

ChatGPT vs Human-authored Text: Insights into Controllable Text Summarization and Sentence Style Transfer

Dongqi Liu and Vera Demberg

Department of Computer Science

Department of Language Science and Technology

Saarland Informatics Campus, Saarland University, Germany

{dongqi, vera}@lst.uni-saarland.de

Abstract

Large-scale language models, like ChatGPT, have garnered significant media attention and stunned the public with their remarkable capacity for generating coherent text from short natural language prompts. In this paper, we aim to conduct a systematic inspection of ChatGPT’s performance in two controllable generation tasks, with respect to ChatGPT’s ability to adapt its output to different target audiences (expert vs. layman) and writing styles (formal vs. informal). Additionally, we evaluate the faithfulness of the generated text, and compare the model’s performance with human-authored texts. Our findings indicate that the stylistic variations produced by humans are considerably larger than those demonstrated by ChatGPT, and the generated texts diverge from human samples in several characteristics, such as the distribution of word types. Moreover, we observe that ChatGPT sometimes incorporates factual errors or hallucinations when adapting the text to suit a specific style.¹

1 Introduction

Generative **P**re-trained **T**ransformer (GPT; *e.g.*, ChatGPT) models, which produce results from given conditional input prompts, have exhibited exceptional performance on various natural language understanding (NLU) and generation (NLG) tasks (Jiao et al., 2023; Wang et al., 2023a; Bang et al., 2023b; Zhou et al., 2023; Dai et al., 2023). For instance, in NLU tasks, Qin et al. (2023) have proved that ChatGPT is comparable to state-of-the-art fine-tuning models in language reasoning. In NLG tasks, Yang et al. (2023a) assessed four widely used benchmark datasets, such as QMSum, and confirmed ChatGPT’s comparability to traditional fine-tuning methods. Peng et al. (2023) further investigated effective strategies for machine translation using ChatGPT and highlight its strong

translation ability. Additionally, ChatGPT can even facilitate multi-modal tasks (Yang et al., 2023b; Shen et al., 2023), as well as the application of data augmentation (Dai et al., 2023). Although the studies mentioned above have demonstrated notable performance of ChatGPT across different domains, there remains a dearth of qualitative and quantitative evaluation of the texts generated by ChatGPT. Such an evaluation is vital to uncover the behavioral differences, potential limitations, and challenges associated with ChatGPT-generated texts, especially when compared with human-authored texts.

Controllable text generation seems to be a task in which ChatGPT-like models could potentially excel. This task is driven by the desire to tailor text for a diverse array of target users (*e.g.*, experts and laypersons) (Kumar et al., 2022; Cao et al., 2020; Luo et al., 2022), and thereby enhancing the accessibility of textual information. In controllable text generation, one delineates a particular set of parameters or provides a prompt that defines the intended target style. This area has recently received growing interest from researchers in the field (Hu and Li, 2021; Li et al., 2022; Zhang et al., 2022; Dathathri et al., 2019a; August et al., 2022; Carlsson et al., 2022; Gu et al., 2022; Li et al., 2022; Keskar et al., 2019; Dathathri et al., 2019b). The traditional natural language generation task (Pu and Sima’an, 2022), which focuses solely on adequately responding with respect to a given input, can be regarded as a special case of controllable natural language generation, wherein the control setting remains unconditioned. Considering ChatGPT as the most recent language generation capability, the assessment of its language generation proficiency, specifically in the realm of controllable language generation, remains largely uncharted. Therefore, our study delves into two distinct applications of ChatGPT, namely controllable summary generation and sentence style trans-

¹The project information of our study can be accessed at https://dongqi.me/projects/ChatGPT_vs_Human.

fer. In the former, we examine ChatGPT’s ability to generate summaries that cater to two distinct readerships, namely experts and non-experts, for a given academic literature. Concerning sentence style transfer, we investigate ChatGPT’s capability to generate both formal and informal sentences for the task of sentence formality.

The objective of this study is to tackle the research question: **In relation to the human-produced text, to what extent does ChatGPT-generated content demonstrate significant divergence from human behavior and the potential susceptibility to inaccuracies?** Our primary contributions are enumerated below:

- To the best of our knowledge, we are the first to utilize ChatGPT to evaluate its effectiveness in controllable text generation.
- Our findings indicate that there are substantial performance disparities between the text generated by ChatGPT and that generated by humans.
- Our study exposes and quantifies the existence of numerous hard-to-spot errors in the text generated by ChatGPT, which have a tendency to amplify with successive transformations of the text.

2 Related Work

2.1 Controllable Text Summarization

Controllable text summarization is a rapidly evolving field that aims to produce summaries with specific characteristics, such as length, style, or content (Shen et al., 2022b; Chan et al., 2021; Sarkhel et al., 2020; Shen et al., 2022a; Goldsack et al., 2022; Keskar et al., 2019; Dathathri et al., 2019b; He et al., 2022; Earle et al., 2021; Liu et al., 2022b). A range of approaches has been proposed for this task, including the use of sequence-to-sequence models such as the Transformer model (Vaswani et al., 2017). These models have demonstrated promising progress in producing high-quality summaries that can be modulated according to specific requirements (Fan et al., 2018; Wu et al., 2021; Amplayo et al., 2021). Additionally, other techniques also have been proposed to enhance the controllability of the summaries, such as conditional generation (He et al., 2022; Luo et al., 2022), prompt-based summarization (Yang et al., 2022; Liu et al., 2022a; Zhang and Song, 2022), and multi-task learning (Cui and Hu, 2021; Gu et al., 2022).

2.2 Text Style Transfer

Text style transfer is a task that involves transforming an input sentence into a desired style while retaining its style-independent semantics (Jin et al., 2022; Zhu et al., 2021; Dai et al., 2019; Li et al., 2020; Babakov et al., 2022; Mir et al., 2019; Ramesh Kashyap et al., 2022; Tokpo and Calders, 2022). To achieve this, prior research has examined sequence-to-sequence learning strategies that utilize parallel corpora with paired source/target sentences in different styles (Cheng et al., 2020; Hu et al., 2021; Nouri, 2022). Owing to the considerable demand for human resources and material investments in data labeling, parallel data across diverse styles are scarce. This has led to an increased interest in exploring more pragmatic situations where only non-parallel stylized corpora are accessible (Malmi et al., 2020; Reif et al., 2022).

2.3 ChatGPT

ChatGPT² is a large language model (LLM), which is built upon the innovations and improvements of its predecessors, such as GPT-3³. In terms of training strategies, ChatGPT employs instruction learning and reinforcement learning from human feedback (RLHF; Ouyang et al., 2022) to enhance its overall performance and adaptability.

Upon its emergence, ChatGPT has garnered considerable attention from researchers, who have undertaken initial studies into the model. Scholars such as Baidoo-Anu and Owusu Ansah (2023); Rudolph et al. (2023); West (2023); Sobania et al. (2023); Gilson et al. (2023); Lai et al. (2023); Wang et al. (2023b) have explored the notable strengths of ChatGPT from the fields of education, science, programming, healthcare, and text generation, respectively. However, Bang et al. (2023a) discovered that ChatGPT suffers from hallucination issues in the context of logical reasoning. Due to its immense and inaccessible training corpus and parameters, and the inability to access external knowledge for reliable sources of support, it is imperative to question whether ChatGPT demonstrates the same hallucination issue as other LLMs when performing sentence generation. Based on these clues, we firmly assert that in-depth analysis of the text generated by ChatGPT and its behavioral patterns are both significant and valuable, and can provide meaningful insights to the readers of this paper.

²<https://openai.com/blog/chatgpt>

³<https://openai.com/research/instruction-following>

3 Study on Controllable Summarization

3.1 Prompt Formulation

In this section, our main objective is to test the zero-shot performance of ChatGPT on controllable summarization, with the goal to generate summaries for laymen vs. experts. To this end, we constructed several prompts as natural language instructions for ChatGPT. The prompts we tested include for the layman style: *Please give me a layman / simple / simplified and understandable / easy-to-comprehend / straightforward / general audience summary of X*, where *X* was replaced by the source text that should be summarized. Similarly, for the expert summary, we experimented with the prompts: *Please give me an expert / a technical / comprehensive and detailed / difficult-to-comprehend / in-depth / complicated summary of X*.

3.2 Experimental Setup

For all experiments, we used ChatGPT *gpt-3.5-turbo*, which was, at the time of experimentation, the best-performing publicly accessible version provided by OpenAI. For the hyper-parameter setting, we set temperature = 0, top p = 1, frequency penalty = 0.2, and presence penalty = 0.2. For summary generation, we configured the maximum number of generated tokens to 512. The remaining hyper-parameters were set to their default values as recommended by OpenAI. It is noteworthy that ChatGPT has the potential to generate empty responses (i.e., empty strings) as the result of network transmission timeouts or API request overloads. Should this arise, we adhere to the established practice of resubmitting the request until ChatGPT provides non-empty responses.

All of our experiments were conducted on the version of ChatGPT between 15 Feb 2023 and 30 Apr 2023 by using the OpenAI’s ChatGPT API.⁴ We should emphasize that to prevent any potential interference from the prior responses, we cleared the conversation history each time we submit a new query to ChatGPT. Unless otherwise specified, we refrained from engaging in any further conversation with ChatGPT to modify its responses.

3.3 Dataset

We selected ELIFE (Goldsack et al., 2022) dataset for our experiments. It contains summaries of aca-

demical literature that exhibit varying levels of readability, tailored to suit either expert or non-expert audiences. By means of this dataset, we can examine to what extent ChatGPT can regulate the summary generation process in accordance with the intended target users, and compare its summaries to human summaries.

3.4 Metrics

In order to assess automatically whether ChatGPT summaries substantially differ in terms of their audience design based on the given prompt, we opted for a set of three automatic readability metrics: Flesch Reading Ease (FRE; Kincaid et al., 1975), Coleman-Liau Index (CLI; Coleman and Liau, 1975), and Dale-Chall Readability Score (DCR; Chall and Dale, 1995).

The Flesch Reading Ease (Kincaid et al., 1975) is a metric that gauges the comprehensibility of a given text. This index relies on the average number of syllables per word and the average number of words per sentence. A higher score signifies an easier-to-understand text. Additionally, the Coleman-Liau Index (Coleman and Liau, 1975) is a measure of the text’s difficulty level, which considers the average number of characters per sentence and the average number of sentences per 100 words. A higher score indicates a more challenging text. The Dale-Chall Readability Score (Chall and Dale, 1995) is computed by comparing the number of complex words in the text with a list of common words. A higher score denotes a more challenging text.

We also employed Rouge scores (Lin, 2004) to evaluate the performance of ChatGPT in the task of text summarization, with the aim of comparing its efficacy against the state-of-the-art model. In order to assess the extent to which the summaries re-use word sequences from the original text, we furthermore evaluated N-gram novelty (See et al., 2017; Gehrmann et al., 2019; Pu et al., 2022). Finally, we quantified inconsistency based on factual consistency checking metric SummaC (Laban et al., 2022), as well as hallucination checking metric (Cao et al., 2022; Fischer et al., 2021). SummaC (Laban et al., 2022) uses sentence compression and summarization techniques to extract important information and improve the detection of inconsistencies in NLI models by segmenting documents and aggregating scores. Named entity hallucination (Fischer et al., 2021) flags potential hallucinations

⁴<https://platform.openai.com/overview>

in named entities if they do not match the original sources. We here used BERT semantic similarity, rather than exact matching, when computing the named entities matching.

3.5 Results on Controllable Summarization

3.5.1 Effect of Prompt Formulation

Table 1 illustrates that different prompt versions are somewhat consistent regarding whether the instructions asking for layman summaries actually lead to more readable texts than those asking for expert summaries, with FRE ranging between scores of 31 and 38 for automatically generated layman summaries, and between 28 and 37 for automatically generated expert summaries. Conversely, human-written summaries exhibit very large differences according to the automatic metrics, with FRE of 53.1 for layman summaries and 22.5 for expert summaries. Similar effects are observed for the CLI and DCR measures. This preliminary test was conducted on a subset of the ELIFE dataset, containing merely 500 random samples; for the rest of the tests, we proceeded to the entire dataset, selecting the prompts asking for “layman” and “expert” summaries, as responses for these prompts seemed to align with the right direction wrt. the readability measures.

Prompt version	FRE	CLI	DCR
layman	37.26 [†]	14.82 [†]	11.21 [†]
simple	31.92 [†]	15.70 [†]	11.54 [†]
simplified and understand.	35.48 [†]	15.17 [†]	11.21 [†]
easy-to-comprehend	36.59 [†]	14.93 [†]	11.32 [†]
straightforward	31.74 [†]	15.58 [†]	11.42 [†]
general audience	35.86 [†]	14.98 [†]	10.96 [†]
human answer (for layman)	53.06	12.36	8.90
expert	29.89 [†]	15.91 [†]	11.88 [†]
technical	36.65 [†]	13.76 [†]	12.20 [†]
comprehensive and detailed	31.62 [†]	15.47 [†]	11.15 [†]
difficult-to-comprehend	28.95 [†]	16.14 [†]	11.71 [†]
in-depth	34.37 [†]	14.93 [†]	10.82 [†]
complicated	29.05 [†]	15.76 [†]	11.40 [†]
human answer (for expert)	22.54	17.65	11.79

Table 1: Reading difficulty on different prompts, tested on a set of 500 randomly selected items from the dataset. [†] indicates statistical significance ($p < 0.05$) against corresponding human answers via paired t-test.

3.5.2 Reading Difficulty Control

Table 2 corroborates that the results of the whole dataset are consistent with the findings from the smaller sample. We conclude that ChatGPT can

produce summaries with different levels of reading difficulty to a certain extent based on the provided prompts. Notably, ChatGPT-generated sentences for expert-style summaries show greater complexity than those for layman-style summaries. However, the magnitude of the difference in the reading difficulty scores between the two types of summaries is considerably smaller than that observed in human-written summaries.

Candidate	FRE	CLI	DCR
Human Layman	52.42	12.46	8.93
Human Expert	23.20	17.62	11.78
ChatGPT Layman	37.38 ^{†‡}	14.78 ^{†‡}	11.17 ^{†‡}
ChatGPT Expert	30.38 ^{†‡}	15.82 ^{†‡}	11.85 ^{†‡}

Table 2: Reading difficulty scores by automatic metrics; [†] and [‡] indicate statistical significance ($p < 0.05$) against same-style human answers, and opposite-style ChatGPT answers via paired t-test, respectively.

3.5.3 Comparison to Previous SOTA Model

We also compared summaries generated by ChatGPT to a previous state-of-the-art (SOTA) neural fine-tuned summarization model (Pu et al., 2023). On the same test split, the summaries produced by ChatGPT reached Rouge-1=25.53, Rouge-2=5.48, Rouge-L=13.30 under unsupervised learning, and Rouge-1=47.88, Rouge-2=13.75, Rouge-L=42.44 in few-shot learning use the training samples from the same subset of Section 3.5.1, while the model by Pu et al. (2023) reached Rouge-1=48.70, Rouge-2=14.84, and Rouge-L=46.13.

3.5.4 Disparities in Summarization Behavior

We next examined whether ChatGPT and Humans are consistent with each other regarding the readability of summarization with respect to different items – it could be possible, that some texts simply lead to less readable summaries than others. However, we discovered that Pearson correlations of FRE scores for summaries by humans and ChatGPT were only 0.31 for expert summaries, and 0.2 for layman summaries. (Scores were similarly low for the CLI and DCR metrics.) In addition, the statistical significance test elucidates the noteworthy divergence between the distinctive response styles produced by ChatGPT and the analogous styles of human-generated answers.

Following this, we contrasted the n-gram novelty of human vs. ChatGPT summaries wrt. the original texts. Figure 1 reveals that a significantly higher

number of novel 4-grams are present in human-written summaries, particularly those aimed at laymen. This suggests that ChatGPT summaries are slightly more extractive compared to human summaries.

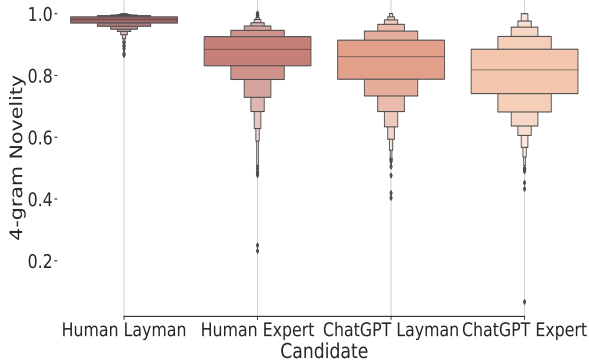


Figure 1: Comparison of abstractiveness between ChatGPT and human-generated summaries

3.5.5 Inconsistencies and Hallucinations

Given that ChatGPT has previously been reported to generate misinformation, we sought to evaluate its risk of hallucinating on our specific task. Figure 2 demonstrates that the SummaC consistency scores are lower for ChatGPT-generated summaries than for human-written summaries. A corresponding phenomenon is verified in the hallucination assessment. Precision scores provided in Table 3 demonstrates the extent to which ChatGPT-generated text contains named entities that are absent in the source text. A lower precision score suggests that the generated text has more named entities that lack support in the source text. The recall scores reflect the ability of ChatGPT to capture named entities from the source text. A lower recall score implies that ChatGPT has missed a considerable number of named entities from the source text. F1 score represents the harmonic mean of the precision and recall scores. By examining Table 3, our findings demonstrate that ChatGPT generates a greater number of named entities that are not present in the source text after undergoing multiple iterations of text conversions and modification. For example, in an expert summary, ChatGPT misinterpreted the meaning of “Geocode” as “regional regulations”.

3.6 Intermediary Discussion

Our experiments show that ChatGPT-generated summaries do not adapt as strongly to the target audience as human-authored summaries. One pos-




Figure 2: Summary consistency detection. L stands for layman, E for expert.

Candidate	Precision	Recall	F1
Human Layman	0.78	0.63	0.70
Human Expert	0.92	0.61	0.73
ChatGPT Layman	0.75 [‡]	0.47 [†]	0.58 [†]
ChatGPT Expert	0.90 [‡]	0.49 [†]	0.63 [†]
ChatGPT L2E2L	0.74 [‡]	0.39 ^{†‡}	0.51 ^{†‡}
ChatGPT E2L2E	0.88 [‡]	0.47 ^{†‡}	0.62 ^{†‡}

Table 3: Named entity hallucination on Elife dataset. [†] and [‡] indicate statistical significance ($p < 0.05$) against same-style human answers, and opposite-style ChatGPT answers via paired t-test, respectively. L stands for layman, E for expert.

sible reason could be that ChatGPT, given the zero-shot setting, had no way to “know” how strongly the texts should be adapted to the target style. Furthermore, we identified evidence for potential hallucinations generated during summarization. We, therefore, carried out two post-hoc experiments: (1) We modified the prompt to include an example from the dataset, so ChatGPT would have a chance to know the expected level of text adaptation. (2) We subjected the resulting summaries to several re-writing steps and test whether this further intensifies the occurrence of hallucinations.

3.6.1 Follow-up Experiment: Example Inclusion in Prompt

We experimented with prompts that also include a human summary example. Unlike the previous few-shot learning experiment, we do not adjust the parameters of the ChatGPT, but just let the model perform unsupervised reasoning through the contents of the prompt. We observe (see Appendix Table 7) that when guided by a human example from the dataset, the summaries generated by ChatGPT indeed tend to be more aligned with human

performance, particularly on the Flesch Reading Ease metric (49.23 for layman, 28.88 for expert summaries). However, no significant changes are detected in other metrics. The degree of control over the summarization style has increased, yet it remains inferior to human capabilities.

3.6.2 Follow-up Experiment: Repeated Re-writing

Summaries are further re-written based on the prompt *Please give me a **layman/expert** style version of X*, where *X* was the previously generated summary. Figure 2 and Table 3 display the performance of ChatGPT after re-writing in the entries “ChatGPT L2E2L” and “ChatGPT E2L2E” which stand for the order in which instructions were given (L stands for layman, and E for expert). The examinations point out that misinformation and hallucinations may be further increased during subsequent rewriting (lower SummaC scores, lower values in the named entity hallucination metric).

4 Study on Text Formality Transfer

4.1 Prompt Formulation and Experimental Setup

Our subsequent set of experiments investigates ChatGPT’s capacity for style transfer concerning language formality. Our prompt for this task was formulated as *Please give me a **formal** / **an informal** version of X*. We utilized the same experimental setup as for the summarization task; however, we restricted the maximum number of generated tokens to 32. We again experimented with various prompts, as shown in Table 4 below. Unless otherwise specified, all experiments used the same configuration.

4.2 Dataset

We investigated whether ChatGPT can proficiently execute style transfer on sentences using data from the GYAFC (Rao and Tetreault, 2018) dataset. The dataset has two branches, Entertainment & Music (EM) and Family & Relationships (FR). With the aid of this dataset, we aim to evaluate ChatGPT’s ability for sentence style transfer, examine the differences in vocabulary selection and syntactic structures between ChatGPT and human performance, and identify the limitations of ChatGPT.

4.3 Metrics

To evaluate the level of formality in the generated text, we utilized Text Formality Score (Heylighen

and Dewaele, 1999) and MTLD Lexical Diversity (McCarthy and Jarvis, 2010) metric. The Text Formality Score (Heylighen and Dewaele, 1999) is a metric that quantifies the degree of formality in language usage within a text, based on the adherence to formal linguistic norms. Another measure that evaluates language formality is the MTLD Lexical Diversity metric (McCarthy and Jarvis, 2010). This index measures the diversity and richness of the vocabulary used in the text, based on the frequency and number of unique words. A higher MTLD score indicates a greater variety of vocabulary, which typically corresponds to a more formal language style. We also utilized BLEU (Papineni et al., 2002) score to draw a comparison between ChatGPT and SOTA approach. We additionally assessed the distribution of POS tags in the generated different styles, as well as the distribution of dependency labels⁵. For quantifying misinformation and hallucinations, we used DAE and named entity hallucination checking. The DAE algorithm (Goyal and Durrett, 2020) utilizes dependency arcs to identify entailment relationships between propositions and identify inconsistencies in factual information based on syntactic and semantic structures.

4.4 Results on Formality Control

4.4.1 Effect of Prompt Formulation

Table 4 presents the results for a set of 500 random samples from the GYAFC dataset. We observe that the Formality scores are very similar for ChatGPT formal vs. informal texts. We note however that the difference in ratings for human-written texts is also small for this metric. The MTLD metric on the other hand shows higher values for ChatGPT-generated formal texts; in fact, the scores are substantially larger than those of human-written texts, but differ not much from each other. We therefore proceed with the prompts using the formulation formal/informal for the rest of the experiments on the whole dataset.

4.4.2 Sentence Formality Control

Table 5 offers supplementary evidence from the full dataset supporting ChatGPT’s capacity to modify the formality level of sentences. By employing the Formality indicator (Heylighen and Dewaele, 1999), it is apparent that the generated text tends to manifest a higher level of formality overall. A primary factor contributing to this result is the pre-

⁵<https://spacy.io/>

Prompt version	Formality	MTLD
informal	51.09	13.22 [†]
unprofessional	51.20	16.23 [†]
spoken version	51.30 [†]	14.47 [†]
easygoing	51.43 [†]	14.11 [†]
casual	51.00	16.30 [†]
laid-back	51.27	13.94 [†]
human answer (for informal)	50.76	11.42
formal	52.22 [†]	31.23 [†]
professional	51.96 [†]	31.98 [†]
written	51.62 [†]	29.69 [†]
stately	51.30 [†]	34.43 [†]
grandiose	52.85 [†]	30.71 [†]
majestic	52.23 [†]	33.49 [†]
human answer (for formal)	53.92	14.99

Table 4: Text formality on different prompts, tested on a set of 500 randomly selected items from the dataset. [†] indicates statistical significance ($p < 0.05$) against corresponding human answers via paired t-test.

disposition of ChatGPT’s training corpus towards written sources, encompassing materials such as books and news articles, as opposed to spoken language corpora (OpenAI, 2023). This perspective is further corroborated by an examination of the generated sentence samples. The MTL D metric underscores that ChatGPT’s lexical diversity is considerably lower when generating informal sentences, but shows a marked increase when generating formal sentences.

Dataset	Candidate	Formality	MTLD
GYAFC-FR	Human Informal	49.87	15.20
	Human Formal	53.57	18.70
	ChatGPT Informal	50.77 ^{†‡}	14.60 [‡]
	ChatGPT Formal	52.06 ^{†‡}	31.68 ^{†‡}
GYAFC-EM	Human Informal	50.11	12.11
	Human Formal	53.76	15.82
	ChatGPT Informal	51.02 ^{†‡}	12.01 [‡]
	ChatGPT Formal	51.98 ^{†‡}	29.80 ^{†‡}

Table 5: Text formality scores by automatic metrics; [†] and [‡] indicate statistical significance ($p < 0.05$) against same-style human answers, and opposite-style ChatGPT answers via paired t-test, respectively.

4.4.3 Comparison to Previous SOTA Model

We also find that ChatGPT outperforms the previous supervised SOTA model (Nouri, 2022) by training on the same subset at Section 4.4.1 for few-shot learning, as evident from the higher BLEU score. Specifically, ChatGPT yields superior scores of

0.711 and 0.697 in the EM and FR branches, as compared to the SOTA model’s scores of 0.671 and 0.652. However, ChatGPT achieved only 0.07 and 0.06 BLEU scores on the EM and FR branches, respectively, in the unsupervised setting.

4.4.4 Effect of Example Inclusion in Prompt

We again examined the impact of including an example of the dataset into the prompt and find that this again helps ChatGPT slightly with matching the dataset style (with details provided in Table 8). Specifically, the formality score for the informal style is 50.67, while it climbs to 52.13 for the formal style, with the MTL D score also displaying an increase from 14.81 for informal texts to 19.22 for formal texts.

4.4.5 Disparities in Style Transfer Behavior

In terms of controlling the formality of sentence style, ChatGPT’s performance still exhibits significant differences compared to human behavior. While the by-item correlation is slightly higher for this dataset than for the summary task (Pearson correlation of around 0.4 for formal style and 0.5 for informal style on the Formality metric; 0.3 for MTL D measure), there are interesting disparities between the distributions of POS tags between ChatGPT and humans. The examination of statistical significance further substantiates our antecedent observation, indicating a substantial disparity between the different response styles engendered by the model, as well as between the answers conforming to the same styles exhibited by humans.

Figure 3 illustrates the absolute differences in the distribution of Part-of-Speech (POS) tags. Based on this figure, it is evident that ChatGPT employs a higher frequency of adjectives, adpositions, determiners, and nouns in the generation of formal sentences when compared to those produced by human writers. Conversely, in the generation of informal sentences, ChatGPT tends to utilize more auxiliary words and punctuation marks. These variances in word choice between formal and informal styles, as exemplified by ChatGPT, are indicative of differences in its selected vocabulary for distinct stylistic modes compare with humans.

By analyzing the distribution of dependency labels (Appendix Figures 5, 6, 7, 8), it is also clear that, in comparison to human-authored sentences, ChatGPT utilizes a greater frequency of adjectival modifiers, auxiliaries, determiners, objects of the preposition, and prepositional modifiers for formal

sentences. Contrarily, compounds and dependents are infrequently employed in the generation of informal sentences by ChatGPT.




Figure 3: Absolute differences in POS tags distribution of ChatGPT and human-generated sentences: GYAFC - EM

4.4.6 Inconsistencies and Hallucinations

In order to assess the risk of introducing erroneous information when ChatGPT performs sentence style transformation, we employed DAE (Goyal and Durrett, 2020) at the sentence level to examine the factuality after text style transformation, and compare again the effect of multiple re-writes. Similar to before, F denotes formal style, I signifies informal style, and X2X2X ($X \in \{F, I\}$) represents multiple rewriting transformations of the text. The outcomes of our inquiry are depicted in Figure 4, and Appendix Figure 14. We also again scrutinized the potential incorporation of hallucinatory information regarding named entities in the ChatGPT-generated text, and the findings are presented in Appendix Table 9.

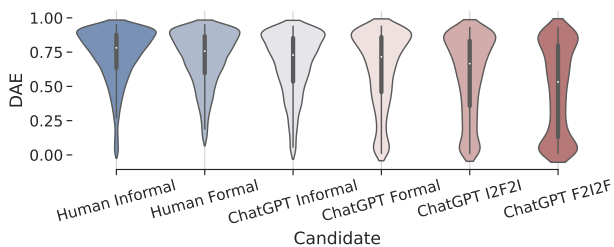


Figure 4: Dependency arc entailment: GYAFC - EM. Data points $>0.95 \approx$ Accurate. To clarify discrepancies, cutoff point = 0.95.

Upon conducting factuality checking (see Figure 4, and Appendix Figure 14), it is discovered that ChatGPT’s performance is inferior to that of humans in sentence-style rewriting. Interestingly, with the increase in the number of text conversions and rewritings, ChatGPT’s tendency to commit factual errors escalates while the output increasingly deviates from the original text, compromising the fidelity of the final result. In a particular instance, the human-generated formal expression states “She

is a poor vocalist”, whereas the formal rendition provided by ChatGPT articulates “She does not possess the ability to sing”. This discrepancy represents a significant semantic alteration. The degree of dependency arc entailment is low in this case. Similarly, Appendix Table 9 reveals that recall scores on the named entity hallucination metric are lower in ChatGPT sentences than in human sentences.

4.4.7 Qualitative Examples

To explore whether ChatGPT-generated sentences significantly alter the original semantics of the input text, we conducted a case study by randomly selecting 15 samples from each branch of the GYAFC dataset. Our findings indicate that ChatGPT poses a relatively severe risk of modifying the original semantics during sentence style transformation, with approximately 18% of the samples exhibiting noticeable semantic inconsistencies. The examples in Table 6 reveal that during the process of sentence style transfer, ChatGPT erroneously modifies the content words, resulting in significant semantic alterations.

Formal to Informal	
It is such a waste of TV space.	(Reference)
Yes, because it’s such a waste of TV space!	(Human)
What a total waste of TV time!	(ChatGPT)
The other boy isn’t that great.	(Reference)
The other boy is not that good.	(Human)
The other kid’s not so hot.	(ChatGPT)
I really enjoy how the composition has the tec...	(Reference)
I really like how they do like the whole techn...	(Human)
I’m diggin’ how the techno beat slows down in ...	(ChatGPT)
Informal to Formal	
Fatboy Slim - Right Here, Right Now Or any oth...	(Reference)
Fatboy Slim is right here and now. He Rocks!	(Human)
Fatboy Slim’s "Right Here, Right Now" is an ex...	(ChatGPT)
loved them since their first album.	(Reference)
I have loved them since their first album.	(Human)
I have held a fondness for them since the rele...	(ChatGPT)
if u occasionally doing it then u alrady r add...	(Reference)
If you occasionally do it, then you are already...	(Human)
If you are engaging in the activity on a regul...	(ChatGPT)

Table 6: Case study of ChatGPT generated output

Furthermore, our examination of the visualized dependency tree (see Appendix Figures 11, 12, and 13), which relies primarily on the dependency arc entailment (DAE) algorithm for fact-checking, reveals that the text generated by ChatGPT contains a higher number of dependency arcs lacking support from the original text, when compared to human responses.

5 Conclusion

This paper presents a broad assessment of ChatGPT’s proficiency in generating controllable text. We conducted quantitative and qualitative examinations at the document level (summarization task) and sentence level (text style transfer). The empirical findings show that ChatGPT outperforms the previous state-of-the-art models on automatic metrics, but that there are substantial disparities between its generated texts and human-written texts. These disparities are reduced by providing a target example of the human writing style. Furthermore, our investigations also confirm the previously reported problems of hallucinations and inaccuracies in text generated by ChatGPT.

6 Limitations

The primary limitations of the current study pertain to the selection of prompts and evaluation metrics. The experimental cost of requesting API responses from OpenAI to assess ChatGPT’s text generation abilities imposes significant constraints on our choice of datasets. Therefore, we have to limit our experimentation to only two related controllable text generation datasets. While we have evaluated ChatGPT’s performance at both the document and sentence levels, we cannot extrapolate that ChatGPT has similar performance for other text generation datasets. Additionally, the experimental cost prohibits us from conducting traversal experiments on the selection of hyperparameters. We relied on the default configuration recommended by OpenAI, and we maintain consistency in all hyperparameters to ensure the fairness of the experiments.

Secondly, although we have studied the impact of prompt engineering on ChatGPT, the selection of prompts is mainly affected by human understanding, and the number of potential prompts is infinite. Hence, we cannot guarantee whether other prompts that we did not select will yield the same conclusions as our experiment. Furthermore, ChatGPT is subject to continuous updates and iterations, which may lead to improved performance, making it difficult to predict if future versions of ChatGPT will have similar results to our experiments.

Finally, to select appropriate evaluation metrics, we have included both domain-related evaluation metrics (such as reading difficulty and text formality) and domain-independent evaluation indicators (such as fact-checking and hallucination detection). However, we acknowledge that the automatic met-

rics may sometimes not capture all aspects of the intended construct correctly.

7 Ethics Considerations

All datasets utilized in this study are publicly available, and we have adhered to ethical considerations by not introducing any additional information into ChatGPT’s inputs.

Acknowledgements

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 Research and Innovation Programme (Grant Agreement No. 948878).



References

- Reinald Kim Amplayo, Stefanos Angelidis, and Mirella Lapata. 2021. [Aspect-controllable opinion summarization](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6578–6593, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Tal August, Katharina Reinecke, and Noah A. Smith. 2022. [Generating scientific definitions with controllable complexity](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8298–8317, Dublin, Ireland. Association for Computational Linguistics.
- Nikolay Babakov, David Dale, Varvara Logacheva, and Alexander Panchenko. 2022. [A large-scale computational study of content preservation measures for text style transfer and paraphrase generation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 300–321, Dublin, Ireland. Association for Computational Linguistics.
- David Baidoo-Anu and Leticia Owusu Ansah. 2023. Education in the era of generative artificial intelligence (ai): Understanding the potential benefits of chatgpt in promoting teaching and learning. *Available at SSRN 4337484*.
- Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Lovenia, Ziwei Ji, Tiezheng Yu, Willy Chung, Quyet V. Do, Yan Xu, and Pascale Fung. 2023a. A multitask, multilingual, multimodal evaluation of chatgpt on reasoning, hallucination, and interactivity. *ArXiv*, abs/2302.04023.

- Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Lovenia, Ziwei Ji, Tiezheng Yu, Willy Chung, et al. 2023b. A multi-task, multilingual, multimodal evaluation of chatgpt on reasoning, hallucination, and interactivity. *arXiv preprint arXiv:2302.04023*.
- Meng Cao, Yue Dong, and Jackie Cheung. 2022. **Hallucinated but factual! inspecting the factuality of hallucinations in abstractive summarization**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3340–3354, Dublin, Ireland. Association for Computational Linguistics.
- Yixin Cao, Ruihao Shui, Liangming Pan, Min-Yen Kan, Zhiyuan Liu, and Tat-Seng Chua. 2020. **Expertise style transfer: A new task towards better communication between experts and laymen**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1061–1071, Online. Association for Computational Linguistics.
- Fredrik Carlsson, Joey Öhman, Fangyu Liu, Severine Verlinden, Joakim Nivre, and Magnus Sahlgren. 2022. **Fine-grained controllable text generation using non-residual prompting**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6837–6857, Dublin, Ireland. Association for Computational Linguistics.
- Jeanne Sternlicht Chall and Edgar Dale. 1995. *Readability revisited: The new Dale-Chall readability formula*. Brookline Books.
- Hou Pong Chan, Lu Wang, and Irwin King. 2021. **Controllable summarization with constrained Markov decision process**. *Transactions of the Association for Computational Linguistics*, 9:1213–1232.
- Yu Cheng, Zhe Gan, Yizhe Zhang, Oussama Elachqar, Dianqi Li, and Jingjing Liu. 2020. **Contextual text style transfer**. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2915–2924, Online. Association for Computational Linguistics.
- Meri Coleman and Ta Lin Liau. 1975. A computer readability formula designed for machine scoring. *Journal of Applied Psychology*, 60(2):283.
- Peng Cui and Le Hu. 2021. **Topic-guided abstractive multi-document summarization**. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1463–1472, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Haixing Dai, Zhengliang Liu, Wenxiong Liao, Xiaoke Huang, Zihao Wu, Lin Zhao, Wei Liu, Ninghao Liu, Sheng Li, Dajiang Zhu, et al. 2023. Chataug: Leveraging chatgpt for text data augmentation. *arXiv preprint arXiv:2302.13007*.
- Ning Dai, Jianze Liang, Xipeng Qiu, and Xuanjing Huang. 2019. **Style transformer: Unpaired text style transfer without disentangled latent representation**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5997–6007, Florence, Italy. Association for Computational Linguistics.
- Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. 2019a. **Plug and play language models: A simple approach to controlled text generation**. *arXiv preprint arXiv:1912.02164*.
- Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. 2019b. **Plug and play language models: A simple approach to controlled text generation**. *ArXiv*, abs/1912.02164.
- Sam Earle, Maria Edwards, Ahmed Khalifa, Philip Bontrager, and Julian Togelius. 2021. **Learning controllable content generators**. In *2021 IEEE Conference on Games (CoG)*, pages 1–9. IEEE.
- Angela Fan, David Grangier, and Michael Auli. 2018. **Controllable abstractive summarization**. In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 45–54, Melbourne, Australia. Association for Computational Linguistics.
- T. Fischer, C. Biemann, Informatik und Naturwissenschaften Universität Hamburg Fakultät für Mathematik, and Universität Hamburg Fachbereich Informatik. 2021. *Finding Factual Inconsistencies in Abstractive Summaries*. Universität Hamburg.
- Sebastian Gehrmann, Zachary Ziegler, and Alexander Rush. 2019. **Generating abstractive summaries with finetuned language models**. In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 516–522, Tokyo, Japan. Association for Computational Linguistics.
- Aidan Gilson, Conrad W Safranek, Thomas Huang, Vimig Socrates, Ling Chi, Richard Andrew Taylor, David Chartash, et al. 2023. **How does chatgpt perform on the united states medical licensing examination? the implications of large language models for medical education and knowledge assessment**. *JMIR Medical Education*, 9(1):e45312.
- Tomas Goldsack, Zhihao Zhang, Chenghua Lin, and Carolina Scarton. 2022. **Making science simple: Corpora for the lay summarisation of scientific literature**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10589–10604, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Tanya Goyal and Greg Durrett. 2020. **Evaluating factuality in generation with dependency-level entailment**. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3592–3603, Online. Association for Computational Linguistics.

- Yuxuan Gu, Xiaocheng Feng, Sicheng Ma, Lingyuan Zhang, Heng Gong, and Bing Qin. 2022. [A distributional lens for multi-aspect controllable text generation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 1023–1043, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Junxian He, Wojciech Kryscinski, Bryan McCann, Nazneen Rajani, and Caiming Xiong. 2022. [CTRL-sum: Towards generic controllable text summarization](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5879–5915, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Francis Heylighen and Jean-Marc Dewaele. 1999. Formality of language: definition, measurement and behavioral determinants. *Interne Bericht, Center “Leo Apostel”, Vrije Universiteit Brüssel*, 4.
- Zhiqiang Hu, Roy Ka-Wei Lee, and Charu C. Aggarwal. 2021. [Syntax matters! syntax-controlled in text style transfer](#). In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2021)*, pages 566–575, Held Online. INCOMA Ltd.
- Zhiting Hu and Li Erran Li. 2021. A causal lens for controllable text generation. *Advances in Neural Information Processing Systems*, 34:24941–24955.
- Wenxiang Jiao, Wenxuan Wang, Jen-tse Huang, Xing Wang, and Zhaopeng Tu. 2023. Is chatgpt a good translator? a preliminary study. *arXiv preprint arXiv:2301.08745*.
- Di Jin, Zhijing Jin, Zhiting Hu, Olga Vechtomova, and Rada Mihalcea. 2022. [Deep learning for text style transfer: A survey](#). *Computational Linguistics*, 48(1):155–205.
- Nitish Shirish Keskar, Bryan McCann, Lav R. Varshney, Caiming Xiong, and Richard Socher. 2019. Ctrl: A conditional transformer language model for controllable generation. *ArXiv*, abs/1909.05858.
- J Peter Kincaid, Robert P Fishburne Jr, Richard L Rogers, and Brad S Chissom. 1975. Derivation of new readability formulas (automated readability index, fog count and flesch reading ease formula) for navy enlisted personnel. Technical report, Naval Technical Training Command Millington TN Research Branch.
- Sachin Kumar, Biswajit Paria, and Yulia Tsvetkov. 2022. [Gradient-based constrained sampling from language models](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2251–2277, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Philippe Laban, Tobias Schnabel, Paul N. Bennett, and Marti A. Hearst. 2022. [SummaC: Re-visiting NLI-based models for inconsistency detection in summarization](#). *Transactions of the Association for Computational Linguistics*, 10:163–177.
- Viet Dac Lai, Nghia Trung Ngo, Amir Pouran Ben Veyseh, Hieu Man, Franck Dernoncourt, Trung Bui, and Thien Huu Nguyen. 2023. Chatgpt beyond english: Towards a comprehensive evaluation of large language models in multilingual learning. *arXiv preprint arXiv:2304.05613*.
- Xiang Li, John Thickstun, Ishaan Gulrajani, Percy S Liang, and Tatsunori B Hashimoto. 2022. Diffusion-lm improves controllable text generation. *Advances in Neural Information Processing Systems*, 35:4328–4343.
- Xiao Li, Guanyi Chen, Chenghua Lin, and Ruizhe Li. 2020. [DGST: a dual-generator network for text style transfer](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7131–7136, Online. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Xiaochen Liu, Yang Gao, Yu Bai, Jiawei Li, Yinan Hu, Heyan Huang, and Boxing Chen. 2022a. [PSP: Pre-trained soft prompts for few-shot abstractive summarization](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 6355–6368, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Yizhu Liu, Qi Jia, and Kenny Zhu. 2022b. [Length control in abstractive summarization by pretraining information selection](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6885–6895, Dublin, Ireland. Association for Computational Linguistics.
- Zheheng Luo, Qianqian Xie, and Sophia Ananiadou. 2022. [Readability controllable biomedical document summarization](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 4667–4680, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Eric Malmi, Aliaksei Severyn, and Sascha Rothe. 2020. [Unsupervised text style transfer with padded masked language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8671–8680, Online. Association for Computational Linguistics.
- Philip M McCarthy and Scott Jarvis. 2010. Mtl-d, vocd-d, and hd-d: A validation study of sophisticated approaches to lexical diversity assessment. *Behavior research methods*, 42(2):381–392.
- Remi Mir, Bjarke Felbo, Nick Obradovich, and Iyad Rahwan. 2019. [Evaluating style transfer for text](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*,

- Volume 1 (Long and Short Papers)*, pages 495–504, Minneapolis, Minnesota. Association for Computational Linguistics.
- Nasim Nouri. 2022. [Text style transfer via optimal transport](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2532–2541, Seattle, United States. Association for Computational Linguistics.
- OpenAI. 2023. Gpt-4 technical report. *ArXiv*, abs/2303.08774.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- 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*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Keqin Peng, Liang Ding, Qihuang Zhong, Li Shen, Xuebo Liu, Min Zhang, Yuanxin Ouyang, and Dacheng Tao. 2023. Towards making the most of chatgpt for machine translation. *Available at SSRN 4390455*.
- Dongqi Pu, Xudong Hong, Pin-Jie Lin, Ernie Chang, and Vera Demberg. 2022. [Two-stage movie script summarization: An efficient method for low-resource long document summarization](#). In *Proceedings of The Workshop on Automatic Summarization for Creative Writing*, pages 57–66, Gyeongju, Republic of Korea. Association for Computational Linguistics.
- Dongqi Pu and Khalil Sima’an. 2022. [Passing parser uncertainty to the transformer: Labeled dependency distributions for neural machine translation](#). In *Proceedings of the 23rd Annual Conference of the European Association for Machine Translation*, pages 41–50, Ghent, Belgium. European Association for Machine Translation.
- Dongqi Pu, Yifan Wang, and Vera Demberg. 2023. Incorporating distributions of discourse structure for long document abstractive summarization. *arXiv preprint arXiv:2305.16784*.
- Chengwei Qin, Aston Zhang, Zhuosheng Zhang, Jiao Chen, Michihiro Yasunaga, and Diyi Yang. 2023. Is chatgpt a general-purpose natural language processing task solver? *arXiv preprint arXiv:2302.06476*.
- Abhinav Ramesh Kashyap, Devamanyu Hazarika, Min-Yen Kan, Roger Zimmermann, and Soujanya Poria. 2022. [So different yet so alike! constrained unsupervised text style transfer](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 416–431, Dublin, Ireland. Association for Computational Linguistics.
- Sudha Rao and Joel Tetreault. 2018. [Dear sir or madam, may I introduce the GYAFC dataset: Corpus, benchmarks and metrics for formality style transfer](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 129–140, New Orleans, Louisiana. Association for Computational Linguistics.
- Emily Reif, Daphne Ippolito, Ann Yuan, Andy Coenen, Chris Callison-Burch, and Jason Wei. 2022. [A recipe for arbitrary text style transfer with large language models](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 837–848, Dublin, Ireland. Association for Computational Linguistics.
- Jürgen Rudolph, Samson Tan, and Shannon Tan. 2023. Chatgpt: Bullshit spewer or the end of traditional assessments in higher education? *Journal of Applied Learning and Teaching*, 6(1).
- Ritesh Sarkhel, Moniba Keymanesh, Arnab Nandi, and Srinivasan Parthasarathy. 2020. [Interpretable multi-headed attention for abstractive summarization at controllable lengths](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6871–6882, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. [Get to the point: Summarization with pointer-generator networks](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada. Association for Computational Linguistics.
- Chenhui Shen, Liying Cheng, Lidong Bing, Yang You, and Luo Si. 2022a. [SentBS: Sentence-level beam search for controllable summarization](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10256–10265, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Chenhui Shen, Liying Cheng, Ran Zhou, Lidong Bing, Yang You, and Luo Si. 2022b. [MRd: A meta-review dataset for structure-controllable text generation](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2521–2535, Dublin, Ireland. Association for Computational Linguistics.
- Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. 2023. [Hugging-gpt: Solving ai tasks with chatgpt and its friends in huggingface](#). *arXiv preprint arXiv:2303.17580*.
- Dominik Sobania, Martin Briesch, Carol Hanna, and Justyna Petke. 2023. An analysis of the automatic

- bug fixing performance of chatgpt. *arXiv preprint arXiv:2301.08653*.
- Ewoenam Kwaku Tokpo and Toon Calders. 2022. [Text style transfer for bias mitigation using masked language modeling](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Student Research Workshop*, pages 163–171, Hybrid: Seattle, Washington + Online. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Jiaan Wang, Yunlong Liang, Fandong Meng, Zhixu Li, Jianfeng Qu, and Jie Zhou. 2023a. Cross-lingual summarization via chatgpt. *arXiv preprint arXiv:2302.14229*.
- Jiaan Wang, Yunlong Liang, Fandong Meng, Beiqi Zou, Zhixu Li, Jianfeng Qu, and Jie Zhou. 2023b. [Zero-shot cross-lingual summarization via large language models](#).
- Colin G West. 2023. Ai and the fci: Can chatgpt project an understanding of introductory physics? *arXiv preprint arXiv:2303.01067*.
- Chien-Sheng Wu, Linqing Liu, Wenhao Liu, Pontus Stenetorp, and Caiming Xiong. 2021. [Controllable abstractive dialogue summarization with sketch supervision](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 5108–5122, Online. Association for Computational Linguistics.
- Kexin Yang, Dayiheng Liu, Wenqiang Lei, Baosong Yang, Mingfeng Xue, Boxing Chen, and Jun Xie. 2022. Tailor: A prompt-based approach to attribute-based controlled text generation. *arXiv preprint arXiv:2204.13362*.
- Xianjun Yang, Yan Li, Xinlu Zhang, Haifeng Chen, and Wei Cheng. 2023a. Exploring the limits of chatgpt for query or aspect-based text summarization. *arXiv preprint arXiv:2302.08081*.
- Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Ehsan Azarnasab, Faisal Ahmed, Zicheng Liu, Ce Liu, Michael Zeng, and Lijuan Wang. 2023b. Mm-react: Prompting chatgpt for multimodal reasoning and action. *arXiv preprint arXiv:2303.11381*.
- Hanqing Zhang and Dawei Song. 2022. [DisCup: Discriminator cooperative unlikelihood prompt-tuning for controllable text generation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 3392–3406, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Hanqing Zhang, Haolin Song, Shaoyu Li, Ming Zhou, and Dawei Song. 2022. A survey of controllable text generation using transformer-based pre-trained language models. *arXiv preprint arXiv:2201.05337*.
- Ce Zhou, Qian Li, Chen Li, Jun Yu, Yixin Liu, Guangjing Wang, Kai Zhang, Cheng Ji, Qiben Yan, Lifang He, et al. 2023. A comprehensive survey on pretrained foundation models: A history from bert to chatgpt. *arXiv preprint arXiv:2302.09419*.
- Wanrong Zhu, Xin Wang, Tsu-Jui Fu, An Yan, Pradyumna Narayana, Kazoo Sone, Sugato Basu, and William Yang Wang. 2021. [Multimodal text style transfer for outdoor vision-and-language navigation](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1207–1221, Online. Association for Computational Linguistics.

A Appendix: One-shot Guidance

B Appendix: Absolute Differences in POS and Dependency Label Distributions

Candidate	FRE	CLI	DCR
Document: {Original Document}, Layman Summary: {Human Layman Summary}.			
Please learn the way of summarization from the previous example, and give me a layman-style summary of X	49.23 [†]	13.26 [†]	10.45 [†]
Human Answer	52.42	12.46	8.93
Document: {Original Document}, Expert Summary: {Human Expert Summary}.			
Please learn the way of summarization from the previous example, and give me an expert-style summary of X	28.88 [†]	15.92 [†]	11.82
Human Answer	23.20	17.62	11.78

Table 7: Reading difficulty of one-shot guidance. [†] indicates statistical significance ($p < 0.05$) against corresponding human answers via paired t-test.

Candidate	Formality	MTLD
Formal: {Formal Sentence}, Informal: {Informal Sentence}.		
Please learn the way of formality conversion from the previous example, and give me an informal version of X	50.67 [†]	14.81
Human Answer	49.87	15.20
Informal: {Informal Sentence}, Formal: {Formal Sentence}.		
Please learn the way of formality conversion from the previous example, and give me a formal version of X	52.13 [†]	19.22
Human Answer	53.57	18.70

Table 8: Text formality of one-shot guidance on GYAFC-FR branch. [†] indicates statistical significance ($p < 0.05$) against corresponding human answers via paired t-test.




Figure 5: Absolute differences in dependency labels distribution of ChatGPT and human-generated formal style sentences: GYAFC - EM




Figure 6: Absolute differences in dependency labels distribution of ChatGPT and human-generated informal style sentences: GYAFC - EM




Figure 7: Absolute differences in dependency labels distribution of ChatGPT and human-generated formal sentences: GYAFC - FR

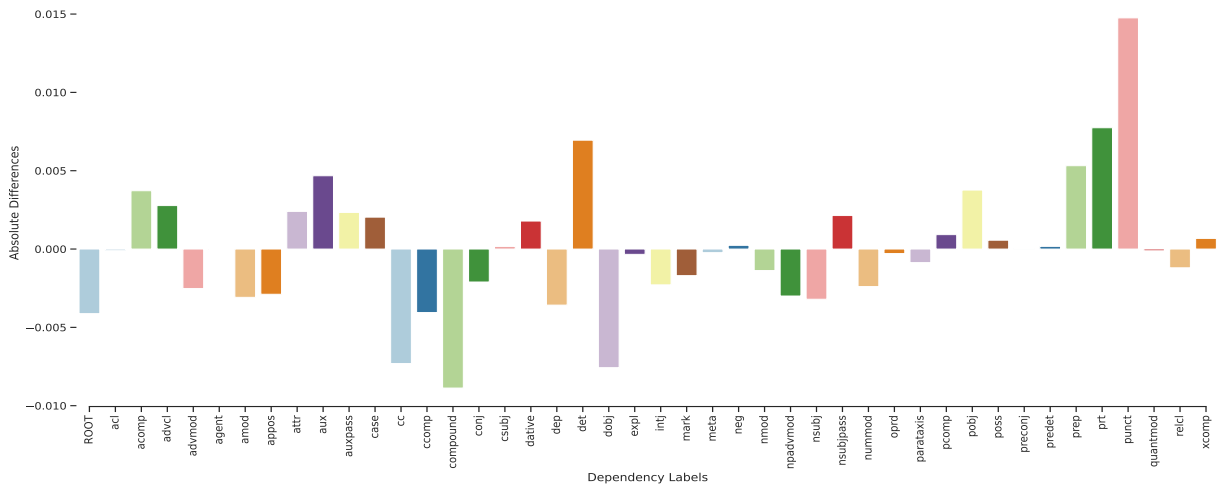


Figure 8: Absolute differences in dependency labels distribution of ChatGPT and human-generated informal sentences: GYAFC - FR




Figure 9: Absolute differences in POS tags distribution of ChatGPT and human-generated formal sentences: GYAFC - FR




Figure 10: Absolute differences in POS tags distribution of ChatGPT and human-generated informal sentences: GYAFC - FR

C Appendix: Dependency Arc Entailment




Figure 11: Case study of dependency tree visualization (Reference)




Figure 12: Case study of dependency tree visualization (Human)




Figure 13: Case study of dependency tree visualization (ChatGPT)




Figure 14: Dependency arc entailment: GYAFC - FR. Data points $>0.95 \approx$ Accurate. To clarify discrepancies, cutoff point = 0.95.

D Appendix: Named Entity Hallucination

Dataset	Candidate	Precision	Recall	F1
GYAFC-FR	Human Informal	0.989	0.988	0.988
	Human Formal	0.988	0.989	0.988
	ChatGPT Informal	0.986	0.985	0.986
	ChatGPT Formal	0.974	0.974	0.974
	ChatGPT I2F2I	0.982	0.982	0.982
	ChatGPT F2I2F	0.974	0.973	0.973
GYAFC-EM	Human Informal	0.979	0.987	0.983
	Human Formal	0.977	0.989	0.982
	ChatGPT Informal	0.975	0.974	0.974
	ChatGPT Formal	0.950	0.952	0.951
	ChatGPT I2F2I	0.970	0.969	0.970
	ChatGPT F2I2F	0.945	0.946	0.945

Table 9: Named entity hallucination - GYAFC