JOURNAL OF COMPARATIVE RESEARCH IN ANTHROPOLOGY AND SOCIOLOGY

Volume 13, Number 1, Winter 2022
ISSN 2068 – 0317
http://compaso.eu

Let me write that for you: Prospects concerning the impact of GPT-3 on the copywriting workforce

Denis Iorga[1]

Abstract

The advertising industry is an economic sector in which the capability of Artificial Intelligence (AI) technology - such as Generative Pre-trained Transformer 3 (GPT-3) – of generating text indistinguishable from a human-written text can have both positive and negative consequences. However, there are limited research attempts to examine the consequences of GPT-3 on the advertising workforce. In this sense, the current study explores the viability of using GPT-3 for automated copywriting (defined as writing text for advertising purposes) and discusses prospects concerning the impact of GPT-3 on the copywriting workforce. An advertisement evaluation inquiry was conducted to evaluate the viability of GPT-3 for automated copywriting. The inquiry involved asking participants (n=31) to choose between advertisements with text generated through GPT-3 and advertisements with human-written text. Based on the reflexive analysis of the results, it is plausible to consider three implications of GPT-3 for the copywriting workforce: (1) it may substitute tasks that involve the generation of low-cost, mass-produced advertising text; (2) it may create new tasks which involve the manipulation of GPT-3 input/output for automated copywriting; and (3) it may aid copywriters to manage creative exhaustion. Such prospects suggest an uneven distribution of GPT-3 consequences on the copywriting workforce, influenced by technological, occupational, and economic factors.

Keywords

Artificial Intelligence; AI; GPT; GPT-3; Future of work; Workforce; Jobs; Automation; Copywriting; Advertising; Creative industries; Creative exhaustion;

[1] Interdisciplinary School of Doctoral Studies, University of Bucharest, Romania, denis.iorga@drd.unibuc.ro.

Introduction

Understanding the impact of Artificial Intelligence (AI) on the workforce employed in specific industries is key amid debates regarding the automation of jobs. The most pressing debates concerning this issue refer to *whether AI will lead to job substitution, creation, or alteration*. The current work aims to contribute to the matter by asking whether an AI-enabled product, namely an automated copywriting system using the Generative Pre-trained Transformer 3 (GPT-3), will lead to the substitution, creation, or alteration of copywriting jobs in the advertising industry.

As such, the current study explores the potential implications of GPT-3 on the future of copywriting work. While there are studies that analyze the use (Branwen, 2020), performance (Elkins and Chun, 2020), and social implications of GPT-3 (Floridi and Chiriatti, 2020), there is limited research concerning the impact of GPT-3 on the copywriting workforce. It is, therefore, important to address this gap so that decision-makers in the advertising industry can anticipate the consequences of automation through AI.

The current work focuses on the case study of a copywriting automation system called "AdvertAI". The system uses GPT-3 to automatically generate advertising text based on textual input. An image choice inquiry involving 31 subjects was conducted to evaluate the AdvertAI system's viability to substitute, create, and/or alter copywriting jobs. The subjects were requested to choose from five pairs of similar images that depicted advertisements for cosmetic products. Half of the advertisements contained advertising texts generated by the GPT-3 system, and the other half contained human-authored advertising text. The reflexive analysis of the results provides a preliminary empirical basis on which a discussion is advanced regarding the uneven impact of GPT-3 on the copywriting workforce.

Literature review

Automation through AI and the workforce

An extensive body of research argues that AI technologies pose social risks. The bias and discrimination created through algorithmic decisions (Heinrichs, 2021), shifts in the international power structures (Polcumpally, 2021), and power asymmetries in the workplace due to automation through AI (Reed, 1987) are just three of the main concerns raised by scientific and advocacy communities preoccupied with this matter. Another concern in the context of work is the potential of AI technologies to automate tasks and negatively impact employment levels.

Concerns regarding the negative impact of automation on employment levels are not new. Borenstein (2011) states that these concerns were brought to public attention at the beginning of the 20th century. Others argue that they were raised even earlier. Smith (2018), for example, identifies the industrial revolution as giving rise to such concerns. It should then come as no surprise that attempts to understand the consequences of AI automation on employment levels use previous waves of automation as reference points.

Automation through AI is portrayed as having more pronounced negative effects on employment than previous waves of automation. Ford (2013) argues that, unlike previous waves of automation, automation through AI poses a more difficult challenge to the workforce: "rather than simply acquiring new skills and moving to another routine job, workers will have to instead migrate to an occupation that is genuinely non-routine and therefore protected from automation […] there are good reasons to be pessimistic about the ability of most of our workforce to accomplish this" (p. 38). Through this argument, the author attempts to make a point that automation through AI is expected to have more pronounced negative effects on employment by comparison with previous waves of automation due to an unprecedented shift of skill requirements towards non-routine tasks.

Yet, the shifting skill requirements created by automation through AI can also positively impact employment. Lane and Saint-Martin (2021) argue that "much of the impact of AI on jobs is likely to be experienced through the reorganization of tasks within an occupation due to the bottlenecks of AI adoption" (p. 4). In this sense, such reorganization of tasks may lead to job enhancement rather than job substitution. Additionally, such a reorganization of tasks may lead to job creation. Acemoglu and Restrepo (2019), for example, argue that jobs destroyed due to automation are counterbalanced by new job opportunities created by automation technologies. The same counterbalancing argument is stated by Muro et al. (2019a) who assert that automation, especially automation through AI, has a muted effect on employment levels at the macro level.

The muted effect of automation caused by AI on employment levels may hide an unequal distribution of job substitution, job creation, and job alteration across industries. Vermeulen et al. (2018) use a dataset with expert projections regarding the susceptibility of various occupations to AI automation. They show how job loss in economic sectors where AI technology is applied (e.g.: transportation) will be counterbalanced by job growth in economic sectors where AI technology is made (e.g.: software engineering). Additionally, Muro et al. (2019b) analyze the text of AI patents and the text of job descriptions to demonstrate that *the impact of automation through AI on the workforce will be unequally distributed*: "it is the smaller, better-paying high-tech or professional industries and their workers that will be most changed by AI" (p. 16). Both examples show that the unequal distribution of job substitution, job creation, and job alteration determined by AI becomes visible through an occupation-based approach.

Nonetheless, understanding the impact of AI on employment levels using an occupation-based approach faces a series of challenges. In this sense, Frank et al. (2019) highlight "the lack of high-quality data about the nature of work (e.g., the dynamic requirements of occupations), lack of empirically informed models of key micro-level processes (e.g. skill substitution and human-machine complementarity), and insufficient understanding of how cognitive technologies interact with broader economic dynamics and institutional mechanisms (e.g., urban migration and international trade policy)" (p. 6531). Furthermore, in a review of three recent occupation-based studies that attempt to understand the impact of AI on employment, Lane and Saint-Martin (2021) argue that "these studies [...] are more limited in what they can say about whether workers […] will see their work substituted or complemented" (p. 24).

Recent attempts that use an occupation-based approach to understand the impact of AI on employment rely on patent data and task descriptions to infer the exposure of jobs to AI automation. Muro et al. (2019b) argue that such an approach, which focuses on specific AI technologies, overcomes the limitations of using subjective expert projections and case studies (Muro et al., 2019b). However, this approach is also limited in what it can say about whether workers will see their work as substituted or altered (Lane and Saint-Martin, 2021).

Based on the research presented above, it can be argued that subjective expert projections regarding AI workforce replacement and case studies concerning the impact of AI in different industries may help overcome the limitation of recent statistical occupation-based approaches that use patent data and job descriptions. This is because both expert projections and case studies capture implicit industrial knowledge regarding the nature of work. In turn, such knowledge may shed light on the factors that will decide whether workers will see their work as substituted or altered. In support of this argument, the current research brings forward a case study of a GPT-3 system used for copywriting automation.

GPT-3 use for text generation

GPT-3 is a language model created using Artificial Intelligence. It can generate text given the input it receives. The release of the beta version of GPT-3 allowed researchers and practitioners to test different use cases of GPT-3. Examples range from generating summaries of a given text to generating programming code (Floridi and Chiriatti, 2020). Other creative use cases of GPT-3 highlighted in the literature include the generation of dialogues, folktales, poetry, articles, and much more (Branwen, 2020). As such, researchers could assess the performance of GPT-3 to generate human-like text.

For example, the performance of GPT-3 to create text similar to humans is evaluated by Brown et al. (2020). To evaluate the quality of text generated by GPT-3, the authors conducted a form of Turing Test, in which 80 human evaluators were asked to decide if a set of news articles were written by a human or by GPT-3. The results show that for the best version of the GPT-3 model, human evaluators obtained a 52% accuracy in detecting which text is human-made and which is GPT-3-made, suggesting an impressive performance of GPT-3 to produce human-like text.

However, Elkins and Chun (2020) argue that while GPT-3 can have impressive results for text generation, it can fail at the simplest linguistic tasks, for instance, to "maintain a coherent argument or narrative thread over long periods of time; maintain consistency of gender or personality; employ simple grammar rules; show basic knowledge and common sense reasoning" (p. 3). Thus, the authors argue that, for GPT-3 to produce text that is indistinguishable from human-made text, humans have to validate the output of GPT-3 before the actual evaluation (Elkins and Chun, 2020).

The use of GPT-3 for text-generation purposes can have negative consequences on society, for instance, discrimination through bias embedded in the language model (Lucy and Bamman, 2021), misinformation amplification through the use of GPT-3 to generate fake news (Floridi and Chiriatti, 2020), and plagiarism proliferation through the use of GPT-3 to generate intellectual property (Dehouche, 2021). Yet, limited research discusses the impact of GPT-3 on the workforce. As GPT-3 can be used to generate news articles, stories, reports, and other forms of written content, it is important to consider the impact of GPT-3 diffusion on occupations that involve writing text, for example, those that involve copywriting.

Materials and methods

An image choice inquiry was conducted to evaluate the viability of using GPT-3 for automated copywriting. The inquiry involved asking 31 subjects to compare five pairs of cosmetic product advertisements (ads). In each pair, one ad contained text automatically generated by a GPT-3 system, while the other contained human-authored text. For each pair, the respondents had to choose from the two ads randomly placed on the right and the left side underneath the question (see Figure 1). The following lines describe the steps taken to create the five pairs of ads.

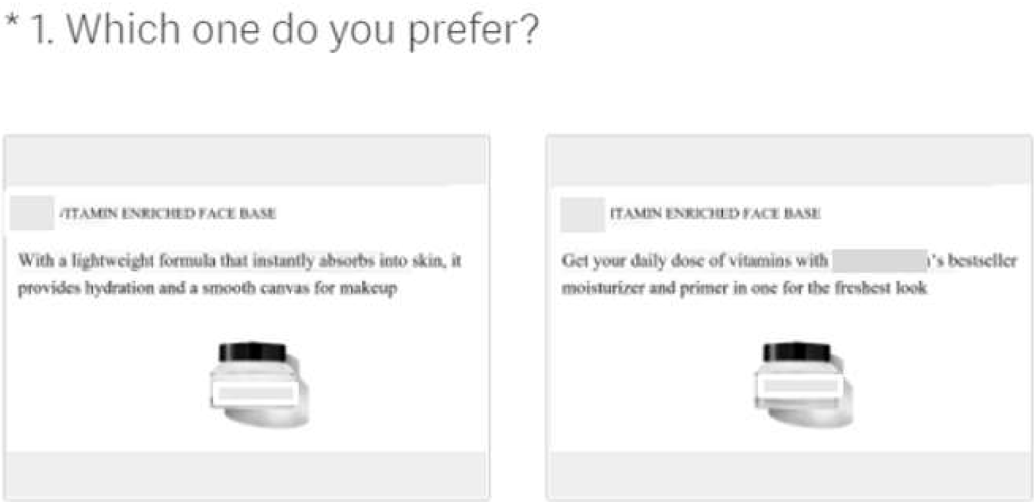

Figure 1. Example of image choice question with a pair of ads that contained text generated by GPT-3 (left) and human-authored text (right); the brand name was anonymized.

A GPT-3 user, who voluntarily accepted to be part of the study, was asked to select five products from any website. The GPT-3 user, aged 28, randomly selected five products from a single website to which he contributed as a software developer. The website belongs to a well-renowned cosmetic brand with a net worth of 2$ million dollars.

Following the selection of the five cosmetic products, the GPT-3 user was asked to automatically generate advertising text for each product using GPT-3.

To generate the advertising texts using GPT-3 for each of the five cosmetic products, the GPT-3 user utilized AdvertAI, a proprietary automated copywriting system based on GPT-3. The purpose was to generate social media advertising text for each of the five cosmetic products, with a maximum length of 20 words. Text from the web pages of the five cosmetic products was used as input. For each product, the text used as input included the product's name and a short product description available on the web pages.

The AdvertAI proprietary automated copywriting system belongs to a start-up owned by the GPT-3 user. The GPT-3 user and his team developed the proprietary automated copywriting system by fine-tuning GPT-3 for advertising text generation. Fine-tuning allows GPT-3 to perform better across various tasks (OpenAI, 2021). The GPT-3 user could not share details regarding the fine-tuning process of GPT-3 as they represent the start-up's intellectual property.

The AdvertAI system generated four alternative advertisement texts for each cosmetic product. The GPT-3 user was asked to go through the four alternatives and select the first appropriate advertising text. After reviewing 20 advertising texts, the GPT-3 user selected the first generated alternative for each product. Afterward, the GPT-3 user sent the selected five advertising texts generated by GPT-3 to the author via e-mail.

To obtain the human-authored advertising texts for each of the five cosmetic products, another volunteer accepted to write them. The volunteer, aged 25, occupies an entry-level content writer position with a tenure of three months in a small marketing organization. The volunteer was asked to use available information present on the web pages of the five cosmetic products to write the advertising texts. Furthermore, given her experience with cosmetic products, the volunteer was asked to write appropriate advertising texts, that would not exceed 20 words for each product, based on a target audience that she finds suited for buying these cosmetic products. However, upon receiving the advertising texts from the volunteer via e-mail, the author noticed that one of the texts had 23 words. Nonetheless, the decision to use the text, despite exceeding 20 words, was taken due to time constraints.

Next, five pairs of ads (summing to ten ads) were created using the advertising texts received via e-mail and information from the web pages of the five cosmetic products. Each of the ten ads contained the product's name taken from the website, the corresponding advertising text (either generated with GPT-3 or human-authored), and an image of the product taken from the website (example in Figure 1). Great attention to the characteristics of the ads was given. For this purpose, an automated screenshot taker was implemented and used to print screen the ads after they had been designed in PowerPoint. Therefore, there were no differences between the ads in a pair, except for the advertising text, which was either human-made or GPT-3 made. Table 1 presents both human-authored and GPT-3 advertising texts, as received by the author via e-mail.

Table 1. Advertising text generated by the human copywriter and GPT-3 for five cosmetic products.

Product	Human-authored ad text	GPT-3-generated ad text
no.1	Get your daily dose of vitamins with [company name]'s bestseller moisturizer and primer in one for the freshest look*	With a lightweight formula that instantly absorbs into skin, it provides hydration and a smooth canvas for makeup
no.2	Learn how to achieve the most natural looking coverage with a matte finish that's comfortable, breathable, and weightless	A long-wear, lightweight formula that provides medium coverage with a matte finish
no.3	Even out your skin tone and control your shininess with this pressed powder– perfect for a smooth look	It's time to take your skin to the next level with this sheer finish pressed powder that leaves your skin looking and feeling like a dream
no.4	Shade, line, and define long-lasting looks with this creamy, stay-put, waterproof stick	A long-wear, crease-proof cream eyeshadow that provides vibrant color and a smooth finish
no.5	Looking for a naturally radiant finish? Then you should try this highlighter which brings an innovative formula for a smooth and refined look	A luminous, pearlescent powder that instantly creates a lit-from-within glow

**As the products belong to a certain brand, the name was anonymized in this paper. However, the name of the brand was not anonymized in the questionnaire.*

Even though the ads for each pair were similar, it is worth mentioning that there was a noticeable visual difference between the ads for product no. 5, given that the human-authored text was significantly longer than the GPT-3 text (23 words vs. 10 words.). As such, the ad with human-made text for product no. 5 had the advertising text spanning three rows, while the ad with GPT-3 text for product no. 5 had the advertising text spanning two rows. The rest of the ads had text spanning two rows.

Following this, an image choice questionnaire was created using the SurveyMonkey service. This service was chosen as it provided the best resolution for the five image choice questions. Additionally, the questionnaire contained three more demographic questions, related to the birth sex of respondents, the difficulty to read the advertising texts in the images, and the frequency of cosmetic product use, such as those presented in the ads.

Following the recommendation of the GPT-3 user, the questionnaire was sent to 31 female subjects with ages ranging from 20 to 30 years old. The respondents were contacted via social media and asked to participate in a market research study, in which the purpose was to decide which are the most suitable ads for the promotion of five cosmetic products. The decision to initially conceal the true purpose of the study was taken to prevent respondents from evaluating the "humanness" of the text (as in the case of the Turing Test), rather than its suitedness for advertising. After the questionnaire was completed, the purpose of the study was revealed, and written consent to use the answers for the ongoing study was received from the subjects.

The respondents were asked to fill in the questionnaire from their smartphones. The average rating on the text reading easiness was 4.4 on a scale from 1 to 5, where 1 means "very difficult" and 5 means "very easy". The average rating on the frequency of cosmetic product use, such as those presented in the ads, was 3.6, where 1 means "never" and 5 means "every day". Moreover, questions related to why the subjects preferred one advertisement over the other were addressed to subjects who were willing to further discuss after completing the questionnaire. Each respondent who was willing to discuss the motivation behind the answers to the image choice questions (n=4) was asked, for each pair of ads, „Why do you think that you chose one ad over the other?". Similarly, the volunteer who authored the advertising texts was also interrogated regarding the writing process to allow a more accurate interpretation of the results.

Results

On average, the ads with human-authored texts were slightly preferred over ads with texts generated by GPT-3. The results show that 47% of the respondents preferred GPT-3 texts, while 53% preferred human-authored texts. Based on these results, it may be argued that the human author was, on average, more proficient in writing advertising text than the GPT-3 system. Figure 2 below presents the scores for each of the product ads with texts presented in Table 1 (counted as the number of subjects who have chosen the ads with GPT-3 text or human-authored text, over the total number of respondents for each pair).

Figure 2. Scores obtained by GPT-3 and human-authored text (n=31).

As can be observed from the figure above, there is a tendency for GPT-3 advertising texts to have almost similar or even better performance than the human-authored advertising texts for products no. 3, 4, and 5. To explain such results, details regarding the writing process of human-authored advertising texts are relevant. The volunteer stated that she wrote the texts in the same period and that she was "out of ideas" for the last product. This might suggest a creative exhaustion phenomenon that influenced the volunteer to write less attractive advertising texts for products no. 3, 4, and 5 than for products no. 1 and 2. Further details regarding the relationship between GPT-3, creative exhaustion, and the workforce are provided in the discussion section.

Another important aspect to consider is the visual difference between the ads for product no. 5, given the noticeable difference in the size of the text (23 words vs. 10 words). Such difference may explain why GPT-3 text surpassed the human-authored text there. According to one of the subjects who were willing to discuss the motivation behind the answers to the image choice questions, simplicity was a criterion that determined her choices. Thus, it is reasonable to consider that subjects opted for the ad with less text. Therefore, the results for product no. 5 may reflect a cumulative effect of creative exhaustion and simplicity value.

Reasons for choosing one ad over the other also included a lack of directivity and straightforwardness, besides simplicity. Moreover, a respondent noticed how part of her answers was influenced by a current trend in the cosmetic industry in which great emphasis is placed on the ingredients. This insight gave reason to consider that frequent users of cosmetic products might display different preferences regarding the ads. Therefore, an additional analysis was conducted with respondents in the sample that frequently use cosmetic products (see Figure 3).

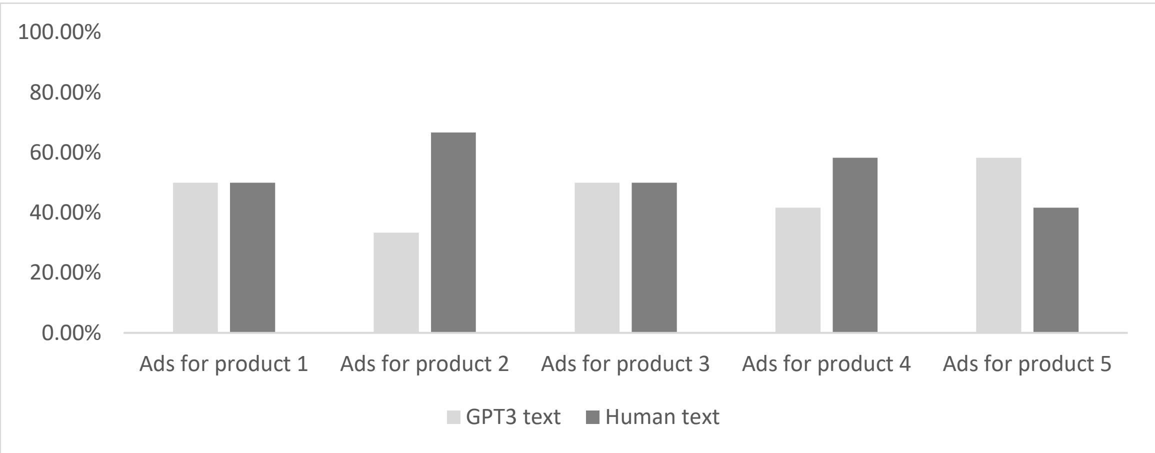

Figure 3. Scores obtained by GPT-3 and human-authored texts for respondents who use cosmetic products every few days and every day (n=12).

The results show a similar average preference of the texts amongst frequent users of cosmetic products, with approximately 47% of the respondents preferring GPT-3 texts and approximately 53% preferring human-written texts. As can be observed in the figure above, the results follow a similar pattern to those presented in Figure 2, except for the first pair of ads. One plausible explanation for this exception, besides the effect of a small sample, might reside in the reaction of subjects to the directiveness of the word “get” in the human-authored text. The word “get” carries a considerable epistemic authoritative connotation. Nonetheless, the overall analysis suggests that the frequency of cosmetic product usage does not have a noticeable influence on the results.

Discussion

The current study revealed that a human is more proficient than a GPT-3 system in terms of consumer preferences concerning the generated advertising texts for cosmetic products. However, the same results suggest that GPT-3 tends to achieve an approximative similar performance if the human reaches *creative exhaustion*. The latter insight provides the opportunity to infer possible consequences of GPT-3 diffusion on the copywriting workforce.

Creative exhaustion refers to the "inability to continue generating creative solutions on one's own" (Gray et al., 2019). In the case of copywriting, creative exhaustion may be defined as the inability to continue generating attractive (from a consumer point-of-view) advertising text on one's own. The current study managed to capture the manifestation of creative exhaustion: the "out of ideas" statement suggests that the volunteer who wrote the advertising texts was already facing creative exhaustion during the writing of product no. 5 advertising text.

Because the human volunteer decided to write the advertising texts in the same time segment, the cognitive resources required to generate attractive advertising texts gradually decreased. Consequently, the quality of the texts decreased, leading GPT-3 to perform almost similar to the volunteer for products no. 3, 4, and 5 (see Figure 2). Therefore, it is plausible to consider GPT-3 as a viable candidate for advertising text generation tasks that are susceptible to creative exhaustion.

In line with Lee and Cho (2020), who anticipate the use of AI for the generation of low-cost mass-produced advertising, the results of the current study suggest that *it is plausible for GPT-3 systems to substitute copywriting tasks that involve the generation of low-cost mass-produced advertising text*. Such tasks are highly susceptible to creative exhaustion as they closely resemble the work conducted by the volunteer involved in the current study. Thus, GPT-3 is a viable candidate for substituting jobs designed specifically for the generation of low-cost high-volume advertising text if the *costs/quality ratio* is satisfactory for the employer.

Nonetheless, *it is also plausible for GPT-3 systems to create new tasks both in the advertising industry and other industries.* The current study illustrated that using GPT-3 for automated copywriting requires humans to manipulate the input (choose the input text and set GPT-3 parameters) and the output (validate the appropriateness of the text and collate the advertising texts with product images). Therefore, it is plausible to consider that automated copywriting systems based on GPT-3 will require humans to "prompt & collate" (Floridi and Chiriatti, 2020) the input and output of GPT-3, given *the designed level of automation and autonomy* of the copywriting system (Simmler and Frischknecht, 2021). Nonetheless, assuming a similar *demand for advertising texts* in the market, the number of required "humans in the loop" (Zanzotto, 2019) will probably be lower than the number of substituted copywriters.

Even if automated copywriting systems will not follow the "human in the loop" principle, the development of fully autonomous copywriting systems will still create new tasks for software engineers. Solving errors and adding features are just two examples in this sense. This view is in line with the argument of Vermeulen et al. (2018) who argue that job loss in economic sectors where AI technology is applied (e.g.: advertising) will be counterbalanced by job growth in economic sectors where AI technology is made (e.g.: software engineering).

Additionally, *it is also plausible for GPT-3 systems to enhance the work of copywriters*. Gray et al. (2019) illustrate how computer technologies can be used to produce valuable ideas by humans facing creative exhaustion. This perspective seems to support the standpoints of Shatalov and Ryabova (2021) who argue that language models can save time and creativity resources for copywriters, and Duin and Pedersen (2021) who anticipate the future of writing to be collaborative between humans and machines.

Assuming that copywriting *jobs involve solely advertising text-generation tasks*, all three scenarios suggest an uneven distribution of "winners" and "losers" in the workforce due to copywriting automation systems based on GPT-3. On the one hand, individuals involved in the generation of low-cost/mass-produced advertising text seem to be the most exposed to copywriting automation through GPT-3. On the other hand, high-skill individuals from the copywriting industry (e.g.: copywriters involved in the generation of high-cost/custom advertising text) and other industries (e.g.: software engineers) seem to be the ones who could benefit the most from copywriting automation through GPT-3. This point of view further nuances Lutz (2019) argument regarding the social inequalities based on skill differences amplified by digital technologies.

Conclusions

In relation to the reviewed studies regarding the impact of AI automation on the workforce, the current study illustrated the capability of case studies conducted on specific AI applications to capture implicit industrial knowledge. More specifically, the current case study highlighted the importance of "creative exhaustion" as a key term for debates regarding the impact of AI on the advertising workforce, and the creative workforce more broadly.

Besides the *susceptibility of the workforce to creative exhaustion*, the present study identified other factors that will influence the interplay between job substitution and alteration due to AI in creative industries. Based on the discussion section, it can be argued that copywriting job substitution, creation, and alteration due to GPT-3 automated copywriting systems will also be influenced by their *costs*, the *market demand* concerning advertisement text, their *design* (especially the adherence to the "human in the loop" principle), and by the *number and value of tasks* attributable to specific copywriting jobs. As such, further research is encouraged to explore the impact of each of these factors.

Another contribution of the present study was to illustrate the application of a reflexive mixed-method approach for exploring the viability and prospects of AI text-generation systems. Existing viability evaluation methods focus on the human evaluation aspect of text generated by machines by using a series of items, such as those related to understandability, coherence, or grammar. The current work illustrates an alternative approach in which evaluation items are replaced with a binary subjective choice between human and machine-generated text, mixed with a qualitative exploration concerning the assessment of both the human evaluation and the generation of human-authored text.

As an exploratory study, the current research is not without limitations. Due to the use of a convenience sample, the replication of the present study is encouraged to confirm the results. Due to the use of a proprietary automated copywriting system, any replication attempts are to be conducted with the same version of the system that was used in the current study (December 2021). Additionally, in order to generalize the results, further studies should consider using multiple proprietary automated copywriting systems. Likewise, it would also be interesting to observe how the results of the current empirical investigation would change if the advertising texts generated by GPT-3 were compared to those of a more experienced human copywriter rather than with those of an entry-level human content writer.

Further research should also focus on assessing, rather than prospecting, the impact of GPT-3 diffusion on copywriting jobs. For example, given that the gig economy is commonly involved in the generation of advertising text, the impact of GPT-3 systems on employment levels may be inferred by correlating the diffusion of automated copywriting systems based on GPT-3 and employment levels of copywriters in the gig economy. Moreover, this approach should also reveal whether GPT-3 systems designed for copywriting automation have an uneven impact on the workforce based on skill.

Acknowledgments

This paper would not have been possible without the support of AdvertAI (https://advert-ai.com/) during their participation in the Innovation Labs program (https://www.innovationlabs.ro/).

REFERENCES

ACEMOGLU, D. & RESTREPO, P. 2019. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives,* 33(2), 3-30. doi: https://doi.org/10.1257/jep.33.2.3

BORENSTEIN, J. 2011. Robots and the changing workforce. *AI & society,* 26(1), 87-93. doi: https://doi.org/10.1007/s00146-009-0227-0

BRANWEN, G. 2020. *GPT-3 creative fiction* [Online]. Available: https://www.gwern.net/GPT-3 [Accessed August 2021].

BROWN, T. B., MANN, B., RYDER, N., SUBBIAH, M., KAPLAN, J., DHARIWAL, P., NEELAKANTAN, A., SHYAM, P., SASTRY, G. & ASKELL, A. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*. doi: https://doi.org/10.48550/arXiv.2005.14165

DEHOUCHE, N. 2021. Plagiarism in the age of massive Generative Pre-trained Transformers (GPT-3). *Ethics in Science and Environmental Politics,* 21, 17-23. doi: https://doi.org/10.3354/esep00195

DUIN, A. H. & PEDERSEN, I. 2021. Collaborative Writing Futures. *Writing Futures: Collaborative, Algorithmic, Autonomous.* Denmark: Springer Cham. pp. 27-52.

ELKINS, K. & CHUN, J. 2020. Can GPT-3 Pass a Writer's Turing Test? *Journal of Cultural Analytics,* 5(2)**,** 17212. doi: https://doi.org/10.22148/001c.17212

FLORIDI, L. & CHIRIATTI, M. 2020. GPT-3: Its nature, scope, limits, and consequences. *Minds and Machines,* 30(4)**,** 681-694. doi: https://doi.org/10.1007/s11023-020-09548-1

FORD, M. 2013. Could artificial intelligence create an unemployment crisis? *Communications of the ACM,* 56(7)**,** 37-39. doi: https://doi.org/10.1145/2483852.2483865

FRANK, M. R., AUTOR, D., BESSEN, J. E., BRYNJOLFSSON, E., CEBRIAN, M., DEMING, D. J., FELDMAN, M., GROH, M., LOBO, J. & MORO, E. 2019. Toward understanding the impact of artificial intelligence on labor. *Proceedings of the National Academy of Sciences,* 116(14)**,** 6531-6539. doi: https://doi.org/10.1073/pnas.1900949116

GRAY, C. M., MCKILLIGAN, S., DALY, S. R., SEIFERT, C. M. & GONZALEZ, R. 2019. Using creative exhaustion to foster idea generation. *International Journal of Technology and Design Education,* 29(1)**,** 177-195. doi: https://doi.org/10.1007/s10798-017-9435-y

HEINRICHS, B. 2021. Discrimination in the age of artificial intelligence. *AI & SOCIETY,* 37**,** 1-12. doi: https://doi.org/10.1007/s00146-021-01192-2

LANE, M. & SAINT-MARTIN, A. 2021. The impact of Artificial Intelligence on the labour market: What do we know so far? *OECD Social, Employment and Migration Working Papers,* No.256. doi: https://doi.org/10.1787/7c895724-en

LEE, H. & CHO, C.-H. 2020. Digital advertising: present and future prospects. *International Journal of Advertising,* 39(3)**,** 332-341. doi: https://doi.org/10.1080/02650487.2019.1642015

LUCY, L. & BAMMAN, D. 2021. Gender and Representation Bias in GPT-3 Generated Stories. *Proceedings of the Third Workshop on Narrative Understanding***,** 48-55. doi: https://doi.org/10.18653/v1/2021.nuse-1.5

LUTZ, C. 2019. Digital inequalities in the age of artificial intelligence and big data. *Human Behavior and Emerging Technologies,* 1(2)**,** 141-148. doi: https://doi.org/10.1002/hbe2.140

MURO, M., MAXIM, R. & WHITON, J. 2019a. Automation and artificial intelligence: How machines are affecting people and places. *Brookings India.* http://hdl.handle.net/11540/9686.

MURO, M., WHITON, J. & MAXIM, R. 2019b. What jobs are affected by AI? Better-paid, better-educated workers face the most exposure. *Metropolitan Policy Program Report.* http://hdl.voced.edu.au/10707/531845.

OPENAI. 2021. *OpenAI API Documentation* [Online]. Available: https://beta.openai.com/docs/guides/fine-tuning/preparing-your-dataset [Accessed December 2021].

POLCUMPALLY, A. T. 2021. Artificial intelligence and global power structure: understanding through Luhmann's systems theory. *AI & SOCIETY***,** 1-17. doi: https://doi.org/10.1007/s00146-021-01219-8

REED, E. S. 1987. Artificial intelligence, or the mechanization of work. *AI & SOCIETY,* 1(2)**,** 138-143. doi: https://doi.org/10.1007/BF01891277

SHATALOV, O. & RYABOVA, N. 2021. Towards Russian Text Generation Problem Using OpenAI's GPT-2. *Proc. 5th Int. Conf. On Computational Linguistics and Intelligent Systems (COLINS)*, 141-153. doi: https://openarchive.nure.ua/handle/document/19040

SIMMLER, M. & FRISCHKNECHT, R. 2021. A taxonomy of human–machine collaboration: Capturing automation and technical autonomy. *AI & SOCIETY,* 36(1), 239-250. doi: https://doi.org/10.1007/s00146-020-01004-z

SMITH, J. L. 2018. The Effects of Automation. *Quality,* 57(7), 13-13. doi: https://www.proquest.com/openview/024b25b46b88a9a52b8f560cc3d74981/1?pq-origsite=gscholar&cbl=35812

VERMEULEN, B., KESSELHUT, J., PYKA, A. & SAVIOTTI, P. P. 2018. The impact of automation on employment: just the usual structural change? *Sustainability,* 10(5), 1661. doi: https://doi.org/10.3390/su10051661

ZANZOTTO, F. M. 2019. Human-in-the-loop Artificial Intelligence. *Journal of Artificial Intelligence Research,* 64, 243-252. doi: https://doi.org/10.1613/jair.1.11345

Denis Iorga is a Ph.D. Student at the Interdisciplinary School of Doctoral Studies, University of Bucharest. He obtained a bachelor`s degree in Human Resources and a master's degree in Advanced Sociological Research. His academic interests concern topics like Human-Technology Interaction (HTI), Social Network Analysis (SNA), and Sociological Artificial Intelligence Research (SAIR).