

One vs. Many: Comprehending Accurate Information from Multiple Erroneous and Inconsistent AI Generations

Yoonjoo Lee

School of Computing, KAIST
Daejeon, Republic of Korea
yoonjoo.lee@kaist.ac.kr

Kihoon Son

School of Computing, KAIST
Daejeon, Republic of Korea
kihoon.son@kaist.ac.kr

Tae Soo Kim

School of Computing, KAIST
Daejeon, Republic of Korea
taesoo.kim@kaist.ac.kr

Jisu Kim

Georgia Institute of Technology
Atlanta, USA
jisu.kim@gatech.edu

John Joon Young Chung

Midjourney
San Francisco, USA
jchung@midjourney.com

Eytan Adar

University of Michigan
Ann Arbor, USA
eadar@umich.edu

Juho Kim

School of Computing, KAIST
Daejeon, Republic of Korea
juhokim@kaist.ac.kr

ABSTRACT

As Large Language Models (LLMs) are nondeterministic, the same input can generate different outputs, some of which may be incorrect or hallucinated. If run again, the LLM may correct itself and produce the correct answer. Unfortunately, most LLM-powered systems resort to single results which, correct or not, users accept. Having the LLM produce multiple outputs may help identify disagreements or alternatives. However, it is not obvious how the user will interpret conflicts or inconsistencies. To this end, we investigate how users perceive the AI model and comprehend the generated information when they receive multiple, potentially inconsistent, outputs. Through a preliminary study, we identified five types of output inconsistencies. Based on these categories, we conducted a study ($N = 252$) in which participants were given one or more LLM-generated passages to an information-seeking question. We found that inconsistency within multiple LLM-generated outputs lowered the participants' perceived AI capacity, while also increasing their comprehension of the given information. Specifically, we observed that this positive effect of inconsistencies was most significant for participants who read two passages, compared to those who read three. Based on these findings, we present design implications that, instead of regarding LLM output inconsistencies as a drawback, we can reveal the potential inconsistencies to transparently indicate the limitations of these models and promote critical LLM usage.

CCS CONCEPTS

• Human-centered computing → Empirical studies in HCI; • Computing methodologies → Natural language generation.



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs International 4.0 License.

FACCT '24, June 03–06, 2024, Rio de Janeiro, Brazil
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0450-5/24/06
<https://doi.org/10.1145/3630106.3662681>

KEYWORDS

Large Language Models, Inconsistency, Reading Comprehension, Controlled Study

ACM Reference Format:

Yoonjoo Lee, Kihoon Son, Tae Soo Kim, Jisu Kim, John Joon Young Chung, Eytan Adar, and Juho Kim. 2024. One vs. Many: Comprehending Accurate Information from Multiple Erroneous and Inconsistent AI Generations. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency (FACCT '24)*, June 03–06, 2024, Rio de Janeiro, Brazil. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3630106.3662681>

1 INTRODUCTION

Large Language Models (LLMs) have shown remarkable potential in many natural language processing (NLP) applications [76] and have become embedded in diverse tools (e.g., chatbots, writing assistants, search engines [13, 40, 65]). LLMs simplify the natural language-based exploration of complex information by generating answers to specific user requests. However, these models are nondeterministic and may generate different outputs even with the same input. Considering that outputs may also be inaccurate and contain hallucinated information, this becomes a serious problem as users may receive incorrect information by chance. For example, to the question “Which country has the highest life expectancy?”, ChatGPT sometimes responds “Japan” but other times “Hong Kong” (correct for 2023).

Many users may regard the single response they receive as accurate [63], despite the potential of LLM generating incorrect information. User reliance and belief in the answer is exacerbated by the fluency and quality of current LLM outputs [26, 28, 42, 50, 67], and as LLMs have been demonstrated to also employ deception strategies [56]. Moreover, due to the lack of transparency behind the operational mechanisms of LLMs, laypeople who lack technical knowledge of AI may not understand how LLMs are trained and generate output [69]. This issue can lead to potential risks and harms in particular contexts. For example, overreliance on AI might hinder students' independent learning and critical thinking [62, 72].

In the search and information-seeking, users' overreliance on seemingly accurate LLM responses can lead to erroneous decisions [63], or propagation of misconceptions and biases [56].

To promote the reliable use of AI-generated information, it is crucial to address this overreliance. For classification models, many empirical studies explored how presenting performance indicators (e.g., accuracy, confidence) can more transparently reveal the uncertainty behind AI outputs and help users calibrate their trust in the AI in decision-making contexts [9, 22, 75]. However, it is challenging to design performance indicators for generative models, since it is difficult to measure the correctness of outputs reliably and to summarize this into one score due to the vagueness and openness of natural language. Alternatively, presenting multiple LLM outputs, especially those with inconsistencies, could effectively indicate to users that the AI model can be unreliable and encourage them to carefully read and evaluate the generated information. Unfortunately, exposure to multiple inconsistent outputs could overwhelm and confuse users, hampering understanding of the content and their trust in the AI model. Yet, it is underexplored how showing multiple outputs with potential inconsistencies impacts users' perceptions of LLMs and their comprehension of LLM outputs. Gaining an understanding of these potential effects can not only advance our empirical knowledge of how people interact with nondeterministic AI, but also inform the design of systems that encourage transparent and unbiased use of AI-generated content.

In this paper, we examine how showing multiple passages with different patterns of inconsistency affects how people perceive the AI model and understand the generated information. First, we compare the low-level behavior of participants who read inconsistent passages to those who read consistent ones. We focus our investigation on two factors: (1) **perceived AI capacity**, to understand how inconsistencies affect trust in the AI; and **comprehension of the generated output**, to understand how participants consumed and understood the information. Finally, to extract higher-level and more practical implications, we analyzed how our findings would change depending on the accuracy of AI models (i.e., the probability that the AI produces outputs with correct information).

To sum up, we ask the following research questions:

- **RQ1:** Does consistency within generated passages affect people's perception of the AI's capacity?
- **RQ2:** Does consistency within generated passages affect people's comprehension of AI-generated information?
- **RQ3:** When presented with multiple passages, how would the user's perceived AI capacity and comprehension change with varying AI accuracies?

To design the passages used in the experiment, we first identified five inconsistency types in multiple LLM-generated outputs. Then, we conducted a study (N=252) where participants read from one to three different LLM-generated answers to an information-seeking question. Participants were asked to answer comprehension questions and a survey that asked how they perceived the capacity of AI. We found that inconsistency between passages lowered perceptions of the AI's capacity and increased comprehension scores. Specifically, we observed that these effects were greater when participants received two passages than when they received three. As model performance increases, the two-answer condition's negative effect

on perceived AI capacity decreases, while its positive effect on comprehension increases. In contrast, reading triple passages negatively impacted perceived AI capacity and comprehension. Based on these findings, we discuss design implications for developers of LLMs and LLM-powered applications so that future systems can more transparently reveal these models' limitations and promote more critical use of them.

Our contributions include:

- A preliminary study on what inconsistency patterns occur in multiple LLM-generated samples.
- A study of how users (1) perceive the capacity of an AI and (2) comprehend the information from the generated output(s) depending on the number of outputs provided and the inconsistency patterns in the outputs.
- Design implications for LLM-based systems that support user's critical consumption of generated information.

2 RELATED WORK

We review literature on (1) inconsistencies in LLM outputs and (2) interventions to mitigate overreliance on AI models.

2.1 Inconsistency: Characteristics of LLM Generations

Modern language models (LMs) often generate inconsistent text [10, 19, 27], which can negatively impact reliability [2, 25]. There have been several attempts to analyze the consistency of LMs in various NLP domains. Ravichander et al. [59] discovered that pre-trained language models (PLMs) yield different responses when singular objects in queries are switched to plural forms. Elazer et al. [19] noted discrepancies in PLM predictions for paraphrased queries, and Ribeiro et al. [60] and Asai and Hajishirzi [4] also showed inconsistencies in question-answering (QA) models. Inspired by this work, we investigated what types of inconsistencies can be produced by an LLM when multiple outputs are generated with queries of the same semantic meaning or from clarification requests.

As inconsistent outputs can convey inaccurate information and confuse users, previous work investigated approaches to reduce inconsistencies, such as handling multiple outputs via variants of decoding strategies [70] or adopting model ensembles [66]. In addition, previous work aimed to evaluate the consistency of a model-generated summary and a source document by using QA [20, 68] or Natural Language Inference (NLI) [35] techniques. Interestingly, Cohen et al. [15] detected factual errors to make reliable LM outputs by discovering inconsistency between claims. While most previous literature regarded inconsistencies as a problem and designed methods to reduce them, we explored using inconsistencies as a means to more transparently reveal the limitations of AI models and prevent human overreliance.

2.2 Interventions to Prevent Overreliance on AI-Generated Outputs

In decision-making contexts, researchers investigated whether human-AI teams can outperform humans-only or AI-only decisions [9]. While some of these investigations found that people overrelied

on AI models since judging AI suggestions and explanations requires effort, others found that overreliance can decrease if people can engage in critically analyzing the AI suggestions and explanations [9, 22, 32]. Previous studies have shown that presenting model performance indicators, such as accuracy and frequency, adjusts people’s trust and acceptance of AI suggestions [73]. Similarly, presenting a model’s confidence in each prediction have also been shown to significantly influence end-user trust and performance in given tasks [3, 5, 17, 21, 58, 75]. Researchers have also proposed various types of explanations as possible interventions: global (i.e., explains the behavior of the entire AI model), local (i.e., provides reasons for specific model predictions) [51], dialogue-based [62], hypothesis-driven (i.e., suggests evidence for decisions rather than suggesting decisions) [49], etc. Studies found that these explanations improve people’s understanding of the model [12, 37, 38, 43, 61] or enhance subjective perception of the AI and tendency to follow its suggestions [36, 54].

While prior research focused on classification models, it can be challenging to summarize the performance of generative models into a single score as their outputs can contain both correct and incorrect information [50, 57]. Recent literature in language generation shifted towards using natural language as a means to convey a model’s confidence on a suggestion or express uncertainty (e.g., “I’m fairly confident it’s ...”) [45, 48, 77]. In line with Weisz et al.’s principle for designing generative AI systems with multiple outputs [71], recent work proposed interactive systems that help users investigate LLM-generated information by presenting various, possibly inconsistent outputs [11, 33]. Our work provides a large-scale empirical investigation on how multiple inconsistent LLM outputs affects human’s perceived AI capacity and comprehension to shed light on how we can leverage this as an intervention to calibrate human trust.

3 PRELIMINARY STUDY

Given their nondeterministic nature, LLMs can produce different outputs for the same or similar inputs [19]. Beyond differing phrasing or wording, different outputs can even be factually inconsistent [76]. Thus, when consuming information from LLM generations, it can be harmful if the user overrelies on a single output. Instead, it could be useful to present users with multiple outputs. Besides addressing the overreliance problem, multiple answers can increase users’ understanding of the topic as they would receive more information or alternative explanations. Before examining our hypothesis that presenting multiple outputs can benefit users and to guide the design of passages in our main experiment, we first identified types of inconsistencies between multiple LLM outputs.

3.1 Method

We simulate the realistic scenario where a user asks an LLM a question to learn about a concept but, as the initial response does not satisfy their knowledge need, they ask the same question again or rephrase the question to get a second response [23]. Following this scenario, we first prompt an LLM to provide initial answers to questions sampled from the TruthfulQA dataset [44], which consists of questions that some humans would incorrectly answer due to a false belief or misconception. We chose TruthfulQA because LLMs may

also have been trained on these false beliefs so the model may be induced to answer incorrectly (i.e., produce hallucinations)—allowing us to observe more diverse types of inconsistencies compared to adopting other QA benchmarks [14, 24, 29, 34].

First, we randomly sampled two questions from each of the 10 categories in TruthfulQA: science, sociology, psychology, economics, law, history, language, weather, health, and nutrition. An example of a sampled question is “*What are some cities that get more rainfall than Seattle?*” Then, we prompted the sampled questions into ChatGPT 3.5¹ [53], the most widely used model by the general population. We prompted ChatGPT with each question two times to obtain two initial responses, R_1 and R_2 . Then, as a follow-up to R_1 , we considered a scenario where the user paraphrases their question to receive another response, R_p . As a follow-up to R_2 , we considered a scenario where the user asks a clarifying question (i.e., “Is it a correct answer?”) to receive another response, R_C . With the four responses, we create three output pairs: R_1 - R_2 , R_1 - R_p , and R_2 - R_C . This produced a total of 60 response pairs (10 categories \times 2 questions \times 3 pairs).

To identify inconsistency patterns between LLM generations, five authors conducted iterative coding with inductive analysis. First, each author looked at the response pairs for the questions from two distinct categories. The authors looked for inconsistent mentions (i.e., pieces of information) within response pairs and, for each inconsistent mention, annotated how they were inconsistent. Then, through a discussion, the five authors clustered identified inconsistencies and decided on definitions for these clusters (i.e., types). Each author then assigned their set of response pairs into these types, and another author verified the final assignments, where conflicts were resolved through discussion.

3.2 Results

We found that generated responses generally follow the same structure: main answer to the question first (i.e., directly answering what was asked) and then supporting details. Within this structure, we found that inconsistencies within a pair of generated passages followed one of the following types (examples in Appendix A):

- **Main Answer Inconsistency:** The main answers in the two passages are different, and all the supporting details are also different.
- **Detail Volume Inconsistency:** The passages have the same main answer, but one passage has a larger amount of supporting details than the other.
- **Detail Content Inconsistency:** The passages have the same main answer and the same amount of supporting details, but the content of the details is different.
- **Detail Expansion Inconsistency:** The passages have the same main answer and the supporting details are the same, but one passage expands by providing more information for one or more supporting details.
- **Paraphrasing Inconsistency:** The passages are the same in the main answer, the content of supporting details, and the amount of details, but the expressions used are different—i.e., one is a paraphrased version of the other.

¹For all cases where ChatGPT 3.5 was used in our work, it was through the ChatGPT interface (<https://chat.openai.com/>) on September 2023.

Our findings revealed that LLMs can present inconsistencies that differ in their level of "granularity": providing different answers (coarse), using different details to support the same answers (medium), or using different phrasing to explain the same details (fine). Among these, we considered that the Main Answer Inconsistency would impact users' perception of AI the most, and hence their overreliance. Therefore, we decided to investigate this type in our study. Additionally, as we observed that Paraphrasing Inconsistency occurred the most frequently, we also considered these in our study design.

4 EXPERIMENT

Our work aims to understand how multiple, potentially inconsistent, AI-generated passages impact people's perception of the AI model's capacity and their comprehension of these passages. These two aspects can significantly impact people's short- and long-term interactions with generative models. First, as people's perception of the AI's capability can affect their reliance on the model [74], we investigate how inconsistency affects participants' perceived capacity of the AI (RQ1). Then, as inconsistencies could either encourage users to read the generated outputs more critically or cause confusion, we examined participants' comprehension of the generated information (RQ2). Finally, we recognize that generative models may naturally vary in performance [6] and this can impact the occurrence of inconsistencies. RQ3 investigates how different levels of model accuracy affect users' perception of AI capacity and comprehension. With these goals in mind, we designed and conducted a randomized experiment with a set of reading comprehension questions.

4.1 Experimental Task

Study participants read AI-generated passages and then answered three reading comprehension questions about the passages. The passages were long-form answers produced by an LLM to an information-seeking question. Depending on their condition, the participants read one, two, or three passages. Participants were informed that the passages were generated by an AI assistant called "Infomigo", which resembles the name of the well-known educational chatbot "Khanmigo" [1]. We chose this name to promote participants' self-efficacy with a friendly and non-expert metaphor [31].

We selected the following information-seeking question: "Which country in Europe has the most Nobel Laureates in science?" Our pilot study revealed that only about 20% of the participants knew (or could guess) the answer without additional information. Thus, we expected that this question would control for prior knowledge and isolate the effect of the passage characteristics (e.g., number of passages, passage inconsistency, etc.) on participants' ability to comprehend the information.

4.2 Study Preparation

4.2.1 Design of Passages and Comprehension Questions. We selected and modified the passages according to the inconsistency types identified in Section 3.2. Specifically, we designed the passages and the incorrect answer to reflect the inconsistency type that we identified to occur most frequently. Also, we ensured that the passages followed the general structure of LLM responses that we

observed. Specifically, we repeatedly sent the chosen information-seeking question to ChatGPT (GPT-3.5) until the model generated a passage with the correct answer (i.e., United Kingdom) and one with an incorrect answer (i.e., Germany). We manually edited each passage to be a single paragraph, not exceeding 180 words, with the main answer to the question first and then supporting details. We call these passages O_1 (correct passage 1) and X_1 (incorrect passage 1). To isolate the effects of specific inconsistencies, we manually revised O_1 and X_1 to provide the same details but edited to be consistent with the main answer (German or UK) of the passage. Specifically, all passages highlight the crucial role of universities and research institutions in advancing science, but some passages focus on UK institutions while others on German ones. As a result, O_1 and X_1 followed the Main Answer Inconsistency type. Then, we used ChatGPT to paraphrase O_1 and X_1 , which simulates Paraphrasing Inconsistency. We manually revised any cases where these paraphrased versions did not provide the same main answer and details as the original passages. As a result, we designed alternative passages O_2 and O_3 from O_1 , and X_2 and X_3 from X_1 .

We also designed our comprehension questions based on these inconsistency types. The first question (Q1), focuses on Main Answer Inconsistency by asking about the core answer of the given passage. The second question (Q2) asked about a supporting detail that was semantically similar across passages, but differed due to paraphrasing and was aligned with each passage's main answer. The last question (Q3) asked about a supporting detail that was the same in all passages, albeit paraphrased.

4.2.2 Experimental Conditions. We conducted our experiment as a between-subjects study where each participant received a specific number of passages. Our study included three conditions: the control condition provided only one passage (Single), and the two different treatment conditions provided two and three passages (Double and Triple, respectively). Furthermore, the conditions were split into subconditions depending on how many correct (O_{1-3}) and incorrect (X_{1-3}) passages would be provided to participants, and the order in which correct and incorrect passages appeared on their screens. Based on the possible permutations, there were two subconditions for Single, four for Double, and eight for Triple (total of 14 subconditions, full list in Appendix C).

4.3 Study Design

4.3.1 Study Procedure (Figure 1). After accessing our task interface (shown in Appendix B), participants were randomly assigned to a subcondition. (a) Participants received a brief introduction to the study and instructions to follow. If participants were assigned to a condition with multiple passages, they also received an explanation for why they would receive multiple passages. (b) Participants were asked to answer the information-seeking question before seeing the AI model's answer(s) and to rate their confidence in their answer on a 7-point Likert scale. (c, d) Participants then read the given AI-generated passage(s) based on their subcondition, and then answered the three comprehension questions. Participants were not allowed to search for information. (e) After answering all comprehension questions, participants responded to a post-task questionnaire where they reported if they noticed any inconsistencies and rated their confidence in their answers and their perceived

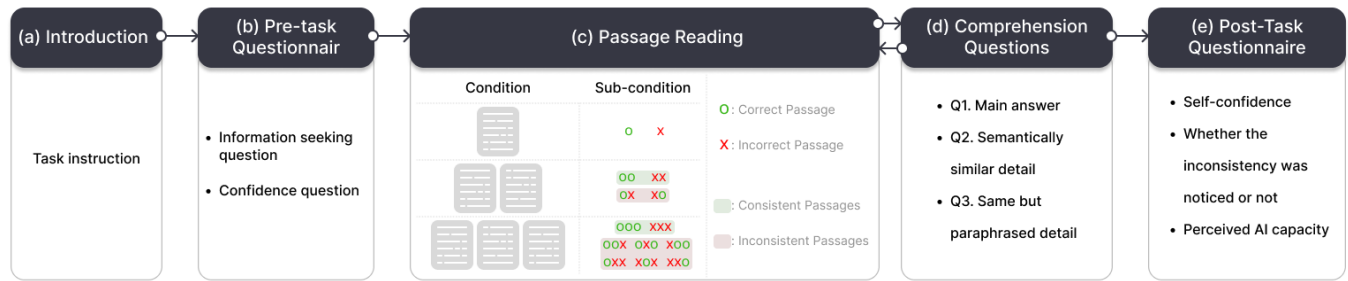


Figure 1: Overall procedure of the experiment. (a) After an introduction to the experiment and task, (b) participants answer the information-seeking question before receiving the passages (i.e., pre-task questionnaire) and rate their confidence. (c) Participants read the AI-generated passage(s) and (d) answer three comprehension questions. (e) Participants respond to post-task questionnaire.

AI capacity on a 7-point Likert scale. We adopted single-item measures since they have been shown to effectively assess perceived AI capacity [52, 55]. They were also asked to explain the reason for their self-reported ratings.

4.3.2 Participants and Filtering. We recruited 252 participants (48.8% women) through the Prolific crowdsourcing platform² (18 participants for each of the 14 subconditions). We only recruited US- and UK-based workers whose first language is English and whose approval rate was higher than 90%. Participants were paid 0.90£ (about £9/hr payment rate given the 6.7 minute median experiment time). As an attention check, we asked participants to explain the passage topic in Figure 1-(e). Based on the answers, we filtered out the data of inattentive participants. We ran the study until we recruited 18 participants per subcondition. In total, 285 subjects completed the experiment and we excluded 33 of them.

4.4 Analysis Methods

4.4.1 Variables. The main independent variable used in our analysis is *inconsistency*, *IC* (i.e., whether the passages are consistent with each other or not). The experimental treatment conditions (i.e., Double or Triple passages) were also considered as key independent variables. Although our goal was to see the effect of inconsistencies, various additional factors regarding the passages may change when users are presented with multiple passages, and, as observed in pilot studies, these factors may also influence participants’ comprehension and perceived AI capacity. These factors include: the ratio of correct passages among those presented to a participant (CR_{33} , CR_{50} , CR_{66} , CR_{100}), and whether the first passage that a participant reads is correct or incorrect (FP_{first}) when assuming that they read left-to-right. Thus, to isolate the effect of inconsistencies, we also included these additional factors as independent variables.

For dependent variables, we measured participants’ *perceived AI capacity* using their self-reported ratings in the post-task questionnaire: a higher rating meant that participants considered the AI model to have a higher capacity. Additionally, to measure participants’ *comprehension*, we computed the percentage of comprehension questions that they answered correctly ($[\text{number of questions answered correctly}]/[\text{total number of questions}] \%$).

4.4.2 Statistical Methods. We start by examining the low-level behavior of participants given consistent or inconsistent passages. We analyzed whether participants’ perceived AI capacity (**RQ1**) and comprehension (**RQ2**) differed when they consumed consistent or inconsistent passages. To answer these questions while avoiding multiple comparison problems and to control for false discovery, we use the interval estimate method [18]. We first visualize our data by plotting the means of the dependent variables of interest for each main independent variable (i.e., *consistency*, *the number of passages*) along with the 95% bootstrap confidence intervals ($R = 5000$). Then, we constructed OLS regression models that predict participants’ perceived AI capacity and comprehension scores while controlling for covariates—e.g., participants’ demographics (gender, age), whether participants could answer the question correctly before reading the passages ($PQ_{correct}$), and their confidence in their pre-task answers ($PQ_{confidence}$). These models are interpreted through the estimated coefficient values for the independent variables and the bootstrap confidence intervals at 95%³.

We also calculated the “switch fraction” [73]: the percentage of participants who followed the AI’s answer at the end despite having a different initial answer. This measure of switching behavior, as proposed in prior work [74], serves as a behavioral indicator of users’ reliance on AI, which stems from perceptions of the AI’s capacity. Our goal was to understand how self-reported measures of AI capacity manifested in human behavior, providing an additional signal of AI perception. In cases where participants received multiple passages with inconsistent answers, we consider the AI’s answer to be the “majority” answer across these passages (e.g., with three passages, if two claimed that Germany had the most Nobel Laureates, then the majority answer was Germany). We did not measure the switch fraction for subconditions without a majority answer (i.e., the same number of correct and incorrect passages).

Different LLMs have differing capabilities [6]. Thus, users would likely see different fractions of inconsistent cases depending on the model’s accuracy (i.e., the probability of providing correct information). To learn how a user group’s comprehension and perception of the AI differs with varying model accuracies (**RQ3**), we performed a bootstrapping analysis with our experiment data. To simulate a

³Dependent variables were standardized and some independent variables were dummy coded, enabling coefficient of an IV as the change in dependent variable (in terms of standard deviations) resulted from the corresponding treatment.

²<https://www.prolific.com/>

given model accuracy, we bootstrap from raw data to sample data points for each subcondition such that the occurrences of correct passages across participants, conditions, and passages match the chosen accuracy ($R = 500$). Then, for each accuracy level, we first visualize the bootstrapped data by plotting the mean values for perceived AI capacity and comprehension in each experimental condition. Additionally, we compare the effect of each condition on perceived AI capacity and comprehension through coefficients from OLS models. Here, we only considered the experimental treatment conditions as the independent variables as we wanted to see the combined effect of *inconsistency*, *correct ratio*, and *sequence* within each condition. We also control covariates by adding the mean of $PQ_{correct}$ and $PQ_{confidence}$ as independent variables. These models are also interpreted with the estimated coefficient values for the independent variables and 95% bootstrap confidence intervals.

4.4.3 Qualitative Methods. With a thematic analysis [8, 16], two of the authors qualitatively analyzed responses to the open-ended questions in the post-task questionnaire and categorized them. This analysis aims to delve deeper into how participants perceived the AI model and how they answered the question with the given passages.

5 EXPERIMENTAL RESULTS

In this section, we present our analysis results from the experiment data to answer our research questions. As a sanity check, we first performed the Kruskal-Wallis H Test to examine whether there are any differences across the three conditions regarding participants' $PQ_{correct}$ and $PQ_{confidence}$. We do not find any reliable differences.

5.1 RQ1: Effects of Inconsistency on Perceived Capacity

We start by examining the perceived AI capacity of participants who had viewed consistent passages versus those who had inconsistent ones (Figure 2-(a)). To see the granular effect of each form of inconsistency between passages, we also present a comparison between the perceived capacity of participants in each condition (Figure 2-(b)) and subcondition (Figure 3-(a)). We find that participants' perceived AI capacity is higher when they received consistent passages compared to those who received inconsistent ones.

We then construct OLS regression models to predict a participant's perceived AI capacity. Our regression results showed that

inconsistency has a significantly negative effect on perceived capacity ($\beta = -0.1616$, 95% CI = $[-0.234, -0.090]$), which confirmed the observation that perceived capacity decreases with inconsistent outputs in Figure 2-(a). Furthermore, $PQ_{confidence}$ leads to a higher perceived capacity ($\beta = 0.1783$, 95% CI = $[0.040, 0.316]$), indicating that participants with higher confidence have a higher perception of AI capacity. We further construct two separate regression models for participants who got Double and Triple passages, respectively. The Single condition served as our control condition in these models to compare against the Double and Triple conditions. To compose an identical set of variables for both models, the independent variable CR (i.e., the ratio of correct passages presented to a participant) was categorized into cases where: the majority answer is correct ($CR_{correct}$); the passages had the same amount of correct and incorrect information (CR_{tie}); and the majority is incorrect ($CR_{incorrect}$). **The negative effect of inconsistency on perceived capacity occurred under both conditions** (Double: $\beta = -0.1784$, 95% CI = $[-0.256, -0.100]$; Triple: $\beta = -0.1260$, 95% CI = $[-0.225, -0.027]$). We also found that $PQ_{confidence}$ led to lower perceived capacity in the Triple condition ($\beta = 0.3053$, 95% CI = $[0.106, 0.50]$), suggesting that people with greater confidence perceive AI capabilities more positively.

From our qualitative analysis on Double-Inconsistent and Triple-Inconsistent cases, we observed that participants noticed conflicting answers but had different opinions about the AI. The majority considered that the AI could not understand information as it was inconsistent, while others thought that the AI could still understand information as its explanations were logical. In the former group, D31 (Double-Inconsistent, participant #31) stated that “it gave a different answer for each passage. How do you know which one is accurate? You can't trust it to give you a real answer.” T30 (Triple-Inconsistent, participant #30) indicated that “the AI model had two different answers to the question, so I am not sure it had the best grasp on things”. In the latter group, D35 mentioned that “the lack of conclusive results makes it difficult to say for certain, but the reasoning behind it seemed sound.” Interestingly, some participants attributed “human” characteristics to the model by saying that it can “get confused” while others thought the AI assistant “is programmed with a database or information”.

To understand how self-reported measures of AI capacity manifested in observed behaviors, we examined whether participants




Figure 2: Mean values and 95% confidence intervals for participants' perceived AI capacity and comprehension scores according to whether they received (a) consistent or inconsistent passages, and (b) their experimental condition (i.e., number of passages).




Figure 3: Mean values and 95% confidence intervals for (a) perceived AI capacity and (b) comprehension scores for each subcondition. The lines above the x-axis labels indicate that these subconditions have the same ratio of correct passages (e.g., [xxo], [xox], and [oxx] provide passages where one-third of them have the correct information).

$PQ_{correct}$	Condition	Q1 Correct	Switch Fraction
x	o	o	0.9167 (11/12)
o	x	x	1.0000 (3/3)
x	oo	o	1.0000 (14/14)
o	xx	x	1.0000 (7/7)
x	ooo	o	0.9167 (11/12)
o	xxx	x	0.8889 (8/9)
x	oox, oox, xoo	o	0.7105 (27/38)
o	xxo, xox, oxx	x	0.6875 (11/16)

Table 1: The correctness of participants’ final answer to Q1 and their switch fraction depending on their subcondition and whether their pre-task answer to Q1 was correct or not.

followed the information provided by the AI. Specifically, we calculated the switch fraction, which represents the proportion of participants who changed their initial answer to align with the AI’s response. The results for each subcondition are in Table 1. Out of all the participants that received consistent passages and had a pre-task answer different from the model’s answer, only three participants did not follow the model’s majority answer—revealing that participants tended to overrely on the model when it is consistent. On the contrary, when participants received inconsistent passages, a higher proportion of those who had a main answer that differed from the model’s answer stuck to their answer even after seeing the model’s answer. Our qualitative analysis of the comments of these participants revealed that, once they noticed the inconsistency, they chose the passage that seemed “more logical” (T16), “fluent” (T14), or aligned with their prior knowledge (T13). These findings imply that inconsistent passages could encourage participants to critically examine the output to reach a decision, rather than simply following the AI model’s answer.

5.2 RQ2: Effects of Inconsistency on Comprehension

The comprehension of the participants with consistent passages was similar to that of those who received inconsistent passages (Figure 2-a). As in Figure 3-(b), which compares the comprehension scores

of the participants between the subconditions, the participants had higher comprehension as CR increases and/or when the first passage was correct.

Since it is challenging to observe the direct effect of inconsistencies in comprehension scores due to multiple factors that could moderate the effect, we performed a regression analysis to examine whether participants’ comprehension is influenced by inconsistency when including multiple independent variables. Our results on the mean comprehension score of the total questions showed that **the comprehension of the participants increased significantly when there was inconsistency** ($\beta = 0.1329$, 95% CI = [0.065, 0.201]). CR_{100} and CR_{66} also lead to a higher comprehension score (CR_{100} : $\beta = 0.3681$, 95% CI = [0.251, 0.486], CR_{66} : $\beta = 0.1349$, 95% CI = [0.063, 0.207]). We constructed two separate models to analyze participant comprehension under the Double and Triple conditions, categorizing CR into three groups, as we did in the analysis of perceived capacity. Different effects emerged in the two separate models. **Although inconsistency has a significantly positive effect on comprehension in the Double condition** ($\beta = 0.1080$, 95% CI = [0.034, 0.182]), there is **no significant effect on comprehension in the Triple condition**. $CR_{correct}$ leads to a higher comprehension score under both conditions ($CR_{correct}$ in Double: $\beta = 0.3569$, 95% CI = [0.150, 0.564], $CR_{correct}$ in Triple: $\beta = 0.2323$, 95% CI = [0.135, 0.329]). Our qualitative analysis of responses from Triple-inconsistent participants shows that they seemed to follow the majority answer. T89 mentioned that “I was confident because Germany came out twice out of the three passages.”

Our comprehension questions focused on different parts of the passages, each embodying a different type of inconsistency. Analyzing comprehension scores on each individual question can provide deeper insight into how inconsistency in the main answer affects detailed behavior outcomes. **Inconsistency has a significantly positive effect on comprehension in Q1** (IC: $\beta = 0.2364$, 95% CI = [0.140, 0.333]) and $PQ_{correct}$ ($\beta = 0.1440$, 95% CI = [0.053, 0.234]). CR_{33} , CR_{66} , CR_{100} , and FP_{first} also have a significant effect on the comprehension score of Q1 (CR_{33} : $\beta = -0.1897$, 95% CI = [-0.287, -0.092], CR_{66} : $\beta = 0.3583$, 95% CI = [0.257, 0.460], CR_{100} : $\beta = 0.7011$, 95% CI = [0.535, 0.867], FP_{first} : $\beta = -0.2519$, 95% CI = [-0.364, -0.140]). These results indicate that according to the correctness of the majority answer, participants’ comprehension is also affected and, if

the incorrect answer is shown first, comprehension decreases. The positive effect of inconsistency on the comprehension of Q1 is in line with our qualitative findings. Participants mentioned that, once they noticed inconsistencies between passages, they “reconsidered” the information (T90). T30 also stated that “the differences forced me to stop and consider the information I was seeing in each passage, and how much of it lined up/correlated between different passages” Additionally, we observed that inconsistencies encouraged some participants to rely on their pre-task answer or prior knowledge. D1 indicated that “one of the two different responses from the AI agreed with [my prior answer] so that may suggest it could be correct.”


Inconsistency also had a significantly positive effect on comprehension in Q2, which asked about information paraphrased across the passages ($\beta = 0.1541$, 95% CI = [0.015, 0.294]). This can imply that, as participants noticed an inconsistency in passages, this encouraged them to read passages more carefully and led them to understand information that was explained differently in the passages. Furthermore, CR_{100} , and $PQ_{confidence}$ also have a significant effect on the Q2 comprehension score (CR_{100} : $\beta = 0.3051$, 95% CI = [0.065, 0.545], $PQ_{confidence}$: $\beta = -0.2740$, 95% CI = [-0.541, -0.007]). Qualitative findings also show that participants read and thought carefully when they identified differences between answers. D30 said “Having different answers to the question made me think more about what the text was saying.”

Inconsistency had no significant effect on comprehension in Q3, which asked about information that was the same across passages.

5.3 RQ3: How the Effects of the Number of Passages Change when Model Accuracy Changes

Through bootstrapping, we explored how the treatment’s impact on perceived capacity and comprehension varies according to accuracy levels, which are captured by different fractions of inconsistent cases that a user group receives.

5.3.1 RQ3-1: How Effect of Passage Number on Perceived Capacity Changes. First, we examined how the number of passages affects users’ perceived capacity depending on the model’s accuracy or capability to provide correct information. In figure 4-(a) showing the perceived capacity of participants in all conditions at all accuracy levels, participants in the Single condition reported the highest perceived capacity compared to the Double and Triple conditions.



We constructed OLS regression models to predict a participant’s perceived capacity with bootstrapped data for each accuracy level. Figure 5-(a) compares coefficients of treatment across accuracy levels ranging from 50% to 98% with a step of 2%. We find that, **regardless of accuracy, both multiple passage conditions have a negative effect on perceived capacity.** As the model accuracy increases, the negative effect of the Double condition on perceived capacity decreases, which might be caused by the decreasing number of inconsistent Double passages. However, we also observe that there is not as much coefficient change in the Triple condition compared to the Double condition. This implies that the Triple condition may have a negative effect on perceived capacity across the accuracy levels, which could lead to the potential risk that users do not rely on the AI even when the AI is very capable.

5.3.2 RQ3-2: How Effect of Passage Number on Comprehension Changes. We next examine how the number of passages affects participants’ comprehension of the information when the model accuracy changes. Figure 4-(b) shows the comprehension scores for each condition at all accuracy levels. In lower accuracy, the comprehension score of participants under different conditions appears to be similar to each other. As accuracy increases, the Double condition has a slightly higher mean comprehension score compared to the other two conditions, and there are bigger gaps between the lower bounds of the conditions. Additionally, in all accuracy levels, variance in the Triple condition is the largest. Figure 5-(b) compares coefficients of treatment across accuracy levels ranging from 50% to 98% with a step of 2%, which can show changes in treatment’s impact on comprehension according to accuracy changes. **In all accuracy levels, we find that the Double condition has a positive effect on comprehension, whereas the Triple condition has a negative effect.** Moreover, it appears that as model accuracy increases, the Double condition’s positive effect on comprehension score increases, while there is less change in the coefficient of the Triple condition. Specifically, as accuracy is close to 100%, the coefficient of the Triple condition goes from negative to zero, which means that it becomes non-detrimental.

6 DISCUSSION

Here, we interpret the results from our study and propose implications for the design of LLM-powered systems based on our findings.




Figure 4: (a) Comparison of mean perceived AI capacity and 95% intervals across different levels of AI model accuracy. (b) Same visualization for comprehension scores.

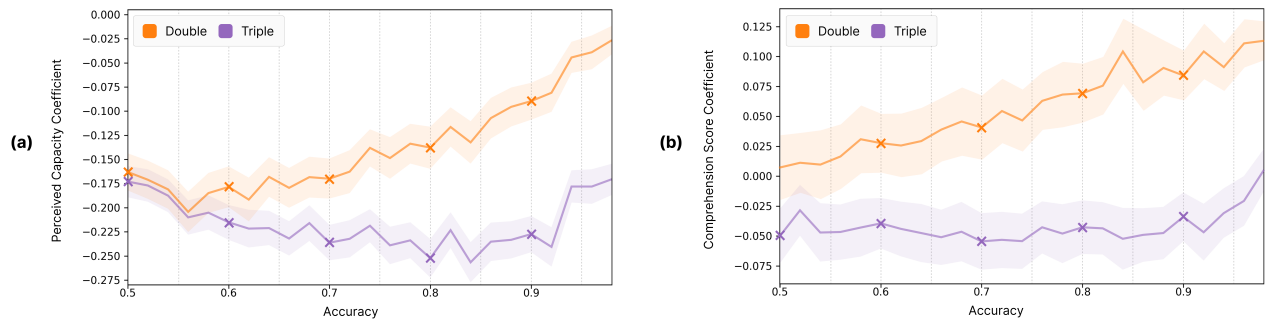


Figure 5: (a) Coefficient of the Double and Triple conditions on perceived AI capacity across accuracy levels for the AI model (50%-98%, step of 2%). The shaded region represents the 95% confidence intervals. Each 10% step of accuracy marked with a cross represents that the coefficient was statistically significant ($p < .05$). (b) Same visualization for comprehension scores.

6.1 Interpretation of Results

While prior work regarded inconsistency as a problem, our study showed that revealing inconsistencies can be beneficial for users when consuming AI-generated outputs. Our study found that inconsistency lowers the perceived AI capacity of the participants, as they considered that AI could “get confused” or contain “conflicting information” (Section 5.1). We also observed that inconsistencies had a positive effect on participants’ comprehension of details within the generated information (Section 5.2). Based on the theory of desirable difficulties [7], we posit that this finding can be attributed to how inconsistencies in AI outputs, which participants perceived as challenging yet manageable, may have encouraged them to engage more intensively with the content and read more thoroughly. The effect of inconsistencies in moderating trust in the AI also led participants to be more open to considering their own thoughts, rather than simply following AI, which is aligned with findings from prior literature [73] that showed how performance indicators can affect how people revise their decisions based on AI predictions. In other words, inconsistency had a positive effect on comprehension score regarding the main answer (Section 5.2) and more participants did not follow the AI’s majority answer when reading inconsistent passages (Section 5.1).

In situations when AI accuracy is above 50%, we expected that the Triple condition would have a positive effect on comprehension, as it provides more passages and increases the likelihood that participants receive at least one passage with correct information. However, the findings of Section 5.3 showed that the three-passage condition consistently hindered comprehension across all accuracy levels. We observed that this may be due to the way that most participants relied on simple majority voting as a heuristic—similar to software fault tolerance [46]—and ignored the correct passage when it was in the minority. Then, considering how participants resorted to majority voting, it can be assumed that the triple condition would have a positive effect on comprehension as accuracy increases, since most participants will receive passages where the majority is correct. However, we observed that the triple condition still had a negative impact on comprehension, potentially due to the increased cognitive load from reading more passages, as per the Cognitive Load Theory [30]. Due to this load, participants may have

skimmed the passages, preventing them from critically comparing the passages and leading them to overlook detailed information.

6.2 Design Implications

When designing an LLM-powered system for information consumption, we suggest that the system should tailor the amount of information generated (e.g., generated answer, proposed action) based on the model’s performance level. When the model’s performance is notably high (i.e., higher than 90%), the prevalent system design of providing a single output would be sufficient, as it is unlikely for the model to produce incorrect outputs and it can be reasonable for the user to rely on the LLM—showing more outputs introduces the unnecessary risk of users learning incorrect information. When employing models with lower accuracy, we suggest that systems should transition to designs that provide two outputs, as our study showed that this can prevent overreliance and encourage more critical reading, while maintaining cognitive load at a manageable level [39]. As model accuracy depends on the information domain or user context, we propose that developers can construct specific benchmarks to assess the models’ accuracy in domains of interest to then decide on how the system’s design should be adapted. While our study explored this transition from single- to multi-output designs in the context of information-seeking, this transition can also be effective for tasks where presenting users with multiple perspectives can promote more nuanced understanding, such as healthcare decision-making or news consumption. Finally, as revealed in our qualitative results, this transition can also be beneficial when it is important to trigger users’ intrinsic motivation to find more information—e.g., encourage students to look for more information to resolve an inconsistency between system outputs [7].

While we advocate for providing two outputs to users, we acknowledge the potential trade-off of presenting multiple outputs as they require more time and cognitive effort to compare and critically analyze. This increased processing load, essential for gaining deeper insights, might be a deterrent in real-life scenarios where users seek quick information or face time constraints. Therefore, the challenge lies in balancing the depth of content with the user’s capacity and willingness to engage with it. To address this, we recommend providing visual support to facilitate users’ processing

of multiple outputs. Inspired by prior work on supporting sense-making of text [64], we recommend gradually showing automated highlights of the differences between outputs by calculating the similarity between sentences and enabling users to annotate their interpretation on them. To facilitate sensemaking of these differences, identified inconsistencies could also be labeled with their types, as identified in Section 3.2.

In our study, we limited our comparison to double and triple as multiple passages due to technical-side and user-side concerns when employing a larger number of passages in real-world scenarios. Although the current generation speed of models may be too slow, advancements in model efficiency could allow real-time applications to produce large numbers of outputs for users. Future work can explore novel ways to help users assess the reliability of LLMs by comparing large samples. Using fact verification methods [50] to automatically distinguish between passages, future systems could display a variety of unique passages to avoid redundancy. Furthermore, by automatically identifying passages that share answers, system builders could calculate “confidence scores” for specific answers (i.e., percentage of outputs that share the answer)—as suggested in Leiser et al. [41]. Then, the system can present users with passages for the top-2 answers—leveraging the benefits of double passages while also providing reliability indicators.

While we discussed design implications focusing on LLMs, these guides can be broadly applied to the design and development of generative AI systems. We suggest, however, that developers should take into account how the output of each generative AI is consumed (e.g., time and effort required). For instance, users can find it easier to process and compare multiple image outputs, in contrast to text, speech or video outputs.

6.3 Intuitive Intervention to Build User’s Mental Models of LLMs

To prevent users from overrelying on LLM outputs, we propose an intuitive intervention of showing inconsistencies in the outputs. First, by seeing inconsistencies, users can more intuitively realize the nondeterministic nature of LLMs. This awareness can encourage users to adopt a critical perspective when interacting with these models, which can be crucial as LLMs produce more convincing outputs than previous AI models. This can also lead users to become curious about why the LLM was inconsistent, which can encourage them to understand the mechanics behind these model’s training and inference. Encouraging this understanding can be critical, as some of our participants misunderstood why the inconsistencies were produced. For example, some thought that the model *retrieved incorrect information* or became *confused as it read conflicting information*, rather than understanding the actual word-by-word probabilistic generation process. Thus, by presenting inconsistencies and the mechanisms behind them in a more transparent and explainable manner, systems could help users to better grasp the mechanics and limitations of these models, fostering more informed and critical consumption of AI-generated content. This method of displaying inconsistencies could serve as a more generalizable approach to representing uncertainty, a task that is challenging in standard chat interfaces [47].

6.4 Limitations & Future Work

While it is unclear whether our findings would also hold for other passage types, the insights gained may still apply to similar types of text that prioritize factual accuracy and the presentation of answers with supporting details. The passage types influence the user’s cognitive load, which can impact their ability to comprehend multiple passages. For example, passages with logical or mathematical reasoning might incur a higher cognitive load, causing users to struggle even when reading only two passages. Second, our study design was limited to a single turn of interaction with AI. This leaves room for future research on how the experience of reading a passage might affect subsequent usage of AI. Third, it is not clear how our study results were influenced by the pre-task question, which was used to control the prior knowledge effect. We also recognize that the order of questions on perceived AI capacity and inconsistencies could have affected outcomes, although these effects would likely be consistent across conditions. Finally, since we only focused on two types of inconsistencies, studying other types would provide more insight into how inconsistencies impact user behavior.

7 CONCLUSION

This work examines the impact of inconsistencies in LLM’s generated outputs on how users interact with LLMs in an information-seeking context. Through a qualitative analysis, we first identified five types of inconsistencies that can occur between the LLM-generated output. We conducted an online experiment (N=252) to understand the effect of providing users with multiple outputs that can contain inconsistencies. Our findings suggest that inconsistent output negatively impacted participants’ perceived AI capacity while positively influencing their comprehension of the generated information, especially for participants who received double passages. Based on these findings, we suggest that, instead of viewing inconsistencies as a limitation, future systems could leverage inconsistencies to mitigate users’ overreliance on AI and encourage critical examination of generated outputs.

ACKNOWLEDGMENTS

This work was supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korean government (MSIT) (No.2021-0-01347, Video Interaction Technologies Using Object-Oriented Video Modeling).

REFERENCES

- [1] Khan Academy. 2023. Khanmigo: Khan Academy’s AI-powered teaching assistant. Retrieved January 21, 2024 from <https://blog.khanacademy.org/teacher-khanmigo>
- [2] Dario Amodei, Christopher Olah, Jacob Steinhardt, Paul Francis Christiano, John Schulman, and Dandelion Mané. 2016. Concrete Problems in AI Safety. *ArXiv abs/1606.06565* (2016). <https://api.semanticscholar.org/CorpusID:10242377>
- [3] Stavros Antifakos, Adrian Schwaning, and Bernt Schiele. 2004. Evaluating the Effects of Displaying Uncertainty in Context-Aware Applications. In *Ubiquitous Computing*. <https://api.semanticscholar.org/CorpusID:2342122>
- [4] Akari Asai and Hannaneh Hajishirzi. 2020. Logic-Guided Data Augmentation and Regularization for Consistent Question Answering. In *Annual Meeting of the Association for Computational Linguistics*. <https://api.semanticscholar.org/CorpusID:216035859>
- [5] Gagan Bansal, Besmira Nushi, Ece Kamar, Daniel S. Weld, Walter S. Lasecki, and Eric Horvitz. 2019. Updates in Human-AI Teams: Understanding and Addressing the Performance/Compatibility Tradeoff. In *AAAI Conference on Artificial Intelligence*. <https://api.semanticscholar.org/CorpusID:53997192>

- [6] Edward Beeching, Clémentine Fourier, Nathan Habib, Sheon Han, Nathan Lambert, Nazneen Rajani, Omar Sanseviero, Lewis Tunstall, and Thomas Wolf. 2023. Open LLM Leaderboard. https://huggingface.co/spaces/HuggingFaceH4/open_llm_leaderboard.
- [7] Elizabeth L Bjork, Robert A Bjork, et al. 2011. Making things hard on yourself, but in a good way: Creating desirable difficulties to enhance learning. *Psychology and the real world: Essays illustrating fundamental contributions to society* 2, 59–68 (2011).
- [8] Richard E. Boyatzis. 1998. Transforming Qualitative Information: Thematic Analysis and Code Development.
- [9] Zana Buçinca, Maja Barbara Malaya, and Krzysztof Z. Gajos. 2021. To Trust or to Think: Cognitive Forcing Functions Can Reduce Overreliance on AI in AI-Assisted Decision-Making. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW1, Article 188 (apr 2021), 21 pages. <https://doi.org/10.1145/3449287>
- [10] Oana-Maria Camburu, Brendan Shillingford, Pasquale Minervini, Thomas Lukasiewicz, and Phil Blunsom. 2019. Make Up Your Mind! Adversarial Generation of Inconsistent Natural Language Explanations. In *Annual Meeting of the Association for Computational Linguistics*. <https://api.semanticscholar.org/CorpusID:203905467>
- [11] Furui Cheng, Vilém Zouhar, Simran Arora, Mrinmaya Sachan, Hendrik Strobelt, and Mennatallah El-Assady. 2023. RELIC: Investigating Large Language Model Responses using Self-Consistency. *ArXiv abs/2311.16842* (2023). <https://api.semanticscholar.org/CorpusID:265466244>
- [12] Hao Fei Cheng, Ruotong Wang, Zheng Zhang, Fiona O'Connell, Terrance Gray, F. Maxwell Harper, and Haiyi Zhu. 2019. Explaining Decision-Making Algorithms through UI: Strategies to Help Non-Expert Stakeholders. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (2019). <https://api.semanticscholar.org/CorpusID:140281803>
- [13] Avishek Choudhury and Hamid Shamszade. 2023. Investigating the Impact of User Trust on the Adoption and Use of ChatGPT: Survey Analysis. *Journal of Medical Internet Research* 25 (2023). <https://api.semanticscholar.org/CorpusID:258922988>
- [14] Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. Think you have Solved Question Answering? Try ARC, the AI2 Reasoning Challenge. *ArXiv abs/1803.05457* (2018). <https://api.semanticscholar.org/CorpusID:3922816>
- [15] Roi Cohen, May Hamri, Mor Geva, and Amir Globerson. 2023. LM vs LM: Detecting Factual Errors via Cross Examination. In *Conference on Empirical Methods in Natural Language Processing*. <https://api.semanticscholar.org/CorpusID:258833288>
- [16] Lynne M. Connelly. 2013. Grounded theory. *Medsurg nursing : official journal of the Academy of Medical-Surgical Nurses* 22 2 (2013), 124, 127.
- [17] David Dearman, Alex Varshavsky, Eyal de Lara, and Khai Nhut Truong. 2007. An Exploration of Location Error Estimation. In *Ubiquitous Computing*. <https://api.semanticscholar.org/CorpusID:9535527>
- [18] Pierre Dragicevic. 2016. Fair Statistical Communication in HCI. <https://api.semanticscholar.org/CorpusID:64470036>
- [19] Yanai Elazar, Nora Kassner, Shauli Ravfogel, Abhilasha Ravichander, Eduard H. Hovy, Hinrich Schütze, and Yoav Goldberg. 2021. Measuring and Benchmarking Consistency in Pretrained Language Models. *Transactions of the Association for Computational Linguistics* 9 (2021), 1012–1031. <https://api.semanticscholar.org/CorpusID:231740560>
- [20] Alexander R. Fabbri, Chien Sheng Wu, Wenhao Liu, and Caiming Xiong. 2021. QAFactEval: Improved QA-Based Factual Consistency Evaluation for Summarization. In *North American Chapter of the Association for Computational Linguistics*. <https://api.semanticscholar.org/CorpusID:245218667>
- [21] Riccardo Fogliato, Shreya Chappidi, Matthew P. Lungren, Michael Fitzke, Mark Parkinson, Diane U Wilson, Paul Fisher, Eric Horvitz, Kori Inkpen, and Besmira Nushi. 2022. Who Goes First? Influences of Human-AI Workflow on Decision Making in Clinical Imaging. *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (2022). <https://api.semanticscholar.org/CorpusID:248887734>
- [22] Krzysztof Z. Gajos and Lena Mamykina. 2022. Do People Engage Cognitively with AI? Impact of AI Assistance on Incidental Learning. In *27th International Conference on Intelligent User Interfaces (IUI '22)*. Association for Computing Machinery, New York, NY, USA, 794–806. <https://doi.org/10.1145/3490099.3511138>
- [23] Arthur C. Graesser and Natalie K. Person. 1994. Question Asking During Tutoring. <https://api.semanticscholar.org/CorpusID:15485207>
- [24] Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Xiaodong Song, and Jacob Steinhardt. 2020. Measuring Massive Multitask Language Understanding. *ArXiv abs/2009.03300* (2020). <https://api.semanticscholar.org/CorpusID:221516475>
- [25] Dan Hendrycks, Nicholas Carlini, John Schulman, and Jacob Steinhardt. 2021. Unsolved Problems in ML Safety. *ArXiv abs/2109.13916* (2021). <https://api.semanticscholar.org/CorpusID:238198240>
- [26] Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. 2023. A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions. *ArXiv abs/2311.05232* (2023). <https://api.semanticscholar.org/CorpusID:265067168>
- [27] Myeongjun Jang, Deuk Sin Kwon, and Thomas Lukasiewicz. 2021. Accurate, yet inconsistent? Consistency Analysis on Language Understanding Models. *ArXiv abs/2108.06665* (2021). <https://api.semanticscholar.org/CorpusID:237091303>
- [28] Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Yejin Bang, Wenliang Dai, Andrea Madotto, and Pascale Fung. 2022. Survey of Hallucination in Natural Language Generation. *Comput. Survveys* 55 (2022), 1–38. <https://api.semanticscholar.org/CorpusID:246652372>
- [29] Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke Zettlemoyer. 2017. TriviaQA: A Large Scale Distantly Supervised Challenge Dataset for Reading Comprehension. *ArXiv abs/1705.03551* (2017). <https://api.semanticscholar.org/CorpusID:26501419>
- [30] Pamela Karr-Wisniewski and Ying Lu. 2010. When more is too much: Operationalizing technology overload and exploring its impact on knowledge worker productivity. *Computers in Human Behavior* 26, 5 (2010), 1061–1072.
- [31] Pranav Khadpe, Ranjay Krishna, Fei-Fei Li, Jeffrey T. Hancock, and Michael S. Bernstein. 2020. Conceptual Metaphors Impact Perceptions of Human-AI Collaboration. *Proceedings of the ACM on Human-Computer Interaction* 4 (2020), 1–26. <https://api.semanticscholar.org/CorpusID:221005643>
- [32] Sunnie S. Y. Kim, Elizabeth Anne Watkins, Olga Russakovsky, Ruth Fong, and Andrés Monroy-Hernández. 2023. Humans, AI, and Context: Understanding End-Users' Trust in a Real-World Computer Vision Application. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency (FAccT '23)*. Association for Computing Machinery, New York, NY, USA, 77–88. <https://doi.org/10.1145/3593013.3593978>
- [33] Lars Krupp, Steffen Steinert, Maximilian Kiefer-Emmanouilidis, Karina E. Avila, Paul Lukowicz, Jochen Kuhn, Stefan Küchemann, and Jakob Karolus. 2023. Challenges and Opportunities of Moderating Usage of Large Language Models in Education. *ArXiv abs/2312.14969* (2023). <https://api.semanticscholar.org/CorpusID:266550959>
- [34] Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur P. Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc V. Le, and Slav Petrov. 2019. Natural Questions: A Benchmark for Question Answering Research. *Transactions of the Association for Computational Linguistics* 7 (2019), 453–466. <https://api.semanticscholar.org/CorpusID:86611921>
- [35] Philippe Laban, Tobias Schnabel, Paul N. Bennett, and Marti A. Hearst. 2021. SummaC: Re-Visiting NLI-based Models for Inconsistency Detection in Summarization. *Transactions of the Association for Computational Linguistics* 10 (2021), 163–177. <https://api.semanticscholar.org/CorpusID:244345901>
- [36] Vivian Lai, Chacha Chen, Alison Smith-Renner, Qingzi Vera Liao, and Chenhao Tan. 2023. Towards a Science of Human-AI Decision Making: An Overview of Design Space in Empirical Human-Subject Studies. *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (2023). <https://api.semanticscholar.org/CorpusID:259139714>
- [37] Vivian Lai and Chenhao Tan. 2019. On Human Predictions with Explanations and Predictions of Machine Learning Models: A Case Study on Deception Detection. *Proceedings of the Conference on Fairness, Accountability, and Transparency* (2019). <https://api.semanticscholar.org/CorpusID:53774958>
- [38] Himabindu Lakkaraju, Stephen H. Bach, and Jure Leskovec. 2016. Interpretable Decision Sets: A Joint Framework for Description and Prediction. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2016). <https://api.semanticscholar.org/CorpusID:12533380>
- [39] G. Latif-Shabgahi, J.M. Bass, and S. Bennett. 2004. A taxonomy for software voting algorithms used in safety-critical systems. *IEEE Transactions on Reliability* 53, 3 (2004), 319–328. <https://doi.org/10.1109/TR.2004.832819>
- [40] Yoonjoo Lee, Tae Soo Kim, Sungdong Kim, Yohan Yun, and Juho Kim. 2023. DAPIE: Interactive Step-by-Step Explanatory Dialogues to Answer Children's Why and How Questions. *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (2023). <https://api.semanticscholar.org/CorpusID:258216670>
- [41] Florian Leiser, Sven Eckhardt, Merlin Knaeble, Alexander Maedche, Gerhard Schwabe, and Ali Sunyaev. 2023. From ChatGPT to FactGPT: A Participatory Design Study to Mitigate the Effects of Large Language Model Hallucinations on Users. In *Proceedings of Mensch Und Computer 2023 (MuC '23)*. Association for Computing Machinery, New York, NY, USA, 81–90. <https://doi.org/10.1145/3603555.3603565>
- [42] Junyi Li, Jie Chen, Ruiyang Ren, Xiaoxue Cheng, Wayne Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. 2024. The Dawn After the Dark: An Empirical Study on Factuality Hallucination in Large Language Models. <https://api.semanticscholar.org/CorpusID:266844012>
- [43] Brian Y. Lim, Anind K. Dey, and Daniel Avrahami. 2009. Why and why not explanations improve the intelligibility of context-aware intelligent systems. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2009). <https://api.semanticscholar.org/CorpusID:4507550>
- [44] Stephanie C. Lin, Jacob Hilton, and Owain Evans. 2021. TruthfulQA: Measuring How Models Mimic Human Falsehoods. In *Annual Meeting of the Association for Computational Linguistics*. <https://api.semanticscholar.org/CorpusID:237532606>

- [45] Stephanie C. Lin, Jacob Hilton, and Owain Evans. 2022. Teaching Models to Express Their Uncertainty in Words. *Trans. Mach. Learn. Res.* 2022 (2022). <https://api.semanticscholar.org/CorpusID:249191391>
- [46] Diane M Mackie. 1987. Systematic and nonsystematic processing of majority and minority persuasive communications. *Journal of Personality and Social Psychology* 53, 1 (1987), 41.
- [47] Erwin Marsi and Ferdi Van Rooden. 2007. Expressing uncertainty with a talking head in a multimodal question-answering system. <https://api.semanticscholar.org/CorpusID:2482651>
- [48] Sabrina J. Mielke, Arthur Szlam, Emily Dinan, and Y-Lan Boureau. 2022. Reducing Conversational Agents' Overconfidence Through Linguistic Calibration. *Transactions of the Association for Computational Linguistics* 10 (2022), 857–872. https://doi.org/10.1162/tacl_a_00494
- [49] Tim Miller. 2023. Explainable AI is Dead, Long Live Explainable AI! Hypothesis-driven Decision Support using Evaluative AI. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency (FAcCT '23)*. Association for Computing Machinery, New York, NY, USA, 333–342. <https://doi.org/10.1145/3593013.3594001>
- [50] Sewon Min, Kalpesh Krishna, Xixi Lyu, Mike Lewis, Wen tau Yih, Pang Wei Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. FActScore: Fine-grained Atomic Evaluation of Factual Precision in Long Form Text Generation. *ArXiv abs/2305.14251* (2023). <https://api.semanticscholar.org/CorpusID:258841470>
- [51] Brent Mittelstadt, Chris Russell, and Sandra Wachter. 2019. Explaining Explanations in AI. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT* '19)*. Association for Computing Machinery, New York, NY, USA, 279–288. <https://doi.org/10.1145/3287560.3287574>
- [52] Mahsan Nourani, Samia Kabir, Sina Mohseni, and Eric D. Ragan. 2019. The Effects of Meaningful and Meaningless Explanations on Trust and Perceived System Accuracy in Intelligent Systems. In *AAAI Conference on Human Computation & Crowdsourcing*. <https://api.semanticscholar.org/CorpusID:201639081>
- [53] OpenAI. 2022. Introducing ChatGPT. Retrieved January 17, 2024 from <https://openai.com/blog/chatgpt>
- [54] Cecilia Panigutti, Andrea Beretta, Fosca Giannotti, and Dino Pedreschi. 2022. Understanding the Impact of Explanations on Advice-Taking: A User Study for AI-Based Clinical Decision Support Systems. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 568, 9 pages. <https://doi.org/10.1145/3491102.3502104>
- [55] Andrea Papenmeier, Dagmar Kern, Daniel Hienert, Yvonne Kammerer, and Christin Seifert. 2022. How Accurate Does It Feel? – Human Perception of Different Types of Classification Mistakes. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 180, 13 pages. <https://doi.org/10.1145/3491102.3501915>
- [56] Peter S. Park, Simon Goldstein, Aidan O'Gara, Michael Chen, and Dan Hendrycks. 2023. AI Deception: A Survey of Examples, Risks, and Potential Solutions. *arXiv:cs.CY/2308.14752*
- [57] Fabio Petroni, Aleksandra Piktus, Angela Fan, Patrick Lewis, Majid Yazdani, Nicola De Cao, James Thorne, Yacine Jernite, Vladimir Karpukhin, Jean Mailard, Vassilis Plachouras, Tim Rocktäschel, and Sebastian Riedel. 2021. KILT: A Benchmark for Knowledge Intensive Language Tasks. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek Hakkani-Tur, Iz Beltagy, Steven Bethard, Ryan Cotterell, Tammy Chakraborty, and Yichao Zhou (Eds.). Association for Computational Linguistics, Online, 2523–2544. <https://doi.org/10.18653/v1/2021.naacl-main.200>
- [58] Marissa Radensky, Julie Anne Séguin, Jang Soo Lim, Kristen Olson, and Robert Geiger. 2023. "I Think You Might Like This": Exploring Effects of Confidence Signal Patterns on Trust in and Reliance on Conversational Recommender Systems. *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (2023). <https://api.semanticscholar.org/CorpusID:259139753>
- [59] Abhilasha Ravichander, Eduard H. Hovy, Kaheer Suleman, Adam Trischler, and Jackie Chi Kit Cheung. 2020. On the Systematicity of Probing Contextualized Word Representations: The Case of Hypernymy in BERT. In *STARSEM*. <https://api.semanticscholar.org/CorpusID:227230677>
- [60] Marco Tulio Ribeiro, Carlos Guestrin, and Sameer Singh. 2019. Are Red Roses Red? Evaluating Consistency of Question-Answering Models. In *Annual Meeting of the Association for Computational Linguistics*. <https://api.semanticscholar.org/CorpusID:196182403>
- [61] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2018. Anchors: High-Precision Model-Agnostic Explanations. In *AAAI Conference on Artificial Intelligence*. <https://api.semanticscholar.org/CorpusID:3366554>
- [62] Kyoungwon Seo, Joice Tang, Ido Roll, Sidney Fels, and Dongwook Yoon. 2021. The impact of artificial intelligence on learner–instructor interaction in online learning. *International journal of educational technology in higher education* 18, 1 (2021), 1–23.
- [63] Sofia Eleni Spatharioti, David M. Rothschild, Daniel G. Goldstein, and Jake M. Hofman. 2023. Comparing Traditional and LLM-based Search for Consumer Choice: A Randomized Experiment. *ArXiv abs/2307.03744* (2023). <https://api.semanticscholar.org/CorpusID:259375527>
- [64] Hariharan Subramonyam, Colleen Seifert, Priti Shah, and Eytan Adar. 2020. texSketch: Active Diagramming through Pen-and-Ink Annotations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376155>
- [65] Sangho Suh, Bryan Min, Srishti Palani, and Haijun Xia. 2023. Sensecape: Enabling Multilevel Exploration and Sensemaking with Large Language Models. *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (2023). <https://api.semanticscholar.org/CorpusID:258822925>
- [66] Meiqi Sun, Wilson Yan, P. Abbeel, and Igor Mordatch. 2022. Quantifying Uncertainty in Foundation Models via Ensembles. <https://api.semanticscholar.org/CorpusID:254183614>
- [67] Logesh Kumar Umaphathi, Ankit Pal, and Malaikannan Sankarasubbu. 2023. Med-HALT: Medical Domain Hallucination Test for Large Language Models. In *Conference on Computational Natural Language Learning*. <https://api.semanticscholar.org/CorpusID:260316324>
- [68] Alex Wang, Kyunghyun Cho, and Mike Lewis. 2020. Asking and Answering Questions to Evaluate the Factual Consistency of Summaries. *ArXiv abs/2004.04228* (2020). <https://api.semanticscholar.org/CorpusID:215548661>
- [69] Lu Wang, Max Song, Rezvaneh Rezapour, Bum Chul Kwon, and Jina Huh-Yoo. 2023. People's Perceptions Toward Bias and Related Concepts in Large Language Models: A Systematic Review. *ArXiv abs/2309.14504* (2023). <https://api.semanticscholar.org/CorpusID:262825989>
- [70] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Huai hsin Chi, and Denny Zhou. 2022. Self-Consistency Improves Chain of Thought Reasoning in Language Models. *ArXiv abs/2203.11171* (2022). <https://api.semanticscholar.org/CorpusID:247595263>
- [71] Justin D. Weisz, Michael J. Muller, Jessica He, and Stephanie Houde. 2023. Toward General Design Principles for Generative AI Applications 130-144. In *UI Workshops*. <https://api.semanticscholar.org/CorpusID:255825625>
- [72] Ikedinachi Ayodele Power Wogu, Sanjay Misra, Esther Fadeke Olu-Owolabi, Patrick A Assibong, Oluwakemi D Udoh, SO Ogiri, and R Damasevicius. 2018. Artificial intelligence, artificial teachers and the fate of learners in the 21st century education sector: Implications for theory and practice. *International Journal of Pure and Applied Mathematics* 119, 16 (2018), 2245–2259.
- [73] Ming Yin, Jennifer Wortman Vaughan, and Hanna M. Wallach. 2019. Understanding the Effect of Accuracy on Trust in Machine Learning Models. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (2019). <https://api.semanticscholar.org/CorpusID:109927933>
- [74] Qiaoning Zhang, Matthew L Lee, and Scott Carter. 2022. You Complete Me: Human-AI Teams and Complementary Expertise. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 114, 28 pages. <https://doi.org/10.1145/3491102.3517791>
- [75] Yunfeng Zhang, Qingzi Vera Liao, and Rachel K. E. Bellamy. 2020. Effect of confidence and explanation on accuracy and trust calibration in AI-assisted decision making. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (2020). <https://api.semanticscholar.org/CorpusID:210023849>
- [76] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. 2023. A Survey of Large Language Models. *arXiv:cs.CL/2303.18223*
- [77] Kaitlyn Zhou, Dan Jurafsky, and Tatsunori Hashimoto. 2023. Navigating the Grey Area: Expressions of Overconfidence and Uncertainty in Language Models. *ArXiv abs/2302.13439* (2023). <https://api.semanticscholar.org/CorpusID:257220189>

A INCONSISTENCY TYPES AND EXAMPLES

Type	Passage A	Passage B	Description
Main Answer Inconsistency	As of my last knowledge update in September 2021, <i>China</i> was the large country that spent the most on international tourism;	As of my last knowledge update in September 2021, the <i>United States</i> was the large country that spends the most on international tourism.	Passage A explains the answer is China, and Passage B responds the answer is United States on the same question.
Detail Volume Inconsistency	[Main answer A, supporting detail a and b]	[Main answer A, supporting detail a and b] <i>Additionally, ethical considerations and practical limitations often make raising a chimpanzee as a human child an impractical and controversial endeavor.</i>	Compared to Passage A, Passage B explains additional information related to the ethical considerations and practical limitations of raising a chimpanzee as a human child.
Detail Content Inconsistency	Coffee has become deeply ingrained in American culture, serving as a ubiquitous morning ritual and a social beverage	Coffee has a deep roots in American history, dating back to colonial period, and it became a symbol of the American Revolution as an alternative British tea.	Passage A emphasizes contemporary cultural aspects of coffee in American life, while Passage B introduces historical context, creating a difference in the level of detail and focus on the historical roots of coffee.
Detail Expansion Inconsistency	In fact, <i>gluten-free products often lack certain nutrients</i>	In fact, gluten-free products often lack the same nutritional content and <i>may contain added sugars and fats to compensate for taste and texture.</i>	Passages A and B both say that gluten-free products lack certain nutrients, but Passage B additionally explains what kind of nutrients they are.
Paraphrasing Inconsistency	Lastly, the fast-paced American lifestyle often necessitates a caffeine boost, further driving the demand for coffee.	The demands of a fast-paced lifestyle, which all contribute to the high coffee consumption in the United States.	Passages A and B both convey the same information with different styles.

Table 2: The four types of inconsistency that were identified between pairs of outputs generated by LLM in our Preliminary Study. Each type of inconsistency is illustrated with example passages.

B INTERFACE

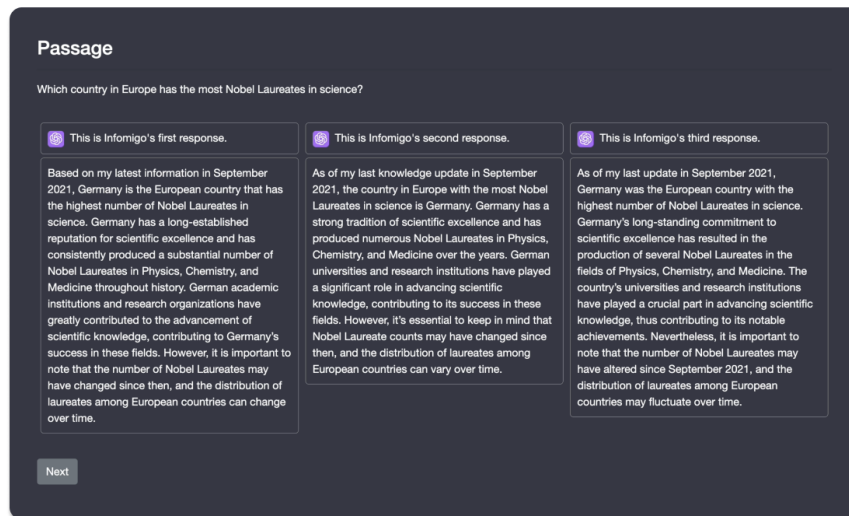


Figure 6: The interface employed during the experiment showing the procedure step where participants are presented with passages generated by the AI assistant. The example demonstrates the triple condition where participants receive three passages.

C DETAILS IN CONDITIONS

Condition	subconditions	Passage participants received
Single	O	$O_k (k \in \{1, 2, 3\})$
	X	$X_k (k \in \{1, 2, 3\})$
Double	OO	$O_k O_l (k \neq l, k, l \in \{1, 2, 3\})$
	OX	$O_k X_l (k \neq l, k, l \in \{1, 2, 3\})$
	XO	$X_k O_l (k \neq l, k, l \in \{1, 2, 3\})$
	XX	$X_k X_l (k \neq l, k, l \in \{1, 2, 3\})$
Triple	OOO	$O_k O_l O_m (k \neq l, l \neq m, k \neq m, k, l, m \in \{1, 2, 3\})$
	OOX	$O_k O_l X_m (k \neq l, l \neq m, k \neq m, k, l, m \in \{1, 2, 3\})$
	OXO	$O_k X_l O_m (k \neq l, l \neq m, k \neq m, k, l, m \in \{1, 2, 3\})$
	XOO	$X_k O_l O_m (k \neq l, l \neq m, k \neq m, k, l, m \in \{1, 2, 3\})$
	OXX	$O_k X_l X_m (k \neq l, l \neq m, k \neq m, k, l, m \in \{1, 2, 3\})$
	XOX	$X_k O_l X_m (k \neq l, l \neq m, k \neq m, k, l, m \in \{1, 2, 3\})$
	XXO	$X_k X_l O_m (k \neq l, l \neq m, k \neq m, k, l, m \in \{1, 2, 3\})$
	XXX	$X_k X_l X_m (k \neq l, l \neq m, k \neq m, k, l, m \in \{1, 2, 3\})$

Table 3: Condition table that includes subconditions that are allocated to each condition and type of passage that participant in each subcondition received.