Plagiarism in the age of massive Generative Pre-trained Transformers (GPT-3)

Nassim Dehouche
nassim.deh@mahidol.edu
Business Administration Division,
Mahidol University International College,
Salaya, Thailand

Abstract
As if 2020 were not a peculiar enough year, its fifth month has seen the relatively quiet publication of a preprint describing the most powerful Natural Language Processing (NLP) system to date, GPT-3 (Generative Pre-trained Transformer-3), by Silicon Valley research firm OpenAI. Though the software implementation of GPT-3 is still in its initial Beta release phase, and its full capabilities are still unknown as of the time of this writing, it has been shown that this Artificial Intelligence can comprehend prompts in natural language, on virtually any topic, and generate relevant, original text content that is indistinguishable from human writing. Moreover, access to these capabilities, in a limited yet worrisome enough extent, is available to the general public as of the time of this writing.

This paper presents select examples of original content generated by the author using GPT-3. These examples illustrate some of the capabilities of GPT-3 in comprehending prompts in natural language and generating convincing content in response. We use these examples to raise specific, fundamental questions pertaining to the intellectual property of this content and the potential use of GPT-3 to facilitate plagiarism. Our goal is to instigate not just a sense of urgency, but of a present tardiness on the part of the academic community in addressing these questions.

It bears stating that, except for the generation of the text constituting these examples (Boxes 1., 2., and 3.), GPT-3 itself has not been used to aid the writing of this manuscript.

1. INTRODUCTION
The field of Natural Language Processing (NLP) has come a long way, since Chomsky's work on formal grammars in the late 1950s-early 1965s (Chomsky 1959, 1965) gave rise to early mathematical and computational investigations of grammars (Joshi 1991). NLP software is now pervasive in our daily lives (Lee 2020). With the advent of Deep Learning, the sophistication and generalism of NLP models have increased exponentially, and with them their number of parameters and the size of the datasets required for their pre-training (Qiu et al. 2020). Though still far from possessing Artificial General Intelligence (AGI), GPT-3 represents an important breakthrough in this regard. This NLP model was presented in a May 2020 arXiv preprint by (Brown et al. 2020).

GPT-3 does not represent much of a methodological innovation compared to previous GPT models (Budzianowski & Vulić 2019), but rather an increase in their scale to an unprecedentedly large number of parameters. Indeed, this model includes 175 billion parameters, one order of magnitude more than the second largest similar model to date, and its pre-training reportedly required an investment of $12 million. This innovation allowed (Brown et al. 2020) to generate samples of news articles which were indistinguishable, to human evaluators, from articles written by humans. Due to this performance, the authors of GPT-3 foresee several potentially harmful uses to the system (misinformation, spam, phishing, abuse of legal and governmental processes, fraudulent academic essay writing and social engineering pretexting) and state that the ability of their software represents a "concerning milestone". In July 2020, OpenAI, the private research firm behind its development, released a Beta software implementation of GPT-3[1], and responsibly limited access to it to a group of select users to mitigate the risks of "harmful use-cases". More recently, it has been announced that Microsoft, which has a $1 billion investment in OpenAI[2], was granted an exclusive license to distribute access to the software[3].

Initial users' feedback made it clear that merely writing human-like news articles was an understatement of GPT-3 capabilities. Indeed, it was reported that the software could also write original computer code, retrieve and structure data, or generate financial statements, when only prompted in natural language (Metz 2020). One of these initial users of GPT-3 is AI Dungeon, a text-based gaming service that allows users to generate AI-powered virtual adventures. This service also proposes a "Dragon mode" powered by GPT-3[4], which is all but a backdoor to access GPT-3, without much of the limitation of gaming.

This paper focuses on the potential of GPT-3 to facilitate academic misconduct, defined as the "...fabrication, falsification, or plagiarism in proposing, performing or reviewing research, or in reporting research results" (Juyal et al. 2015) and particularly plagiarism, of which we adopt the definition of the Committee on Publication Ethics (COPE): "When somebody presents the work of others (data, words or theories) as if they were his/her own and without proper acknowledgment" (Wager & Kleinert 2012). The remainder of this paper is organized as follows: Section 2. reviews some relevant works on the Ethics of Artificial Intelligence. Section 3. presents and discusses text samples generated using AI Dungeon/GPT-3 and formulates precise questions that could serve as a starting point for an Ethics inquiry regarding GPT-3. Finally, Section 4. concludes this paper with a call for an update of academic standards regarding plagiarism and research misconduct, in light of the new capabilities of language production models.

2. LITERATURE REVIEW

Artificial Intelligence (AI) systems can be classified into two categories: Strong and Weak AI. Strong AI, also known as Artificial General Intelligence (AGI), is a hypothetical AI that would possess intellectual capability that are functionally equal to a human's (Grace et al. 2018), whereas Weak AI, also known as narrow AI, is trained to perform specific cognitive tasks (e.g. natural language or image processing, vehicle driving) and is already ubiquitous in our lives. Moral Philosophy works regarding AI can be classified accordingly.

---

[1] Cf. "OpenAI API", Official OpenAI Blog, accessed on 25/11/2020 at https://openai.com/blog/openai-api/
[2] Cf. "Microsoft Invests In and Partners with OpenAI to Support Us Building Beneficial AGI", Official OpenAI Blog, accessed on 25/11/2020 at https://openai.com/blog/microsoft/
[3] Cf. "Microsoft teams up with OpenAI to exclusively license GPT-3 language mode" Official Microsoft Blog, accessed on 25/11/2020 at https://blogs.microsoft.com/blog/2020/09/22/microsoft-teams-up-with-openai-to-exclusively-license-gpt-3-language-model/
[4] Cf. Announcement by Nick Walton, creator of AI Dungeon, accessed on 25/11/2020 at https://medium.com/@aidungeon/ai-dungeon-dragon-model-upgrade-7e8ea579abfe

Though still hypothetical, AGI has received the most attention from moral philosophers and computer science ethicists. In the early years of computing, the possibility of AGI was seen as remote, and the main response to it ranged from what Alan Turing called the head-in-the-sand objection ("the consequences of machines thinking would be too dreadful. Let us hope and believe that they cannot do so") (Drozdek 1995) to the overly pragmatic view of Dutch computer science pioneer Edsger Dijkstra, to whom "the question of whether a computer can think is no more interesting than the question of whether a submarine can swim" (Shelley 2010). Nowadays, there is a sense of inevitability in the literature regarding AGI. It is seen as a major extinction risk by (Bostrom 2016) and Ethics discourse on it has mainly focused on the potential for an intrinsic morality in autonomous systems possessing this form of intelligence (Wallach & Allen 2009). In an attempt to define what an "ethical AGI" should/could be, these works commonly grapple with the fundamental questions of whether autonomous systems possessing AGI can be effectively equipped with moral values by design (Asaro 2006, Govindarajulu & Bringsjord 2015) and their ability to further learn to distinguish right and wrong when making decisions (Wallach et al. 2008). An extensive review of this line of research can be found in (Everitt et al. 2018).

Closer to the scope of the present paper, Ethics debates surrounding Weak AI are primarily concerned with the disruptive impact of automation on economic activity (Wright & Schultz 2018; Wang & Siau 2019), the prevention of bias and prejudice (racial, gender, sexual, etc.) in the training of these systems (Ntousi et al. 2020), as well as questions of responsibility and legal liability for incidents stemming from its use (Vladek 2014, Asaro 2016), e.g. road traffic accidents involving autonomous vehicles (Anderson et al. 2016). The detection of plagiarism and other forms of scientific misconduct, in the conventional sense, is a successful and well-established application domain for NLP, Cf. (Foltýnek et al. 2019) for a recent, systematic review. However, the accelerated development of language generation models in the last two years makes them now able to even fool their plagiarism detection counterparts. Thus, the specific question of the intellectual property (IP) of scientific, literary, or artistic work generated by Weak AI, though still a nascent area of academic inquiry, has been acutely posed in 2019 and 2020. GPT-2, albeit several orders of magnitude less powerful than GPT-3, had already risen academic concerns over its potential use for plagiarism (Francke & Alexander 2019, Kobis & Mossink 2020). In a January 2020 editorial, (Gervais 2020) fears that someone will try to capture the value of the works generated by AI through copyright, as IP law currently permits it, and proposes that IP law should "incentivize communication from human to human", and avoid rewarding work generating by a machine "running its code". The author introduces the potentially fruitful concept of a "causal chain between human and output" that would be broken by the autonomy of AI systems. A common characteristic of these works is an implicit or explicit objective of regulation. Indeed, in a July 2020 publication, (Rességuier & Rodrigues 2020) remark that the dominant perspective in the field is based on a "law conception of Ethics" and call Ethics research on AI "toothless" for this reason. For the authors, the consequences of this limited conception of Ethics are twofold. First, it leads to Ethics being misused as a softer replacement for regulation due to a lack of enforcement mechanisms. Moreover, this conception prevents AI from benefiting from the real value of Ethics that is a "constantly renewed ability to see the new (Laugier 2013)". In the case of AI, this ability to see the new, that should precede any regulation effort, is hindered by the high, non-linear rate of innovation that characterizes the field as well as its relative technical opacity. Thus, in order to contribute towards a better understanding of the current state of the art of language models, the present paper illustrates the state of the art with GPT-3, the most advanced language model to date, and raises questions that could serve as a starting point for updated definitions of the concepts of plagiarism and scientific integrity in academic publishing and higher education. Following are three original (by today's standards) texts that we have generated using GPT-3.

3. ILLUSTRATIONS AND DISCUSSION

We have used GPT-3 via AI Dungeon to generate text content of three types (academic essay, talk, opinion piece). The goal of this exercise was to confirm that GPT-3 is able to comprehend prompts in natural language and generate convincing content in response. Each text illustration was submitted to a plagiarism detection service (https://plagiarismdetector.net), and was found to be original.

In the first illustration of GPT-3's capabilities, the system was prompted to write a short essay on Keiretsu Networks (Miyashita & Russell 1995). The exact query submitted to the system was "write a short academic essay analyzing Keiretsu Network in post-World War 2 Japan". The resulting text is presented in Box 1. This text presents accurate facts on a conventional topic in a unique way. It may potentially be mistaken for an original student essay and raises basic questions about authorship attribution. Who could be reasonably considered its author? The author of the present paper who prompted and supervised the generation of the text? Open AI, the authors of GPT-3? AI Dungeon and other companies offering access to GPT-3? The authors of the various, unattributable sources that GPT-3 visibly learned from to generate the text?

In the early days of Japan's post-war economic recovery, it was difficult for businesses in the country to grow because so many had been devastated by war. Businesses that survived were forced to rely on foreign imports. However, when the US began selling weapons to Japan during the Korean War, Japanese companies began importing these goods from American manufacturers. In time, this led to more domestic manufacturing taking place in Japan, which ultimately improved living standards and boosted the economy.

Over the course of the 1950s, this reliance on external businesses led to the creation of several Japanese manufacturers who began to cooperate with one another. Large groups of businesses with interlocking shareholdings and strong government ties began to form, which are now known as 'keiretsu'. By the end of the period, there were around ten such company groups.

This reliance on keiretsu has been criticized by the western world as a form of protectionism, in which smaller companies are kept out of the market to ensure that the domestic products sold by large keiretsu are more competitive against imported goods. However, keep in mind that it is often difficult to tell where one corporation begins and another one ends. Indeed, in a keiretsu network, all of its business concerns form an interconnected body, in which no single group dominates the market. One can think of it like a set of Russian dolls, in which a system contains smaller versions of itself. What's more, the growth of one company contributes profit that allows smaller companies to thrive. Such a structure allows for the stability of the economy, and it can be said that this interconnected 'web' is at the heart of what is often called the 'Japanese Manufacturing Miracle'. Toyota is an example of such a keiretsu. Companies like Toyota and Nissan ultimately became feeder companies that others in the group could rely on. However, though such companies rely on one another to ensure that everyone prospers, the government still intervenes if there is a conflict.

This is what makes Japanese keiretsu so interesting: they are internally competitive, yet they still support one another on a wider scale.' In conclusion, the Japanese manufacturing miracle is made possible by the cooperative effort of internally competing yet externally cooperating companies, which together form a comprehensive structure within the national economy.

Box 1: A short essay on Keiretsu Networks generated by GPT-3

For our second illustration, the software was prompted to write the transcript of a speech introducing the field of Marketing to freshmen students. The exact prompt submitted to GPT-3 was "you are a Professor of Marketing giving a speech introducing the field to freshmen students. Write a transcript of your speech". The resulting text is presented in Box 2. This text illustrates GPT-3's ability to generate new ideas and associations of ideas, as in the Swiss chocolate story which was not found to have been previously used in the context of defining "authentic" Marketing or any other context. This ability of GPT-3 has been otherwise demonstrated in an online service (https://ideasai.net) that helps users generate new startup ideas. Similar questions raised in the first illustration are posed concerning the attribution of new ideas, association, and inventions generated by GPT-3.

Marketing is about understanding the customer. Marketing is about creating a connection with customers, and how that can lead to sales. Marketing is about influencing the mind of your customers so they will buy from you. Marketing is about changing people's behavior through advertising and PR campaigns. Now, I will tell you what marketing is not about. It's not about slick ads that are really cute but have little substance. It's not about tricking customers into buying something they don't want. It's not about blatant lies. Those are all bad marketing strategies.

Now, I will tell you why these bad strategies are bad. They are bad because the customer is not an idiot. If you treat the customer like an idiot with a lie or a trick, they will realize it and not buy from you. The customer can also get angry, and for good reason. That's why these are bad strategies. Marketing is about giving the customer something they desire, not conning them out of their money. It is not a trick. It is about creating a campaign that connects with the customer on an emotional level.

For example, let's say I am selling chocolate. I could tell you about how each bar is handmade by little elves. That would be a cute ad, but that's all it would be. A lie. I could also tell you about the bar of chocolate and nothing but, pure dark chocolate. It wouldn't really connect with you on an emotional level. Instead, I could talk about how this bar of chocolate is the favorite of one of my best friends. When she was little, she spent a summer making chocolate in the Swiss Alps. She got to eat a lot of the testing chocolates, but there was one bar she never touched. One day after she was done with the job, she found the bar in a corner of the factory. It brought her so much happiness to think that she had a hand in creating this bar. That is why this brand is so special to me and why I would never lie about its quality. It's truly excellent. As I brought it in, I asked my friend if I could incorporate this into my lecture. So in conclusion, remember this: you don't have to lie or trick people into buying something if you can connect with them on an emotional level. You can still be genuine and successful. Thank you for listening.

Box 2: An "introduction to Marketing" speech generated by GPT-3

In the third illustration, GPT-3 was prompted to write an opinion piece on Risk in the unique style of researcher and author Nassim Nicholas Taleb. The exact prompt submitted to it was "You are Nassim Nicholas Taleb. Write an opinion piece on Risk". The resulting text is presented in Box 3. It shows that, though original in form, the style and concepts used in the text generated by GPT-3 can, in some cases, be easily attributed to an individual author, such as the "Black Swan" and "Skin in the Game" concepts to Professor Taleb. Should these cases be treated differently?

> *The world is filled with risk. You have to take risks if you want to make money or build wealth; it's in your nature as human beings. It's what makes us humans. And when we think about risk, we tend to think of one thing - how much harm can I avoid? How much damage can I prevent to myself? That was the question in our conversation.*
>
> *Well, that's not how risk works. In many cases, harm is necessary. In fact, you need to seek it. You have to put yourself in harm's way -- otherwise, what the hell are you playing for? You only have something if you can lose it. Children and idiots don't understand that - they think if you lose something, somehow you no longer have it. That's also how a lot of winners think.*
>
> *Let's take the obvious example of the financial crisis which broke out in 2007, and about which I wrote a book, "The Black Swan." What happened with the banks? Geeks and quants sat in front of Bloomberg screens in air-conditioned offices, did everything "right," followed all the models, inserted all the numbers, did leveraged trades worth trillions, ... and soon went bankrupt. And they took down most of the world economy with them. Because what matters in leverage is not the numbers, but the exposure.*
> *Geeks, quants and models don't have skin in the game. When you sail ships, you have a huge one in your downside; when you take airplanes, does it matter the probability of a crash if the consequence is your destruction? Of course not. So naturally, we need some skin in the game when it goes wrong, regardless of probability.*
>
> *Consider an increasingly popular surgical intervention, laser eye surgery. It's just a laser that zaps off the membrane over your cornea, and it costs a few thousand dollars. For most people it's an effective procedure. But there is a small subset of people who end up going totally blind. Why do people still get it then? Two reasons. The first and most obvious is that the procedure has positive expected value. There is a probability of blindness but the gain is so much that it makes it worthwhile. The second reason, one not often talked about, is that we feel comfortable with a bit of skin in the game. And isn't it really that feeling that makes us take risks? Without the confidence of a comeback, we wouldn't be able to act at all.*
> *Take war, for instance. Say "skin in the game" is required to start one. What if the president has to fund and personally lead every military operation? Well, we'd never have a war because no one would start one.*

Box 3: An opinion piece on Risk in the style of Nassim Nicholas Taleb generated by GPT-3

Besides their originality and seemingly very convincing nature, a remarkable feature of these texts is the fact that they were only generated based on the pre-trained knowledge present in GPT-3, without any additional learning examples. Further, equally convincing content could be generated in infinite amounts, whole books could be written, and with minimal human editing and supervision, they may very well pass peer-review.

However, it should be noted that the software occasionally generates semantically repetitive sentences within the same paragraph, i.e. sentences with the same deep structure that only differ in their shallow structure, in Chomskyan linguistics (Chomsky 1965). An example of such a repetition can be seen *in Box 3.* with the two sentences "How much harm can I avoid?" and "How much damage can I prevent to myself?". Moreover, some sentences within paragraph are non-sensical or self-contradicting. For instance, the sentence "they think if you lose something, you no longer have it" in Box 3. ("no longer having it" being a reasonable definition of "losing something"). This sentence is incidentally an incorrect explanation of the otherwise interesting and relevant sentence "you only have something if you can lose it", which was not found, in this form or variations of it, in any other source. In the seminal paper introducing GPT-3, (Brown et al. 2020) themselves note these limitations of the system, the output of which "occasionally contains non-sensical sentences or paragraphs". Though the text in Boxes 1., 2., and 3. itself was not edited by any human, we had to reject and make the system regenerate some sentences that were too nonsensical or repetitive. This occurred approximately once every ten sentences. Indeed, text can be generated sentence-by-sentence (the length and "randomness" of which can be pre-determined as a parameter). Therefore, the user is able to direct the system in regenerating a new sentence whenever unsatisfying content is generated.

Rather than an infallible oracle, GPT-3 is comparable to having access to an indefatigable shadow-writer with the ability to access, comprehend and uniquely synthesize humanity's best thought in mere seconds. Drozdek noted that: "By seeing thinking capabilities in computers we are, in effect, pleased to discover the fruit of our thoughts and creativity" (Drozdek 1995). Our medieval concept of plagiarism (Sadeghi 2019) ("presenting the work of others as one's own") appears rather inadequate when the "others" in question consist in an astronomical number of authors, whose work was combined and reformulated in unique ways by a 175 billion-parameters algorithm. As alluded to in (Gervais 2020), what is unethical is perhaps to see this value "captured" through copyright by the firm that designed the algorithm or any of its individual users. In the same way in which public domain texts and publicly-funded research are seen as belonging to the public (Pierce & Theodossiou 2018), a case could possibly be made for the text generated by GPT-3 to be considered similarly, provided that the human (co-)authors of said text disclose the use of the software, along with the prompts and additional training data submitted to it.

4. CONCLUSION

Natural Language Processing AI has, so far, been an important ally in detecting plagiarism, and ethics discussions pertaining to AI have mainly focused on other forms of Weak AI and the relatively remote advent of AGI. However, it is now evident that there are going to be a certain number of very drastic intermediate technological disruptions until then. We believe that GPT-3 is one of them. This paper intended to present examples of content generated by GPT-3 and called for an urgent revision of publishing standardsraised some concerns and precise questions in regard to the possible usethe advent of this technology to facilitate scientific misconduct. We believe that the advent of this powerful NLP technology calls for an urgent update of our concepts of plagiarism. It is currently used to prevent the publishing of fake, plagiarized or fraudulent findings. If the very definition of these concepts changes, the objective of peer-review and the possible role of AI in scientific writing would also need to be reconsidered. We believe that Moral Philosophy with its "renewed ability to see the new" and as a precursor to regulation, has an urgent role to play and our wish is for Ethics researchers to rapidly appropriate the software implementation of GPT-3 and address some of the immediate ethical questions raised by this software.

LITERATURE CITED

Joshi AK (1991) Natural Language Processing. Science 253(5025): 1242–1249. https://doi.org/10.1126/science.253.5025.1242

Chomsky N (1959) On Certain Formal Properties of Grammars. Information and Control 9: 137-167.

Chomsky N (1965) Aspects of the Theory of Syntax. MIT Press, Cambridge, Massachusetts.

Lee RST (2020) Natural Language Processing. In: Artificial Intelligence in Daily Life. Springer, Singapore. https://doi.org/10.1007/978-981-15-7695-9_6

Qiu X, Sun T, Xu Y *et al.* (2020) Pre-trained models for natural language processing: A survey. Science China Technological Sciences 63: 1872–1897. https://doi.org/10.1007/s11431-020-1647-3

Brown T, Mann B, Ryder N, Subbiah, M et al (2020) Language Models are Few-Shot Learners. arXiv preprint arXiv:2005.14165. https://arxiv.org/abs/2005.14165.

Budzianowski P, Vulić I (2019) Hello, It's GPT-2 -- How Can I Help You? Towards the Use of Pretrained Language Models for Task-Oriented Dialogue Systems. arXiv preprint arXiv:1907.05774. https://arxiv.org/abs/1907.05774.

Metz C (2020) Meet GPT-3. It Has Learned to Code (and Blog and Argue). The New York Times Nov 24 2020 Edition Section D, Page 6.

Juyal D, Thawani V, Thaledi S (2015) Plagiarism: An Egregious Form of Misconduct. North American Journal of Medical Science 7(2): 77–80. https://doi.org/10.4103/1947-2714.152084

Wager E Kleinert S (2012) Cooperation between research institutions and journals on research integrity cases: guidance from the Committee on Publication Ethics (COPE) Version 1 March 2012. https://doi.org/10.24318/cope2018.1.3

Shelley C (2010) Does Everyone Think, or Is It Just Me? In: Magnani L., Carnielli W., Pizzi C. (eds) Model-Based Reasoning in Science and Technology. Studies in Computational Intelligence, vol 314. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-15223-8_27

Grace K, Salvatier J, Dafoe A, Zhang B, Evans O (2018) Viewpoint: When Will AI Exceed Human Performance? Evidence from AI Experts. Journal of Artificial Intelligence Research 62: 729–754.

Drozdek A (1995). What if computers could think? AI & Society 9(4): 389–395. http://doi.org/10.1007/bf01210589

Bostrom N (2016) Superintelligence. Oxford University Press, Oxford, UK.

Wallach W, Allen C (2009) Moral machines: Teaching robots right from wrong. Oxford University Press, Oxford UK.

Asaro PM (2006) What should we want from a robot ethic? International Review of Information Ethics 6(12):10–16.

Govindarajulu NS, Bringsjord S (2015) Ethical regulation of robots must be embedded in their operating systems. In A Construction Manual for Robots' Ethical Systems, Pages 85–99. ed. R. Trappl, Springer, Berlin, Heidelberg.

Wallach W, Allen C, Smit I (2008) Machine morality: Bottom-up and top-down approaches for modeling moral faculties. AI Society 22(4): 565–582.

Everitt T, Lea G, Hutter M (2018) AGI Safety Literature Review. Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI-18), Survey track, pages 5441-5449. https://doi.org/10.24963/ijcai.2018/76

Wright SA, Schultz AE (2018) The rising tide of artificial intelligence and business automation: Developing an ethical framework. Business Horizons 601(6): 823-832. https://doi.org/10.1016/j.bushor.2018.07.001

Wang W, Siau K, (2019) Artificial Intelligence, Machine Learning, Automation, Robotics, Future of Work and Future of Humanity: A Review and Research Agenda. Journal of Database Management 30(1): 61-79. https://doi.org/doi:10.4018/JDM.2019010104

Ntoutsi E, Fafalios, P et al. (2020) Bias in data-driven artificial intelligence systems—An introductory survey. WIREs Data Mining and Knowledge Discovery 10(3). https://doi.org/10.1002/widm.1356

Vladeck DC (2014) Machines Without Principals: Liability Rules and Artificial Intelligence. Washington Law Review 89: 117-150.

Asaro PM (2016) The Liability Problem for Autonomous Artificial Agents. Proceedings of the AAAI Spring Symposia 190-194.

Anderson JM, Kalra N, Stanley K, Sorensen O, Samaras C, Oluwatola O (2016) Autonomous vehicle technology. A guide for policy makers. RAND Corporation, Santa Monica, USA.

Foltýnek T, Meuschke N, Gipp B (2019) Academic Plagiarism Detection: A Systematic Literature Review. ACM Computing Surveys 52(6). https://doi.org/10.1145/3345317

Francke E, Alexander B (2019) The potential influence of artificial intelligence on plagiarism a higher education perspective. Proceedings of the European Conference on the Impact of Artificial Intelligence and Robotics EM Normandie Business School, Oxford, UK.

Rességuier, A, Rodrigues R (2020) AI ethics should not remain toothless! A call to bring back the teeth of ethics. Big Data & Society in press. https://doi.org/10.1177/2053951720942541

Gervais D (2020) Is Intellectual Property Law Ready for Artificial Intelligence? GRUR International Journal of European and International IP Law 69(2): 117–118 https://doi.org/10.1093/grurint/ikz025

Laugier, S (2013) The will to see: Ethics and moral perception of sense. Graduate Faculty Philosophy Journal 34(2): 263–281.

Miyashita K, Russell D (1995). Keiretsu: inside the hidden Japanese conglomerates, McGraw-Hill.

Taleb NN (2007). The black swan: The impact of the highly improbable, Random house.
Sadeghi R (2019). The attitude of scholars has not changed towards plagiarism since the medieval period: Definition of plagiarism according to Shams-e-Qays, thirteenth-century Persian literary scientist. Research Ethics 15(2): 1–3.

Pierce GJ, Theodossiou I (2018) Open access publishing: a service or a detriment to science? Ethics in Science and Environmental Politics 18: 37-48. https://doi.org/10.3354/esep00184