GENERATIVE AI IN KNOWLEDGE WORK

Jeffrey V. Nickerson[1], Shiyan Zhang[2]

[1]Stevens Institute of Technology, USA, [2] Stevens Institute of Technology, USA

Keywords: generative AI, journalism, knowledge work

Large language models like GPT-4 and Bard are capable of answering questions and composing text (Brown et al., 2020). Knowledge workers, and in particular journalists, are experimenting with these technologies, and find them useful in some ways but not in others. This paper considers how interaction with these models might be improved, especially by augmenting these language models with other technologies. The focus of this work is on journalists, because in and of itself the profession of journalism is important to society. In addition, the processes journalists engage in are not far from the processes of other kinds of knowledge work. For example, the analytic processes of journalism are similar to the processes used by financial analysts and management consultants, so a better understanding of the effects of AI on journalism may also shed light on many related occupations.

Some information systems researchers have described the combination of artificial intelligence and humans as metahuman systems, different from our normal preconceptions of systems in that not only humans but also machines are learning in the interactions, and learning at different rates (Lyytinen et al., 2021). Other related streams of research have discussed humans in the loop (Cranshaw et al., 2017; Yang et al., 2019) and machines in the loop (Clark et al., 2018; Seidel et al., 2019) as alternative ways to think about the collaborations between people and machines (Cranshaw et al., 2017; Yang et al., 2019).

There is also an economic angle to the proliferation of large-scale language models (Luitse & Denkena, 2021). GPT-3 was a collaboration between OpenAI and Microsoft: Microsoft supplied a supercomputer for training with 285,000 CPUs and 10,000 GPUs (Langston 2020). While exact figures are not available, it is reasonable to infer that machine operational costs alone ran into the millions of dollars, and the human effort involved was also costly in terms of wages for highly trained technologists. Recent models have become more efficient to train, as companies are focusing on the environmental impact of training (Patterson et al., 2021). Still, training industry-competitive models is arguably outside the capabilities of most universities and companies.

In the face of these costs, one strategy is to focus on the meta levels associated with these models. Specifically, we can ask what the design space for these models and their ancillary components looks like, and in which ways we can, through resources within our control, study the impacts of different interface arrangements, drawing on scholarship about digital transformation (Kallinikos et al., 2013; Lyytinen et al., 2021; Seidel et al., 2019). The ideas developed here are applied against an illustrative domain, the changing use of AI in journalism, with examples inspired by a participatory design project the authors are engaged with. The project team has produced a forthcoming paper that documents a system based on top of GPT-3 and its evaluation by journalists (Petridis et al., 2023).

Figure 1 (left) represents most current interactions with large language models. Humans talk directly to one system, in this case GPT-3. On the right of Figure 1 is an alternative, where humans use a complementary technology, a knowledge graph.

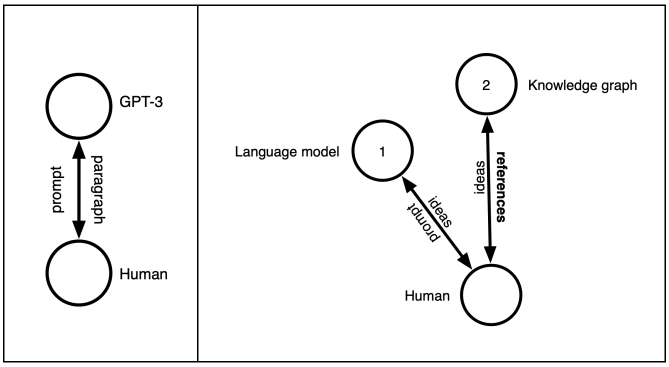

Figure *1. Interaction with a language model. On the left, a direct back and forth with the model. On the right, a kind of mediated interaction, with the human steering the model and then feeding information to a different technology, a knowledge graph.*

Figure 2 shows an extension of Figure 1, incorporating symbolic AI as a third system. In addition, the figure shows that there are many humans involved: those who design, code, and train the systems. Indeed, the effort involved in creating language models, knowledge graphs, and other AI systems is considerable: the cultivation of an AI is an important, difficult, and expensive job (Lyytinen et al., 2020).

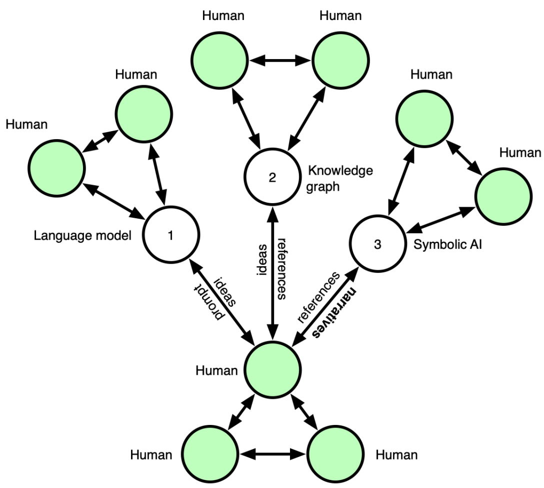

Figure *2. The interaction of the previous figure can be generalized to interaction with three different kinds of machines. Shown in green is human involvement. Not only do humans consume the output of AI, but also teams of humans create and cultivate the technologies.*

With respect to journalism, many different experiments are being tried, exemplified by both academic and industry papers at conferences such as Computing & Journalism (2022). One of the highest impact technologies is relatively simple: template-based composition of news stories using decision trees and related symbolic technologies (Belz, 2019; Diakopoulos, 2019). These technologies are relatively easy to control, and they don't create the off-key stories that language models sometimes do.

News production managers find such tools useful because they allow for a template to be created once in one local market and parameterized with information from hundreds or thousands of other markets, creating stories that are read. The argument in favor of these systems is that they allow local stories to be told that otherwise might not be told. The argument against them is that they can't always capture the context of a location or even the significance of data. The templates take data at face value, and don't usually understand assumptions and changing conditions the same way journalists do.

For journalists, such decision tree technologies are a double-edged sword. They can automate away boring news work, but they can also smooth away the idiosyncratic aspects of local news. The model argued for here would see such decision-tree technologies as one possible component in a larger system. The system would have interface points so that the templates can easily be controlled at different stages, ranging from ideation about what to templatize to recombination with other systems that might forward unusual data for human interpretation and verification.

Verification of news is rarely automatic: it involves reporters who call people up and talk to them, and then call other people up if need be until it becomes clear if there is an angle for readers, and if the information assembled is true. Even something as easy to understand as verification is complex to do, in part because people who make public statements are sometimes motivated to obfuscate facts and intentions. The counter to that obfuscation is probably not machines, but journalists who have time to focus on important stories.

Ideally, journalists get tools that help them dig deeper into stories. Other professions are similar. Analysts of all stripes, in industry, government, and academia, have need to pursue processes of ideation, reference collecting, and narrative construction. Which tools will work for which tasks is not always clear, so either architectures that allow for interaction across systems or architectures that expose enough dials, levers, and intermediate outputs from large systems, are important for the generation of new knowledge.

Interacting with large language models is a different experience than interacting with traditional technology. The natural language interface makes interactions feel close to human, until something jarring is produced. Then it becomes clear that either the models need to drastically improve, or the existing models need to be augmented. This paper has described a kind of architecture that can keep humans in the loop in a news generation task. This use case is important in itself because journalists perform important oversight and explanatory functions in society. The case is potentially generalizable to other knowledge workers, in that many knowledge workers will want to control AI and ground the ideas AI might help generate.

References

Belz, A. (2019). Fully Automatic Journalism: We Need to Talk About Nonfake News Generation. In TTO.

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., … Amodei, D. (2020). Language Models are Few-Shot Learners. In *arXiv [cs.CL]*. arXiv. http://arxiv.org/abs/2005.14165

Clark, E., Ross, A. S., Tan, C., Ji, Y., & Smith, N. A. (2018). Creative Writing with a Machine in the Loop: Case Studies on Slogans and Stories. *23rd International Conference on Intelligent User Interfaces*, 329–340.

Cranshaw, J., Elwany, E., Newman, T., Kocielnik, R., Yu, B., Soni, S., Teevan, J., & Monroy-Hernández, A. (2017). Calendar.Help: Designing a Workflow-Based Scheduling Agent with Humans in the Loop. *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 2382–2393.

Diakopoulos, N. (2019). Towards a design orientation on algorithms and automation in news production. Digital Journalism, 7(8), 1180-1184.

Kallinikos, J., Aaltonen, A., & Marton, A. (2013). The Ambivalent Ontology of Digital Artifacts. *MIS Quarterly*, *37*(2), 357–370.

Luitse, D., & Denkena, W. (2021). The great Transformer: Examining the role of large language models in the political economy of AI. *Big Data & Society*, *8*(2), 20539517211047736.

Lyytinen, K., Nickerson, J. V., & King, J. L. (2021). Metahuman systems = humans+ machines that learn. *Journal of Information Technology*, *36*(4), 427–445.

Patterson, D., Gonzalez, J., Le, Q., Liang, C., Munguia, L.-M., Rothchild, D., So, D., Texier, M., & Dean, J. (2021). Carbon Emissions and Large Neural Network Training. In *arXiv [cs.LG]*. arXiv. http://arxiv.org/abs/2104.10350

Petridis, S., Diakopoulos, N., Crowston, K., Hansen, M., Jastrzebski, Henderson, Nickerson, J. V., & Chilton, L. (2023). AngleKindling: Supporting Journalistic Angle Ideation with Large Language Models. *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*.

Seidel, S., Berente, N., Lindberg, A., Lyytinen, K., & Nickerson, J. V. (2019). Autonomous Tools and Design: A Triple-loop Approach to Human-machine Learning. *Communications of the ACM*, *62*(1), 50–57.

Yang, Y., Kandogan, E., Li, Y., Lasecki, W. S., & Sen, P. (2019). HEIDL: Learning Linguistic Expressions with Deep Learning and Human-in-the-Loop. In *arXiv [cs.CL]*. arXiv. http://arxiv.org/abs/1907.11184