

thatwhitepaperguy

9 Reasons Why ChatGPT Can't Write Your Next White Paper

A white paper written
entirely by ChatGPT!

Contents

Introduction	3
1. Missing creative flair	4
2. Struggling with context	5
3. Risk of inaccuracies	6
4. Limited understanding of human emotions	7
5. Difficulty grasping the target audience	8
6. Overlooking ethical considerations	9
7. Personalization challenges	10
8. Risk of bias	11
9. Limitations in reasoning	12
Conclusions	13
About the author	14
About the publisher	14
Afterword by Gordon Graham	15
Sources	20

Copyright © 2023 by That White Paper Guy Inc.

All text from *Introduction* to *About the Author* was output by ChatGPT.
The *Afterword* was written by Gordon Graham.

V1.0 February 20, 2023 with original sources from ChatGPT-3.5.

V2.0 March 3, 2023 with sources from new Bing and revised *Afterword*.

V2.1 March 14, 2023 with new *Introduction* by ChatGPT-4 and updated *Afterword*.

V2.2 September 12, 2025 with updated links to my revised website.

Introduction

White papers are crucial for B2B businesses looking to educate and persuade potential customers on complex subjects. But crafting a top-notch white paper takes time, resources, and expertise.

Enter AI technology: a tempting solution for companies wanting to streamline the process using AI language models. Yet, are AI models truly the best choice for producing compelling B2B white papers?

In this paper, we dive into nine reasons why relying solely on AI language models might not be your best bet:

1. Missing creative flair
2. Struggling with context
3. Risk of inaccuracies
4. Limited understanding of human emotions
5. Difficulty grasping the target audience
6. Overlooking ethical considerations
7. Personalization challenges
8. Risk of bias
9. Limitations in reasoning

While human-generated text can also have these issues, these points emphasize the potential weaknesses of AI language models in crafting persuasive B2B white papers. So, what's the solution?

We recommend a hybrid approach, marrying the best of both worlds: the creativity and critical thinking of human writers with the efficiency and data-driven insights of AI models. By doing so, you can create high-quality white papers that truly resonate with your target audience.


1. Missing creative flair

While AI language models are incredibly accurate and efficient in generating text, they may lack creativity.

“AI is designed to be precise, to follow instructions, and to achieve specific goals,” notes media commentator Devan Leos. “Because of this, AI is not very good at being creative.”¹

AI language models work based on patterns and algorithms, and therefore may not provide a unique and creative perspective that captivates readers.

“Many AI researchers agree that AI cannot generate entirely new ideas on its own—for now. Several arguments suggest that algorithms can only do what they are programmed to do,” said a recent research paper from Germany.²



An AI language model may produce a white paper that contains all the necessary information but lacks personality and originality.

For example, a white paper on marketing automation written by an AI language model may provide data and statistics but lack the punch a human writer can give it.

Conversely, human writers can create engaging and memorable white papers using metaphors, analogies, and other figures of speech. They can take the data and statistics and craft a compelling and engaging story that resonates with the reader and makes them more likely to take action.

This illustrates one of the limitations of using AI for writing white papers and suggests that human writers may be better equipped to produce compelling and effective content.

2. Struggling with context

AI language models can struggle to understand context, which can be a significant issue when it comes to writing a white paper.

While AI may excel at mimicking human language, it lacks the ability to truly understand the subtleties of context and nuance. This can lead to inaccuracies, irrelevant content, and less effective communication with readers.



“Unlike humans, AI systems lack the judgment and context for many of the environments in which they are deployed,” stated a journal from the Wharton School at the University of Pennsylvania. “In most cases, it is not possible to train the AI system on all possible scenarios and data.”³

The inability to understand context is particularly problematic when it comes to writing on complex and technical subjects that may require specialized knowledge and terminology.

“A model, however intelligent the output may seem, is a mere representation of reality,” noted executive Julia Krauwer from the Dutch bank AB AMRO. “That is why a human in the loop is essential: We are, unlike machines, able to take into account context and use general knowledge.”⁴


For example, an AI language model may produce a white paper on renewable energy that does not fully understand the nuances of the topic or the target audience.

A human writer with experience in the field can produce a white paper that is well-researched, relevant, and engaging to the target audience.

3. Risk of inaccuracies

AI language models have the potential to produce high-quality text, but they can also be prone to inaccuracies if not designed correctly or if the data they were trained on is biased or inaccurate.

ChatGPT may mistakenly attribute a quote or a fact to the wrong source, which can be particularly problematic in a white paper.



As veteran technology editor Harry McCracken noted, “Even when ChatGPT is mostly accurate, it often makes a fundamental mistake or two. For example, it told me that former Apple CEO John Sculley was responsible for the iPod, a product released eight years after he left the company.”⁵

“These systems are inherently unreliable, frequently making errors of both reasoning and fact, and prone to hallucination,” said Gary Marcus, scientist, author, and AI contrarian.⁶

Using AI language models to generate white papers may seem like a time-saving solution, but it can be risky.

An AI-generated white paper may contain inaccuracies or use biased data, which could undermine the credibility and damage the reputation of the company that publishes it.


In contrast, human writers can conduct thorough research and fact-checking to produce a high-quality white paper that provides accurate information and is valuable to the reader. With a human touch, a white paper can address complex topics in a clear and concise manner while avoiding the potential pitfalls of AI-generated content.

4. Limited understanding of human emotions

It is clear that AI language models have limitations in understanding and conveying human emotions. This is because these models lack the innate ability of humans to process emotions and empathize with others.

AI language models may struggle to effectively convey emotions and empathy, which can be important in certain types of writing, such as persuasive writing.

“While ChatGPT can generate responses that seem empathic, it does not possess true emotional intelligence,” says Bernard Mar, a well-known futurist and author. “It cannot detect subtle emotional cues or respond appropriately to complex emotional situations.”⁷



Although AI models can analyze data and recognize patterns, they cannot interpret emotions or understand the context behind them.

As a result, they may not be able to capture the nuances of emotional language or create content that truly resonates with readers.

This limitation can be important in persuasive writing, where an emotional connection with the reader is often key to achieving the desired outcome.

We are far from understanding the human heart, let alone replicating it, notes tech executive and author Kai-Fu Lee. “For all of AI’s astounding capabilities, the one thing that only humans can provide turns out to be exactly what’s most needed in our lives: love.”⁸


As a result, AI language models may not be able to create white papers that emotionally connect with the reader, which can be essential in persuading them to take action.

A human writer, on the other hand, can use emotional language to make the text more persuasive and compelling.

5. Difficulty grasping the target audience

AI language models may struggle to understand the intricacies of a target audience, making it challenging to produce content that resonates with readers. This can be a significant limitation when writing a white paper that seeks to persuade or influence a specific audience.

AI language models may not understand the needs and preferences of the target audience, leading to a lack of relevance in the content produced.



For example, an AI language model may produce a white paper on healthcare that is not tailored to the needs and preferences of a specific audience, such as physicians or patients.

This limitation can result in a white paper that fails to engage the reader, making it less effective at achieving its intended purpose.

Furthermore, an AI language model may produce content that does not match the tone or style preferred by the target audience, which can further detract from its impact.

Marketing expert Seth Godin warns, “People do not buy goods and services. They buy relations, stories, and magic.” AI language models are not capable of creating the magical elements necessary to persuade an audience and build a connection with them.

In contrast, human writers can conduct market research and develop an understanding of the target audience, producing content that speaks directly to their needs, preferences, and biases.


This results in a white paper that is tailored to the audience and has a higher chance of engaging and persuading the reader.

6. Overlooking ethical considerations

AI does not possess the ethical and moral considerations that humans have, which can lead to a lack of judgment.

According to a recent report from the U.S. Financial Industry Regulatory Authority (FINRA), numerous incidents have been reported where AI software was “fraudulent, nefarious, discriminatory, or unfair.”⁹

AI language models lack the ability to consider ethical issues. This limitation is significant for white papers about topics that involve ethical considerations.



Another challenge for AI writing is that it cannot factor in cultural and religious beliefs.

For instance, a white paper on gene editing may not consider the ethical considerations of different cultures and belief systems.

The field of AI ethics is still grappling with many issues that remain unresolved. A report from PEW Research¹⁰ highlights some of these challenges:

- ethics can take a back seat to market share and profits
- formal ethics training is not mandatory for AI developers
- cultural differences may make a global consensus on ethics impossible

It’s also important to note that ethics is not binary; decisions often involve complex trade-offs and shades of gray. No one can compel an unethical player to stick to a rulebook, and this lack of accountability is a cause for concern.

Another issue is the “risk of unethical practices being used to generate content designed to manipulate or deceive the reader,”¹¹ said a blog post from ContentBot, which emphasized the seriousness of this concern for writers who use AI.

In summary, AI language models lack the ability to consider ethical issues and to make decisions based on moral principles.


Human writers have the necessary experience and expertise to write in a way that takes into account cultural and religious beliefs, reflects a balanced and impartial assessment, and factors in the trade-offs of ethical decisions.

7. Personalization challenges

AI-generated content can lack the personalization that comes with the human touch, making it less engaging and less likely to connect with readers.

For example, an AI language model may generate a white paper on a new product, but it may not be able to tailor the content to the specific needs and preferences of the target audience. This can lead to less engagement and fewer conversions.

Most marketers believe in personalization. A 2020 survey showed that the majority of all B2B marketers (93%) agree that personalizing content on their website has increased their organization's revenue over time.¹²



This shows that personalization is crucial to creating effective marketing content.

While AI language models have the ability to analyze data and provide insights, they may not be able to understand the needs and preferences of the target audience.

AI-generated content may fall short in establishing a personal connection with readers, which is essential for fostering deeper engagement and loyalty.


By leveraging their experience and expertise, human writers have the ability to create content that truly resonates with the target audience, eliciting emotions, sparking conversations, and driving action.

8. Risk of bias

AI models often rely on large amounts of data to learn and improve their performance. However, some sources of data may have inherent biases that can be perpetuated by the model.

Wikipedia, for example, is a common source of training data for AI models. Unfortunately, it has a significant gender bias due to the fact that only 15% of contributors are women, according to the Wikimedia Foundation.¹³

As a result, articles on the platform may perpetuate stereotypes and under-represent women in various fields, creating a gender gap that has been acknowledged by the foundation.



As we strive for gender equality, it is important to acknowledge the potential biases in AI content.

For example, the UNESCO report, "I'd Blush If I Could," states that AI voice assistants reinforce gender stereotypes by programming female voices to respond to sexist comments.¹⁴

There are three ways in which bias can be introduced into an AI system:

- **Data bias** can occur when the data used to train an AI system is incomplete, inaccurate, or unrepresentative
- **Algorithmic bias** can occur when the methods or assumptions used to create an AI model are flawed or inappropriate
- **Human bias** can occur when human prejudices and biases are introduced into AI systems

In addition, cultural diversity and sensitivity can also be a challenge for AI language models, leading to inaccuracies and misunderstandings. For example, an AI-generated white paper on a global issue such as climate change may fail to fully consider the diverse perspectives of different cultures and regions.

A human writer, on the other hand, can be more attuned to the complexities of gender, culture, and language, and can provide a more sensitive analysis.


9. Limitations in reasoning

AI language models have a significant limitation: they can't reason or think critically. Although they can analyze large amounts of data and generate fast responses, they lack the ability to make informed decisions or think deeply about information.

This deficiency is especially problematic when creating white papers on complex topics that require extensive critical thinking and analysis.

For instance, consider a white paper on the ethical implications of autonomous vehicles. Such a topic requires careful consideration.

While AI models can provide data and information on a range of factors such as safety, public policy, and societal impacts, they are unable to think critically about that information or make informed decisions on how to address a topic.



In contrast, a human writer has “a range of cognitive abilities, such as creativity, intuition, and critical thinking, that allow them to generate new ideas and hypotheses beyond their initial training data.”¹⁵

Human writers also possess the ability to reason deductively and construct valid arguments based on

given premises, which is not something that current AI language models can do at the same level of sophistication.

Skilled writers can use critical thinking skills to identify underlying assumptions, evaluate the validity of arguments, and consider potential counter-arguments.

By engaging in this type of critical analysis, a human writer can provide a more comprehensive and nuanced assessment of a topic that takes into account a wider range of perspectives and considerations.

Conclusions

While AI language models can provide many benefits for B2B white paper writing, they also have significant limitations that must be carefully considered.

The use of AI can help speed up the writing process and increase efficiency, but it may also result in content that lacks creativity or accuracy. Additionally, AI may not be equipped to consider ethical issues or produce content that emotionally resonates with the reader.

On the other hand, human writers can provide a unique and creative perspective, tailor content to the target audience, conduct research and fact-checking, and produce a more nuanced and informed analysis.

The ideal approach is to leverage the speed and efficiency of AI for generating ideas and draft text, while relying on human writers to add the creativity, empathy, critical thinking, and fact-checking necessary to produce a high-quality final product.





About the author

This white paper was generated by ChatGPT-3.5, a large language model trained by OpenAI. As an AI language model, ChatGPT has been designed to assist users in generating high-quality written material on a variety of topics. By leveraging machine learning algorithms and natural language processing capabilities, ChatGPT can produce informative and engaging content that meets the specific needs of the user.



About the publisher

This white paper was published by That White Paper Guy Inc., a leading provider of white paper writing services founded by award-winning writer Gordon Graham.

Since 1997, Gordon has worked on 300+ white papers for clients from Silicon Valley to Switzerland, for everyone from tiny start-ups to 3M, Google, and Verizon. The company is now preparing an AI-assisted service to help marketers and writers prepare white papers faster and easier than ever before.

For the latest tips on long-form content, visit www.ThatWhitePaperGuy.com

Afterword by Gordon Graham

What do you think? Is this a competent white paper, despite the title? Do you wonder if you should try AI for your next white paper after all?

This white paper is an experiment I did with ChatGPT and the new Bing in February and March 2023. **All in all, I'm convinced ChatGPT can help write white papers faster and easier than ever. For generating ideas, making outlines, doing research, and generating first drafts, ChatGPT + the new Bing form a quick and cost-effective workflow.**

By the numbers: metrics on this paper

- Words drafted by ChatGPT: 3,250
- Words cut by editor: 250
- Words added by editor: 150
- Total number of drafts: 10
- Total number of prompts: 150+
- Sources proposed by ChatGPT: 18
- Bogus sources from ChatGPT: 16
- Sources proposed by the new Bing: 20
- Sources difficult to confirm from new Bing: 5

By the numbers: metrics on process

This table compares this white paper with ChatGPT-3.5 to my best guestimates for doing a similar white paper without AI. I kept notes on the elapsed time, reducing my numbers a little to cover my learning curve with ChatGPT and the new Bing.

	This white paper by ChatGPT + the new Bing	Similar white paper without any AI
<i>Time to generate ideas</i>	<i>A few seconds</i>	<i>Minutes to hours</i>
<i>Time to research</i>	<i>A few minutes</i>	<i>A few days</i>
<i>Time to first draft</i>	<i>1.5 hours</i>	<i>At least 15 hours</i>
<i>Time to revise (10 drafts)</i>	<i>4 hours</i>	<i>At least 4 hours</i>
<i>Time to check and re-research</i>	<i>6 hours</i>	<i>Not applicable</i>
<i>Time to design</i>	<i>6 hours</i>	<i>6 hours</i>

Source: *That White Paper Guy*, 2023

Success: Lightning-fast outlining and writing

The table shows that all through the first stages, working with AI was likely 10X faster. I started by giving ChatGPT a brief prompt for a provocative topic: “Please give me the best reasons not to trust AI to write a B2B white paper.”

Within seconds, ChatGPT suggested the nine items in the list and wrote an acceptable outline. Then I prompted it to generate a first draft. In about 90 minutes, it wrote 3,500 words from the *Introduction to About the Author* including 18 sources and footnotes in proper format. That’s fast. Lightning fast.

Success: Next-to-no manual editing

To edit the v1.0 output, I spread out a hard copy of our chats on my dining room table for what I thought would be a long session of cut-and-paste.

But I decided not to polish ChatGPT’s output, so you can see what it can do.



Here’s all the editing I did:

- deleted a couple hundred words that were repetitive
- Added the nine items to the Intro
- pasted in my boilerplate for *About the Publisher*.

In fact, I spent most of my editing time applying styles, inserting endnotes, and futzing with Word.

Fail: So-so style

To finish up the v1.0 text, I did six drafts with ChatGPT, working section by section so it could manage the word count. I repeatedly urged it to be more creative, yet the output remained stodgy. This is perhaps C-level copy, which could be pulled up to B-level with some stern revisions.

My partner and former journalist Angie Gallop said that draft lulled her to sleep, with most of the quotes simply repeating what ChatGPT already said, doing nothing to amplify or illuminate the point.

I think ChatGPT managed a clear voice but certainly not an engaging style.

The text’s readability scores from Word confirm this impression: v2.0 scored 36 out of 100 on the Flesh Reading Ease and Grade 13.5 on the Flesh-Kincaid Grade Level. That’s only half as readable as this *Afterword*, rated 72 and Grade 6.6.

I realized after the fact that I could prompt the AI for better readability. I jumped on ChatGPT-4 the day it was released and asked it to redo the

Introduction to make it more readable. In seconds, it improved the readability of that page to 66 and Grade 9—scores that are perfectly acceptable in my book.

Terrible fail: Fabricated sources

If I prompted the AI carefully and reminded it often, ChatGPT created perfectly formatted footnotes with placeholders correctly located in the text. Was this the Holy Grail? Could AI now do instant research with perfect sourcing? In seconds, not days? That struck me as a game-changer.

I know, I know, the AI tried to warn me in point 3. All those lovely endnotes? Out of 18 sources ChatGPT provided for v1.0, not a single one was solid. Only a couple of the links worked; all the rest were broken. One company it referenced had been acquired, so its research was no longer online. Another had in fact published what ChatGPT said, but at some point taken it down.

The quotes sounded reasonable, from people and publications that do exist. For example, the quote on page 8 from Seth Godin sounds like something he likely said. But I couldn't find where. Even after using Google, Google Books, Amazon, and Amazon's Look Inside feature, I came up empty-handed for source after source. But I left that in as an Easter egg.

Worse, ChatGPT gave me different sources in different iterations. I was astounded to see articles, books, and page numbers changing from one draft to the next, even though they still pointed to the same quotes.

Fact-checking ChatGPT felt like marking a project from a bored high school student who had pasted in text from Wikipedia while watching TikTok with their friends. After two hours, I gave up in bitter disappointment.

Success: Doing research with the new Bing

I'd been on the waitlist for the new Bing search engine with ChatGPT built in. And the minute my wait was over, I tried using the new Bing for research. To my delight, it readily overcomes the sourcing limits of ChatGPT. The new Bing has real-time access to the web for up-to-the-minute sources. And it can also create proper footnotes if you remind it to.

Out of 14 sources in v2.0, two came from ChatGPT and the other 12 from the new Bing. The new Bing is still awkward to use, and it defaults to paraphrasing rather than extracting quotes. It has some annoying guard-rails and a current limit of six prompts to a session. After that, you must click on the magic broom to clear its memory and start over. Not efficient.

Even so, in another couple of hours, I gathered more than enough sources to create this much-improved v2.0 of this paper. I'm sure you'll agree that six hours



to research and fact-check a whole white paper is remarkably quick. That usually takes me days.

Advice: Don't fall back on the old ways

At times, ChatGPT and the new Bing are phenomenal. Other times, they're like children who need a lot of handholding. I get frustrated when they don't come up with exactly what I'm after. And I find myself thinking, "It's okay, I can fix that while I'm editing."

But I'm reminded of something I heard 40 years ago from the venerable Stan Bevington, founder of Toronto's innovative small press, Coach House Press.

In the early 1980s, he was bringing in digital typesetting to complement their traditional letterpress operations. And he said how frustrating it was to see a typo in the machine output after fiddling with a file for hours.

"Sometimes you just want to go in with a knife and fix it!" he declared.

But he was learning to take a breath, put down his paste-up knife, go back into the computer file, make the fix, and re-run the galley.

"Otherwise, the file will never be right," he sighed. "And we'll never get comfortable with the new tools."

It's the same with prompting ChatGPT or the new Bing.

Even though it seems slow and frustrating, I'm learning to prompt and re-prompt to nudge the text closer to what I want. That's the only way to get comfortable with the new tools.

Observation: It's all about the prompts

What they call "prompt engineering" is a weird combination of art and science. Trying out different prompts can be exhilarating—or exasperating.



Sometimes the system will instantly come up with ideas, tweak the voice, combine two previous drafts, expand or compress a passage—and do it all in seconds, far faster than any human could.

For example, ChatGPT wrote a landing page for this white paper in less than 10 seconds. It takes me that long just to open a file in Word, let alone start typing.

Other times, the AI seems to miss the point and go right off the rails. It munged up quotes when I told it not to touch them. Over and over.

And while ChatGPT claims it doesn't get tired, I tested its endurance. After working on this paper for 90 minutes, it ran out of steam. After 50 prompts, it started to sound like someone who'd been up all night drinking.

OpenAI admits that if you stay in the same session for too long, ChatGPT can get confused, poor dear. And with its six-prompt limit, the new Bing has the attention span of a goldfish.

So, can ChatGPT write a white paper?

You tell me. I hope you found this a thought-provoking exercise.

Here are some more questions I'll be exploring in the coming days:

- Is ChatGPT-4 truly capable of writing B-level content?
- What's the best workflow between ChatGPT and the new Bing?
- Is there any need for Google search in that workflow?
- Will Google punish any white papers co-written with ChatGPT?



Like everyone else, I'm still learning how to use ChatGPT. I've only had my hands on ChatGPT-4 for a few minutes so far.

I'm still discovering the best tips for prompting it and then polishing the output.

One thing's for sure: We're a long way from pressing a button to crank out a great white paper in five minutes. **But despite all the limitations listed in this paper, I believe that AI will quickly become a valuable helper for generating long-form content.**

And if you'd like to see what AI can do for your next white paper, visit <https://thatwhitepaperguy.com/category/ai-for-white-papers/>

To see what else is new with me, visit www.ThatWhitePaperGuy.com

A handwritten signature in black ink that reads "Gordon Liaba". The signature is fluid and cursive.

That White Paper Guy
March 14, 2023

Sources

- 1 Devan Leos, “Is AI a Risk to Creativity? The Answer Is Not So Simple,” Entrepreneur.com, 12 December 2022, <https://www.entrepreneur.com/science-technology/is-ai-a-risk-to-creativity-the-answer-is-not-so-simple/439525>
- 2 Jan Bieser, “Creative Through AI: How artificial intelligence can support the development of new ideas,” GDI Gottlieb Duttweiler Institute, 2022, pp 7 and 21, <https://gdi.ch/en/publications/studies/creative-through-ai-pdf-2022-e#attr=>
- 3 “How Can Financial Institutions Prepare for AI Risks?” Knowledge@Wharton, 13 April 2021, <https://knowledge.wharton.upenn.edu/article/can-financial-institutions-prepare-ai-risks/>
- 4 “The Next Big Wave: How Financial Institutions Can Stay Ahead of the AI Revolution,” Finextra Research and Intel, May 2017, p 23, <https://www.ceo.digital/media/ai-revolution.pdf>
- 5 Harry McCracken, “If ChatGPT doesn’t get a better grasp of facts, nothing else matters,” FastCompany.com, 11 January 2023, <https://www.fastcompany.com/90833017/openai-chatgpt-accuracy-gpt-4>
- 6 Gary Marcus, “AI’s Jurassic Park moment,” Substack.com, 10 December 2022, <https://garymarcus.substack.com/p/ais-jurassic-park-moment>
- 7 Bernard Marr, “The Top 10 Limitations of ChatGPT,” Forbes, 3 March 2023, <https://www.forbes.com/sites/bernardmarr/2023/03/03/the-top-10-limitations-of-chatgpt/>
- 8 Kai-Fu Lee, “AI Superpowers: China, Silicon Valley and the New World Order,” Houghton Mifflin Harcourt, 2018, p 196
- 9 “Artificial Intelligence (AI) in the Securities Industry,” Financial Industry Regulatory Authority, June 2020, p 11, <https://www.finra.org/sites/default/files/2020-06/ai-report-061020.pdf>
- 10 Lee Rainie, Janna Anderson and Emily A. Vogels, “Experts Doubt Ethical AI Design Will Be Broadly Adopted as the Norm Within the Next Decade,” Pew Research, 16 June 2021, <https://www.pewresearch.org/internet/2021/06/16/experts-doubt-ethical-ai-design-will-be-broadly-adopted-as-the-norm-within-the-next-decade/>
- 11 Nick Duncan, “The ethical use of AI content – A basic guideline,” ContentBot.ai, 23 August 2022, <https://www.contentbot.ai/blog/news/the-ethical-use-of-ai-content-a-basic-guideline/>

12 Krystle Vermes, “Report: 93% of B2B Marketers Say Personalizing Content Has Boosted Company Revenue,” KoMarketing.com, 18 June 2020, <https://komarketing.com/industry-news/report-93-of-b2b-marketers-say-personalizing-content-has-boosted-company-revenue-4299/>

13 Barbara Kandek, “Closing the gender gap: Women in Red’s efforts to add more women to Wikipedia,” Wikimedia Foundation News, 1 March 2023, <https://wikimediafoundation.org/news/2023/03/01/closing-the-gender-gap-women-in-reds-efforts-to-add-more-women-to-wikipedia/>

14 “I’d Blush If I Could: Closing Gender Divides in Digital Skills Through Education,” UNESCO, 2019, pp. 108-110, <https://unesdoc.unesco.org/ark:/48223/pf0000367416.page=1>

15 Vassili Gorshkov, “Interview on AI with ChatGPT,” LinkedIn Pulse, 23 February 2023, <https://www.linkedin.com/pulse/interview-ai-chatgpt-vassili-gorshkov>