

At EVIA 2023, the panelists answered the question: “Name one paper involving generative LLM’s (either for a task, or for evaluation) that you think everyone should read.” Here are their answers:

1. Chirag Shah and Emily M. Bender, Situating Search, CHIIR, 2022, <https://dl.acm.org/doi/10.1145/3498366.3505816>
2. Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, et al., Llama 2: Open Foundation and Fine-Tuned Chat Models, arXiv:2307.09288 [cs.CL], 2023, <http://arxiv.org/abs/2307.09288>
3. Information Retrieval meets Large Language Models: A strategic report from Chinese IR community, AI Open, vol 4, 2023, <https://www.sciencedirect.com/science/article/pii/S2666651023000049?via%3Dihub>
4. Two answers:
 - a. **Point:** Aher, Arriaga, and Kalai. 2023. Using large language models to simulate multiple humans and replicate human subject studies. arXiv 2208.10264v5 [cs.CL], 2023, <https://arxiv.org/abs/2208.10264v5>
 - b. **Counterpoint:** Durmus et al., 2023. Towards measuring the representation of subjective global opinions in language models. arXiv 2306.16388v1 [cs.CL], 2023, <https://arxiv.org/abs/2306.16388>