

REFERENCES

- Bergmann, D. (2024) *What is fine-tuning?*, IBM. Available at: <https://www.ibm.com/think/topics/fine-tuning>.
- Cao, L. (2024) "Learn to Refuse: Making Large Language Models More Controllable and Reliable through Knowledge Scope Limitation and Refusal Mechanism." Available at: <https://doi.org/10.48550/arXiv.2311.01041>.
- Chen, A. *et al.* (2020) "MOCHA: A Dataset for Training and Evaluating Generative Reading Comprehension Metrics." Available at: <https://doi.org/10.18653/v1/2020.emnlp-main.528>.
- Darwish, A.M., Rashed, E.A. and Khoriba, G. (2025) "Mitigating LLM Hallucinations Using a Multi-Agent Framework," *Information*, 16(7), p. 517. Available at: <https://doi.org/10.3390/info16070517>.
- Dhuliawala, S. *et al.* (2023) "Chain-of-Verification Reduces Hallucination in Large Language Models." Available at: <https://doi.org/10.48550/arXiv.2309.11495>.
- Gao, Y. *et al.* (2023) "Retrieval-Augmented Generation for Large Language Models: A Survey." Available at: <https://doi.org/10.48550/arXiv.2312.10997>.
- Huang, L. *et al.* (2025) "A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions," *ACM Transactions on Information Systems*, 43(2). Available at: <https://doi.org/10.1145/3703155>.
- Hu, E.J. *et al.* (2021) "LoRA: Low-Rank Adaptation of Large Language Models." Available at: <https://doi.org/10.48550/arXiv.2106.09685>.
- Huang, L. *et al.* (2025) "A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions," *ACM Transactions on Information Systems*, 43(2). Available at: <https://doi.org/10.1145/3703155>.
- Ji, Z. *et al.* (2022) "Survey of Hallucination in Natural Language Generation." Available at: <https://doi.org/10.48550/arXiv.2202.03629>.
- Lewis, P. *et al.* (2021) "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks." Available at: <http://arxiv.org/abs/2005.11401>.
- Lv, Q. *et al.* (2024) "Coarse-to-Fine Highlighting: Reducing Knowledge Hallucination in Large Language Models." Available at: <https://doi.org/10.48550/arXiv.2410.15116>.

- Metropolitansky, D. and Larson, J. (2025) “VeriTrail: Closed-Domain Hallucination Detection with Traceability.” Available at: <https://doi.org/10.48550/arXiv.2505.21786>.
- Mucci, T. (2024) *Conversational AI use cases for enterprises, IBM.com*. Available at: <https://www.ibm.com/think/topics/conversational-ai-use-cases> (Accessed: May 16, 2025).
- Naveed, H. *et al.* (2023) “A Comprehensive Overview of Large Language Models.” Available at: <https://doi.org/10.48550/arXiv.2307.06435>.
- Rajpurkar, P. *et al.* (2016) “SQuAD: 100,000+ Questions for Machine Comprehension of Text.” Available at: <https://doi.org/10.48550/arXiv.1606.05250>.
- Vaswani, A. *et al.* (2017) “Attention Is All You Need.” Available at: <https://doi.org/10.48550/arXiv.1706.03762>.
- Wang, C. *et al.* (2023) *Survey on Factuality in Large Language Models: Knowledge, Retrieval and Domain-Specificity, ACM Computing Surveys*. Article. Available at: <https://doi.org/10.48550/arXiv.2310.07521>.
- Xu, L. *et al.* (2023) “Parameter-Efficient Fine-Tuning Methods for Pretrained Language Models: A Critical Review and Assessment.” Available at: <http://arxiv.org/abs/2312.12148>.
- Yang, J. *et al.* (2023) “Harnessing the Power of LLMs in Practice: A Survey on ChatGPT and Beyond.” Available at: <https://doi.org/10.48550/arXiv.2304.13712>.
- Yang, L. *et al.* (2024) “Can bidirectional encoder become the ultimate winner for downstream applications of foundation models?” Available at: <https://doi.org/10.48550/arXiv.2411.18021>.
- Zhang, T. *et al.* (2019) “BERTScore: Evaluating Text Generation with BERT.” Available at: <https://doi.org/10.48550/arXiv.1904.09675>.