1	212	5.63	0.02
	symbol	1	155	2.20	0.14
	digits	1	15	0.00	1.00
<hr/>					
<i>Dependency tree</i>					
Syntactic features	max depth of dep. tree	1	155	0.28	0.60
	ave. depth of dep. tree	1	170	0.59	0.44
	max length of dep. links	1	201	0.01	0.91
<i>Others</i>					
	noun phrases	1	171	8.05	< 0.01
	perplexity	1	316	14.15	< 0.01
<hr/>					
<i>Emotion</i>					
Stylistic features	joy	1	78	0.00	1.00
	anticipation	1	75	0.65	0.42
	<i>Others</i>				
	textual specificity	1	11	0.75	0.39
	brackets	1	50	14.25	< 0.01

Table 3: Results of the chi-square test for linguistic features. Df and N refer to the degree of freedom and the number of cases for each type of feature, respectively.

man preference judgment? To answer this question, we analyzed the linguistic features influencing human preferences, focusing particularly on cases in which the judges’ preferences were aligned (§4.1). Furthermore, to demonstrate the practical applications of the findings, we explored methods for generating attractive ad texts that incorporated the linguistic features identified in the analysis (§4.2).

4.1 Analyzing Linguistic Features that Influence Human Preferences

4.1.1 Linguistic Features

An ideal ad text is designed to capture attention and generate interest. Therefore, readability, visibility, and informativeness are key factors in improving the attractiveness of ad texts (Wang and Pomplun, 2012; Murakami et al., 2023). We analyzed various linguistic features from the surface-level, such as text length, to deeper factors, such as lexical choice and emotion. Our feature set consists of

four groups: raw text, lexical, syntactic, and stylistic features, including the 24 types listed in Table 3. For the detailed definitions of these features, see Appendix D.

Raw text features Raw text features, such as text length, are the most basic features of ad text that affect readability and informativeness. We measured word and character counts.

Lexical features Lexical features include word-level features such as the number of content words, lexical choice, and character types. We hypothesized that ad texts with more content words would be more informative and attractive. Regarding lexical choices, we posit that ad texts using common words would be more readable and preferable. To measure this, we calculated the average word frequency of each ad text using the Balanced Corpus of Contemporary Written Japanese (Maekawa et al., 2010) and counted the common and proper nouns. Because ADPARAPHRASE comprises Japanese ad texts, we also included character-type features, such as hiragana counts, which are known to affect readability by helping readers distinguish words more clearly (Sato et al., 2008).

Syntactic features Syntactic features relate to the structure of the entire ad text or its parts, such as phrases, including the depth of the dependency tree, length of the dependency links, number of noun phrases, and perplexity (PPL).

Stylistic features Stylistic features comprise features related to the wording and style of ad texts, including high-level linguistic features such as emotions (e.g., joy and anticipation) and the specificity of the ad text. We hypothesized that more positive or specific ad texts would be preferred. To determine these features, we used separate classification models for emotion and specificity, which we had constructed (see Appendix D for details). In addition, we used the presence or absence of bracket symbols such as 「」 or 「」 as a unique linguistic feature in Japanese ad texts because these brackets are often used to highlight important information for visibility. For instance, in the ad text [Official Site] ABC Insurance, the brackets highlight the key information.

4.1.2 Analysis Method

We used the chi-square test of independence to analyze the relationship between the linguistic features and human preferences. This statistical method

tests the independence of two categorical variables. In our study, these variables were (1) the ad text selected by the majority of human judges and (2) the ad text scoring value (higher or lower) for each linguistic feature. For example, when examining the impact of PPL on human preferences, we investigated whether ad texts with lower PPL tended to be selected by the majority of the judges.

To ensure the reliability of our analysis, we used the 316 (out of 725) pairs of ad texts, in which the choices of more than seven human judges for attractiveness were consistent, as this indicated fair agreement by Fleiss' kappa (Landis and Koch, 1977).

4.1.3 Results

Table 3 lists the chi-square test results. Our analysis focused on the differences in features between the two ad texts, considering only the cases in which each feature differed within the pair. Consequently, the number of cases varied for each feature type. For example, the character counts differed in 248 pairs.

In Table 3, linguistic features with high chi-square values and low p-values indicate a strong correlation with human preferences. Therefore, the results demonstrate that the ad texts preferred by the majority of human judges exhibit significantly lower PPL, longer character counts, higher frequencies of nouns, more noun phrases, and greater use of bracket symbols ($p < 0.01$). These findings suggest that to create attractive ad texts, it is crucial to focus on these specific linguistic features.

However, we observed no statistically significant differences in other linguistic features. Several factors may explain these results. We speculated that the complexity of the dependency trees had little impact because of the short length of the ad texts. The quality of emotion labels may be limited by a domain mismatch because the emotion classification model was trained on social media texts, which differ from ad texts. Concerning specificity, because we used paraphrase pairs for attractiveness evaluation, the specificity of the pair differed in a few cases and we could not obtain a sufficient number of cases for analysis. However, a promising result was obtained, in that 10 out of the 11 cases with high specificity were preferred by the majority of human judges.

4.2 Generating Attractive Ad Texts

To demonstrate the practical application of these findings, we conducted an experiment using ATG.

Specifically, we explored methods for generating attractive ad texts based on the findings of our analysis.⁴

4.2.1 Ad Text Generation Methods

In this experiment, we focused on ad text refinement (Youngmann et al., 2020; Mishra et al., 2020). The goal was to generate more attractive ad text by rephrasing its wording and style without adding or deleting information from the input ad text.

Inspired by the success of integrating human preferences into LLMs, we explored methods for generating attractive ad texts using LLMs. Various methods have emerged for integrating human preferences into LLMs, including reinforcement learning, supervised fine-tuning, and in-context learning (ICL) (Ouyang et al., 2022; Fernandes et al., 2023; Kirk et al., 2023). Owing to limited preference data, we focused on ICL, which is a learning paradigm in which an LLM learns a new task from context, including instructions and input–output demonstrations, and is known for its effectiveness in few-shot settings (Brown et al., 2020).

We explored methods for generating more attractive ad texts by learning human preferences using ICL. For instructions, we incorporated the findings of the linguistic features into the prompt. For the demonstrations, we created two types of input–output pairs using the results of attractiveness evaluations: those in which attractiveness improved and those in which it did not. For the former, we selected cases in which the preferences of at least seven out of ten human judges matched, using the less-preferred text as the input and the more-preferred text as the output. For the latter, we selected cases in which no consensus was reached among the judges. These were referred to as positive and negative examples, respectively. Full experimental details, including prompts, are provided in Appendices E and F.

4.2.2 Evaluation Method

We conducted a human evaluation based on two aspects: paraphrase identification⁵ and attractiveness. Specifically, we assessed whether the generated ad texts were semantically equivalent to the input ad texts and which were more attractive, following the same procedures in §2.2 and §3.1.

⁴In this experiment, we tested the findings other than the frequency of nouns and noun phrases.

⁵Although automatic metrics such as ROUGE can be used, Shen et al. (2022) reported that they do not align well with human judgments. Therefore, we opted for a human evaluation.

To verify the effectiveness of the instructions and demonstrations, we compared five variations of GPT-4-based models listed in Table 4, in addition to the zero-shot baselines of Llama2, GPT-3.5, and GPT-4 in §2. For example, GPT-4-fewshot-findings-pos refers to the GPT-4 model that uses the findings and positive examples as few-shot examples. For comparison with the baselines, we used the same 133 ad texts as inputs for the evaluation. To prevent data leakage, the ad texts were not used as few-shot examples.

4.2.3 Results and Discussion

Table 4 summarizes the results. The success rate for each evaluation task was calculated. For paraphrase identification, the success rate represents the proportion of generated texts labeled as paraphrases by the majority of human judges. For attractiveness, we defined the success rate as the proportion of paraphrased texts found to be more attractive by at least seven out of ten judges. We also calculated the overall success rates for both tasks. The following can be observed from Table 4.

Findings significantly improved performance

Explicitly, incorporating the findings into the instructions (e.g., GPT-4-zeroshot-findings) significantly improved both evaluation tasks compared to the baselines (e.g., GPT-4-zeroshot). We found that directly incorporating these findings into the instructions was straightforward and effective. Further improvements would be possible by identifying other factors that influence preferences.

Negative examples are also useful

Introducing only positive examples as few-shot examples improves the paraphrase identification performance. However, contrary to our expectations, this yielded little improvement in attractiveness. Interestingly, when we introduced negative examples as well, the attractiveness evaluation performance improved significantly.⁶ One possible explanation is that negative examples consist of minor edits between the input and output texts compared to positive examples. For example, the average Jaccard similarity between the input and output were 0.560 and 0.602 for the positive and negative examples, respectively. This discrepancy may have influenced the generation model to actively perform editing operations

⁶Fisher’s exact test showed GPT-4-fewshot-both significantly outperformed GPT-4-fewshot-pos ($p < 0.05$), while no significant difference was found between GPT-4-fewshot-findings-pos and -both.

Model	Instruction	Few-shot example		Success Rate \uparrow		
	w/ findings	positive	negative	Paraphrase	Attractive	Overall
Llama2-zeroshot	—	—	—	0.647	0.081	0.053
GPT-3.5-zeroshot	—	—	—	0.737	0.184	0.135
GPT-4-zeroshot	—	—	—	0.609	0.185	0.113
GPT-4-zeroshot-findings	\checkmark	—	—	0.887	0.314	0.278
GPT-4-fewshot-pos	—	\checkmark	—	0.820	0.147	0.120
GPT-4-fewshot-both	—	\checkmark	\checkmark	0.842	0.286	0.241
GPT-4-fewshot-findings-pos	\checkmark	\checkmark	—	0.857	0.263	0.226
GPT-4-fewshot-findings-both	\checkmark	\checkmark	\checkmark	0.842	0.366	0.308
HUMAN	—	—	—	0.940	0.232	0.218

Table 4: Human evaluation results for ad text generation models.

Model	Jaccard similarity
GPT-4-fewshot-pos	0.379
GPT-4-fewshot-both	0.358
GPT-4-fewshot-findings-pos	0.473
GPT-4-fewshot-findings-both	0.375

Table 5: Jaccard similarity between input and generated ad texts.

such as the active use of brackets. To verify this, we measured the Jaccard similarity between the input and generated text for models using only positive examples and those using both positive and negative examples. The results are shown in Table 5. Lower values indicate less lexical similarity between the input and generated text, suggesting that the output has been actively rephrased. These results further support the explanation. These findings suggest that considering both positive and negative examples can help the model learn the characteristics of attractive ad texts.

Attractiveness has room for improvement According to the overall success rate, the model (GPT-4-fewshot-findings-both), which incorporated the findings and both positive and negative examples, performed the best, outperforming HUMAN. However, room for improvement still exists in terms of attractiveness, even though the five models performed well in paraphrase identification. Therefore, it is worth exploring the factors that influence human preferences and the methods for generating more attractive ad texts in future work.

To examine whether the generated texts reflected the linguistic features provided by ICL, we analyzed the linguistic features of the generated texts. Specifically, we calculated the features identified in the analysis, including PPL, number of characters, and presence of brackets. Table 6 presents the

	PPL \downarrow	#Char \uparrow	Brackets \uparrow
GPT-4-zeroshot	288.5	12.9	0.180
GPT-4-zeroshot-findings	245.1	14.2	0.977
GPT-4-fewshot-pos	248.2	13.0	0.105
GPT-4-fewshot-both	284.5	13.6	0.128
GPT-4-fewshot-findings-pos	206.3	14.0	0.511
GPT-4-fewshot-findings-both	206.6	14.8	0.699
HUMAN	158.8	14.3	0.376

Table 6: Analysis results of linguistic features for generated ad texts. PPL and #Char represent the average perplexity and the number of characters of the generated texts, respectively. Brackets represent the proportion of generated texts that contain brackets.

results. These results confirm that GPT-4-fewshot-findings-both, the model with the highest overall success rate, tends to have a lower PPL, longer text length in characters, and a higher proportion of texts containing brackets, indicating that the generated texts successfully reflect the findings. Importantly, the results in Table 4 and Table 6 suggest that the linguistic features identified through the linguistic feature analysis (§4.1) contribute to enhancing the attractiveness of ad texts.

5 Analysis

Through experiments, we explored the linguistic features that influence human preferences and the methods for generating attractive ad texts. Because the primary goal of advertising is to capture users' attention, it is crucial to examine how attractively paraphrasing ad text affects ad performance, such as clicks. To this end, we analyzed the relationship between human preference and pCTR. In addition, we conducted an online evaluation of paraphrased ads in real-world ad-delivery scenarios.

	#Cases	#Aligned \uparrow	Ratio \uparrow
Entire evaluation set	725	355	0.490
high-agreement group	316	173	0.547
low-agreement group	409	182	0.445

Table 7: Alignment of human preferences with pCTR. #Aligned indicates the cases where ad texts favored by the majority of judges exhibited a higher pCTR.

Ad group	Imp \uparrow	Click \uparrow	Cost \uparrow
HR _{adg1}	1.628	1.233	1.531
Education _{adg2}	0.912	0.843	0.740
Education _{adg3}	0.627	0.365	0.401

Table 8: Relative improvements of our paraphrased ads compared to original ads in online evaluation.

5.1 Relationship between Human Preferences and pCTR

Table 7 lists the number of cases in which human preferences align with pCTR in the evaluation set. Following Hughes et al. (2019), we used an in-house CTR prediction model. We identified ad texts that were both preferred and had a higher pCTR. Only 355 of the 725 ad texts exhibited this alignment, suggesting a weak correlation between human preference and pCTR. For a detailed analysis, we divided the evaluation data into two groups based on the level of agreement. The high-agreement group, in which more than seven judges preferred the same ad, showed a 54.7% alignment, while the low-agreement group, which included all other cases, showed a 44.5% alignment. This suggests that the ad texts preferred by many judges tend to align with pCTR. Therefore, attractive rephrasing can positively affect pCTR, although to a limited extent.

5.2 Ad Performance via Online Evaluation

We conducted an online evaluation to verify the extent to which the paraphrasing method (§4.2) affected the ad performance. Specifically, we performed an A/B test to compare the performance of the original ads as a baseline with that of the paraphrased ads. See the link⁷ for the glossary of advertising terms. We used Google Ads’ search ads as the evaluation platform. For the evaluation data, we used three ad groups from two companies in the human resources (HR) and education industries. Each ad group consists of a maximum of 15

⁷<https://support.google.com/google-ads/topic/3121777>

headlines and four descriptions. We applied the paraphrasing methods to the headline ad text. We delivered the original and generated ads for two weeks and compared the number of impressions, clicks, and costs, which is the budget spent. The higher the value, the better for all metrics.

Table 8 presents the results of the online evaluation, indicating a relative improvement in the generated ads compared to the original ads. A value exceeding 1 indicates the superior performance of the generated ad.⁸ For example, if the baseline impressions were 100 and the generated text had 110, the improvement ratio will be $110/100 = 1.1$. According to Table 8, performance varies across ad groups; the ad performance for the HR company improved, whereas the others degraded. Future work should focus on conducting extensive online evaluations and developing paraphrasing methods that consistently enhance performance across diverse ad texts.

6 Discussion for Practical Application

As a potential practical application of this study, the development of a writing assistant tool based on our paraphrasing method could be considered to enhance the attractiveness of input ad texts. In advertising production, there are situations where the advertised content is fixed due to campaign requirements, and only the expression of the ad text needs to be refined. In fact, many tools already exist with the aim of supporting ad production in various ways.⁹ Such a writing assistant tool could further improve the wording and style of ad texts, making them more appealing.

7 Related Work

7.1 Generating Attractive Ad Text

With the increasing demand for online advertising, the manual creation of ad text has reached its capacity limits. Therefore, researchers have focused on ATG (Murakami et al., 2023). The goal of advertising is to attract interest in a product or service and motivate users to take action such as clicking and purchasing. Therefore, generating attractive ads is critical to the success of online advertising.

Various methods have been developed to generate attractive ad texts, ranging from template-based approaches (Bartz et al., 2008; Thomaidou et al.,

⁸Due to confidentiality, absolute values for each metric cannot be disclosed.

⁹<https://cyberagent.ai/products/>

2013) to neural-based techniques (Hughes et al., 2019; Kamigaito et al., 2021). Our work differs from previous studies in that it focuses on how the wording and style of ad text affect attractiveness. Unlike previous studies that evaluated generated ad texts without considering content variations, we ensured semantic equivalence between ad pairs before assessing their attractiveness. This approach allowed us to specifically analyze the impact of expressions on attractiveness. We believe that identifying the key linguistic factors that influence the attractiveness of ad texts is crucial for exploring the potential directions for advancing ATG methods.

7.2 Understanding Attractive Ad Text

Understanding the factors that affect the attractiveness of ad text is crucial to the success of the ad creation process. Various efforts have been made to analyze the factors influencing the attractiveness of ad texts, such as advertising appeal (Murakami et al., 2022), persuasive tactics (Yuan et al., 2023), and emotions (Youngmann et al., 2020). This study investigates the linguistic features of ad texts that affect their attractiveness.

A common approach for studying the factors influencing human preferences is to use log data, which measures attractiveness based on clicks and views. However, these log data are often proprietary and not publicly available, hindering research replication and knowledge advancement. Therefore, we collected human preferences through manual evaluations to make the data publicly available.

In a study closely related to our work, Pryzant et al. (2018) examined the impact of writing style on ad performance, while controlling for potential confounding variables. However, this study only considered the influence of the brand names and neglected other content-related factors. In our study, we constructed a paraphrase dataset varying in writing style to focus on the linguistic differences between ad texts while mitigating content variations, thereby enabling the analysis of human preferences centered on these linguistic features.

8 Conclusion

In this paper, we proposed a dataset called ADPARAPHRASE, which comprises semantically equivalent ad pairs, to analyze the linguistic features that influence human preferences. We collected data on human preferences for these ad pairs

and demonstrated that factors such as fluency, number of characters, and use of brackets significantly affected the attractiveness of the ad text. In addition, we conducted experiments on ATG, showing that considering human preferences can lead to the generation of attractive ad texts. In future work, we plan to expand our dataset to analyze other attractiveness factors and explore ATG methods.

9 Limitations

This study had several limitations that should be addressed in future studies.

Dataset size ADPARAPHRASE consists of 1,238 pairs of ad text, including 725 paraphrases. Although the limited size of the dataset has affected the robustness and generalizability of our findings, we ensured reliable results. Unlike many existing studies that use three human judges per pair to evaluate attractiveness, we asked ten human judges per pair to collect more comprehensive preference data. This approach helped mitigate the limitations of dataset size. In future work, we plan to expand the dataset to include a broader range of linguistic features in ad texts, thus further enhancing the depth and applicability of our research.

Linguistic features In this study, we analyzed a variety of linguistic features, including deep linguistic features that extend beyond surface-level features such as lexical choice and textual specificity. However, other unexplored features such as discourse intent were not considered. Therefore, future studies should consider these additional linguistic features to provide a more comprehensive understanding of the factors that influence human preferences.

Generalizability to other languages ADPARAPHRASE consists of Japanese ad texts, and all linguistic features and preference analyses are based on Japanese ad texts. Therefore, please note that some features, such as character types, are unique to the Japanese language, and we do not intend that our analysis results be applicable to all languages. However, we believe that other languages such as English and Chinese also have their unique linguistic features. Multilingualization of a dataset is an important future direction for revealing linguistic features that can be applied to a wide range of languages. In this study, we focused on Japanese ad texts, but we hope that this will pave the way for the

development of advertisement paraphrasing data in different languages.

Annotator demographics Although human preferences can be influenced by the demographic background of the annotators, we did not collect such information due to constraints within the crowdsourcing platform. Consequently, we did not consider the effects of demographic factors. Therefore, future studies should include an analysis of preferences that considers the demographic attributes of annotators to understand how these factors influence preference judgments.

References

- Kevin Bartz, Cory Barr, and Adil Aijaz. 2008. [Natural language generation for sponsored-search advertisements](#). In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 1–9.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Jan Cegin, Jakub Simko, and Peter Brusilovsky. 2023. [ChatGPT to replace crowdsourcing of paraphrases for intent classification: Higher diversity and comparable model robustness](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1889–1905.
- Patrick Fernandes, Aman Madaan, Emmy Liu, António Farinhas, Pedro Henrique Martins, Amanda Bertsch, José G. C. de Souza, Shuyan Zhou, Tongshuang Wu, Graham Neubig, and André F. T. Martins. 2023. [Bridging the Gap: A Survey on Integrating \(Human\) Feedback for Natural Language Generation](#). *Transactions of the Association for Computational Linguistics*, 11:1643–1668.
- J.L. Fleiss et al. 1971. Measuring nominal scale agreement among many raters. *Psychological Bulletin*, 76(5):378–382.
- Yebowen Hu, Kaiqiang Song, Sangwoo Cho, Xiaoyang Wang, Hassan Foroosh, and Fei Liu. 2023. [Decipher-Pref: Analyzing influential factors in human preference judgments via GPT-4](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 8344–8357.
- J. Weston Hughes, Keng-hao Chang, and Ruofei Zhang. 2019. [Generating better search engine text advertisements with deep reinforcement learning](#). In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 2269–2277.
- Tomoyuki Kajiwaru, Chenhui Chu, Noriko Takemura, Yuta Nakashima, and Hajime Nagahara. 2021. [WRIME: A new dataset for emotional intensity estimation with subjective and objective annotations](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2095–2104.
- Hidetaka Kamigaito, Peinan Zhang, Hiroya Takamura, and Manabu Okumura. 2021. [An empirical study of generating texts for search engine advertising](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Industry Papers*, pages 255–262.
- Hannah Kirk, Andrew Bean, Bertie Vidgen, Paul Rottger, and Scott Hale. 2023. [The past, present and better future of feedback learning in large language models for subjective human preferences and values](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2409–2430.
- J. Richard Landis and Gary G. Koch. 1977. [The measurement of observer agreement for categorical data](#). *Biometrics*, 33(1):159–174.
- Kikuo Maekawa, Makoto Yamazaki, Takehiko Maruyama, Masaya Yamaguchi, Hideki Ogura, Wakako Kashino, Toshinobu Ogiso, Hanae Koiso, and Yasuharu Den. 2010. [Design, compilation, and preliminary analyses of Balanced Corpus of Contemporary Written Japanese](#). In *Proceedings of the Seventh International Conference on Language Resources and Evaluation*, pages 1483–1486.
- Shaunak Mishra, Manisha Verma, Yichao Zhou, Kapil Thadani, and Wei Wang. 2020. [Learning to create better ads: Generation and ranking approaches for ad creative refinement](#). In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management*, pages 2653–2660.
- Masato Mita, Soichiro Murakami, Akihiko Kato, and Peinan Zhang. 2024. [Striking gold in advertising: Standardization and exploration of ad text generation](#). *Preprint*, arXiv:2309.12030.
- Soichiro Murakami, Sho Hoshino, and Peinan Zhang. 2023. [Natural language generation for advertising: A survey](#). *Preprint*, arXiv:2306.12719.
- Soichiro Murakami, Peinan Zhang, Sho Hoshino, Hidetaka Kamigaito, Hiroya Takamura, and Manabu Okumura. 2022. [Aspect-based analysis of advertising appeals for search engine advertising](#). In *Proceedings of the 2022 Conference of the North American*

- Chapter of the Association for Computational Linguistics: Human Language Technologies: Industry Track*, pages 69–78.
- OpenAI. 2024. [GPT-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems 35*, volume 35, pages 27730–27744.
- Reid Pryzant, Sugato Basu, and Kazoo Sone. 2018. [Interpretable neural architectures for attributing an ad’s performance to its writing style](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 125–135.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#). OpenAI Blog.
- Satoshi Sato, Suguru Matsuyoshi, and Yohsuke Kondoh. 2008. [Automatic assessment of Japanese text readability based on a textbook corpus](#). In *Proceedings of the Sixth International Conference on Language Resources and Evaluation*.
- Lingfeng Shen, Lemao 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*, pages 3178–3190.
- Kazuma Takaoka, Sorami Hisamoto, Noriko Kawahara, Miho Sakamoto, Yoshitaka Uchida, and Yuji Matsumoto. 2018. [Sudachi: a japanese tokenizer for business](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation*.
- Stamatina Thomaidou, Ismini Lourentzou, Panagiotis Katsivelis-Perakis, and Michalis Vazirgiannis. 2013. [Automated snippet generation for online advertising](#). In *Proceedings of the 22nd ACM International Conference on Information and Knowledge Management*, pages 1841–1844.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *Preprint*, arXiv:2307.09288.
- Hsueh-Cheng Wang and Marc Pomplun. 2012. [The attraction of visual attention to texts in real-world scenes](#). *Journal of Vision*, 12(6):26–26.
- Xiting Wang, Xinwei Gu, Jie Cao, Zihua Zhao, Yulan Yan, Bhuvan Middha, and Xing Xie. 2021. [Reinforcing pretrained models for generating attractive text advertisements](#). In *Proceedings of the 27th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 3697–3707.
- Penghui Wei, Xuanhua Yang, Shaoguo Liu, Liang Wang, and Bo Zheng. 2022. [CREATER: CTR-driven advertising text generation with controlled pre-training and contrastive fine-tuning](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Industry Track*, pages 9–17.
- Ikuya Yamada, Akari Asai, Hiroyuki Shindo, Hideaki Takeda, and Yuji Matsumoto. 2020. [LUKE: Deep contextualized entity representations with entity-aware self-attention](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6442–6454.
- Brit Youngmann, Elad Yom-Tov, Ran Gilad-Bachrach, and Danny Karmon. 2020. [The automated copywriter: Algorithmic rephrasing of health-related advertisements to improve their performance](#). In *Proceedings of The Web Conference 2020*, pages 1366–1377.
- Yuan Yuan, Fengli Xu, Hancheng Cao, Guozhen Zhang, Pan Hui, Yong Li, and Depeng Jin. 2023. [Persuade to click: Context-aware persuasion model for online textual advertisement](#). *IEEE Transactions on Knowledge and Data Engineering*, 35(2):1938–1951.
- Peinan Zhang, Yusuke Sakai, Masato Mita, Hiroki Ouchi, and Taro Watanabe. 2024. [AdTEC: A unified benchmark for evaluating text quality in search engine advertising](#). *Preprint*, arXiv:2408.05906.

A Details of Paraphrase Creation

We sampled the ad texts from the dev set of CAMERA (Mita et al., 2024), a publicly available dataset for ATG tasks.

The process of creating paraphrased candidates using CAMERA was divided into two steps. In

Please rephrase the provided ad text to make it more attractive according to the following conditions.

Conditions

- An ad text must be within 15 full-width characters.
- Do not add new information or remove existing information from a given ad text.

Figure 4: Instruction for paraphrase creation presented to human experts.

You are a professional copywriter responsible for creating search engine ads. Please rephrase the provided ad text to make it more attractive according to the following conditions

Conditions

- An ad text must be within 15 full-width characters.
- Do not add new information or remove existing information from a given ad text.

Answer

Input: {ad text}

Output: {paraphrased ad text}

Figure 5: Prompt for paraphrase creation presented to LLMs.

the first step, we asked two human experts working as professional ad writers in an advertising agency to rephrase their wording or styles to enhance the attractiveness of the given ad text. We instructed ad writers under the following two conditions: The first was not adding new information or deleting any existing information from the text. Second, the length of the ad text was limited to 15 full-width characters, which is the length limit for advertising delivery platforms such as Google Ads¹⁰. Figure 4 shows the instructions presented to the human experts for creating paraphrases. Due to their limited working hours, we asked them to create approximately 100 paraphrases, which resulted in 133 paraphrases.

In the second step, we used LLMs including GPT-3.5, GPT-4, and Llama2 to generate paraphrased ad texts from the source ad texts. Figure 5 shows the prompt for the paraphrase creation. A total of 133 source ad texts, identical to those presented to human experts, were used as inputs. This was performed to allow further analysis to directly compare the abilities of LLMs and human experts to generate paraphrases.

¹⁰<https://support.google.com/google-ads/answer/1704389>

Please determine whether the following pair of ad texts is a paraphrase.

Criteria

- Does the pair have the same meaning?
- Does the pair convey the same content in different expressions?

Ad Text #1

Ad Text #2

Yes No

Figure 6: User interface for paraphrase identification.

B Details of Paraphrase Identification

For the paraphrase candidates collected in §2.1, we conducted a manual annotation for paraphrase identification as a binary classification task to determine whether a pair was a *paraphrase* or *non-paraphrase*. Our defined criteria for paraphrases involved determining whether each candidate pair of ad texts was semantically equivalent at the sentence level rather than at the word or phrase level. The annotations were performed by five in-house workers with extensive experience in advertising creation. To ensure a consistent understanding of the annotation criteria and maintain quality control, we conducted two practice sessions in advance and provided comprehensive feedback on all the questions. The final gold label was determined by a majority vote of the five workers. The user interface and the instructions presented to the judges are presented in Figure 6, while the annotation examples presented to the judges are shown in Figure 7.

C Details of Attractiveness Evaluation

The attractiveness of each ad pair was evaluated by ten workers. Among them, four were workers with experience in in-house dataset production and the remaining six were workers from a crowdsourcing platform. We used Yahoo! Crowdsourcing¹¹ as the crowdsourcing platform. The compensation was set at 10 yen per 15 tasks, based on the platform’s regulations. They were informed that the annotated results would be used for research purposes. The user interface and instructions presented to judges are shown in Figure 8.

¹¹<https://crowdsourcing.yahoo.co.jp/>

Examples of non-paraphrase pairs

- (a)
 - Ad #1: Select your tax accountant at tax-attorney.com!
 - Ad #2: Find a tax attorney here!
 - Reason: The service name “tax-attorney.com” is not mentioned in Ad #2.
- (b)
 - Ad #1: Light and thin, but your back doesn't sink in
 - Ad #2: Light and thin mattress
 - Reason: The product name “mattress” is explicitly mentioned only in Ad #2, but it is implied in Ad #1.
- (c)
 - Ad #1: 14-day trial enrollment/closing soon
 - Ad #2: Free trial ends soon!
 - Reason: “14-day trial enrollment” and “free trial” refer to different concepts.

Examples of paraphrase pairs

- (d)
 - Ad #1: 60-day money-back guarantee
 - Ad #2: [Full Money Back Guarantee] 60-day trial
- (e)
 - Ad #1: A wide variety of products is here
 - Ad #2: See a wide variety of items
- (f)
 - Ad #1 Tarot reading / partially free
 - Ad #2 Tarot reading / partially free of charge

Figure 7: Annotation examples presented to judges for paraphrase identification.

D Details of Linguistic Features

This study analyzed a wide range of linguistic features, from surface-level features to deeper factors, including lexical choice and textual specificity. Our feature set consists of four groups: raw text, lexical, syntactic, and stylistic features, including the 24 types of linguistic features listed in Table 3.

In the following sections, we explain the definitions and extraction methods for each feature.

D.1 Raw text features

Text length We calculated the number of words and characters as a measure of text length. We used Sudachi (Takaoka et al., 2018), a Japanese morphological analyzer, to split Japanese texts into words.

D.2 Lexical features

Content words We hypothesized that ad texts with more content words would be more informative and attractive. We defined content words as nouns, verbs, adjectives, adjectival nouns, and adverbs and counted their occurrences in each ad text. For morphological analysis of Japanese ad texts, we used Sudachi (Takaoka et al., 2018).

Lexical choice For word frequency, we hypothesized that ad texts containing more common words

Please select more attractive ad text between the two ad texts. Here are some criteria you can consider.

Criteria

- Which do you want to click on?
- Which is more memorable?
- Which is more attractive?
- Which is more eye-catching?
- Which is easier to understand?
- Which is easier to read?

Notes
If the impression of both ads is the same, please select “Skip”.

Ad Text #1

Ad Text #2

Skip

Figure 8: User interface for attractiveness evaluation.

would be preferred; therefore, we calculated the frequency of words in each ad text. Specifically, we calculated the average frequency of occurrence per million words for each word in the ad text using the Balanced Corpus of Contemporary Written Japanese (Maekawa et al., 2010). Furthermore, to verify whether ad texts containing common or proper nouns were preferred, we used Sudachi to extract these nouns and counted their occurrences.

Character type We counted the numbers of hiragana, katakana, kanji, symbols, and numerals in each Japanese ad text.

D.3 Syntactic features

Dependency tree We used spaCy with GiNZA¹² to perform the dependency parsing of ad texts. The depth of the dependency tree refers to the longest path from the roots to leaves in the dependency tree. The length of the dependency link refers to the number of words between the syntactic head and its dependent.

Noun phrases We used spaCy with GiNZA to extract the noun phrases and counted their occurrences.

Perplexity We calculated the perplexity using GPT-2¹³ (Radford et al., 2019) trained on a web-crawled corpus and the Wikipedia dataset.

¹²<https://github.com/megagonlabs/ginza>

¹³<https://huggingface.co/rinna/japanese-gpt2-medium>

D.4 Stylistic features

Emotion To label the emotion of ad texts, we used the LUKE model¹⁴ (Yamada et al., 2020) trained on WRIME (Kajiwara et al., 2021), which is a Japanese emotion analysis dataset based on social media text. This model is an 8-class classifier that determines the most appropriate emotions from the following eight categories: joy, sadness, anticipation, surprise, anger, fear, disgust, and trust. When we applied the classifier to the ad texts, the majority of cases were labeled as either “joy” or “anticipation,” so we used only these two labels in our analysis. The accuracy of the classifier was 68.6%.

Textual specificity Owing to the lack of datasets or existing models for measuring the specificity of Japanese sentences, we created our own specificity classifier using GPT-4 with a few-shot setting. We defined this as a three-class classification problem and constructed a model that compares two ad pairs and judges with higher specificity. If the specificity of both is equivalent, it outputs a label of “equal.” To evaluate the performance of the model, we randomly sampled 100 predictions and conducted a manual evaluation, achieving an accuracy of 88.0%.

Brackets We created a binary label indicating whether the ad text contains bracket symbols 「」 or 「」.

E Details of Ad Text Generation Experiment

Figure 9 shows an example of the prompts used in the ATG experiment. Specifically, this prompt was used for GPT-4-fewshot-findings-both in Table 4. In this prompt, we include few-shot examples and the findings revealed from the analysis of linguistic features that influence human preferences. For the few-shot examples, we created two types of examples, where attractiveness improved and where it did not, from the results of the attractiveness evaluation and used 20 examples for each type. We report the results of a single experiment.

F Implementation Details

Models We used GPT-3.5, GPT-4, and Llama2 to perform the paraphrase candidate generation (§2.1) and an experiment using ATG

¹⁴<https://huggingface.co/Mizuiro-sakura/luke-japanese-large-sentiment-analysis-wrime>

```
You are a professional copywriter responsible for creating search engine ads. Please rephrase the provided ad text to make it more attractive according to the following conditions.

# Conditions
- An ad text must be within 15 full-width characters.
- Do not add new information or remove existing information from a given ad text.

# Tips for Creating Attractive Ad Texts Findings
- Utilize bracket symbols such as 「」 or 「」.
- Use the full character limit (15 full-width characters).
- Make the text more fluent.

# Examples Few-shot examples
## Examples where attractiveness did not improve (20 cases)
Input: Save from 5,000 yen per month
Output: Possible to save from 5,000 yen per month
      :
## Examples where attractiveness improved (20 cases)
Input: Over 100,000 episodes of popular free comics
Output: 【Free】 Over 100,000 episodes of popular comics
      :

Refer to the examples where attractiveness improved and rephrase a given ad text to make it more attractive.

# Answer
Input: {ad text}
Output: {paraphrased ad text}
```

Figure 9: Prompt for ad text generation experiment. For visibility, Japanese prompt is translated into English.

(§4.2). For GPT-3.5 and GPT-4, we used the Azure OpenAI API¹⁵ service with version 2024-03-01-preview and a default temperature parameter of 1.0. In the case of Llama2, we used the ELYZA-japanese-Llama-2-7b model available via Hugging Face¹⁶, which enhances Japanese-language capabilities through additional pre-training based on Llama2. See link¹⁷ for the detailed parameters of the Llama2 model.

Tokenizer We used Sudachi (Takaoka et al., 2018), a Japanese morphological analyzer, to tokenize Japanese text into words. Sudachi offers three splitting modes to provide tokens of different granularities for each application purpose: short, medium, and named entities. In this study, we used the named entity mode.

¹⁵<https://azure.microsoft.com/en-us/products/ai-services/openai-service/>

¹⁶<https://huggingface.co/>

¹⁷<https://huggingface.co/elyza/ELYZA-japanese-Llama-2-7b>