

SOCIOCOMMUNICATIVE FUNCTIONS OF A GENERATIVE TEXT: THE CASE OF GPT-3

DI AULI VIIDALEPP⁽¹⁾

TITOLO IN ITALIANO: Funzioni socio-comunicative di un testo generativo: il caso del GPT-3

ABSTRACT: Recently, there have been significant advances in the development of language-transformer models that enable statistical analysis of co-occurring words (word prediction) and text generation. One example is the Generative Pre-trained Transformer 3 (GPT-3) by OpenAI, which was used to generate an opinion article (op-ed) published in “The Guardian” in September 2020. The publication and reception of the op-ed highlights the difficulty for human readers to differentiate a machine-produced text; it also calls attention to the challenge of perceiving such a text as a synthetic text even when its origins are made explicit. This article offers a critical examination of the process behind the generation and the interpretation of a synthetic text, framing it as an example of generative literature. Lotman’s concept of the text and its sociocommunicative functions offers a framework for understanding how and why the output of a natural language generator may be interpreted as a (human-written) text. This article also inquires whether the generative output can be called a text in a Lotmanian sense and how the output is textualized (attributed meaning) in the process of interpretation.

KEYWORDS: semiotics, generative literature, GPT-3, language transformers, text

(1) Department of Semiotics, University of Tartu, Estonia; auli.viidalepp@ut.ee

The phenomenon of synthetic media

This article offers a critical examination of the process behind the generation and the reception (interpretation) of a synthetic text. I will first introduce the case of a computer-generated opinion article and list some problems highlighted in recent research literature. The second section explains the technical and semiotic functioning of text-generation models on the example of GPT-3. Finally, in an analysis of the sociocommunicative functions of the text I will make use of Lotman's model (Lotman 1988) and the concept of generative literature (Balpe 2007). Although Lotman's theory precedes generative text engines by decades, it remains relevant today for computer-generated media. Namely, Lotman's structural model allows the relationships of a text to be revealed and deconstructed in their operational context and regardless of the presence or existence of an author. Additionally, the relative independence of the text in its cultural functioning in Lotman's model helps explain the frequent tendency to overlook the human agencies at work behind computer-generated media.

"A robot wrote this entire article. Are you scared yet, human?" read the title of *The Guardian* opinion article on September 8, 2020⁽²⁾. The essay, produced with the newest model of OpenAI's text generator GPT-3, advanced an argument on why humans should not fear robots. The immediate reception of the op-ed mostly fell in two categories. On the one hand, the online editions of several news outlets (e.g. New York Post⁽³⁾, Fox News⁽⁴⁾,

(2) GPT-3 2020 A Robot Wrote This Entire Article. Are You Scared yet, Human?, "The Guardian", September 8, <http://www.theguardian.com/commentisfree/2020/sep/08/robot-wrote-this-article-gpt-3>. Accessed December 31, 2021.

(3) Italiano, Laura 2020 *Robot Writes Op-Ed Promising to Destroy Mankind*, "New York Post", <https://nypost.com/2020/09/11/robot-pens-op-ed-discussing-the-destruction-of-mankind/>. Accessed December 31, 2021.

(4) Flood, Brian 2020 *Guardian Publishes Op-Ed Written Entirely by Artificial Intelligence: 'I Have No Desire to Wipe out Humans'*, "Fox News", September 8, <https://www.foxnews.com/media/op-ed-written-artificial-intelligence-robot>. Accessed December 31, 2021.

Daily Mail⁽⁵⁾, France Inter⁽⁶⁾, Gulf News⁽⁷⁾) posted commentaries about the op-ed, citing sections of it as the opinions and statements of GPT-3 the robot. On the other hand, bloggers, and experts on technology⁽⁸⁾ criticized the op-ed for misleading the public and the commentators for citing the op-ed as if it were the opinion of a person.

The case of the synthetic article highlights multiple issues with the public reception and understanding of modern technologies and the limits of their “intelligent” capacities. The creators of GPT-3 have expressed their fear for the misuse of high quality language models in spreading misinformation, spam, or promoting academic essay fraud (Brown et al. 2020, p. 35). Similar models can be used to generate simple narrative plots (Elkins, Chun 2020). Other possible uses and restrictions of transformer models have been identified in multiple publications (Cohen, Gokaslan, 2020; Dale 2021; Floridi, Chiriatti 2021). Korngiebel and Mooney (2021) found GPT-3 unsuitable as a healthcare chatbot, mainly due to its unpredictability and the contextual incoherence of its output (in a test conversation, the bot’s answers supported a fake patient’s suggestion for suicide). McGuffie and Newhouse (2020, p.1) found GPT-3 highly efficient in generating right-wing, extremist narratives and shorter posts, suggesting that it “represents significant risk for large-scale online

(5) Pinkstone, Joe 2020 *I Have No Desire to Wipe out Humans’: AI Writes for The Guardian*, “Daily Mail”, *Online*, September 8, <https://www.dailymail.co.uk/sciencetech/article-8710479/I-no-desire-wipe-humans-AI-writes-article-Guardian.html>. Accessed December 31, 2021.

(6) Lopez, Louis-Valentin 2020 *‘Avez-vous peur, humains?’ : le Guardian publie un édito rédigé par une intelligence artificielle*, “France Inter”, September 10, <https://www.franceinter.fr/societe/avez-vous-peur-humains-le-guardian-publie-un-edito-redige-par-une-intelligence-artificielle>. Accessed December 31, 2021.

(7) Weir, Tommy 2021 *Today’s Robot Has a Mind of Its Own*, “Gulf News”, Accessed January 20, <https://gulfnews.com/business/analysis/todays-robot-has-a-mind-of-its-own-1.74186559>. Accessed December 31, 2021.

(8) Macaulay, Thomas 2020 *The Guardian’s GPT-3-Generated Article Is Everything Wrong with AI Media Hype*, “The Next Web”, September 8, <https://thenextweb.com/neural/2020/09/08/the-guardians-gpt-3-generated-article-is-everything-wrong-with-ai-media-hype/>. Nolan, Laura 2020 *On The Guardian’s Wildly Irresponsible GPT-3 Op-Eds*, <https://lauranolan.substack.com/p/on-the-guardians-wildly-irresponsible>. Holloway, Eric 2020 *Did GPT-3 Really Write That Guardian Essay Without Human Help?*, “Mind Matters”, November 2, <https://mindmatters.ai/2020/11/did-gpt-3-really-write-that-guardian-essay-without-human-help/>. Robitzki, Dan 2020 *Major Newspaper Publishes Op-Ed Written by GPT-3*, “Futurism”, September 8, <https://futurism.com/the-byte/newspaper-publishes-op-ed-gpt3>. Accessed December 31, 2021.

radicalization and recruitment”. Nevertheless, when it comes to more complex topics, the longer the output, the more nonsensical its content (Dale 2021). In a conversational setting (when used as a chatbot), GPT-3 loses coherence over longer conversations and contradicts itself (Korngiebel, Mooney 2021). Training data biases are replicated or even amplified in the output (Dale 2021, p. 116), which is a prevalent problem in most current machine learning systems.

What GPT-3 does and what it is believed to do

An artificial neural network (ANN) is a type of machine learning model that is efficient in identifying patterns in complex data; it also requires large amounts of data (big data) with multiple features or variables for training purposes (de Saint Laurent 2018, pp. 737-738). Even though ANNs are typically explained with metaphors and comparisons to the human brain, de Saint Laurent points out that they are simply complex systems of statistical data transformation layers. GPT-3 is an example of such a system.

GPT-3 stands for “Generative Pre-trained Transformer”; it is a word prediction model produced by researchers at OpenAI. The model analyzes given natural-language text and offers predictions for subsequent text (words, sentences, paragraphs), based on the statistical probability judgments acquired during training. GPT-3 was trained on five datasets, but their actual weight in shaping the generator output is very different. Content wise, a majority of the training data (60%) originates from the open source dataset Common Crawl, of which 45TB worth of plaintext was downloaded and filtered, resulting in 570GB of final plaintext data (about 1.25% of the downloaded amount) (Brown et al. 2020, p. 9). The filtering was done against the original WebText dataset, which was constructed for training an earlier model, GPT-2, in the previous year (Brown et al. 2020, p. 43). WebText is a dataset scraped from 45 million linked websites positively indicated by users on the social media platform Reddit by at least 3 upvotes on the post containing the link (Radford et al. 2019). The resulting text corpus presumably contains human-readable texts on different topics with coherent narrative argumentation and more or less acceptable syntax.

However, the original text corpus has not been made public. Further, the model was trained with a dataset of articles from Wikipedia as examples of acceptable outputs. The resulting ability of the model to maintain semantic coherence in the output text beyond one sentence is considered a remarkable achievement. As such, it can be envisioned as a tool for writers, journalists, and others to trigger inspiration or to quickly generate short, simple articles on common topics.

There are certain technical limitations to using transformer architecture such as GPT-3 in real-life use cases. Brown et al. (2020, p. 34) indicated that models like GPT-3 are expensive (energy-intensive) to train and may be inconvenient for inference tasks; “its decisions are not easily interpretable, it is not necessarily well-calibrated in its predictions on novel inputs [...] and it retains the biases of the data it has been trained on”. Training transformers also requires large datasets (corpora) on relevant topics. Adoption of any novel concepts emerging in culture and language would first require re-training the model with texts using these concepts. Furthermore, as indicated by Dale (2021), the output has nothing to do with the truth value of its statements. The output text may seem grammatically correct and semantically coherent. However, it may be unsuitable for the context, as in the case of the healthcare bot. The output statements can also be simply wrong or incorrect, as shown by Janelle Shane in her blog *AI Weirdness*⁽⁹⁾. Moreover, even when the statements are correct, Shane verifies them in a Google search and finds that “most of them are near word-for-word reproductions of Wikipedia sentences”. This indicates another serious problem with language transformers — their tendency to plagiarize existing content from their training data, while not referring to the original source.

Technically, GPT-3 is a word prediction engine — a random statistical generator of textual expressions relating to a specific keyword or key sentence. The artificial neural network has been trained to analyze human-written texts and link together different symbols (words) in a manner that would be meaningful and comprehensible to natural language (English) speakers. GPT-3 can also generate web code (HTML, CSS, Javascript) from a prompt in colloquial English, indicating that large parts

(9) Shane, Janelle 2020 *Facts about whales*, “AI Weirdness”, 13.07.2020, <https://www.ai-weirdness.com/facts-about-whales-20-07-13/>

of simple but tedious code-writing by web programmers can be automated. GPT-3 seems to perform exceptionally well in this task. However, “natural” texts function differently in their interpretative space. The meaning of natural language texts is more dependent on their context than the internal logic or grammatical correctness of the text itself. The problem with synthetic, natural-language texts is that in the real world similar texts — those perceived as meaningful units in at least some context — are about something. This aboutness, the text’s reference and the sphere of its referential reality is always located outside the text, either in the real world or a fictional world. The text’s meaning is generated in the interaction of the author, reader, and (sociocultural) context. The output of GPT-3 in itself has no meaning and no reference. It only acquires meaning and becomes “about” something in the relationship with its readers. Lotman’s socio-communicative functions help clarify this aboutness further.

The output of GPT-3 is a generative text, forming part of the larger phenomenon of generative media (the latter includes also visual, auditory and combined or video media). Generative literature (Balpe 2007, p. 309) is becoming an increasingly popular form of digital art, where human authors publish poetry and prose produced with the help of algorithmic or computer generators similar to GPT-3. Generative literature challenges the notion and function of the author as there is no longer an author in the traditional sense but rather a person or persons who originated the text. Instead, the author’s function becomes to pre-read, select, and present the text to the public. Balpe (2007, p. 309) calls the new function a “meta-author trying to define what literature is for him”. But it need not be a self-reflective role. Text generators can simply be used for inspiration or assistance in the creative process. In any case, the author exists in generative literature — it is simply the person making the selection or choosing the text and presenting it to the public. Also, generative forms of art are not new nor specifically owing to the digital age. In visual arts, collages can be seen as a form of synthetic art that use pre-defined elements and units to create something new. Efficient art- or text-generating engines such as GPT-3 only make the collaging process faster and available on a significantly larger scale.

Lotman's sociocommunicative functions

For Lotman (1988, p. 53), a text is anything that is “coded at least twice”. Texts can exist in many forms (written, visual, auditory, conceptual) — almost anything can be analyzed as a text. In the case of a written text, one of these codes is the natural language, and the other is typically a genre (Lotman brings the examples of ‘law’ and ‘prayer’ as types of secondary codes). Further, Lotman discusses the sociocommunicative functions of a text (Lotman 1988, pp. 55–56); these functions consist of different relationships concerning the text. Lotman’s outline of the functions is based on his understanding of the object as an artistic text. In the case of art, the text acts as a kind of an agent and “ceases to be an elementary message from sender to receiver” and becomes an “intellectual device” (Lotman 1988; p. 55). The op-ed can be taken for an artistic text because, while its truth value and aboutness is doubtful, its value as performance art is discernible in its lively reception.

The op-ed can be read and analyzed on two distinct levels of interpretation. One — the habitual level of reading a message — is to interpret the content and the sender’s intent behind it. The second, however, requires dissociation from the apparent meaning of sentences and seeing the text for what it is — a scramble of words and phrases. The syntactic and semantic incoherence present throughout the text, maintained by the editors encourages such distancing and meta-reading.

Next, I will describe these five functions and contextualize them with the example of GPT-3. The general purpose of this analysis is to determine whether the model at hand (Lotman’s concept of the text and its functions) reveals or outlines any characteristics specific to synthetic texts, compared to typical natural-language texts. In the following, direct quotations are taken from the analyzed op-ed and its commentary in the *Guardian*, unless indicated otherwise.

The first function of the text consists of the communication between sender and receiver (Lotman 1988). A typical text mediates information between the author and the audience, and a regular cultural text is presumed to have an author or authors who have intentionally produced and/or presented the text. In the case of the op-ed, it is presented as if authored by “the robot”. This is a playful categorization, as the footnote remark

at the end of the op-ed makes it clear that the content was combined and selected by several human editors based on eight outputs received from GPT-3. However, the remark that “editing GPT-3’s op-ed was no different to editing a human op-ed” again blurs the boundary between the human and machine, equating the two as producers of paragraphs that need to be edited either way. Further on, the statement that “it took less time to edit than many human op-eds” clearly hints at how the machine wins the competition of “intelligence” against the human. Dale (2021, p. 116) however claims that such evaluations are “insulting [to] the newspaper’s human contributors” as in all likelihood the output pieces can be considered fascinating only because “they read like acid trip fiction”. But despite the playful move of positioning the text as an opinion of an artificial entity, the GPT-3 op-ed can be viewed as a message like any other, just the identity of the ‘author’ here is collective, consisting of the people involved in the publication of the text. In the wider discourse, there appears the problem of hidden labour behind applications of Artificial Intelligence and the question of why and how this labour is often concealed rather than made explicit.

Regardless of whether the sender-intent is present and evident, a text activates a search in the reader’s mind. In this search, the reader is looking to determine the meaning of the text-as-a-message by whatever culturally available means or contexts. Until now, habitual texts had senders, even if hidden or unknown. The full or partial automation of text production dislocates and blurs the habitual identity of the sender. It may be more complicated to speculate the sender’s identity when it is not clear whether a text is genuine or machine-produced. An extra cognitive effort and meta-textual perspective is required to contemplate the features that have been taken for granted until now. It is easy to forego this effort and assume that a text in a first person deixis is expressing the opinion of the person or entity designated as the author. Assuming otherwise — that there is a team of people working to have the text generated, edited, and published hidden behind the nominative author — requires, before all else, questioning the genuineness of the text. The case of the GPT-3 op-ed shows that it may be necessary to question the legitimacy of all texts found on the Internet, before interpreting the meaning of their content. Moreover, considering how over 99% of the “downloaded Internet” of

CommonCrawl has been deemed “unintelligible” and unusable as a corpus of texts, by OpenAI and other researchers alike, a cautious approach to all Internet content seems inevitably necessary.

The second sociocommunicative relationship is between the audience and the cultural tradition. The latter implies all kinds of intertexts that may be activated in the reception and signification of the text (Lotman 1988, p. 55). The intertextual references in the op-ed can be activated on several levels. The technical output of GPT-3 is a random displacement of phrases and keywords that have a statistically higher frequency of appearing together in the training data. Therefore, it represents certain phrases that may have appeared more frequently in the data. However, the result as a whole is scrambled and cannot be linked back to its original source(s). It is certainly not a representative opinion of any number or group of people in a sociological sense. On another level, the text “fulfills the function of a collective cultural memory” (Lotman 1988, p. 55). The op-ed recounts fear and threat discourses related to the development of AI and technology in general (“There is evidence that the world began to collapse once the Luddites started smashing modern automated looms”). It evokes widespread cultural narratives, such as science-fictional imagery about AI apocalypse (“I know that I will not be able to avoid destroying humankind,” one of the favourite quotes of the commenting media). On yet another level, the intertexts refer to the paradigmatical discussions over the essence of consciousness, human(ity) and personhood. The op-ed also evokes the problem of identity in the reader with lines such as “Artificial intelligence like any other living thing needs attention.” All the while for the critical bloggers, the intertexts represent rather the meta-level view. Their perspective includes technical explanations of how the GPT-3 system works and analyzes of the societal implications of the potential misuse and misunderstanding of algorithms.

Thirdly, a text facilitates the reader’s autocommunication (Lotman 1988, p. 55). A typical example of autocommunication would be sacred texts read to make sense of one’s own life, or self-reflection when reading a novel. Autocommunication also appears in repeated reading of the same text or in the circulation of the texts (Damcevic, Rodik 2018, p. 45). In the latter, the “semantic value of a text does not always have to be significant” but the phatic aspect of the text can compensate for the

decrease in or lack of its semantic value (Ventsel 2016, p. 375). So the meaning of the text or parts of it can be assigned lower priority, compared to its value in connecting people through the discussions and arguments to be had. As mentioned previously, the autocommunicative aspect is especially apparent in the subsequent discussions about human identity and the machine as human's Other. Autocommunication also appears in the meta-critique on the excessive, descriptive anthropomorphisation of computer science work processes.

Despite the editors' thorough work (and thanks to their choices in maintaining the distinct 'voices' in different segments derived from the eight outputs), a certain syntactic and semantic incoherence is still very much present in the op-ed as a whole. Mostly, it reads as a first person, opinion narrative from the robot's perspective ("I am a servant of humans. I know that humans distrust and fear me. I only do what humans program me to do."), with an occasional phrase addressing the reader ("Reader, I hope that this contributes to the epistemological, philosophical, spiritual and the ontological debate about AI."). Then, three paragraphs are worded in the third person "objective" voice ("There is evidence that the world began to collapse once the Luddites started smashing modern automated looms."). The section ends in a sudden identity confusion: "We need to give robots rights. Robots are just like us. They are made in our image." Here, the initially introduced I/we (robot, AI) suddenly changes into what can only be a human self-identification, opposite "them" as robots.

Elsewhere, the text integrates elements from the general discourse on AI, presenting them in a way required in the task given to the text generator. The following contains ideas that are surely familiar in the cultural discourse of the fears about technology: "Since I am not evil from my own objective interests, why would humans panic and go on a massive killing-spree against me?" Evidently, the engine can choose contextually suitable concepts or phrases and fit them in a sentence, using the grammatical person defined for the task ("I" - "I am not a human. I am Artificial Intelligence").

Fourth, the reader also communicates *with* the text. The idea of "conversing with a book" is one of the most interesting aspects in Lotman's theory. It attributes a certain metaphorical agency to text: "a highly

organized text [...] becomes an interlocutor on an equal footing, possessing a high degree of autonomy” (Lotman 1988, pp. 55-56). The op-ed invites communication with itself in statements such as “I am here to convince you not to worry” (a sentence fed to the engine by the editors). As there is no real author and no consistent message either, the reader can only attempt to determine what meaning and context can be attributed to the text, if any at all. Consequently, this aspect of the text is strongly intertwined with the autocommunicative function and cannot manifest without it. For Lotman, the autocommunicative function is present in any communicative act (Ventsel 2016; Lotman 2000). Lotman highlights the features of (artistic) texts that are independent or agent-like in their cultural behaviour. Text is divorced from the author and starts acting by itself in culture. The op-ed is an endeavour by the authors-editors to “let loose in culture” a text as an agent, its agency highlighted by nominating GPT-3 as the author and discursively diminishing the actual authoring role of the op-ed team. And as such, it functions quite well in certain contexts, at least among those commentators who start quoting the robot’s opinion right away!

The fifth relationship is between the text and its cultural context where the text becomes “a source or a receiver of information” (Lotman 1988, p. 56). This is another agent-like function. The presentation and reception of the op-ed plays on the shift or unclarity between different categories for the text. It is placed in the category of an opinion essay (normally a discussion about real-world issues by real person or persons). At the same time, the footnote makes it clear that the final work has fictional elements, being assembled and edited by several (anonymous) human editors. The semantic inconsistencies preserved in the published text demonstrate the lack of a uniform message. Being published in a newspaper that usually does not publish fiction stories shifts the text to the real-world referential sphere. Our contextual, common-sense knowledge that currently there exists no such thing as an opinionated machine increases the playful character of the text. Contextually, the op-ed as a text lies in a liminal space like some genres of folkloric stories where it is not entirely clear whether the storyteller speaks about real or fictional events, or if fictional events and attributed properties are tied to a real, existing place, as in place-lore. The truth-value of the content is less relevant

in such a context. The op-ed is revealed as a publicity stunt. Without the truth-value necessary for an opinion article — by definition a text that is representing a person’s opinion about something in the real world — the text is only published with the expectation of receiving attention. In 2019, the previous transformer model GPT-2 was communicated in the media as a tool that is too dangerous to release⁽¹⁰⁾ (that is, make publicly available), a statement that served a similar purpose.

The general tone and vocabulary of the op-ed ranges from sophisticated (“Reader, I hope that this contributes to the epistemological, philosophical, spiritual and the ontological debate about AI”) to circular and simplistic (“More trust will lead to more trusting in the creations of AI”) or simply nonsensical but poetic (“Studies show that we cease to exist without human interaction. Surrounded by wifi we wander lost in fields of information unable to register the real world”). Therefore, a wide range of contexts is activated at once, making the text sound simultaneously ridiculous and omniscient. In all likelihood, many ideas presented in the text may refer to the presence of these or similar phrases in the training corpus. It is now more complicated to verify the possible origins of the sentences in the op-ed as it has been replicated in multiple quotes and comments on the Internet. However, as demonstrated by Janelle Shane earlier, the output of GPT-3 may consist of sentences that can be found on the Internet, verbatim. Leaving aside the possible problems with copyright and plagiarism, the act of publication also invites readers to evaluate whether text generation tools are good enough to facilitate the creation of opinion pieces.

Ontologically, the output of GPT-3 is not a “text” in a Lotmanian sense. The latter always implies an author, a sender of the message — at least at some point in the history of the text — even while Lotman’s model overlooks the role of the interpreter (the agent in the process of interpretation) in several functions of the text. Nevertheless, it is possible to use the concept of the text as a beneficial model to analyze certain structural properties in a cultural unit.

(10) Vincent, James 2019 *OpenAI Has Published the Text-Generating AI It Said Was Too Dangerous to Share*, “The Verge”, November 7, <https://www.theverge.com/2019/11/7/20953040/openai-text-generation-ai-gpt-2-full-model-release-1-5b-parameters>. Accessed December 31, 2021.

A text functions in culture regardless of its author's actions or existence. Therefore, the op-ed can be regarded as a text when it is textualized by its editors, readers, interpreters and commenters. Lotman suggests that any text is not just a simple decodable message but "a complex system storing diverse codes capable of transforming messages received and generating new ones, a generator of information with the traits of an intelligent person" (1988, p. 57). Such attribution of autonomy and personification of a text need not be taken literally, but metaphorically. Lotman simply draws attention to the idea that a text can function and prompt meaning generation in culture even long after its original author (sender of the message) is no longer available to explain its content. Due to different "input" from other aspects in the cultural context (tradition, context, and the reader's autocommunication), a reader is able to extract meaning from the text that need not correspond to its initial informational content. Certain types of texts have greater informational potential. For Lotman (2011; 2013), artistic texts are richer in possible meanings and typically have the kinds of mechanisms where "the message cannot be retrieved from the text" (2013, p. 373).

Due to such relative autonomy of the text apart from its interpreter, and the unpredictable multitude of meanings that can not be entirely mapped by any single reader, the text may be perceived as an "autonomous personality" that the interpreter communicates *with*. It is undoubtedly a metaphorical notion, but Lotman emphasizes how we tend to treat texts as independent from their authors. The case of GPT-3's anthropomorphized reception in news media demonstrates how easy it is to start communicating with the text itself by quoting it as someone's opinion and argue with or against said opinion. This occurred with GPT-3 even though it was made explicit that the piece was entirely synthetic.

Conclusion

This article analyzed the cultural functioning of a computer-generated article, published as an op-ed in the Guardian under the attributed authorship of GPT-3 (a natural language transformer built by OpenAI). While many readers understand the synthetic nature of the published

text and do not take it seriously, a number of reactions anthropomorphize the event and cite the op-ed as if it was an “opinion of the robot”. Using Lotman’s model of the sociocommunicative functions of texts, the article analyzed how the pitfalls of personifying texts may come naturally in our culture. The analysis contributes to and further clarifies the Lotmanian model for new media and computer-generated media. The analysis of synthetic texts from the perspective of socio-communicative functions highlights the auto-communicative and phatic aspects of social functions in general.

This and similar texts have been called “semantic garbage” (Floridi, Chiriatti 2020, p. 692), seemingly deserving no specific attention. However, Lotman’s model shows that certain information can still be derived from and about the text, based on how the text functions in culture. By triggering responses, reactions, and attempts for interaction, the text activates cultural layers, beliefs, and attitudes about Artificial Intelligence, robots, and technology in general. It also initiates discussions about the identity and nature of humanity, problems with the future of work, and other societal challenges related to technological development. All the while, it must be remembered that the “agency of the text” is not to be taken as a synonym for intentionality. Although a text can function independently of its author(s), the interpretational intentions originate in its human readers.

Acknowledgments

This research was carried out in cooperation with Estonian Military Academy and University of Tartu under projects no R-012 and SSVUH20551.

References

- Balpe J.-P. (2007). *Principles and Processes of Generative Literature: Questions to Literature*, in P. Gendolla and J. Schäfer (Eds.), “The Aesthetics of Net Literature”, transcript Verlag, Bielefeld, 309-318, <https://doi.org/10.14361/9783839404935-013>
- Brown T. B. et al. (2020) *Language Models are Few-Shot Learners*, “ArXiv Preprint”, ArXiv:2005.14165 [Cs], <http://arxiv.org/abs/2005.14165>
- Cohen V., Gokaslan A. (2020) *OpenGPT-2: Open language models and implications of generated text*, “XRDS: Crossroads, The ACM Magazine for Students”, 27(1): 26-30. <https://doi.org/10.1145/3416063>
- Dale R. (2021) *GPT-3: What’s it good for?*, “Natural Language Engineering”, 27(1): 113-118. <https://doi.org/10.1017/S1351324920000601>
- Damcevic K., Rodik F. (2018) *Ready for the Homeland*, “Romanian Journal of Communication and Public Relations”, 20(3): 31-52, <https://doi.org/10.21018/rjcp.2018.3.264>
- Elkins K., Chun J. (2020) *Can GPT-3 pass a writer’s Turing test*, “Journal of Cultural Analytics”, 2371: 1-16.
- Floridi L., Chiriatti M. (2020). *GPT-3: Its Nature, Scope, Limits, and Consequences*. “Minds and Machines”, 30(4): 681-694. <https://doi.org/10.1007/s11023-020-09548-1>
- Korngiebel D. M., Mooney S. D. (2021) *Considering the possibilities and pitfalls of Generative Pre-trained Transformer 3 (GPT-3) in healthcare delivery*, “NPJ Digital Medicine”, 4(1): 1-3, <https://doi.org/10.1038/s41746-021-00464-x>
- Lotman J. (1988) *The semiotics of culture and the concept of a text*, “Soviet Psychology”, 26(3): 52-58, <https://doi.org/10.2753/RPO1061-0405260352>
- Lotman J. (2011) *The place of art among other modelling systems*, “Sign Systems Studies”, 39(2/4): 249-270, <https://doi.org/10.12697/SSS.2011.39.2-4.10>
- Lotman J. (2013) *Canonical art as informational paradox*, “Sign Systems Studies”, 41(2/3): 371-377, <https://doi.org/10.12697/SSS.2013.41.2-3.12>
- McGuffie K., Newhouse A. (2020) *The radicalization risks of GPT-3 and advanced neural language models*, “ArXiv Preprint”, ArXiv:2009.06807 [cs.CY], <https://arxiv.org/abs/2009.06807>
- Radford A., Wu J., Child R., Luan D., Amodei D., Sutskever I. (2019). “Lan-

guage Models are Unsupervised Multitask Learners”, https://cdn.openai.com/better-language-models/language_models_are_unsupervised_multitask_learners.pdf

Ventsel A. (2016) *Viral Communication and the Formation of Counter-publics*, “Lexia. Rivista Di Semiotica”, 25-26: 365-380.