

Hey, ChatGPT, how does CSS research need to change?

Research questions in CSS in the wake of the perceived LLM singularity

Olga Zagovora¹ and Ralf Lämmel²

¹GESIS – Leibniz Institute for the Social Sciences

²University of Koblenz, CS Faculty, SoftLang Team

¹olga.zagovora@gesis.org, ²laemmel@uni-koblenz.de

Abstract

CSS and other research fields are affected by LLMs entering common usage and with AGI possibly coming upon us soon enough. Accordingly, we sketch a few new *types* of research questions in CSS. Our sketch is informed by a review of past CSS research challenges to see what role, if any, the shift from genuine human behavior to AI/LLM-based or -enriched behavior has played in CSS research so far. (We are much less concerned here with understanding how CSS research is possibly disrupted, for good and bad, by AI, in how it enables new forms of automation.)

Keywords: Computational Social Science, CSS, Large Language Model, LLM, Chatbot, ChatGPT, AI, AGI

Disruption of CSS research by ChatGPT et al.

Right after November 30, 2022 (ChatGPT launch), grown-up researchers found themselves dumbfounded, e.g., David Maslach (Assoc. Prof. of Strategy, Innovation, and Entrepreneurship, College of Business, Florida State University), certainly no stranger to innovation, put it — perhaps somewhat dramatically — like this: “*My Academic Research Career Is Dead. ChatGPT Did It.*” (Maslach 2022).

Closer to the heart of CSS, Maslach specifically singles out quantitative research, as Eager, a colleague, summarizes Maslach’s video message: “*quantitative researchers will be the first hit. AI will become more and more sophisticated and able to replicate core skills involved in the processes of collecting and analyzing quantitative data and writing up findings in the form of journal articles*” (Eager 2023).

Different ‘coping strategies’ have emerged among researchers since ChatGPT launch. As far as CSS is concerned, reducing it to quantitative research would not be appropriate — given its powerful interaction with social science and its theory. As a kind of technical response to LLMs, AI guru LeCun (LeCun 2023) suggests sensory grounding for LLMs’ meaning and understanding because “*autoregressive LLMs are doomed. They cannot be made factual, non-toxic, etc. They are not controllable.*” and he adds that this is not fixable (because of the exponential increase of the probability of ‘incorrect’ answers with token length). See also elsewhere (Mazzocchi 2015) for a less technical and

more philosophical rejection of the idea of AI killing science, as we know it. Furthermore, ‘advice’ has come along on how research, generally, needs to respond to LLM et al. (van Dis et al. 2023) — much of it is also applicable to CSS.

Importantly, a good part of the potential disruption of CSS research by LLMs is likely a welcome transformation towards more automation. For instance, CSS researchers may benefit from LLM usage in labeling tasks (Gilardi, Alizadeh, and Kubli 2023; Ziems et al. 2023), where they would previously rely on human annotators, who are potentially costly and are hard to get by. The aforementioned ‘advice’ (van Dis et al. 2023) also calls out ‘Embrace the benefits of AI’.

In the rest of the paper, we move our focus from the disruption of CSS research to *CSS research challenges and questions*, as they might relate to the expected shift from genuine human behavior to AI/LLM-based or -enriched behavior.

Past CSS research challenges

CSS is a relatively young research field that was originally defined or identified as an emerging field that “*leverages the capacity to collect and analyze data at a scale that may reveal patterns of individual and group behaviors*” (Lazer et al. 2009).

CSS challenges were called out by some scholars over the years. Watts’ challenges (Watts 2016) were about the social super-collider for combining data sources, the need to expand virtual labs, and paying more attention to social science. Lazer et al.’s challenges (Lazer et al. 2020) were about data sharing, research ethics, and improved incentives.

The potential effect of AI (LLMs) on CSS research — other than perhaps just allowing us to *use* AI (NLP, ML, data science, etc.) in CSS methodologies — was not an issue in those aforementioned ‘challenges papers’.

In a recent paper (Lindgren and Holmström 2020), a social science perspective on AI is presented; the focus is here on social science really, without much, if any, engagement with the adjective ‘computational’ in CSS. Still, the building blocks for the social science perspective on AI proxy for CSS research challenges; the authors suggest: “*i) the interaction between humans and machines must be studied in its broader societal context; ii) technological and human actors must be seen as social actors on equal terms; iii) we*

must consider the broader discursive settings in which AI is socially constructed as a phenomenon [...]; iv) [...] AI, algorithms and datafication affect social science research objects and methods”.

In another recent paper (Wagner et al. 2021), the effects of ‘algorithms’ (e.g., recommendation systems) and the associated platforms for social, political, economic, and scientific processes are discussed with regard to the challenges that social science and CSS face; the challenges focus on ‘measurements’ in the sense of using observable data for reasoning about higher-level social constructs.

New types of CSS research questions

Let us now sketch a few new types of CSS-specific research questions — in response to the arrival of LLMs, as LLM usage is starting to permeate human behavior and with AGI possibly coming upon us soon enough.

We designed types I–III in a way that type I is the most obvious and welcoming type (‘How to find evidence of humans to excel with the help of AI?’), while type II admits that ‘confusion’ may be part of the emerging disruption, and type III covers the fearful singularity prospects and how CSS, if at all, can help.

None of our types I–III of CSS research questions aims at specificity comparable to actual research questions, as we would expect them to underlie specific studies. Instead, these ‘types’ are blueprints (‘themes’) for authoring proper research questions.

I: Study the agency of humans and AIs

There is a new type of research questions targeting the understanding of how humans and AIs complement each other in utterances or other data accessible to CSS research.

If you ask ChatGPT or your colleague, you typically witness encouragement regarding the beneficial agency of humans and AIs. It requires genuine CSS research to reason about and measure that cooperation properly. For whom is it beneficiary under which conditions? How to tell apart those (not) using LLMs? What are the tangible benefits?

If CSS, until now, favored studying mostly genuine human behavior, CSS may now need to embrace ‘modern social machines’ = modern humans as engaging in collaboration with machines (AIs); see also the notion of ‘social machine’, i.e., “*processes in which the people do the creative work and the machine does the administration*” (Berners-Lee and Fischetti 1999; dos Santos Brito et al. 2020), which likely needs an update. (We also mention the actor-network theory (Lindgren and Holmström 2020) as being relevant here in passing.) In the broader context, CSS research needs to further its interconnection with the HCI/CSCW (human-computer interaction and computer-supported cooperative work) research communities (Fede et al. 2022; Zamfirescu-Pereira et al. 2023).

II: Study how humans are challenged by AIs

If the previous type focused on *agency*, then let the present type focus on *challenges* (uncertainty, confusion, other

forms of problematic disruption). The disruption of our society due to the strength of available AI is significant and not all of it is smoothly transformational: job profiles will be eradicated; individuals will experience distress in terms of self-confidence, career and family planning. It requires genuine CSS research to properly reason about and measure the effects on humans being challenged by AIs. Such research is crucial to allow society to respond in a data-driven manner to the mounting challenges. Thereby, CSS would have an immediate impact on the developing AI-related ethics and policy (Zhang et al. 2021; Ryan et al. 2021; van Dis et al. 2023).

While many specific research questions may fit this pattern, the corresponding methodologies are likely to be challenged by some common circumstances. In particular, we need to be able to *tell apart genuine human behavior and AI (-enriched) behavior and to measure the impact of the AI on the human behavior*.

Prior research on bots and trolls can serve as a pattern to be re-instantiated more powerfully for AI ‘intrusion’ or ‘confusion’ (e.g., social bot detection (Färber, Qurdina, and Ahmedi 2019; Yang et al. 2020; Cresci 2020), troll detection (Falher et al. 2017; Im et al. 2020), removal of organizational users, when aiming at individual behavior (Sen et al. 2021), and impact on social bots (Ross et al. 2019), such as the bots succeeding in manipulating humans).

Much of the new ‘AI detection capability’ is likely to be developed outside the CSS community, but the CSS community is in a leading position to put it to good use so that the effects of AIs on human behavior can be thoroughly studied.

III: Study the imminent singularity

Let us now address the AI singularity (or any sort of precursor thereof), which can be seen as a potentially ‘existential’ threat to mankind. Aiming at a type of research question here should come with an expectation of trying to envisage how CSS research could possibly be of help with regard to the threat. Our basic assumption is here that we will continue to see a lot of efforts on AI-related ethics and policy (Zhang et al. 2021; Ryan et al. 2021; van Dis et al. 2023), but it is the genuine responsibility of CSS research to help inform society about the threat, as it continues to build up.

We make our point by referring to the slightly obscure term ‘dark infinities’ (Cai 2021) — a cousin of the emerging singularity. Dark infinities are defined as “*routines of thinking or actions of infinite possibilities which appear common but cannot be completed with limited resources*” (Cai 2021) and thus are potentially very destructive. The interesting twist is that humans (especially with LLMs at their hands) could be as prone to dark infinities as machines. An early and human- and research-centric form of a dark infinity would be this situation:

Researchers end up ‘generating’ (with the so-called help of an LLM) an unmanageable amount of submissions, thereby failing the system of peer review and clogging the necessary process of research informing society.

In fact, think of CSS researchers, in how they focus on quantitative research (Eager 2023; Maslach 2022); they are particularly prone to dark infinities.

While CSS research is not the right field to prevent us — technically — from AI catastrophes to happen, CSS research is best equipped to detect such catastrophes in the making — very well assuming that the AI catastrophe arrives slowly enough to see it coming and to respond in a timely manner. To conclude, we need CSS research that warns us about dark infinities we may be going down into.

Acknowledgment: The authors are grateful to Katrin Weller for feedback on an earlier version of this paper and to Claudia Wagner for input on coverage for prior art regarding the topic.

References

- Berners-Lee, T., and Fischetti, M. 1999. *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web by Its Inventor*. Harper San Francisco, 1st edition.
- Cai, H. 2021. Reaching consensus with human beings through blockchain as an ethical rule of strong artificial intelligence. *AI Ethics* 1(1):55–59. The same author apparently introduced “dark infinities” in [8] of this paper, but this resource cannot be located on the internet, as of writing.
- Cresci, S. 2020. A decade of social bot detection. *Commun. ACM* 63(10):72–83.
- dos Santos Brito, K.; de Lima, A. A.; Ferreira, S. E.; de Aruda Burégio, V. A.; Garcia, V. C.; and de Lemos Meira, S. R. 2020. Evolution of the web of social machines: A systematic review and research challenges. *IEEE Trans. Comput. Soc. Syst.* 7(2):373–388.
- Eager, B. 2023. Will artificial intelligence (AI) take your academic job in 2023? Blog post by Dr. Bronwyn Eager (Senior Lecturer in the Management Discipline in the College of Business and Economics, University of Tasmania) available at <https://broneager.com/ai-and-academic-jobs>.
- Falher, G. L.; Cesa-Bianchi, N.; Gentile, C.; and Vitale, F. 2017. On the troll-trust model for edge sign prediction in social networks. In Singh, A., and Zhu, X. J., eds., *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, AISTATS 2017*, volume 54 of *Proceedings of Machine Learning Research*, 402–411. PMLR.
- Färber, M.; Qurdina, A.; and Ahmedi, L. 2019. Identifying twitter bots using a convolutional neural network. In Cappellato, L.; Ferro, N.; Losada, D. E.; and Müller, H., eds., *Working Notes of CLEF 2019 - Conference and Labs of the Evaluation Forum*, volume 2380 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Fede, G. D.; Rocchesso, D.; Dow, S. P.; and Andolina, S. 2022. The idea machine: Llm-based expansion, rewriting, combination, and suggestion of ideas. In *C&C '22: Creativity and Cognition*, 623–627. ACM.
- Gilardi, F.; Alizadeh, M.; and Kubli, M. 2023. Chatgpt outperforms crowd-workers for text-annotation tasks.
- Im, J.; Chandrasekharan, E.; Sargent, J.; Lighthammer, P.; Denby, T.; Bhargava, A.; Hemphill, L.; Jurgens, D.; and Gilbert, E. 2020. Still out there: Modeling and identifying russian troll accounts on twitter. In Ferrara, E.; Leonard, P.; and Hall, W., eds., *WebSci '20: 12th ACM Conference on Web Science*, 1–10. ACM.
- Lazer, D.; Pentland, A.; Adamic, L.; Aral, S.; Barabási, A.-L.; Brewer, D.; Christakis, N.; Contractor, N.; Fowler, J.; Gutmann, M.; Jebara, T.; King, G.; Macy, M.; Roy, D.; and Alstynne, M. V. 2009. Computational social science. *Science* 323(5915):721–723.
- Lazer, D. M. J.; Pentland, A.; Watts, D. J.; Aral, S.; Athey, S.; Contractor, N.; Freelon, D.; Gonzalez-Bailon, S.; King, G.; Margetts, H.; Nelson, A.; Salganik, M. J.; Strohmaier, M.; Vespignani, A.; and Wagner, C. 2020. Computational social science: Obstacles and opportunities. *Science* 369(6507):1060–1062.
- LeCun, Y. 2023. Do large language models need sensory grounding for meaning and understanding? spoiler: Yes! Presentation slides at Philosophy of Deep Learning Conference and Pre-Conference Debate available at <https://drive.google.com/file/d/1BU5bV3X5w65DwSMapKcsr0ZvrMRUNbi/view>.
- Lindgren, S., and Holmström, J. 2020. A social science perspective on artificial intelligence: Building blocks for a research agenda. *Journal of Digital Research* 2(3):1–15. Part of a special issue <https://jdsr.se/ojs/index.php/jdsr/issue/view/5>.
- Maslach, D. 2022. My academic research career is dead. chatgpt did it. Youtube video by David Maslach (Associate Professor of Strategy, Innovation, and Entrepreneurship, College of Business, Florida State University) available at <https://youtu.be/wMqg-c15Udo>.
- Mazzocchi, F. 2015. Could big data be the end of theory in science? *EMBO reports* 16(10):1250–1255.
- Ross, B.; Pilz, L.; Cabrera, B.; Brachten, F.; Neubaum, G.; and Stieglitz, S. 2019. Are social bots a real threat? an agent-based model of the spiral of silence to analyse the impact of manipulative actors in social networks. *Eur. J. Inf. Syst.* 28(4):394–412.
- Ryan, M.; Antoniou, J.; Brooks, L.; Jiya, T.; Macnish, K.; and Stahl, B. C. 2021. Research and practice of AI ethics: A case study approach juxtaposing academic discourse with organisational reality. *Sci. Eng. Ethics* 27(2):16.
- Sen, I.; Flöck, F.; Weller, K.; Weiß, B.; and Wagner, C. 2021. A Total Error Framework for Digital Traces of Human Behavior on Online Platforms. *Public Opinion Quarterly* 85(S1):399–422.
- van Dis, E. A. M.; Bollen, J.; van Rooij, R.; Zuidema, W.; and Bockting, C. L. 2023. Chatgpt: five priorities for research. *Nature* 614. Comment.
- Wagner, C.; Strohmaier, M.; Olteanu, A.; Kıcıman, E.; Contractor, N.; and Eliassi-Rad, T. 2021. Measuring algorithmically infused societies. *Nature* 595. Perspective.
- Watts, D. 2016. Computational social science: Exciting progress and future challenges. In Krishnapuram, B.; Shah, M.; Smola, A. J.; Aggarwal, C. C.; Shen, D.; and Rastogi, R., eds., *Proc. of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2016*, 419. ACM.
- Yang, K.; Varol, O.; Hui, P.; and Menczer, F. 2020. Scalable and generalizable social bot detection through data selection. In *AAAI 2020, IAAI 2020, EAAI 2020*, 1096–1103. AAAI Press.
- Zamfirescu-Pereira, J. D.; Wong, R. Y.; Hartmann, B.; and Yang, Q. 2023. Why johnny can't prompt: How non-ai experts try (and fail) to design LLM prompts. In Schmidt, A.; Väänänen, K.; Goyal, T.; Kristensson, P. O.; Peters, A.; Mueller, S.; Williamson, J. R.; and Wilson, M. L., eds., *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems, CHI 2023*, 437:1–437:21. ACM.
- Zhang, B.; Anderljung, M.; Kahn, L.; Dreksler, N.; Horowitz, M. C.; and Dafoe, A. 2021. Ethics and governance of artificial intelligence: Evidence from a survey of machine learning researchers. *J. Artif. Intell. Res.* 71:591–666.
- Ziems, C.; Held, W.; Shaikh, O.; Chen, J.; Zhang, Z.; and Yang, D. 2023. Can large language models transform computational social science? <https://arxiv.org/abs/2305.03514>.