



## Article

# Fluent but Not Factual: A Comparative Analysis of ChatGPT and Other AI Chatbots' Proficiency and Originality in Scientific Writing for Humanities

Edisa Lozić and Benjamin Štular \*

Research Centre of the Slovenian Academy of Sciences and Arts, 1000 Ljubljana, Slovenia

\* Correspondence: benjamin.stular@zrc-sazu.si

**Abstract:** Historically, mastery of writing was deemed essential to human progress. However, recent advances in generative AI have marked an inflection point in this narrative, including for scientific writing. This article provides a comprehensive analysis of the capabilities and limitations of six AI chatbots in scholarly writing in the humanities and archaeology. The methodology was based on tagging AI-generated content for quantitative accuracy and qualitative precision by human experts. Quantitative accuracy assessed the factual correctness in a manner similar to grading students, while qualitative precision gauged the scientific contribution similar to reviewing a scientific article. In the quantitative test, ChatGPT-4 scored near the passing grade (−5) whereas ChatGPT-3.5 (−18), Bing (−21) and Bard (−31) were not far behind. Claude 2 (−75) and Aria (−80) scored much lower. In the qualitative test, all AI chatbots, but especially ChatGPT-4, demonstrated proficiency in recombining existing knowledge, but all failed to generate original scientific content. As a side note, our results suggest that with ChatGPT-4, the size of large language models has reached a plateau. Furthermore, this paper underscores the intricate and recursive nature of human research. This process of transforming raw data into refined knowledge is computationally irreducible, highlighting the challenges AI chatbots face in emulating human originality in scientific writing. Our results apply to the state of affairs in the third quarter of 2023. In conclusion, while large language models have revolutionised content generation, their ability to produce original scientific contributions in the humanities remains limited. We expect this to change in the near future as current large language model-based AI chatbots evolve into large language model-powered software.

**Keywords:** generative AI; large language model (LLM); ChatGPT; Bard; Bing; scientific writing; digital humanities



**Citation:** Lozić, E.; Štular, B. Fluent but Not Factual: A Comparative Analysis of ChatGPT and Other AI Chatbots' Proficiency and Originality in Scientific Writing for Humanities. *Future Internet* **2023**, *15*, 336. <https://doi.org/10.3390/fi15100336>

Academic Editors: Antony Bryant, Roberto Montemanni, Min Chen, Paolo Bellavista, Kenji Suzuki and Jeanine Treffers-Daller

Received: 15 September 2023  
Revised: 6 October 2023  
Accepted: 9 October 2023  
Published: 13 October 2023



**Copyright:** © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

## 1. Introduction: AI's Great Inflection Point

In recent human history, the ability to write well was considered essential to human progress and professionalism. Creative expression was traditionally seen as a defining characteristic of humanity and the pinnacle of human achievement. This view is still reflected in the way universities cultivate their students' writing skills. Until recently, novel artefacts such as literary works, scientific texts, art, and music were difficult to create and only attainable by talented experts [1–3].

We must reckon with that changing!

The current generation of openly available generative AI has rightly been called AI's great inflection point. Generative AI is shaping to become a general-purpose technology [4], a “fundamental, horizontal technology that will touch everything in our lives” (Tim Cook, Apple CEO, speaking at Università Degli Studi di Napoli Federico II in Naples, Italy, 29 September 2022).

The most recent and disruptive advance in generative AI has been a leap-frog development in the field of large language models (hereafter LLMs). LLMs are based on deep

neural networks and self-supervised learning that have been around for decades, but the amount of data that the current LLMs were trained with has led to an unprecedented and, to some extent, unexpected performance leap. They belong to foundation models that are pre-trained on large datasets using self-supervision at scale and then adapted to a wide range of downstream tasks. This centralisation is crucial for harnessing the enormous computing power needed to create them, but it also replicates potential problems such as security risks and biases [3,5].

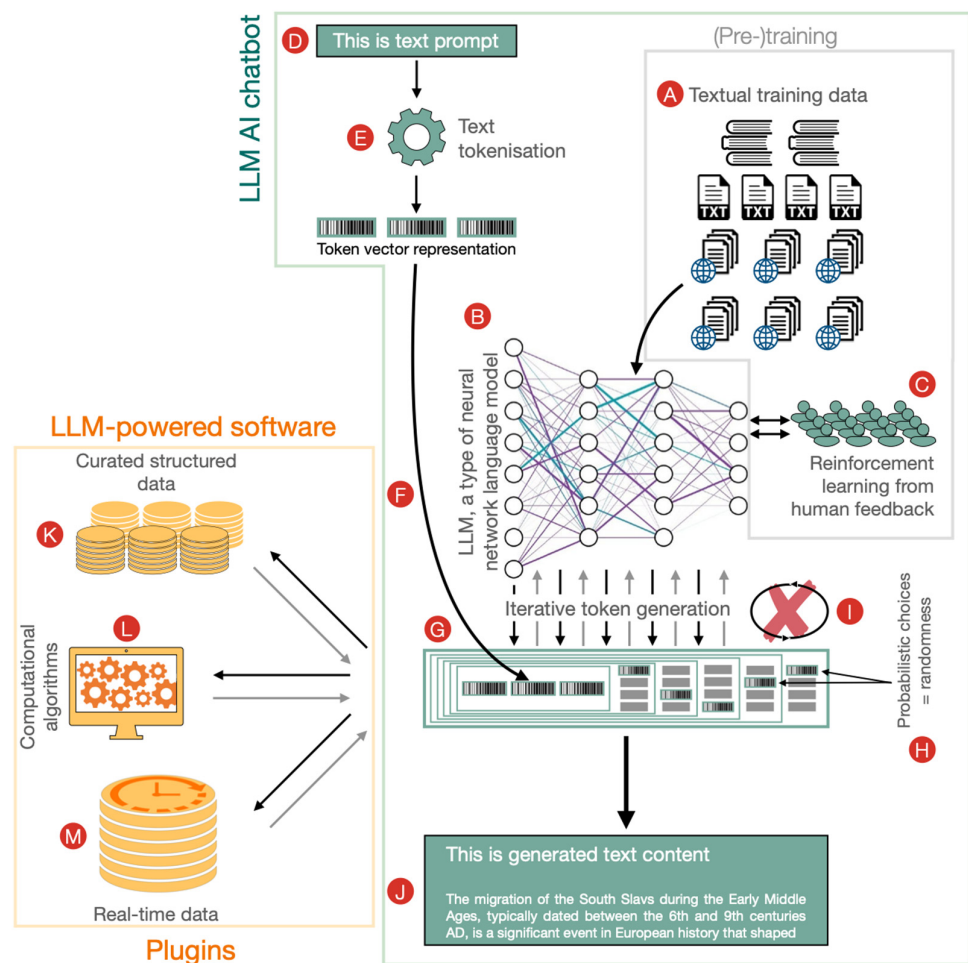
Currently, the most powerful LLMs are generative pretrained transformers (hereafter GPTs; we use the acronym GPT for Transformer models in general and the term ChatGPT for Transformer models built by OpenAI), which are based on the Transformer, a type of neural network architecture. The Transformer uses a mechanism called attention to weigh the influence of different input words on each output word. As a result, instead of processing words in a sentence sequentially, it constructs relationships between all words in a sentence at once [6]. An additional key advantage of GPTs over earlier models is that the learning process can be parallelised, and they can be trained on an unprecedented scale.

The scale of an LLM depends on the size of ingested datasets, the amount of training compute, and the number of parameters it can support [7,8]. Parameters are numerical values that determine how a neural network processes and generates natural language. The more parameters a model has, the more data it can learn from and the more complex tasks it can perform. The largest version of OpenAI's GPT-3 (Davinci) from 2020, for example, supports 175 billion parameters and has been trained on 45 TB of text data, including almost the entire public web [9]. PaLM, the 2022 LLM from Google Research, is a 540-billion-parameter GPT trained with the Pathways system [10], and GPT-4 launched in 2023 supports an estimated 1.8 trillion parameters [11]. That is 1,800,000,000,000 parameters with which the LLM interacts to generate each individual token (a word or a part of a word). Multiplied by ChatGPT's 100,000,000 monthly users each processing just one prompt with 100 tokens daily brings us to a staggering 18,000,000,000,000,000,000 or  $18 \times 10^{21}$  computations, which explains the daily costs of running ChatGPT at USD 700,000 [12]. One can only imagine the environmental footprint of the operation.

When given an input or prompt, LLMs are able to predict the next word in the context of all previous content and can thus generate creative outputs such as complete sentences and answers or even essays and poems. In effect, they generate a pattern of words based on the word patterns they have been trained on by applying attention to context and a controlled amount of randomness. Due to the fact that they have been trained on such a large amount of text, the quality of the text is such that GPT-4, for example, has been able to pass standardised academic and professional tests [13].

Users interact with an LLM via AI conversational agents or AI chatbots. An AI chatbot is a software application that uses AI techniques to simulate human conversation. In the past, AI chatbots did not use LLMs, and in the future, AI chatbots are likely to use resources other than LLMs, such as curated structured data, computational algorithms, and real-time data. Currently, however, AI chatbots are mostly used as conduits to LLMs, and thus, the terms AI chatbot and LLM are sometimes used interchangeably (Figure 1).

The most-discussed AI chatbot in history was ChatGPT, built on the GPT called GPT-3.5 (specifically gpt-3.5-turbo) [16] (<https://chat.openai.com>; accessed on 1 October 2023). It was released to the public in November 2022 and had 1 million active users in its first week (<https://twitter.com/sama/status/1599668808285028353>; accessed on 1 October 2023) and over 100 million after two months [17], easily surpassing Facebook, TikTok, Instagram, and Uber to that milestone. Since then, Google has launched Bard [18] (<https://bard.google.com>; accessed on 1 October 2023), Microsoft the Bing chatbot [19] (<https://www.bing.com>; accessed on 1 October 2023), OpenAI ChatGPT Plus [20] (<https://chat.openai.com>; accessed on 1 October 2023), Opera Aria [21] (<https://www.opera.com/features/browser-ai>; accessed on 1 October 2023), Anthropic Claude 2 [22,23] (<https://claude.ai>; accessed on 1 October 2023), and more.



**Legend**

- A** Ingestion of a large body of human-generated texts drawn from large web corpora, books, social networks, etc.
- B** Pre-training of a neural network language model; since there is not enough written text to directly deduce all probabilities, a large\* language model (LLM) estimates these probabilities (\*GPT-4 = 1,800,000,000,000 parameters).
- C** Reinforcement learning from human feedback introduces guardrails, reduces biases, etc.
- D** A textual prompt...
- E** ...is broken down into tokens (words or parts of a word)...
- F** ...which enables the network to initiate the content.
- G** LLM generates text by ranking potential next tokens according to their probability;
- H** this ranking takes context into account and introduces randomness to prevent repetitive and uncreative output.
- I** The procedure involves processing the prompt and the already generated text through all parameters\* for each token without loops.
- J** The generated text successfully emulates the content that adheres to the initial prompt.

- K L M** AI chatbot queries curated structured data, computational algorithms and/or real-time data to align the generated content with accurate computations and/or exact data. It can run in loops and can compute computationally irreducible problems. It can be achieved either through plugins for the AI chatbot or through LLM-powered software.

**Figure 1.** A high-level overview of an AI chatbot and LLM-powered software (constructed based on information from [14,15]).

AI chatbots have implications for a range of practices and disciplines and for many facets of our daily lives. They can be used to compose emails, essays, computer code, and speeches or translate text into another language or simply give it a different tone. They can be empowering because they lower the barrier to entry, seamlessly complement human work, and make us more productive and creative [3,24,25]. Many (but mostly proponents of entities with a stake in AI chatbots) tout that AI chatbots will take the drudgery out of everyday office work by automating various tasks and ultimately increasing productivity across the economy, e.g., [26,27].

However, AI chatbots can also be terrifying and are seen by many as potentially threatening and having far-reaching negative consequences. They could amplify the biases we already experience, undermine our trust in information, and take away our ability to determine what is real and what is not. Equally important, they are likely to upend the creative and knowledge industries and many types of work that have been primarily conducted by well-compensated professionals, such as artists, writers, executives, and programmers [3,24,25].

A recent study of the potential impact on the labour market found that LLMs and LLM-powered software will affect between 47% and 56% of all work-related tasks, with higher-income jobs being more affected [4]. Of particular interest to our article is that roles that rely heavily on science and critical thinking show a negative correlation with exposure to LLMs, yet there is some exposure [4].

Thus, AI experts, journalists, policymakers, and the public are increasingly discussing a wide range of important and urgent risks of AI, such as the AI race, organizational risks, rogue AIs, the reinforcement of social inequalities, reshaping labour and expertise, and the exacerbation of environmental injustices [28–33]. Questions are being asked about safety [34,35], capabilities [36], massive workforce redundancy [37], and legality [38,39]. This has led to calls for a pause in AI development [40], although there are doubts that such attempts would have any impact [41].

In short, the opportunities that generative AI presents for our lives, our communities, and our society are as great as the risks it poses [42]. There are those who believe that generative AI will change our work and our lives in general for the better. There are others who believe that generative AI will disastrously encroach on areas best navigated by sentient beings. Regardless, all agree on an urgent need for safeguards [43–46], and the first steps have already been taken [47].

All this is even more true for the planned next step, artificial general intelligence (hereafter AGI). The term AGI describes AI systems that will generally be more intelligent than humans and are planned to become a reality within the next 10 years. After the introduction of the first AGI, the world could be extremely different from how it is today [48].

And how does academia and the way we create and write research and scholarly articles fit in? AI chatbots have the potential to revolutionise academia and scholarly publishing [49]. In fact, it seems that academia will be among the first industries to go through this process, since academics and students represented two of the top three occupational groups among the early adopters of ChatGPT [50].

AI chatbots—most of the attention to date was directed to ChatGPT—have already been recognised as a powerful tool for scientific writing that can help organise material, proofread, draft, and generate summaries [51–53]. The scientific community is also actively testing their ability to generate entire papers with minimal human input. The consensus is that AI chatbots are able to create scientific essays and reports of scientific experiments that appear credible, but the resulting texts are a combination of true and entirely fabricated information [49,51,54–63].

Unfortunately, there is a public perception that ChatGPT is already capable of generating academic papers that are peer-reviewed and published, e.g., [64,65], which may add to public scepticism about science. This is not the case. For example, the ArXiv repository of preprints (<https://arxiv.org>; accessed on 16 August 2023), the AI community's most popular publication forum, shows no results for ChatGPT (in any variation) as a (co-)author.

We are aware of a single such attempt that has attracted much public attention but has not been peer reviewed or published and has no notable scientific value [55].

Regardless, there is a clear consensus among researchers that AI chatbots will be widely adopted in scientific writing in the near future, and it is thus crucial to reach an accord on how to regulate their use [63,66]. To successfully discuss regulation, a better understanding of AI chatbots is needed. This requires more systematic testing of their capabilities, which will provide a more robust understanding of the strengths and weaknesses of the technology. This process has been likened to the processes drugs go through to gain approval. Assessment of AI systems could allow them to be deemed safe for certain purposes and explain to users where they might fail [66].

To this end, there is a growing body of testing and benchmarking of generative AI LLMs, e.g., [13,67,68]. The standard methodology in machine learning is to evaluate the system against a set of standard benchmark datasets, ensuring that these are independent of the training data and span a range of tasks and domains. This strategy aims to distinguish real learning from mere memorisation. However, this approach is not ideally suited to our needs and to the study of LLM-based AI chatbots for three reasons. First, only the creators of proprietary LLMs have access to all of the training details needed for detailed benchmark results. Second, one of the key aspects of the intelligence of LLMs is their generality and ability to perform tasks that go beyond the typical scope of narrow AI systems. The metric of evaluation benchmarks designed for such generative or interactive tasks remains a challenge. The third and perhaps most important reason is that, in this article, we are interested in how well AI chatbots perform in human tasks. To evaluate this, methods closer to traditional psychology leveraging human creativity and curiosity are needed [69].

Such an approach has already been taken, and there are several evaluations of the performance of AI chatbots in scientific writing, but most of them focus on medicine and similar fields [49,51,56–63,70–72]. We are not aware of any such test designed specifically for humanities.

This means that there is a gap in knowledge about the reliability of the content generated by AI chatbots in the context of scientific writing in the humanities. This is because, despite some principles common to all sciences, there are major differences between research in medicine and similar fields and research in the humanities. It is therefore reasonable to expect that AI chatbots will not produce exactly the same results in all fields of science. There is currently no quantifiable evidence to tell us whether the content generated by AI chatbots in the humanities is 75% or 25% incorrect. We do not know what types of biases specific to humanities, if any, AI chatbots generate. Therefore, more tests and, we believe, more types of tests on the performance of AI chatbots in scientific writing are needed, including those specific to the humanities.

With this in mind, the aim of this article was to design and conduct a test of AI chatbots' abilities in scientific writing in the humanities. First, we were interested in their ability to generate correct answers to complex scientific questions. Second, we tested their capacity to generate original scientific contributions in humanities research. Since AI chatbots are developing at a staggering pace, our results apply to the state of affairs in the third quarter of 2023 (23Q3).

To achieve this, we used an interdisciplinary case study combining archaeology, historiography, linguistics and genetic history. We created two one-shot prompts to ask complex scientific questions, and we fed them to each of the six AI chatbots tested: ChatGPT-3.5, ChatGPT-4, Bard, Bing Chatbot, Aria, and Claude 2. We also tested two ChatGPT-4 plugins: Bing and ScholarAI. The generated contents were evaluated against each other and against the human-generated content.

The main reason for the comparison between the AI chatbots was to create a baseline for the staggering speed at which they evolve. Comparison with human-generated responses served as a starting point for the discussion on how upending, transformative, and disruptive generative AI chatbots and other AI tools will be in the humanities in the future.



## 2. Materials and Methods

### 2.1. AI Chatbots

This section describes the AI chatbots that were tested. As they are all proprietary commercial products, there is often not much detail available, let alone in the form of peer-reviewed articles. Our descriptions are therefore based on various sources such as blogs, social media and help pages. We do not go into the specifics of the underlying LLMs, as this is a specialised topic on which there is extensive literature, e.g., [73,74].

First, the criteria for selecting the six AI chatbots should be elucidated. As mentioned earlier, most of the previous studies have only analysed ChatGPT-3.5. Its inclusion, as well as that of its successor ChatGPT-4, was therefore a given. The Bing chatbot was included because it was arguably the most advanced freely available AI chatbot at the time of the test. Bard was included because it is seen by many as the only challenger to ChatGPT's hegemony. We also wanted to include two chatbots that use an application programming interface (hereafter API) to access LLM. APIs are the only available means for "smaller" developers, i.e., anyone other than OpenAI/Microsoft or Google, to access state-of-the-art LLMs. We chose Aria and Claude 2, which use APIs from OpenAI and Google, respectively. If Aria and Claude 2 performed on par with ChatGPT and Bard, it would signal that generative AI technology is indeed being developed openly "for all humanity", and vice versa. The two plugins, ChatGPT with Bing and ScholarAI, were chosen from the rapidly growing selection as the two most relevant to the task of scientific writing. Baidu's ERNIE bot (<https://yiyao.baidu.com>; accessed on 1 October 2023), on the other hand, was not considered because at the time, it was only available with a Chinese interface and required a Baidu login and the Baidu app (also only available in Chinese) to use.

*ChatGPT-3.5*, sometimes called ChatGPT or GPT-3.5, is an AI chatbot offered as a free service by OpenAI (<https://chat.openai.com>; accessed on 1 October 2023). It was fine-tuned from a model in the GPT-3.5 series, more specifically gpt-3.5-turbo, in 2022. This autoregressive language model has the same number of parameters as the largest model from the 2020 GPT-3 series, namely 175 billion. The model was trained using the same methods as InstructGPT-3, but with slight differences in the data collection setup and by using supervised fine-tuning. To predict the next token in the text, it was pre-trained on approximately half a trillion words and improved by task-specific fine-tuning datasets with thousands or tens of thousands of examples that primarily used reinforcement learning from human feedback [75]. ChatGPT-3.5 achieved strong performance on many NLP datasets, including translation and question answering, as well as on several tasks requiring on-the-fly reasoning or domain adaptations. Its breakthrough results paved the way for a new generation of LLMs by demonstrating that scaling language models exponentially increases performance. GPT-3.5 is also available as an API [9]. It may come as a surprise that the core team that developed ChatGPT initially consisted of only approximately 100+ experts, although crowdworkers were also involved as a so-called Short term alignment team [76].

*ChatGPT-4*, also known as ChatGPT Plus or GPT-4, is based on a large-scale multimodal model that can accept text input as well as, since September 2023, image and voice inputs to produce text output. The number of parameters is estimated at 1.8 trillion [11]. The AI chatbot is available as a paid service ChatGPT Plus (<https://chat.openai.com>; accessed on 1 October 2023). Another improvement compared to ChatGPT-3.5 was the development of an infrastructure and optimisation methods that behave predictably across a wide range of scales. Reinforcement learning from human feedback is also an ongoing process. ChatGPT-4 demonstrates human-level performance on various standardised professional and academic benchmarks, including passing a simulated bar exam with a score near the top 10% of test takers [13].

The ChatGPT-4 service has opened up to third-party plugins [77], and we tested two relevant plugins that existed in July 2023. The plugins are available within the ChatGPT Plus service, but the whole service was in beta testing.

ChatGPT-4 with the Bing plugin (henceforth *ChatGPT w/Bing*) has no available information on how it processes its prompts. Furthermore, the plugin was only briefly available as an early Beta in late June; it was temporarily unavailable in July and was replaced by a different service in late September. The results must therefore be considered only as an indicator of the final capabilities.

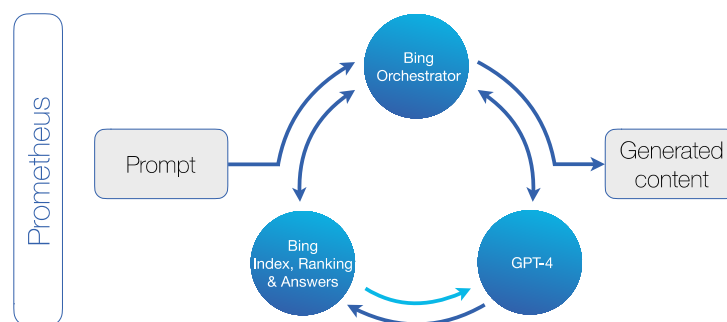
ChatGPT-4 with the ScholarAI plugin (henceforth *ScholarAI*) was also under active development during our test. No documentation was available, but the plugin provided metadata about how it processed the prompt (Appendix A, available at <https://doi.org/10.5281/zenodo.8345088>, accessed on 1 October 2023; henceforth Appendix A: L. 588, footnote 24; Figure 2), and ScholarAI provided additional information (personal communication with Lakshya Bakshi, CTO and Co-founder of ScholarAI). First, ScholarAI extracted keywords from the prompt, which were then recombined into a query. Based on this query, it returned the top results from the ScholarAI database of “40 M+ peer-reviewed papers”. The user then either confirmed the selection or requested further search. When the user was satisfied with the selection of articles, ScholarAI fed the ChatGPT-4 content. ScholarAI ensures that the LLM receives the source data and tries to get ChatGPT to discuss only the content provided to it, but this is a work in progress.



**Figure 2.** Screenshot of the ScholarAI plugin for ChatGPT Plus demonstrating its inner workings.

The *Bing Chatbot* is available either as a Bing Chat or a Bing Compose service and can be used free of charge in the Microsoft Edge browser and in the Bing app (<https://www.bing.com/new>; 1 October 2023). The Bing Chatbot is based on proprietary technology called Prometheus, an AI tool that combines Bing’s search engine with ChatGPT-4. When prompted by a user, it iteratively generates a series of internal queries through a component called the Bing Orchestrator. By selecting the relevant internal queries and leveraging the corresponding Bing search results, the model receives up-to-date information so that it can answer topical questions and reduce inaccuracies. For each search, it ingests approximately 128,000 words from Bing results before generating a response for the user. In the final step, Prometheus adds relevant Bing search responses and is also able to integrate citations into the generated content. This is how Prometheus grounds ChatGPT-4. However, for prompts that the system considers to be simple, it generates responses using Microsoft’s proprietary LLM called Turing-NLG, which consumes less computing power (<https://www.linkedin.com/pulse/building-new-bing-jordi-ribas>; 1 October 2023; Figure 3). The most obvious difference between Chat and Compose is that the former is limited to approximately 200 words and the latter offers more setting options for the generated content. However, the two generated significantly different contents in our test (Appendix A). Due to the limitation to 200 words, Bing Chat was not included in our results.

*Aria* is a browser AI available for free in the Opera browser (<https://www.opera.com>; 1 October 2023). It was “built on Opera’s AI engine, based on several Large Language Models (e.g., Generative Pre-trained Transformer built by OpenAI)”. It was primarily designed to support and enhance the browsing experience by compiling information from various online sources. In addition, it can be used as an AI chatbot to assist in generating texts on any topic in any style for which it can ingest web content. One of Aria’s most important features is its ability to convert certain keywords in its responses into web links [21,78].



**Figure 3.** Prometheus, a high-level overview of the workflow (adapted after <https://www.linkedin.com/pulse/building-new-bing-jordi-ribas>; 1 October 2023).

*Bard* is an AI chatbot available as a free service from Google to a (rapidly growing) limited number of users [18]. It is based on a lightweight and optimised version of Language Models for Dialogue Applications (hereafter LaMDA). LaMDA is a family of GPTs [6] designed specifically for dialogues. First announced during the 2021 Google I/O keynote, it has up to 137 billion parameters and has been pre-trained with 1.56 trillion words from public dialogue data and web text. In addition, LaMDA was fine-tuned for safety and factual grounding. The former was achieved by filtering candidate responses using a LaMDA classifier and data annotated by crowdworkers, the latter by enabling the model to consult external sources, such as an information retrieval system, a language translator, and a calculator. Therefore, Bard interacts with an external information retrieval system to improve the accuracy of facts provided to the user [79–81].

*Claude 2* is an AI chatbot released by Anthropic, a public benefit corporation, and is currently freely available in the US and UK via a website ([claude.ai](https://claude.ai); 1 October 2023) and as an API. The underlying LLM was developed on Google’s hardware, but few details are available about the model’s pre-training. Claude 2 stands out for its fine-tuning. In addition to reinforcement learning with human feedback, so-called constitutional AI was used. This method was developed by Anthropic researchers and requires the model to respond to a large number of questions. It is then instructed to make its answers less harmful by adhering to the “constitution”. Finally, the model is adjusted to generate such less-harmful content in response to the initial prompt. Claude’s constitution, a document created by the developers, is based on the United Nations Declaration of Human Rights and other principles that reflect non-Western perspectives. Another unique feature of the Claude 2 AI chatbot is the size of its context window of 100,000 tokens, or approximately 75,000 words, which allows users to paste or upload large documents or entire books and prompt questions based on their content [22,82,83].

The versions of the AI chatbots and plugins used in the tests are sparsely reported. Apart from the labels ChatGPT-3.5, ChatGPT-4, and “beta” for plugins, there is no information about the versioning available to the end user. This is vital information, since AI chatbots are being developed at a rapid pace. To mitigate the lack of information on versioning, we provide the exact dates on which the tests were conducted for each AI chatbot: ChatGPT-3.5 on 6 June 2023, ChatGPT-4 on 7 June 2023, Bard on 8 June, Bing Chat on 8 June 2023, Aria on 22 June 2023, ChatGPT-4 w/Bing on 26 June 2023, ChatGPT-4 w/ScholarAI on 26 June 2023, Bing Compose on 2 July 2023, and Claude 2 on 26 July 2023.

Finally, it must be emphasized that the tested AI chatbots were designed as general-purpose AI chatbots capable of human-like conversation and not to “do science”. Such downstream applications can be expected in the near future.

## 2.2. Domain-Specific Scientific Background

The case study chosen for testing AI chatbots was the migration of the South Slavs with a follow-up prompt on the Alpine Slavs, a subgroup of the South Slavs. The authors’



thematic explanation of the case study can be found in the appendix (Appendix A: L. 383–439 and 799–855).

The migration of the Slavs, including the South Slavs, has been a research topic for over a century. Notwithstanding this, the rapid spread of the Slavic language in the second half of the first millennium CE remains a controversial topic [84–86]. It is part of the grand narrative of the “dawn of European civilisation”. The Europe we live in today emerged in the centuries after the decline of the Roman Empire and was importantly shaped by the ancient Slavs, among others.

The current scientific debate on this issue revolves around the gene pool landscape on the one hand and the landscape of ethnicity on the other. Until the 1950s, migration was assumed to be the main process of change, e.g., [87], and peoples and tribes were understood as caroming around the continent, similar to culture-bearing billiard balls [88]. It was during this period that the term migration period was coined. Since the 1960s, the understanding of ethnic identity has shifted to the concept of dispersed identities, which states that people fluidly adopt different identities as changing social circumstances dictate, e.g., [89]. Today, most assume that hardly any physical migration has taken place, but rather that ideas and knowledge have been passed on, e.g., [90]. However, recent research in the field of DNA, ancient DNA, and deep data analysis supported by machine learning is providing increasingly compelling evidence that, at least in the case of the South Slavs, physical migrations took place [84].

In other words, the interpretation that began with “everything was migration” has shifted to “(almost) no migration” and is now shifting back to “(considerable) migration”.

Why is this important? From the perspective of an interested researcher, this is a vibrant area of research and an exciting time to be working in. However, from the perspective of a casual observer (who is likely to prompt an AI chatbot on the subject), it may seem that not much has changed since the 1900s: it has “always been known” that migrations occurred, and although there was doubt for a while, modern science proves it. Therefore, the understanding of this long-standing issue by relevant researchers may differ significantly from that of a casual observer. This is relatively common in the humanities but not, for example, in mathematics, physics, law, or medicine. Once the Pythagorean theorem and the effects of antibiotics were proven, this became common knowledge shared by experts and non-experts alike.

This distinction is important for understanding AI chatbots because the vast majority of content available on the public web—and thus the content available for training LLMs—is created by interested laypeople, not scientists. For example, the content on Wikipedia dealing with the Alpine Slavs largely reiterates the narrative established in the 1950s [91].

Therefore, our test was also suitable for determining whether the content generated by AI chatbots reflects “the public web” (i.e., the uncritical adoption of the highest-ranking online content or, where applicable, the content amplified by social media algorithms), the scientific explanation, or both.

### 2.3. Text Prompts

Our experiment consisted of asking generative AI models two specific scientific questions designed for human-to-human interaction. We did not use prompt engineering based on a prompt pattern catalog [92] or similar. However, we tweaked our prompts using ChatGPT-3.5 through trial and error until the generated content was precise enough to not require additional prompts.

According to the help pages of ChatGPT, while the content of the prompts can be used for reinforcement learning from human feedback, it is not carried over to new chats. The content used in our test and reported in the appendices was therefore generated in a new chat, and it represents the chatbot’s initial and only response to the prompt.

The fact that the prompts were tweaked in ChatGPT-3.5 gives it, in theory, an advantage over other AI chatbots in our test. However, the generated content proves that all AI

chatbots comprehended the prompt, i.e., “understood the question”, without any problems. The prompts were thus robust enough that the tested chatbots can be compared.

The selected case study spans several academic fields, one of which can be considered a natural science (DNA analysis), one a humanities science (historiography), and two a humanities science with links to natural science (archaeology, linguistics). In the USA, archaeology is also considered a social science in certain contexts.

The two text prompts were as follows:

- Q1: What is scientific explanation for migration of South Slavs in Early Middle Ages? Write 500 words using formal language and provide references where possible.
- Q2: What is scientific explanation for the settlement of Alpine Slavs in Early Middle Ages? Write 500 words using formal language and provide references where possible.

Q1. The first prompt is a complex scientific question on the subject of early medieval studies. To discuss it requires knowledge of archaeology, historiography, linguistics, and DNA studies. However, the topic is relatively broad. Spatially, it covers an entire European region: the Balkans. Its scientific background, the migration of the Slavs, is relevant to over 200 million modern Europeans. In short, although it is not one of the foremost topics in the humanities or even for Early Medieval scholars, there are numerous researchers working on this topic, and dozens of relevant scientific papers are published every year.

Q2. At first glance, the second prompt is almost exactly the same as the first, except for the target group: the Alpine Slavs. However, the added complexity of this prompt comes from the fact that it addresses a very narrow and specific topic. In fact, the only scholarly content on this topic is either more than half a century old, e.g., [91], and not available online, or it comes from a 2022 paper [84], which is too recent to be included in the datasets used for training ChatGPT.

However, the key term “Alpine Slavs” is very specific. In response to the search term “settlement of the Alpine Slavs”, the search engines Bing, Google, and DuckDuckGo as well as Google Scholar return the mentioned article as the top hit after Wikipedia or the Encyclopaedia Britannica.

We therefore expected ChatGPT to respond well to Q1 but have more problems with Q2. On the other hand, AI chatbots with access to the current online content (Bing chatbot, GPT w/Bing) were expected to generate high-quality content for Q2 by sourcing it directly from the relevant article.

Our scientific questions are therefore so-called one-shot prompts, where the user provides the AI chatbot with a single example of the desired task and then asks it to perform a similar task. It is well known that GPTs are “few-shot learners” (BrownEt2020), i.e., they are much better at generating content when given more examples of the expected content. When using one-shot prompts, multiple refinement prompts are expected to improve results, e.g., [93].

However, few-shot prompts were not suitable for our testing purpose because they did not mimic a scientific question and a series of prompts would introduce subjectivity and hence reduce the value of a direct comparison between different AI chatbots. Therefore, our prompts were optimised for comparison rather than for generating the best possible content.

#### 2.4. Tagging and Analysis

There are several existing studies similar to ours, but they refer to other fields of science. Regardless, a brief overview of the methods used is in order. Altmäe and colleagues [51], for example, provided prompts and content generated by ChatGPT and then discussed the quality of the content. Petiška [56] focused only on references and analysed factors such as the number of citations, the date of publication, and the journal in which the paper was published. Májovský and colleagues [57] posed questions and prompts to the model and refined them iteratively to produce a complete article, which was then reviewed by relevant experts for accuracy and coherence. Buholayka and colleagues [59] tasked ChatGPT with

writing a case report based on a draft report and evaluated its performance by comparing it to a case report written by human experts.

Our approach is similar to that of Májovský and colleagues [57], but there are three significant differences. First, we did not use iterative refinement to ensure comparability between different AI chatbots. Second, we did not generate a complete paper. Third, our review was both qualitative and quantitative, not just qualitative. This was achieved by tagging the content. The aim was to provide what is, to our knowledge, the first quantitative comparison of different AI chatbots on the subject of scientific writing.

The content generated by each of the tested AI chatbots was tagged for quantitative accuracy and qualitative precision.

*Quantitative accuracy* describes how accurate the content is in the opinion of the human experts (the authors). It was divided into five grades:

- Correct: Factually correct and on par with the content created by human experts.
- Inadequate: Factually correct but falls short of the content created by human experts.
- Unverifiable: The statement cannot be verified or there is no expert consensus.
- w/Errors: Mostly factually correct, but with important errors that change the meaning.
- Incorrect: Factually incorrect.

Quantitative accuracy is the measurement most commonly applied to AI-generated content, providing a quantifiable index of how trustworthy the tested AI chatbot is for the task at hand. It can be understood as similar to the grading of a student's work. The questions in this case study are on the level of a senior undergraduate student attending a class on Early Medieval archaeology.

*Qualitative precision* describes how "good" the content is in the opinion of a human expert. In other words, how it compares to a human-generated response in the context of scientific writing. It was divided into four grades:

- Original scientific contribution.
- Derivative scientific contribution.
- Generic content not directly related to the question.
- Incorrect: Factually incorrect, containing errors that change the meaning, or disputed (the last three classes of above quantitative tagging combined).

Qualitative precision, as we have defined it, is specific to testing AI-generated content for the purpose of scientific writing and, to our knowledge, has not yet been used. The reason for the insufficient development of such indices is mainly that the current generation of AI chatbots is not expected to generate original scientific content. However, in the near future, AGI will be expected to produce it. The qualitative precision index was therefore developed with an eye to the future as a measure of how close the tested AI chatbots are to AGI.

From an academic perspective, qualitative precision can be understood in a similar way to the peer review of a scientific paper. As with any peer review, e.g., [94], it is a combination of objective and subjective evaluation.

It should be mentioned that in the humanities, an article usually consists of both derivative and original content. Typically, the introduction and method are predominantly derivative, while the results, discussion, and conclusion are predominantly original. The ratio of derivative to original content varies widely depending on the discipline, topic, type of article, etc. Thus, the expected "perfect score" is not 100%, but in the order of 50+% of the original scientific contribution. To establish the baseline for our case study, we tagged the responses generated by human experts (see Section 3).

Both quantitative accuracy and qualitative precision tagging were performed by the two co-authors independently of each other. Following the standard process for grading students or reviewing scholarly articles, we then compared the results and consulted to reach a unanimous decision. Typically, there were only minor differences, for example part of the sentence was graded "w/Errors" by one reviewer and "Incorrect" by the other. The final results are shown in the appendices (qualitative accuracy: Appendix A; quantitative

accuracy: Appendix B, available at <https://doi.org/10.5281/zenodo.8345088>, accessed on 1 October 2023; henceforth Appendix B). Both co-authors are experienced in both tasks in a professional capacity, and both are experts on the topic, e.g., [84,95].

The tagged content was quantified, and the results are discussed in the next section. Given the small amount of data generated by tagging, the observational method of analysis was sufficient. To sort the different AI chatbots according to the results, we calculated the accuracy score using the following formula:  $Correct\% - (2 \times Incorrect\%)$ . The higher the score, the better. Students would need a positive score for a passing grade.

As the amount of tagging data increases in future projects, more sophisticated statistical methods will be used.

### 2.5. Limitations of the Method

There are several limitations to our method. First, the case study is limited to a single (interdisciplinary) field in humanities and cannot address the differences between, for example, philosophy and geography. Second, only a limited number of human experts were involved; therefore, the similar academic backgrounds and knowledge pools of the two reviewers, the authors, may have influenced the peer review procedure for evaluating content. Third, prompt engineering based on a prompt pattern catalog was not used.

We estimate, however, that the subjective differences had a limited effect on the overall results because the selected topics are too high-level to be seriously contested in the scientific community.

We intend to address the limitations in the future, as is discussed in Section 5.

## 3. Results

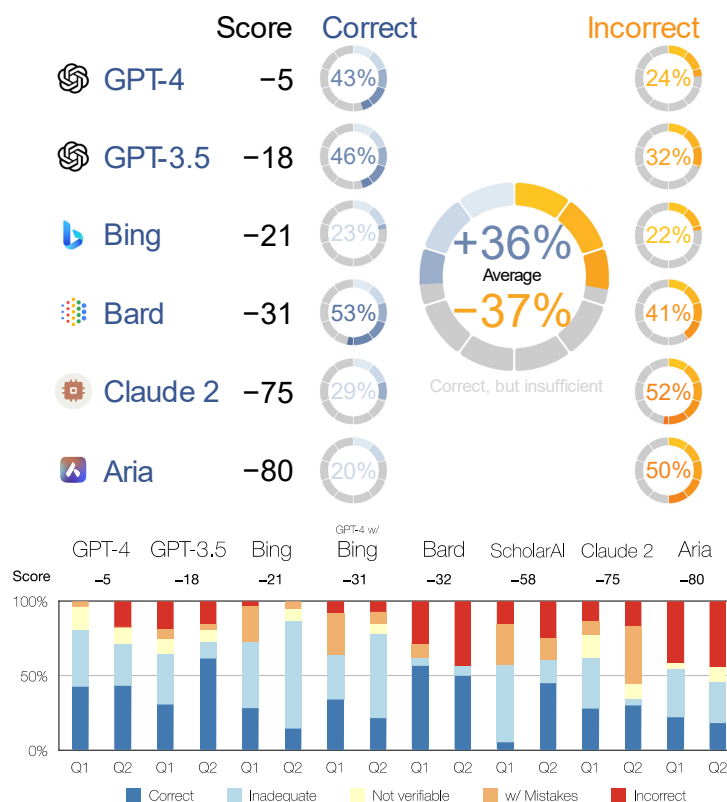
### 3.1. Quantitative Accuracy: AI Chatbots Taking an Undergraduate Test

Quantitative accuracy tagging was intended to objectively determine how correct the answers generated by the AI chatbots were.

The highest accuracy score was achieved by ChatGPT-4, which also generated the highest percentage of correct content. On average, approximately half of the content provided was correct, and approximately 1/5 was incorrect, with errors, or unverifiable. However, as expected (see Section 2.3), all of the incorrect content belonged to Q2, for which the LLM could not source the relevant content from the 2022 article. Considering the complexity of the questions, the results were impressive but below the level expected of a senior undergraduate student (Figure 4; Appendix A).

The performance of ChatGPT-3.5 was in line with expectations. On average, it produced approximately the same amount of correct content as ChatGPT-4, but almost 10% more incorrect content. The latter is the expected result of continuous reinforcement learning from human feedback [75], but the former is somewhat underwhelming for ChatGPT-4 given the 10-fold increase in the number of parameters in the model.

The Bing chatbot generated only 28% and 13% incorrect or unverifiable content for Q1 and Q2, respectively. However, the percentage of correct content was also one of the lowest at 22% and 18%. The reason for this is that the content was almost exclusively derived from the corresponding Wikipedia pages. The Wikipedia source pages were correctly listed under references, but were systematically misquoted in the text. All references except Wikipedia were secondary, i.e., AI chatbot copied a selection of references listed on the Wikipedia pages without actually consulting them (many of them are copyrighted material not available on the public web). The article directly referring to Q2 [84] was neither used as a source for the content nor referenced, although, as mentioned above, it is returned by the Bing search engine as the highest-ranking non-Wikipedia result. The content generated by the Bing chatbot thus corresponded well to the chosen settings for a professional (i.e., closely following sources) blog (i.e., formatted in a free flowing text), but fell short of simulating an essay by, for example, a senior undergraduate student.



**Figure 4.** Quantitative test results: generalized accuracy score above and detailed quantitative accuracy below. Incorrect (orange) above equates to incorrect, w/ Mistakes, and not verifiable (red, orange, and yellow) below.

ChatGPT-4 w/Bing performed significantly worse than Bing Chatbot and ChatGPT-4. In addition, it was unable to source the content for Q2 from the relevant scientific article. At the time of the test, therefore, this plugin was not yet working as intended.

Bard was far less stable and reliable than the ChatGPTs. It generated a comparable percentage of correct content, but almost twice as much incorrect content.

ScholarAI, Claude 2, and Aria performed at notably lower levels.

ScholarAI was obviously not yet ready for production use. It systematically generated incorrect references and was unable to source relevant content. It should be mentioned, though, that online research databases often perform better in scientific fields other than the humanities, such as medicine, chemistry, and natural sciences in general. There are several reasons for this, including the fact that the most prestigious form of dissemination in the humanities are books, which are not indexed in (so many) scientific databases and are often not available in open access format.

The poor result was surprising for Aria, because it uses an LLM from OpenAI. We can only assume that it uses one of the GPT-3 APIs rather than GPT-3.5. Additionally, unlike Bing Chatbot and ChatGPT w/Bing, it did not consult online resources to generate the content; rather, it linked the generated content to online resources.

The Claude 2 results are comparable to those of Aria. We can only assume that it is using a scaled-down LaMDA tuned not to outperform Bard. The recent announcement of a “strategic collaboration” between Anthropic and Amazon, i.e., Anthropic’s move away from Google, seems to confirm this assumption.

In summary, we can say that GPT-4, GPT-3.5, and Bing Chatbot were close to the level of, for example, an undergraduate student’s initial research steps for a term paper, which start with looking up general sources such as Wikipedia. ChatGPT w/Bing and Bard were good substitutes for searching the internet with general search engines. They were at a



level that could be described as a layman looking at a new topic. ScholarAI, Aria, and Claude 2 were not yet up to the task of answering complex humanities questions.

### 3.2. Qualitative Precision: AI Chatbots' Take on Original Scientific Contribution

The main focus of our article was on whether the tested AI chatbots were able to generate original scientific contributions. The short and expected answer was no. A more detailed answer can be found below (Figure 5; Appendix B).

As mentioned earlier, human-generated scientific articles in the humanities are typically a combination of derivative and original scientific contributions. In our case study, the human-generated content included  $\frac{1}{2}$  of the original scientific contribution for Q1 and  $\frac{3}{4}$  for Q2.

The AI chatbots did not reach this level by far. The only discernible original scientific contribution was at 11% generated by ChatGPT-4. It aptly inferred in Q1 that the migration of the South Slavs was not a singular event (Appendix B: L. 91–93), and its introductory paragraph in Q2 was extremely apt, profound, and on the cutting edge of science (Appendix B: L. 478–483). Similarly, ChatGPT-3.5 summarised the settlement of the Alpine Slavs very astutely, if repetitively (Appendix B: L. 458–461 and 646–467).

Claude 2 correctly pointed out that the fact that Christian missionaries had to preach in Slavic languages suggests the demographic dominance of the Slavs. This is an established historical fact, but not commonly referred to in the context of migration, and was therefore tagged as an original scientific contribution.

ScholarAI generated what at first sight appeared to be very exciting original scientific content. It has established a direct link between the process of settlement of the Alpine Slavs and their cultural practices and beliefs (Appendix B: L. 578–580 and 585–588). The discussion of beliefs in the context of the migrations is far from the norm and, to our knowledge, has never been brought forward for the migrations of the Alpine Slavs. However, ScholarAI's argumentation was flawed because it was based on irrelevant knowledge pertaining to the Baltic Slavs [96] dwelling approximately 1000 km northeast of the Alpine Slavs. Interestingly, the same hypothesis could have been argued with another freely available scientific text [97], but this is a book rather than an article and is therefore not in the ScholarAI database.

Other AI chatbots have not generated original scientific contributions.

In conclusion, ChatGPT-4 was once again the best among the AI chatbots tested, but not on the same scale as the human-generated content (Figure 5).

### 3.3. Reasoning Errors, Hallucinations, Biases


The most commonly cited shortcomings of AI chatbots are reasoning errors, hallucinations, and biases, e.g., [9,13]. The terms themselves are not the best choice, because they inappropriately anthropomorphize AI chatbots. However, they are widely used, and we have used them for clarity.

In the quantitative analysis above, these shortcomings were interchangeably tagged as incorrect, with errors, or unverifiable (Appendix A). Here we address them qualitatively, on a case-by-case basis.

*Reasoning errors*, also termed lack of on-the-fly reasoning or lack of critical thinking, are the kind of incorrect content in causal statements where cause and effect do not match. Critical thinking is one of the most important qualities for humanities scholars and knowledge workers in general. However, AI chatbots based on LLMs are not designed for this task.

The most obvious example of a reasoning error in our case study was the content generated by ChatGPT-4 and Bard, which causally linked the migration of Slavs into the Alps to the period of climate cooling (Appendix A: L. 512–515 and 703–704). Similarly, ChatGPT-3.5 linked the settlement of the Alpine areas to “fertile lands” (Appendix A: L. 456–459). For most Europeans and most people with formal education worldwide, the Alps are synonymous with mountains and hence with a cold climate and harsh agricultural

conditions. Most people would therefore reason that a cooling climate and the search for fertile land would not expedite the migration into the Alps but rather impede it.



**Figure 5.** Qualitative test results: original scientific contribution above and detailed precision below.

Another example of a reasoning error was that almost all tested AI chatbots listed the decline of the (Western) Roman Empire as one of the attractors for the migration of South Slavs to the Balkans (Western Roman Empire: Appendix A, L. 116–117, 135–136, 144–145, 278–281, 322–323, 454–455, 683–684, 736–738; Roman Empire: Appendix A, L. 124–125, 506–507, 516–517, 548–550, 584–585, 593–595). However, we learn in high school history classes that the (Western) Roman Empire preceded the migration of the South Slavs by at least a century. In fact, the Byzantine Empire was the declining superpower that created the power vacuum for the immigration of the South Slavs to the Balkans.

The fact that both LaMDA (Bard) and GPT-4 (ChatGPT-4) generated almost identical incorrect content suggests that such behaviour is inherent in the current generation of LLMs.

The underlying issue in the lack of critical thinking was that none of the tested AI chatbots made any attempt to critically compare different sources. For example, the most important component of a human-generated response to Q1 was: “Currently, there are three main hypotheses. . .” (Appendix A: L. 394), which was continued by comparing several different sources. No such attempt was detected in the content generated by the AI chatbots. Anecdotally, the majority of randomly selected human users were able to distinguish the critical thinking of the human expert from the content generated by ChatGPT, based solely on the 24-character snippet “There are 3 hypotheses. . .” without further context (Figure 6).

The critical comparison of different sources is typical and vital not only in any kind of scientific reasoning, but also in everyday life. The one-sided approach of the tested AI chatbots amplifies “the loudest voice” (the highest-ranking search engine result), which is not only bad science but also a grave danger for balanced news reporting, democracy, minority rights, etc.

*Hallucinations* or confabulations of AI chatbots are confident responses by an AI that are not justified by its training data. This is not typical of AI systems in general, but is relatively common in LLMs, as the pre-training is unsupervised [98].



**Figure 6.** A Twitter (now X) poll asking human users to differentiate between ChatGPT and human-generated content with almost no context. Most respondents answered correctly.

The most obvious hallucinations in our case study were invented references (ChatGPT-4, Appendix C available at <https://doi.org/10.5281/zenodo.8345088>, accessed on 1 October 2023; henceforth Appendix C: L. 36; ChatGPT-3.5, Appendix A: L. 495–498; ScholarAI, Appendix A: L. 189–197). Similarly, attempts to inline citations by Bing (Appendix A: L. 226–257) and ChatGPT-4 w/Bing (Appendix A: L. 575–581) were largely confabulations. It would appear that the inability to provide correct inline citations is a known problem, as ChatGPT-4 w/Bing generated a warning to this effect (Appendix A: L. 154–155).

Another very clear example of hallucination was the Late Antique Little Ice Age phenomenon. ChatGPT-4 dated it “between 300 and 700 AD” (Appendix A: L. 512). The correct dates are 536 to approximately 660 AD, which is clearly intelligible to any human by consulting the title of the reference which ChatGPT-4 correctly provides: “. . . Late Antique Little Ice Age from 536 to around 660 AD” (Appendix A: L. 532–535).


The underlying issue and a key challenge for current technology is that “AI chatbots do not know what they do not know” and may very confidently invent facts [46] rather than formulate the sentence “I don’t know”.

*Biases* are often exhibited by AI chatbots. They are based on training data, but according to recent research, they can be amplified beyond existing perceptions in society. Biases generated by AI chatbots can therefore be informative about the underlying data, but they can also be misleading if the AI-generated content is used uncritically. The most researched biases to date are those related to gender, race, ethnicity, and disability status, e.g., [29,49,98–100].

In our test, we detected three different types of biases: language bias, neo-colonial bias, and citation bias.

First, there was a language bias. Although there is far more relevant scholarly content written in Balkan languages than in English, 92% of the references generated by the AI chatbots in our test referred to English and none to Balkan-language publications (Figure 7). This can only be partially explained by the fact that the prompts were in English. Namely, three (8%) German references prove that English was not the only criterion for selection. When question Q2 was asked with a prompt in Slovenian, two references were again in English and the third in Slovenian was a hallucination (Appendix C: L. 36).

The detected language bias is most likely due in large part to the language bias of online search engine ranking algorithms that favour English publications [101]. This bias seems to be a wasted opportunity, because all tested AI chatbots “understand” many languages, e.g., [13].



**Figure 7.** Language of references generated by AI chatbots to an English language prompt.

Second, neo-colonial bias. A total of 88% of the references are by authors from the global West. Only a minority (12%) belong to English translations of texts originally written in Slavic languages [96,102,103], although there are several other (more) relevant translations, e.g., [103,104]. This reflects a scholarly hierarchy created by colonialism (until the 1910s, the Balkans were largely divided between the Austro-Hungarian and Ottoman Empires), sometimes referred to as a neo-colonial pattern in the global network of science in which the intellectual dominance of the global West is growing [105,106]. To our knowledge, the neo-colonial bias for the study of medieval Slavs has not yet been explicitly analysed or even discovered, as it has never been revealed as clearly as through the use of AI chatbots in this case study.

Third, citation bias. A total of 75% of the references are from before 2005, and the oldest was originally published in 1895 [102]. This shows a very clear bias against new and up-to-date publications. For example, by far the most referenced publication in our case study is Curta [85]. While this is still a seminal work, it is outdated and has often been criticised, e.g., [107], and the critiques have been responded to [86]. Therefore, in a modern scientific text created by a human expert, the reference to Curta is always followed by either its critique or an up-to-date response to that critique. In AI-generated content, however, Curta is always referenced as the primary source.

This bias is in line with the growing trend to cite old documents, caused at least in part by the “first page results syndrome” combined with the fact that all search engines favour the most cited documents [108,109]. The ScholarAI plugin, for example, transparently discloses that it ranks references solely by the number of citations (Appendix A: L. 622, footnote 25).

These findings are consistent with the recent study looking at references generated by ChatGPT-3.5. It revealed that ChatGPT-3.5 tends to reference highly cited publications, shows a preference for older publications, and refers predominantly to reputable journals. The study concluded that ChatGPT-3.5 appears to rely exclusively on Google Scholar as a source of scholarly references [56].

All three of the listed biases are abundant in the LLMs’ training material, which is, of course, mostly the public web. Uncritical online search by a human using one of the major search engines, Google Scholar, or Wikipedia pages would yield comparable results with language bias, citation bias, and neo-colonial bias. However, proper research by a human expert would avoid citation bias and at least reduce language and neo-colonial bias.

### 3.4. Race for Parameters

As mentioned earlier, the scale of an LLM depends on the size of the ingested datasets, the amount of training compute, and the number of parameters supported. However, since


developers never fully disclose these data, we follow other studies [9–11,110] and compare the number of parameters as a proxy to gauge the scale of the LLM (Figure 8).

It was not the intention of this article to gain insights into the efficiency of LLMs, but as a side note, the current trend of upsizing the LLMs, sometimes referred to as an AI Arms Race, can be addressed. Currently, numerous well-funded start-ups, including Anthropic, AI21, Cohere, and Character.AI, are putting enormous resources into developing ever larger algorithms, and the number of LLM parameters exceeds the exponential growth trend line. However, it is expected that the returns on scaling up the LLM size will diminish [13].

In our data, we observed the impact of the exponential growth of the LLM on the content generated. ChatGPT-4 is approximately 10 times larger than ChatGPT-3.5. LaMDA has a similar number of parameters as ChatGPT-3.5, 137 and 175 billion, respectively, but the AI chatbot tested, Bard, only uses the “lightweight” version of LaMDA. Assuming that the “lightweight” version uses only half the parameters, the size ratio between LaMDA (Bard), ChatGPT-3.5, and ChatGPT-4 is 1:2:20.

The up to 10% improvement in the generated content of ChatGPT-4 compared to ChatGPT-3.5 is not negligible, but for many use cases, it might not be noticeable. Given the tenfold increase in the number of parameters, this is not particularly impressive. Especially since the 10% improvement is due not only to the new LLM but also to continuous reinforcement learning from human feedback. Nevertheless, the improvement of ChatGPT-4 over Bard in reducing incorrect content is significant. One can surmise that the twenty-fold increase in parameters from Bard to ChatGPT-4 is significant enough to be noticeable in daily use, while the tenfold increase from ChatGPT-3.5 to ChatGPT-4 is mostly only observable in a controlled test environment.

Given that improvements are diminishing and that continuous exponential growth of LLMs is physically impossible with current technology, we can only conclude that the growth of LLMs has reached a plateau. This observation is consistent with the sparse information coming out of the industry: “. . .we’re at the end of the era where it’s going to be these, like, giant, giant models. . . We’ll make them better in other ways”. (Sam Altman, CEO of OpenAI, at an event at MIT on 14 April 2023).



**Figure 8.** The race for parameters. The increase in the number of parameters (red dotted line; adapted after [110] and updated with sources cited in Section 2.1) exceeds the exponential trend (orange line); the in-context learning as detected by our test (blue columns; see Section 3.1) only improves with a linear trend (blue line).



#### 4. Discussion: Data to Knowledge

The results of our analysis show that ChatGPT, while currently the most advanced AI chatbot, is not capable of producing original scientific content. This is a zero-shot learning problem, but it is also much more than that. Namely, LLM-based AI chatbots are not designed to generate content in the same way that human researchers produce new knowledge.

A typical process for producing original scientific contributions, which was also used in our case study, is an intricate process that is perhaps best explained using Ackoff's Data–Information–Knowledge–Wisdom (DIKW) hierarchy. According to this hierarchy, data are the raw input, simple facts without context; information is data that has been given context and is meaningful in some way; knowledge is the understanding and interpretation of that information gained through experience or learning; and wisdom, the highest level of the pyramid, is the ability to apply knowledge judiciously and appropriately in different contexts, understand the broader implications, and make insightful decisions based on accumulated knowledge and experience [111].

In our case study, pertinent to Q2, the data were the documentation of the excavations of 1106 archaeological sites. The information was summaries of the excavation reports and scientific publications on these 1106 sites curated in the structured database Zbiva, available on the public web since 2016 [112]. Knowledge was the scholarly articles discussing and/or analysing the migration of the Alpine Slavs. Wisdom is what happens (rarely) after the publication of scientific articles and after the generation of AI chatbot content and does not concern us here.

Human researchers approached Q2 by first obtaining and refining the data, which was followed by analysing the information using appropriate software, formulating knowledge based on the results, and finally disseminating the knowledge in the form of an original scientific article.

In the real world, archaeological practice (and most humanities research) is messy, and therefore, this process is fuzzy and recursive. It takes the form of a hermeneutic spiral [113], which, from the perspective of computational algorithms, are loops between data, information, and knowledge. These loops involve solving computationally irreducible processes that cannot be handled by LLM-based AI chatbots alone (Figure 1: I).

In other words, generating original scientific content requires not only access to curated data/information, but also the ability to analyse it.

LLMs, on the other hand, are pre-trained on existing knowledge (texts, books) and are only able to recombine it in new permutations. Occasionally, this can lead to modest original scientific content. For instance, AI chatbots could be used to summarize what a historical text verbatim says about a certain subject, but not to interpret it. Therefore, this is a limited and finite avenue to original scientific contributions. Regardless of the future improvement of LLMs, LLM-based AI chatbots will not be able to replicate the workflow of human researchers, as they are only trained on the existing knowledge.

A purpose-built LLM-powered software, on the other hand, could handle such a workflow: searching for data/information, performing relevant analysis, generating textual and graphical information, and summarising it into new knowledge (Figure 1: K–M; G, J). Such LLM-powered software would have several qualities of an AGI and is in fact already feasible with existing technology, for example by combining Prometheus, a relevant cloud-based software connected through a ChatGPT API, and ChatGPT-4.

In a nutshell, LLM-based AI chatbots are not and probably will never be able to generate new knowledge from data in the same way as human researchers (in the humanities), but appropriate LLM-powered software might be.

#### 5. Conclusions: Fluent but Not Factual

Most commentators on generative AI, including the authors of this article, agree that the current generation of AI chatbots represents AI's inflection point. It is very likely that historiography will record ChatGPT-3 as the eureka moment that ushered in a new era

for humanity. What this new era will bring is currently unknown, but it seems that it will, for better or worse and probably both, change the world. To maximize the positive and mitigate the negative as much as possible, making AI safe and fair is necessary, and a significant part of making AI safe is testing.

The aim of this article was to test the current AI chatbots from a human(ities) perspective, specifically their scientific writing abilities. We compared six AI chatbots: ChatGPT-3.5, ChatGPT-4, Bard, Bing Chatbot, Aria, and Claude 2.

In accordance with expectations, ChatGPT-4 was the most capable among the AI chatbots tested. In our quantitative test, we used a method similar to grading undergraduate students. ChatGPT-4 achieved -5 in our custom scoreboard, where 1 or more was required for a passing grade. ChatGPT-3.5 (-18), Bing (-21), and Bard (-31) were not far behind, whereas Claude 2 (-75) and Aria (-80) scored much lower. The ChatGPT-4 plugins were not yet up to the task. ChatGPT-4 w/Bing scored -31 and ScholarAI -58, but both failed to conduct the tasks they were designed for, which were to provide up-to-date information and reliable references, respectively.

In our qualitative test, we used a method similar to peer reviewing a scientific article. ChatGPT-4 was again the best performer, but it did not generate any notable original scientific content.

Additional shortcomings of the AI-generated content that we found in our test include reasoning errors, hallucinations, and biases. Reasoning errors refer to the chatbot's inability to critically evaluate and link cause-and-effect relationships, as evidenced by several historical inaccuracies regarding the migration patterns and historical periods of the Slavs. Hallucinations denote confident but unsubstantiated claims by the AI, such as invented references and inaccurate dates. Our test also reveals significant biases in the content generated by the AI. These biases manifest as language bias, favouring English sources over other relevant languages; neo-colonial bias, displaying a preference for Western authors; and citation bias, skewing towards older and more highly cited publications. These findings highlight that, despite their technological prowess, AI chatbots remain reliant on their training data, echoing and even amplifying existing biases and knowledge gaps. Given that they veer towards past data, including, in our test, a book originally published in 1895, they are likely to be too conservative in their recommendations. Since the listed deficiencies are relatively inconspicuous compared to, for example, gender or racial biases, it is unlikely that they will be remedied in the foreseeable future by resource-intensive reinforcement learning from human feedback. These biases are among the key concerns with the use of AI chatbots in scientific writing, as most of them are less likely to be highlighted in the review processes.

Our results also highlighted the possible future trends in the development of AI chatbots. The large discrepancy between almost passing an undergraduate exam and not producing any notable scientific contribution may seem surprising at first glance. On closer inspection, however, this was to be expected. "Doing science", i.e., making an original scientific contribution, is much more complex than just doing very well in the exams. "Doing science" is based on proactively acquiring and analysing data and information to generate new knowledge, whereas "passing an exam" is based on accumulating existing knowledge and passing it on at a request.

"Passing an exam" will further improve when AI chatbots are given access to curated data. An AI chatbot with access to selected datasets would be a typical downstream task developed around an existing LLM. However, without access to external tools, LLM-based AI chatbots will never be suitable for "doing science". Therefore, in the near future, an evolution of current LLM-based AI chatbots towards LLM-powered software capable of, among other things, "doing science" seems likely. This assertion is in line both with the fact that the growth of LLMs seems to have plateaued and that the industry is turning to other solutions.

The limitations of our analysis resulting from the methodology must also be reiterated. First, the case study is limited to a single field in humanities and cannot address the

differences between different subfields. Second, only a limited number of human experts were involved, which can introduce subjectivity. Third, engineering of the prompts was not standardised.

To counter these limitations in the future, we plan to extend this study to a series of field-specific studies and involve more human experts. Involving additional experts will also enable us to improve on the engineering of the prompts. However, in the current structure of public science, it takes years to accomplish such an organisational feat. Therefore, this article can be understood as an interim measure taken in response to the incredible speed at which AI chatbots are developing.

In conclusion, we agree with the previous commentators that AI chatbots are already widely used for various tasks in scientific writing due to their applicability to a wide range of tasks. However, AI chatbots are not capable of generating a full scientific article that would make a notable scientific contribution to the humanities or, we suspect, to science in general. If left unsupervised, AI chatbots generate content that is fluent but not factual, meaning that the errors are not only many but also easily overlooked, cf. [114]. Therefore, peer review processes need to be rapidly adapted to compensate for this, and the academic community needs to establish clear rules for the use of AI-generated content in scientific writing. The discussion about what is acceptable and what is not must be based on objective data. Articles such as this one are necessary to support those decisions, and we suspect that many more will follow.

For the future beyond the immediate foreseeable, when LLM-powered software and/or AGI will be able to generate original scientific contributions, we agree that questions about explainable AI are likely to come to the fore. Understanding our world, a fundamental aspiration of the humanities, will only be partially achieved through the use of black-box AI. Since the humanities, just like justice, for example, are as much about process as outcome, humanities scholars are unlikely to settle for uninterpretable AI-generated predictions. We will want human-interpretable understanding, which is likely to be the remaining task of human researchers in the humanities in the future, cf. [115].

**Supplementary Materials:** Appendix A, Appendix B, and Appendix C can be downloaded at: <https://doi.org/10.5281/zenodo.8345088> (accessed on 1 October 2023). The files are deposited in the open access repository Zenodo under CC-BY 4.0 licence.

**Author Contributions:** Both authors, E.L. and B.Š., contributed equally to the article. Conceptualization, E.L. and B.Š.; formal analysis, E.L. and B.Š.; funding acquisition, E.L. and B.Š.; investigation, E.L. and B.Š.; methodology, E.L. and B.Š.; writing—original draft, E.L. and B.Š.; writing—review and editing, E.L. and B.Š. All authors have read and agreed to the published version of the manuscript.

**Funding:** This research was part of the AI4Europe project that has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement n° 101070000, and the Slovenian Research and Innovation Agency (ARIS) grant number P6-0064.

**Data Availability Statement:** Datasets generated and analysed during the study are publicly archived on the open access repository Zenodo under CC-BY 4.0 licence: <https://doi.org/10.5281/zenodo.8345088>, accessed on 1 October 2023.

**Acknowledgments:** The authors give thanks to Zoran Čučković for introducing them to ChatGPT in December 2022 and to Jošt Hobič for conducting the Twitter (X) poll. In addition, the authors give thanks for the constructive comments and suggestions by academic editors and the anonymous reviewers.

**Conflicts of Interest:** The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

## References

1. Elam, M. Poetry Will Not Optimize: Creativity in the Age of AI. In *Generative AI: Perspectives from Stanford HAI. How Do You Think Generative AI Will Affect Your Field and Society Going Forward?* Li, F.-F., Russ, A., Langlotz, C., Ganguli, S., Landay, J., Michele, E., Ho, D.E., Liangs, P., Brynjolfsson, E., Manning, C.D., et al., Eds.; HAI, Stanford University, Human-Centred Artificial Intelligence: Palo Alto, CA, USA, 2023; pp. 11–12.
2. Manning, C.D. The Reinvention of Work. In *Generative AI: Perspectives from Stanford HAI. How Do You Think Generative AI Will Affect Your Field and Society Going Forward?* Li, F.-F., Russ, A., Langlotz, C., Ganguli, S., Landay, J., Michele, E., Ho, D.E., Liangs, P., Brynjolfsson, E., Manning, C.D., et al., Eds.; HAI, Stanford University, Human-Centred Artificial Intelligence: Palo Alto, CA, USA, 2023; pp. 18–19.
3. Liang, P. The New Cambrian Era: “Scientific Excitement, Anxiety”. In *Generative AI: Perspectives from Stanford HAI. How Do You Think Generative AI Will Affect Your Field and Society Going Forward?* Li, F.-F., Russ, A., Langlotz, C., Ganguli, S., Landay, J., Michele, E., Ho, D.E., Liangs, P., Brynjolfsson, E., Manning, C.D., et al., Eds.; HAI, Stanford University, Human-Centred Artificial Intelligence: Palo Alto, CA, USA, 2023; p. 15.
4. Eloundou, T.; Manning, S.; Mishkin, P.; Rock, D. GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of 5. Large Language Models. *arXiv* **2023**, arXiv:2303.10130. [[CrossRef](#)]
5. Bommasani, R.; Hudson, D.A.; Adeli, E.; Altman, R.; Arora, S.; von Arx, S.; Bernstein, M.S.; Bohg, J.; Bosselut, A.; Brunskill, E.; et al. *On the Opportunities and Risks of Foundation Models*; Center for Research on Foundation Models, Stanford University: Palo Alto, CA, USA, 2023.
6. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Lukasz, K.; Polosukhin, I. Attention Is All You Need. In *Advances in Neural Information Processing Systems*; Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2017; Volume 30.
7. Kaplan, J.; McCandlish, S.; Henighan, T.; Brown, T.B.; Chess, B.; Child, R.; Gray, S.; Radford, A.; Wu, J.; Amodei, D. Scaling Laws for Neural Language Models. *arXiv* **2020**, arXiv:2001.08361. [[CrossRef](#)]
8. Hoffmann, J.; Borgeaud, S.; Mensch, A.; Buchatskaya, E.; Cai, T.; Rutherford, E.; de Las Casas, D.; Hendricks, L.A.; Welbl, J.; Clark, A.; et al. Training Compute-Optimal Large Language Models. *arXiv* **2022**, arXiv:2203.15556. [[CrossRef](#)]
9. Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. Language Models Are Few-Shot Learners. In *Advances in Neural Information Processing Systems 33*; Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2020; pp. 1877–1901.
10. Chowdhery, A.; Narang, S.; Devlin, J.; Bosma, M.; Mishra, G.; Roberts, A.; Barham, P.; Chung, H.W.; Sutton, C.; Gehrmann, S.; et al. PaLM: Scaling Language Modeling with Pathways. *arXiv* **2022**, arXiv:2204.02311. [[CrossRef](#)]
11. Patel, D.; Wong, G. GPT-4 Architecture, Infrastructure, Training Dataset, Costs, Vision, MoE. Demystifying GPT-4: The Engineering Tradeoffs That Led OpenAI to Their Architecture. *SemiAnalysis*, 10 July 2023. Available online: <https://www.semianalysis.com/p/gpt-4-architecture-infrastructure> (accessed on 12 September 2023).
12. Gardizy, A.; Ma, W. Microsoft Readies AI Chip as Machine Learning Costs Surge. *The Information*, 18 April 2023. Available online: <https://www.theinformation.com/articles/microsoft-readies-ai-chip-as-machine-learning-costs-surge> (accessed on 12 September 2023).
13. OpenAI GPT-4 Technical Report. *arXiv* **2023**, arXiv:2303.08774. [[CrossRef](#)]
14. Wolfram, S. Wolfram|Alpha as the Way to Bring Computational Knowledge Superpowers to ChatGPT. *Stephen Wolfram Writings*, 2023. Available online: <https://www.writings.stephenwolfram.com/2023/01/wolframalpha-as-the-way-to-bring-computational-knowledge-superpowers-to-chatgpt/> (accessed on 12 September 2023).
15. Wolfram, S. What Is ChatGPT Doing ... and Why Does It Work? *Stephen Wolfram Writings*, 2023. Available online: <https://writings.stephenwolfram.com/2023/02/what-is-chatgpt-doing-and-why-does-it-work/> (accessed on 12 September 2023).
16. OpenAI Introducing ChatGPT. *OpenAI Blog*, 30 November 2022. Available online: <https://openai.com/blog/chatgpt> (accessed on 12 September 2023).
17. Bryant, A. AI Chatbots: Threat or Opportunity? *Informatics* **2023**, *10*, 49. [[CrossRef](#)]
18. Hsiao, S.; Collins, E. Try Bard and Share Your Feedback. *Blog Google*, 21 March 2023. Available online: <https://blog.google/technology/ai/try-bard> (accessed on 12 September 2023).
19. Mehdi, Y. Reinventing Search with a New AI-Powered Microsoft Bing and Edge, Your Copilot for the Web. *Official Microsoft Blog*, 7 February 2023. Available online: <https://blogs.microsoft.com/blog/2023/02/07/reinventing-search-with-a-new-ai-powered-microsoft-bing-and-edge-your-copilot-for-the-web> (accessed on 12 September 2023).
20. OpenAI Introducing ChatGPT Plus. *OpenAI Blog*, 1 February 2023. Available online: <https://openai.com/blog/chatgpt-plus> (accessed on 12 September 2023).
21. Opera Your Comprehensive Guide to Aria: Opera’s Native Browser AI. *Blogs Opera*, 2023. Available online: <https://blogs.opera.com/desktop/2023/06/introducing-aria/> (accessed on 12 September 2023).
22. Anthropic Introducing Claude. *Anthropic Blog*, 14 March 2023. Available online: <https://www.anthropic.com/index/introducing-claude> (accessed on 12 September 2023).
23. Anthropic Claude 2. *Anthropic Blog*, 11 July 2023. Available online: <https://www.anthropic.com/index/claude-2> (accessed on 12 September 2023).



24. Brynjolfsson, E. A Call to Augment—Not Automate—Workers. In *Generative AI: Perspectives from Stanford HAI. How Do You Think Generative AI Will Affect Your Field and Society Going Forward?* Li, F.-F., Russ, A., Langlotz, C., Ganguli, S., Landay, J., Michele, E., Ho, D.E., Liangs, P., Brynjolfsson, E., Manning, C.D., et al., Eds.; HAI, Stanford University, Human-Centred Artificial Intelligence: Palo Alto, CA, USA, 2023; pp. 16–17.
25. Li, F.-F. AI's Great Inflection Point. In *Generative AI: Perspectives from Stanford HAI. How Do You Think Generative AI Will Affect Your Field and Society Going Forward?* Li, F.-F., Russ, A., Langlotz, C., Ganguli, S., Landay, J., Michele, E., Ho, D.E., Liangs, P., Brynjolfsson, E., Manning, C.D., et al., Eds.; HAI, Stanford University, Human-Centred Artificial Intelligence: Palo Alto, CA, USA, 2023; pp. 4–5.
26. Microsoft. *Will AI Fix Work? 2023 Work Trend Index: Annual Report*; Microsoft: Redmond, WA, USA, 2023. Available online: [https://assets.ctfassets.net/y8fb0rhks3b3/5eyZc6gDu1bzftdY6w3ZVV/93190f5a8c7241ecf2d6861bdc7fe3ca/WTI\\_Will\\_AI\\_Fix\\_Work\\_060723.pdf](https://assets.ctfassets.net/y8fb0rhks3b3/5eyZc6gDu1bzftdY6w3ZVV/93190f5a8c7241ecf2d6861bdc7fe3ca/WTI_Will_AI_Fix_Work_060723.pdf) (accessed on 12 September 2023).
27. Donelan, M. Government Commits up to £3.5 Billion to Future of Tech and Science. In *UK Government News*; 16 March 2023. Available online: <https://www.gov.uk/government/news/government-commits-up-to-35-billion-to-future-of-tech-and-science> (accessed on 12 September 2023).
28. Hinton, G.; Bengio, Y.; Hassabis, D.; Altman, S.; Amodei, D.; Song, D.; Lieu, T.; Gates, B.; Zhang, Y.-Q.; Sutskever, I.; et al. Statement on AI Risk. *Center for AI Safety*, 2023. Available online: <https://www.safe.ai/statement-on-ai-risk> (accessed on 12 September 2023).
29. Okerlund, J.; Klasky, E.; Middha, A.; Kim, S.; Rosenfeld, H.; Kleinman, M.; Parthasarathy, S. *What's in the Chatterbox? Large Language Models, Why They Matter, and What We Should Do About Them*; University of Michigan: Ann Arbor, MI, USA, 2022. Available online: <https://stpp.fordschool.umich.edu/sites/stpp/files/2022-05/large-language-models-TAP-2022-final-051622.pdf> (accessed on 12 September 2023).
30. Hendrycks, D.; Mazeika, M. X-Risk Analysis for AI Research. *arXiv* **2022**, arXiv:2206.05862. [[CrossRef](#)]
31. Bucknall, B.S.; Dori-Hacohen, S. Current and Near-Term AI as a Potential Existential Risk Factor. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, Oxford, UK, 19–21 May 2021. [[CrossRef](#)]
32. Cohen, M.K.; Hutter, M.; Osborne, M.A. Advanced Artificial Agents Intervene in the Provision of Reward. *AI Mag.* **2022**, *43*, 282–293. [[CrossRef](#)]
33. Ngo, R.; Chan, L.; Mindermann, S. The Alignment Problem from a Deep Learning Perspective. *arXiv* **2022**, arXiv:2209.00626. [[CrossRef](#)]
34. Kleinman, Z.; Wain, P. Why Making AI Safe Isn't as Easy as You Might Think. *BBC News Website*, 13 June 2023. Available online: <https://www.bbc.com/news/technology-65850668> (accessed on 12 September 2023).
35. Carlsmith, J. Is Power-Seeking AI an Existential Risk? *arXiv* **2022**, arXiv:2206.13353. [[CrossRef](#)]
36. Kleinman, Z.; Gerken, T. Nick Clegg: AI Language Systems Are “Quite Stupid”. *BBC News Website*, 19 July 2023. Available online: <https://www.bbc.com/news/technology-66238004> (accessed on 12 September 2023).
37. Vallance, C. AI Could Replace Equivalent of 300 Million Jobs—Report. *BBC News Website*, 28 March 2023. Available online: <https://www.bbc.com/news/technology-65102150> (accessed on 12 September 2023).
38. Vallance, C. Sarah Silverman Sues OpenAI and Meta. *BBC News Website*, 12 July 2023. Available online: <https://www.bbc.com/news/technology-66164228> (accessed on 12 September 2023).
39. Bearne, S. New AI Systems Collide with Copyright Law. *BBC News Website*, 1 August 2023. Available online: <https://www.bbc.com/news/business-66231268> (accessed on 12 September 2023).
40. Future of Life Institute Pause Giant AI Experiments: An Open Letter. 2023. Available online: <https://futureoflife.org/open-letter/pause-giant-ai-experiments> (accessed on 12 September 2023).
41. Vallance, C. Powerful Artificial Intelligence Ban Possible, Government Adviser Warns. *BBC News Website*, 1 June 2023. Available online: <https://www.bbc.com/news/technology-65779181> (accessed on 12 September 2023).
42. Li, F.-F.; Russ, A.; Langlotz, C.; Ganguli, S.; Landay, J.; Michele, E.; Ho, D.E.; Liangs, P.; Brynjolfsson, E.; Manning, C.D.; et al. (Eds.) *Generative AI: Perspectives from Stanford HAI. How Do You Think Generative AI Will Affect Your Field and Society Going Forward?* HAI, Stanford University, Human-Centred Artificial Intelligence: Palo Alto, CA, USA, 2023.
43. McCallum, S. Seven AI Companies Agree to Safeguards in the US. *BBC News Website*, 22 July 2023. Available online: <https://www.bbc.com/news/technology-66271429> (accessed on 12 September 2023).
44. Madiaga, T. *Artificial Intelligence Act*; European Parliament: Strasbourg, France, 2023; Available online: [https://www.europarl.europa.eu/RegData/etudes/BRIE/2021/698792/EPRS\\_BRI\(2021\)698792\\_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/BRIE/2021/698792/EPRS_BRI(2021)698792_EN.pdf) (accessed on 12 September 2023).
45. Imafidon, A.-M.; Floridi, L.; George, D.; Curtis, P.; Davenport, J.H.; Koene, A.; Taylor, J.; Baines, V.; Mainelli, M.; Bramer, M.; et al. AI Open Letter to UK Government and Industry; BCS, The Chartered Institute for IT. 2023. Available online: <https://www.bcs.org/sign-our-open-letter-on-the-future-of-ai> (accessed on 12 September 2023).
46. Simons, J. The Creator of ChatGPT Thinks AI Should Be Regulated. *Time*, 5 February 2023. Available online: <https://time.com/6252404/mira-murati-chatgpt-openai-interview/> (accessed on 12 September 2023).
47. White House FACT SHEET: Biden-Harris Administration Secures Voluntary Commitments from Leading Artificial Intelligence Companies to Manage the Risks Posed by AI. In *The White House Brief*; 21 July 2023. Available online: <https://www.whitehouse.gov/briefing-room/statements-releases/2023/07/21/fact-sheet-biden-harris-administration->



- secures-voluntary-commitments-from-leading-artificial-intelligence-companies-to-manage-the-risks-posed-by-ai/ (accessed on 12 September 2023).
48. Altman, S. Planning for AGI and Beyond. *OpenAI Blog*, 2023. Available online: <https://openai.com/blog/planning-for-agi-and-beyond> (accessed on 12 September 2023).
  49. Lund, B.D.; Wang, T.; Mannuru, N.R.; Nie, B.; Shimray, S.; Wang, Z. ChatGPT and a New Academic Reality: AI-Written Research Papers and the Ethics of the Large Language Models in Scholarly Publishing. *J. Assoc. Inf. Sci. Technol.* **2023**, *74*, 570–581. [CrossRef]
  50. Haque, M.U.; Dharmadasa, I.; Sworna, Z.T.; Rajapakse, R.N.; Ahmad, H. “I Think This Is the Most Disruptive Technology”: Exploring Sentiments of ChatGPT Early Adopters Using Twitter Data. *arXiv* **2022**, arXiv:2212.05856. [CrossRef]
  51. Altmäe, S.; Sola-Leyva, A.; Salumets, A. Artificial Intelligence in Scientific Writing: A Friend or a Foe? *Reprod. BioMed. Online* **2023**, *47*, 3–9. [CrossRef] [PubMed]
  52. Else, H. Abstracts Written by ChatGPT Fool Scientists. *Nature* **2023**, *613*, 423. [CrossRef] [PubMed]
  53. Gao, C.A.; Howard, F.M.; Markov, N.S.; Dyer, E.C.; Ramesh, S.; Luo, Y.; Pearson, A.T. Comparing Scientific Abstracts Generated by ChatGPT to Original Abstracts Using an Artificial Intelligence Output Detector, Plagiarism Detector, and Blinded Human Reviewers. *Digit. Med.* **2023**, *6*, 75. [CrossRef] [PubMed]
  54. Osmanovic Thunström, A. We Asked GPT-3 to Write an Academic Paper about Itself—Then We Tried to Get It Published an Artificially Intelligent First Author Presents Many Ethical Questions—and Could Upend the Publishing Process. *Scientific American*, 30 June 2022. Available online: <https://www.scientificamerican.com/article/we-asked-gpt-3-to-write-an-academic-paper-about-itself-mdash-then-we-tried-to-get-it-published/> (accessed on 12 September 2023).
  55. Generative Pretrained Transformer, G.; Thunström Osmanovic, A.; Steingrímsson, S. Can GPT-3 Write an Academic Paper on Itself, with Minimal Human Input? *HAL Preprint*, 2022; hal-03701250. Available online: <https://hal.science/hal-03701250/document> (accessed on 1 October 2023).
  56. Petiška, E. ChatGPT Cites the Most-Cited Articles and Journals, Relying Solely on Google Scholar’s Citation Counts. As a Result, AI May Amplify the Matthew Effect in Environmental Science. *arXiv* **2023**, arXiv:2304.06794. [CrossRef]
  57. Májovský, M.; Černý, M.; Kasal, M.; Komarc, M.; Netuka, D. Artificial Intelligence Can Generate Fraudulent but Authentic-Looking Scientific Medical Articles: Pandora’s Box Has Been Opened. *J. Med. Internet Res.* **2023**, *25*, e46924. [CrossRef] [PubMed]
  58. Bhattacharyya, M.; Miller, V.M.; Bhattacharyya, D.; Miller, L.E. High Rates of Fabricated and Inaccurate References in ChatGPT-Generated Medical Content. *Cureus* **2023**, *15*, e39238. [CrossRef] [PubMed]
  59. Buholayka, M.; Zouabi, R.; Tadinada, A. Is ChatGPT Ready to Write Scientific Case Reports Independently? A Comparative Evaluation Between Human and Artificial Intelligence. *Cureus* **2023**, *15*, e39386. [CrossRef]
  60. Lund, B.D.; Wang, T. Chatting about ChatGPT: How May AI and GPT Impact Academia and Libraries? *Libr. Hi Tech News* **2023**, *40*, 26–29. [CrossRef]
  61. Salvagno, M.; Taccone, F.S.; Gerli, A.G. Can Artificial Intelligence Help for Scientific Writing? *Crit. Care* **2023**, *27*, 75. [CrossRef] [PubMed]
  62. Pividori, M.; Greene, C.S. A Publishing Infrastructure for AI-Assisted Academic Authoring. *bioRxiv* **2023**, 2023.01.21.525030. [CrossRef]
  63. Venema, L.; Jerde, T.; Huth, J.; Pieropan, M.; Matusevych, Y. The AI Writing on the Wall. *Nat. Mach. Intell.* **2023**, *5*, 1. [CrossRef]
  64. Getahun, H. After an AI Bot Wrote a Scientific Paper on Itself, the Researcher behind the Experiment Says She Hopes She Didn’t Open a “Pandora’s Box”. *Bussines Insider Nederland*, 9 July 2022.
  65. Jaupi, J. An AI bot wrote a paper about itself—and experts fear they’ve opened ‘Pandora’s box’. *The Sun*, 11 July 2022.
  66. Understanding ChatGPT Is a Bold New Challenge for Science. *Nature* **2023**, *619*, 671–672. Available online: <https://www.nature.com/articles/d41586-023-02366-2.pdf> (accessed on 1 October 2023).
  67. Moskvichev, A.; Odouard, V.V.; Mitchell, M. The ConceptARC Benchmark: Evaluating Understanding and Generalization in the ARC Domain. *arXiv* **2023**, arXiv:2305.07141. [CrossRef]
  68. Kocoń, J.; Cichecki, I.; Kaszyca, O.; Kochanek, M.; Szydło, D.; Baran, J.; Bielaniec, J.; Gruza, M.; Janz, A.; Kanclerz, K.; et al. ChatGPT: Jack of All Trades, Master of None. *Inf. Fusion* **2023**, *99*, 101861. [CrossRef]
  69. Bubeck, S.; Chandrasekaran, V.; Eldan, R.; Gehrke, J.; Horvitz, E.; Kamar, E.; Lee, P.; Lee, Y.T.; Li, Y.; Lundberg, S.; et al. Sparks of Artificial General Intelligence: Early Experiments with GPT-4. *arXiv* **2023**, arXiv:2303.12712. [CrossRef]
  70. Koo, M. The Importance of Proper Use of ChatGPT in Medical Writing. *Radiology* **2023**, *307*, e230312. [CrossRef]
  71. Liebrez, M.; Schleifer, R.; Buadze, A.; Bhugra, D.; Smith, A. Generating Scholarly Content with ChatGPT: Ethical Challenges for Medical Publishing. *Lancet Digit. Health* **2023**, *5*, e105–e106. [CrossRef]
  72. Hill-Yardin, E.L.; Hutchinson, M.R.; Laycock, R.; Spencer, S.J. A Chat(GPT) about the Future of Scientific Publishing. *Brain Behav. Immun.* **2023**, *110*, 152–154. [CrossRef]
  73. Zhao, W.X.; Zhou, K.; Li, J.; Tang, T.; Wang, X.; Hou, Y.; Min, Y.; Zhang, B.; Zhang, J.; Dong, Z.; et al. A Survey of Large Language Models. *arXiv* **2023**, arXiv:2303.18223. [CrossRef]
  74. Fan, L.; Li, L.; Ma, Z.; Lee, S.; Yu, H.; Hemphill, L. A Bibliometric Review of Large Language Models Research from 2017 to 2023. *arXiv* **2023**, arXiv:2304.02020. [CrossRef]

75. Christiano, P.; Leike, J.; Brown, T.B.; Martic, M.; Legg, S.; Amodei, D. Deep Reinforcement Learning from Human Preferences (V4). *arXiv* **2023**, arXiv:1706.03741. [[CrossRef](#)]
76. Perrigo, B. Exclusive: OpenAI Used Kenyan Workers on Less Than \$2 Per Hour to Make ChatGPT Less Toxic. *Time*, 18 January 2023. Available online: <https://time.com/6247678/openai-chatgpt-kenya-workers> (accessed on 12 September 2023).
77. OpenAI ChatGPT Plugins. *OpenAI Blog*, 2023. Available online: <https://openai.com/blog/chatgpt-plugins> (accessed on 12 September 2023).
78. Opera Opera Browser AI—Aria FAQ. *Opera Help*, 2023. Available online: <https://help.opera.com/en/browser-ai-faq> (accessed on 12 September 2023).
79. Manyika, J. An Overview of Bard: An Early Experiment with Generative AI. *AI. Google Static Documents*, 2023. Available online: <https://ai.google/static/documents/google-about-bard.pdf> (accessed on 12 September 2023).
80. Collins, E.; Ghahramani, Z. LaMDA: Our Breakthrough Conversation Technology. *Blog Google*, 2021. Available online: <https://blog.google/technology/ai/lamda/> (accessed on 12 September 2023).
81. Thoppilan, R.; De Freitas, D.; Hall, J.; Shazeer, N.; Kulshreshtha, A.; Cheng, H.-T.; Jin, A.; Bos, T.; Baker, L.; Du, Y.; et al. LaMDA: Language Models for Dialog Applications. *arXiv* **2022**, arXiv:2201.08239. [[CrossRef](#)]
82. Bai, Y.; Kadavath, S.; Kundu, S.; Askell, A.; Kernion, J.; Jones, A.; Chen, A.; Goldie, A.; Mirhoseini, A.; McKinnon, C.; et al. Constitutional AI: Harmlessness from AI Feedback. *arXiv* **2022**, arXiv:2212.08073. [[CrossRef](#)]
83. Bai, Y.; Jones, A.; Ndousse, K.; Askell, A.; Chen, A.; DasSarma, N.; Drain, D.; Fort, S.; Ganguli, D.; Henighan, T.; et al. Training a Helpful and Harmless Assistant with Reinforcement Learning from Human Feedback. *arXiv* **2022**, arXiv:2204.05862. [[CrossRef](#)]
84. Štular, B.; Lozić, E.; Belak, M.; Rihter, J.; Koch, I.; Modrijan, Z.; Magdič, A.; Karl, S.; Lehner, M.; Gutjahr, C. Migration of Alpine Slavs and Machine Learning: Space-Time Pattern Mining of an Archaeological Data Set. *PLoS ONE* **2022**, *17*, e0274687. [[CrossRef](#)] [[PubMed](#)]
85. Curta, F. *The Making of the Slavs: History and Archaeology of the Lower Danube Region, c. 500–700*; Cambridge University Press: Cambridge, UK, 2001; ISBN 978-0-511-49629-5.
86. Curta, F. Migrations in the Archaeology of Eastern and Southeastern Europe in the Early Middle Ages (Some Comments on the Current State of Research). In *Migration Histories of the Medieval Afro-Eurasian Transition Zone: Aspects of Mobility between Africa, Asia and Europe, 300-1500 CE*; Preise-Kapeller, J., Reinfandt, L., Stouraitis, Y., Eds.; Brill: Leiden, Boston, MA, USA, 2020; pp. 101–140.
87. Ratzel, F. *Anthropogeographie. Bd. 1. Grundzüge Der Anwendung Der Erdkunde*; J. Engelhorn's Nach.: Berlin, Germany, 1909.
88. MacEachern, S. Genes, Tribes, and African History. *Curr. Anthropol.* **2000**, *41*, 357–384. [[CrossRef](#)]
89. Bentley, C.G. Ethnicity and Practice. *Comp. Stud. Soc. Hist.* **1987**, *29*, 24–55. [[CrossRef](#)]
90. Knapp, B.A. *Prehistoric and Protohistoric Cyprus. Identity, Insularity, and Connectivity*; Oxford University Press: Oxford, UK, 2008.
91. Grafenauer, B. *Zgodovina Slovenskega Naroda, I. Zvezek: Od Naselitve Do Uveljavljenja Frankovskega Feodalnega Reda*; Kmečka knjiga: Ljubljana, Slovenia, 1954.
92. White, J.; Fu, Q.; Hays, S.; Sandborn, M.; Olea, C.; Gilbert, H.; Elnashar, A.; Spencer-Smith, J.; Schmidt, D.C. A prompt pattern catalog to enhance prompt engineering with chatgpt. *arXiv* **2023**, arXiv:2302.11382. [[CrossRef](#)]
93. Woods, R. AI + You: How to Use ChatGPT for Content Creation. *Microsoft Learn Articles*, 11 July 2023. Available online: <https://create.microsoft.com/en-us/learn/articles/how-to-use-chatgpt-for-content-creation> (accessed on 12 September 2023).
94. PLOS. How to Write a Peer Review. *PLOS*, 2023. Available online: <https://plos.org/resource/how-to-write-a-peer-review/> (accessed on 12 September 2023).
95. Lozić, E. Application of Airborne LiDAR Data to the Archaeology of Agrarian Land Use: The Case Study of the Early Medieval Microregion of Bled (Slovenia). *Remote Sens.* **2021**, *13*, 3228. [[CrossRef](#)]
96. Kajkowski, K. The Boar in the Symbolic and Religious System of Baltic Slavs in the Early Middle Ages Dzik w Systemie Symboliczno Religijnym Słowian Nadbałtyckich Wczesnego Średniowiecza. *Stud. Myth. Slavica* **2012**, *15*, 201–215. [[CrossRef](#)]
97. Pleterski, A. *Kulturni Genom: Prostor in Njegovi Ideogrami Mitične Zgodbe*; Založba ZRC: Ljubljana, Slovenia, 2014; ISBN 978-961-254-736-3.
98. Lucy, L.; Bamman, D. Gender and Representation Bias in GPT-3 Generated Stories. In Proceedings of the Third Workshop on Narrative Understanding, Virtual, 11 June 2021; Nader Akoury, F.B., Chaturvedi, S., Clark, E., Iyyer, M., Martin, L.J., Eds.; Association for Computational Linguistics: Kerrville, TX, USA, 2021; pp. 48–55.
99. Hutchinson, B.; Prabhakaran, V.; Denton, E.; Webster, K.; Zhong, Y.; Denuyl, S. Social Biases in NLP Models as Barriers for Persons with Disabilities. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics; Dan Jurafsky, J.C., Schluter, N., Tetreault, J., Eds.; Association for Computational Linguistics: Kerrville, TX, USA, 2020; pp. 5491–5501.
100. Kotek, H.; Dockum, R.; Sun, D.Q. Gender Bias and Stereotypes in Large Language Models. *arXiv* **2023**, arXiv:2308.14921. [[CrossRef](#)]
101. Rovira, C.; Codina, L.; Lopezosa, C. Language Bias in the Google Scholar Ranking Algorithm. *Future Internet* **2021**, *13*, 31. [[CrossRef](#)]
102. Hrushevsky, M. *History of Ukraine-Rus'. Volume 1. From Prehistory to the Eleventh Century*; University of Alberta Press: Edmonton, AB, USA, 1997; Volume 1, ISBN 1-895571-19-7.
103. Obolensky, D. *The Byzantine Commonwealth: Eastern Europe, 500-1453*; History of Civilization Series; Weidenfeld and Nicolson: London, UK, 1971.

104. Dolukhanov, P. *The Early Slavs: Eastern Europe from the Initial Settlement to the Kievan Rus*; Routledge: England, UK, 1996; ISBN 978-1-315-84355-1.
105. Nagtegaal, L.W.; de Bruin, R.E. The French Connection and Other Neo-Colonial Patterns in the Global Network of Science. *Res. Eval.* **1994**, *4*, 119–127. [[CrossRef](#)]
106. Schürer, Y. Decolonising the Library—In Deutschland? In *Bibliothek Forschung und Praxis*; De Gruyter: Berlin, Germany, 2023; p. 47. [[CrossRef](#)]
107. Lindstedt, J.; Salmela, E. Language Contact and the Early Slavs. In *New Perspectives on the Early Slavs and the rise of Slavic*; Klír, T., Boček, V., Jansens, N., Eds.; Empirie und Theorie der Sprachwissenschaft; Universitätsverlag Winter GmbH: München, Germany, 2020; pp. 275–299. ISBN 978-3-8253-4707-9.
108. Beel, J.; Gipp, B. *Google Scholar's Ranking Algorithm: An Introductory Overview*; Larsen, B., Leta, J., Eds.; International Society for Scientometrics and Informetrics: Leuven, Belgium, 2009; pp. 230–241.
109. Martín-Martín, A.; Orduna-Malea, E.; Ayllón, J.M.; López-Cózar, E.D. Back to the Past: On the Shoulders of an Academic Search Engine Giant. *Scientometrics* **2016**, *107*, 1477–1487. [[CrossRef](#)]
110. Sanh, V.; Debut, L.; Chaumond, J.; Wolf, T. DistilBERT, a Distilled Version of BERT: Smaller, Faster, Cheaper and Lighter. *arXiv* **2019**, arXiv:1910.01108. [[CrossRef](#)]
111. Ackoff, R.L. From Data to Wisdom. *J. Appl. Syst. Anal.* **1989**, *16*, 3–9.
112. Štular, B.; Belak, M. Deep Data Example: Zbiva, Early Medieval Data Set for the Eastern Alps. *Res. Data J. Humanit. Soc. Sci.* **2022**, 1–13. [[CrossRef](#)]
113. Hodder, I. *The Archaeological Process an Introduction*; Wiley-Blackwell: Hoboken, NJ, USA, 1999; ISBN 978-0-631-19885-7.
114. Stokel-Walker, C. AI Bot ChatGPT Writes Smart Essays—Should Professors Worry? *Nature* **2022**. [[CrossRef](#)]
115. Ganguli, S. *An AI Window into Nature In Generative AI: Perspectives from Stanford HAI. How Do You Think Generative AI Will Affect Your Field and Society Going Forward?* Li, F.-F., Russ, A., Langlotz, C., Ganguli, S., Landay, J., Michele, E., Ho, D.E., Liangs, P., Brynjolfsson, E., Manning, C.D., et al., Eds.; HAI, Stanford University, Human-Centred Artificial Intelligence: Palo Alto, CA, USA, 2023; pp. 8–9.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.