

Opportunities for Generative Artificial Intelligence to Accelerate Deployment of Human-Supervised Autonomous Robots

Jason M. Gregory¹, Satyandra K. Gupta²

¹ DEVCOM Army Research Laboratory, Adelphi, MD, USA

² Viterbi School of Engineering, University of Southern California, CA, USA
guptask@usc.edu, jason.m.gregory1.civ@army.mil

Abstract

Autonomous robots have the potential to supplement human capabilities while reducing cognitive and physical burden. However, deploying such systems in natural settings is currently a time-consuming process that revolves around a human’s ability to research, design, test, and evaluate the robot – thereby introducing unnecessary bottlenecks and significant delays to technology adoption. The current work in the field of human-robot interaction (HRI) has historically focused on robot use even though humans play a critical role during autonomous system design and deployment. We argue that the scope of HRI must be expanded beyond that of the current views within the scientific community, to include all phases of system development, deployment, and use. Furthermore, to facilitate the pursuit of this new expanded scope, we present eight opportunities for technological advances in HRI and autonomy using Generative AI that, if realized, could have transformational impact on the fielding of human-supervised autonomous robots. Broadly speaking, our identified opportunities relate to interaction and trustworthiness, collaboration and cooperation, robot motion, robot perception, synthetic scenario generation, testing and evaluation, failure detection, and robot design.

Introduction

Generative Artificial Intelligence refers to a field of artificial intelligence that uses models and techniques for creating new content. These methods are designed to understand and learn the underlying patterns, structures, and characteristics of the training data, which enables the generation of new content. Generative AI utilizes techniques such as deep neural networks to capture and imitate the statistical distribution of the training data. When these models are trained on large datasets they can generate new examples that exhibit similar characteristics to the training data. They have been successfully used to generate a variety of outputs, including images, texts, audio, and videos (Kim et al. 2018; Liu et al. 2021; Aldausari et al. 2022). Recent advances in Generative AI, such as ChatGPT (OpenAI 2023) and DALL-E 2 (Ramesh et al. 2022), are expected to revolutionize many industry sectors and we’ve already seen unprecedented improvements and impact in some fields, including manufac-

turing (Kusiak 2020), education (Pavlik 2023), and medical research (Miljković, Rodríguez-Pérez, and Bajorath 2021).

With the introduction of Generative AI, the future of work is expected to change. To this end, labor shortages are requiring robotic solutions to be deployed rapidly and supplement traditional human work. Unfortunately, the deployment of autonomous robots takes a significant amount of human effort due to the time required to write software, test systems, and build trust and safety assurances. The increasing complexity of robotic systems is complicating this problem and the availability of human expertise can create bottlenecks in robot deployment and delays in technology adoption. Once robotic systems are deployed, they can require significant oversight to ensure safe and proper operations. Generative AI marks a new frontier in human-robot integration and offers opportunities to remove burden from the human so that they can focus more on supervision while maintaining necessary levels of safety.

Historically, the HRI community has focused on robot use and recently is considering design, development, testing, evaluation, and deployment stages. In the context of robotic deployments alongside humans, the field of HRI plays an increasingly important role and requires a new expanded scope to account for an autonomous robot’s full life cycle. Rather than narrowly considering robot use given a productized robotic system, the scientific community must investigate challenges pertaining to the human, robot, and the combined team from target problem conceptualization through fielding and maturation. Since the robotics community has generally been at the forefront of leveraging the latest AI advances, a natural question explored herein is: *how will Generative AI concepts and technologies be used by the robotics community to expand HRI and accelerate the deployment of the next generation of intelligent, human-supervised robots?* There are significant investments in this area due to the enormous amount of interest in Generative AI ideas. Recent efforts are showing early signs of success in using Generative AI in robotics applications (Kumra, Joshi, and Sahin 2020; Ren and Ben-Tzvi 2020; Ha, Agrawal, and Song 2021) and the breadth and applications are expected to grow.

In this work, we present eight opportunities for using Generative AI to realize technological advances in HRI and autonomy that could appreciably accelerate robot deployment and have transformational impact in both the future of

work and expansion of HRI as a field. The authors have more than 40 years of combined experience in designing, testing, and evaluating robotic systems, and the presented opportunities are informed by the author’s most recent works in light of the literature (Gregory et al. 2016, 2019; Al-Hussaini et al. 2020, 2021; Gregory et al. 2022a,b; Al-Hussaini et al. 2022). If the technological advances described herein are realized, we believe that highly-capable, resilient autonomous robots could be deployed alongside humans at increasingly accelerated rates and significantly reduced costs in terms of money and human burden.

Developing Human-Like Responses for Enhanced Interactions and Trustworthiness

The nature and effectiveness of the interaction between the human and robot is central to HRI applications. In many collaborative tasks, robots that exhibit human-like responses and motions can lower training costs because interactions can be more intuitive in nature, and, in some cases, promote trust building. Consider the example of a robot that delivers parts and tools to a human performing a complex maintenance task. Humans may be more comfortable collaborating with this robot if its motions are similar to the movement of a human co-worker because human-like movement aligns with the human’s expectations and experiences from working with other humans. If robots physically operate more like humans and humans gain sufficient trust we posit that humans will also feel more comfortable assuming supervisory roles over robots. Generative AI has already demonstrated capabilities to generate paintings and writings based on the style of a particular artist, such as Dali or Shakespeare (Mihailova 2021; Pataranutaporn et al. 2023). Therefore, we expect that Generative AI can facilitate high-quality imitation learning, where robots learn to mimic human motions and respond accordingly to human interactions. Depending on the target application and robotic platform, humans and robots can have different morphology so exactly imitating human motion will not always be possible or necessary. However, by training models on human-motion data or recorded demonstrations, robots will be able to generate their own sequences of actions that produce human-like behaviors. Therefore, Generative AI can enhance HRI by enabling robots to exhibit more natural and contextually-appropriate responses. By training on human language data, generative models will also be able to generate robot speech responses that are coherent, fluent, and aligned with human communication norms. This has the potential to make robot interactions more engaging, intuitive, and effective. In turn, trustworthy robotic co-workers will likely be deployed ubiquitously once humans grow accustomed to human-like responses. We anticipate that human-like responses will be especially important in social and assistive robotic applications where robots interact with humans of varying age, abilities, and needs. To this end, explainability of the Generative AI and robotic capabilities will be paramount. This has already been identified as a necessary technology for non-experts (Anjomshoae et al. 2019) and will likely improve the quality of interaction, accessibility, and trustworthiness.

Improving Task Planning and Decomposition in Collaborative and Cooperative Settings

Many human-agent teaming applications require robots to plan and perform a complex task or series of tasks that are directed by a human. For example, consider the task of replacing a motor of a cooling fan in a control box. This requires the top-level task to be decomposed into much simpler subtasks and to determine the sequence of tasks for the robot or human-agent team to complete. Once the task sequence is determined, a robot can generate motions for the simple tasks and execute this task. Traditionally, a task planner needs to be developed for each specialized domain and the addition of new objects or processes requires an update to the planner. Tasks planners often struggle to deal with new failure modes and therefore recovering from failures becomes challenging during task planning. Large Language Models (LLM) have become the foundation of many Generative AI techniques and can be used to identify sequences of atomic tasks needed to perform complex tasks (Zhao et al. 2023). With the latest advancements in LLMs a human can pose a query such as, “Provide step-by-step directions to obtain a tool from a locked shelf.” and generate a sequence of various subtasks necessary to perform the overall task. An exemplary set of resulting subtasks includes locating the shelf, reaching the shelf, unlocking the shelf, extracting the tool, locking the shelf, and transporting the tool. Once atomic tasks have been identified, the robot can use a motion planner to generate the motion to execute the task. LLMs can be extremely useful in automatically generating task sequences based on common sense knowledge and using them can eliminate the need for developing domain-specific task planners. Already, there exists some work on leveraging Generative AI to generate plans using embodied, multimodal LLMs (Driess et al. 2023), using symbolic planning (Pallagani et al. 2022), and corrective re-prompting (Raman et al. 2022). However, research suggests that current methods lack fundamental formal and physical world reasoning and are sensitive to prompts (Xie et al. 2023), which limits their applicability to real-world problems. Succinctly put, Generative AI, specifically LLMs, currently excel at tasks requiring formal linguistic competence, but fail on many tests requiring functional competence (Mahowald et al. 2023). While LLMs are still significantly worse than humans in simple common-sense planning domains (Valmeekam et al. 2022), there is tremendous opportunity through the development of assessment frameworks in addition to new techniques and models. Once the scientific community realizes sufficiently performant systems using Generative AI, humans will be able to issue complex requests to robotic teammates and either supervise the task planning and decomposition or have greater control and flexibility over the level of autonomous execution, i.e., shared control (Li et al. 2022).

Generating Robot Motion from Natural Language Description

Robots often need to perform complex motions to successfully execute a task or subtask. Consider the example of

sanding where the robot needs to move the sanding tool in a complex motion pattern to produce a scratch-free surface finish. In the past, if a human expert needed a robot to follow a particular type of sanding tool motion, they would use one of the following methods. The first option would be to teleoperate the robot using a tech pendant and specify the tool motion. The second option would be to demonstrate the tool motion using a handheld tool and use a motion capture system to record the tool motion. The robot then attempts to imitate human motion by analyzing the motion capture data. Third, they can program the motion using the motion commands. Unfortunately, all of these approaches are labor intensive and take significant time for complex parts. This can delay the deployment of robots in new applications. Many process experts would prefer a new modality to generate robot motion based on natural language description of the motion - such an interaction modality aligns more closely with human supervision by allowing the human to only provide intent and desire rather than explicit execution. Generative AI now offers the capability to generate code from the text description, which enables humans to communicate with robots in a more natural, time-efficient manner and automatically create robot motion. This means human experts with no programming experience can interact with robots to perform desired motions. The elimination of specialized robot programming is expected to remove bottlenecks and can significantly speed up task execution. This fundamental capability of generating code from text enables future opportunities in collaborative and cooperative HRI where dialog systems and task learning could be introduced for more intuitive and efficient teaming.

Improving Expressiveness and Prediction of Robotic Perception Capabilities

Robots use perception to build models of their operational environment from which they make decisions and execute autonomous behaviors to perform tasks. Complex environments often create challenges in building complete models due to occlusions and sensor errors. For example, consider the task of a robot scanning a large part to remove the paint. Scanning the part at high speed to obtain fast cycle times may produce some phantom holes in the surface model. These holes can pose challenges for decision-making due to a lack of information. Generative AI can be employed to complete occluded or hidden parts of an object through a technique similar to image inpainting. This involves filling in missing or occluded regions of an image with plausible content that blends seamlessly with the surrounding area (Lugmayr et al. 2022; Rombach et al. 2022). Once the Generative AI model is trained, it can be used to complete occluded or hidden parts of the objects. The incomplete object model is fed into the trained model, which then generates predictions for the missing regions based on the learned patterns and context from the training data. This approach can create a visually coherent result by filling in the missing details and enabling a robot to perform autonomous tasks that might otherwise be halted by poor perception. Without the ability to perceive the world sufficiently well and per-

form autonomous tasks, greater burden and unwanted interactions will be required from the human to tend to the robot.

Generating Synthetic Scenarios for Machine Learning

Machine learning (ML), specifically Reinforcement Learning (RL), has emerged as a useful tool for robots to acquire new skills, which can directly supplement or enhance the human's capabilities. This type of learning most commonly involves the robot training a policy using a simulator and trial-and-error approach. Manually generating a large number of scenarios for RL-based techniques is highly time-consuming and may still not cover all the relevant cases or provide adequate diversity depending on the complexity of environments and tasks to be performed. To illustrate this, consider the case of a legged robot that seeks to learn locomotion on challenging terrains to perform inspection tasks. Robust operations in real-world deployments will typically require a wide variety of challenging terrains to learn robust and efficient gaits.

Generative AI technology can be used to generate distinct, synthetic simulation scenarios that are of interest to human teammates and improve the learning efficiency of RL-based algorithms. By enabling the robots to train on a diverse set of examples and scenarios, Generative AI can enable robots to learn new behaviors, strategies, or responses based on previously-learned patterns and context. This alleviates the human of having to create simulation environments - instead allowing the human to supervise the applicability and deployment of the robot - and accelerates the rate at which robots can develop robust and adaptive behaviors to overcome new situations more effectively.

Generating Test Plans for Testing and Evaluation of Autonomous Robots

Autonomous robots need to be rigorously tested before field deployment alongside humans to ensure safe operations in a wide variety of challenging environments. Consider the case of a mobile manipulator performing wind turbine finishing. This robot needs to be carefully tested to make sure that it will not damage the wind turbine. Manually developing test plans is time-consuming for autonomous robots operating in complex environments. Testing and evaluation is further complicated when autonomous robots make use of sophisticated AI and ML-based techniques.

Based on what is currently being explored in the medical literature for drug discovery (Jain et al. 2022), we believe Generative AI models could learn the structure and logic of test cases for autonomous robots from prior test plans of similar systems and generate new ones. This approach can help researchers, engineers, practitioners, and operators to supervise testing and evaluation by automatically generating diverse and valuable test cases to evaluate different aspects of autonomous robots, such as perception, planning, decision-making, state estimation, control, and coordination.

Simulation has proved to be a useful tool to evaluate autonomous robots and Generative AI could further be em-

ployed to enhance the capabilities of simulation environments used for testing autonomous robots. By generating realistic and diverse synthetic data, such as sensor inputs or environmental conditions, the AI model can augment the existing simulation data and create more challenging and representative test cases to aid in the testing and evaluation of autonomous robot capabilities. Finally, Generative AI models can be used to generate realistic synthetic faults in the behavior of autonomous systems. This enables the evaluation process to assess whether or not the system responds appropriately to faults and remains safe.

Autonomously Detecting and Recovering from Failure for Resilient Human-Agent Teaming

The ability to recover from failure is needed in the field deployment of human-supervised autonomous robotic teams in challenging applications. Whenever testing reveals that the system is not able to recover from failure, the system needs to be improved. This necessitates an ability to efficiently detect failures and take recovery actions. Take, for example, an assembly operation where a previously-installed component is damaged due to a subsequent assembly operation. To proceed with the assembly, the damaged part will need to be replaced through a series of many disassembly steps. Manually coding all possible contingency actions is intractable. By learning the patterns and characteristics of normal operation, the AI model can generate synthetic data that deviates from the norm. These generated anomalies can be used to train a failure detection and recovery system to autonomously detect and respond to unexpected or abnormal situations, helping to improve its robustness and fault tolerance.

Generative models can also learn to generate synthetic sensor data streams that represent typical behavior during normal operation. During robot operation, deviations from the expected patterns may be detected by comparing the generated sensor data stream with the real-time sensor data. This allows robots to identify faults, anomalies, or malfunctions and take appropriate actions to recover or mitigate the impact. Ultimately, autonomous failure detection and recovery can alleviate the human supervisor from having to monitor operations and allow them to focus their effort on higher priority tasks.

Automating Robot Design

As autonomous robots are sought for new and greater roles in real-world deployments with humans, their designs will be constantly evolving to enable more sophisticated behaviors and capabilities. To date, robot design is a time-consuming process that can introduce significant non-recurring engineering costs. For example, many material handling tasks in manufacturing require the use of end-effectors and, depending on the task, these end-effectors need to be customized for optimal performance. In the case of picking and transporting a highly-compliant large sheet, designing a complex end-effector can take multiple days. If significant human effort is needed during the end-effector

design process, the availability of humans can become a bottleneck. Generative AI could eliminate these challenges by designing or aiding in the design of end-effectors so that a human expert only has to approve the design and supervise its use. Generative AI is showing promise in a wide variety of design tasks and therefore it can be utilized to rapidly automate the design of an end-effector for new tasks. By reducing the human's time required to design new robots and parts tailored to a specific task, autonomous robots can be deployed more swiftly.

Conclusions

There are significant opportunities for technological advancements in HRI using Generative AI that could revolutionize the deployment of autonomous robots. Here, we identified eight opportunities that can be described as interaction and trustworthiness, collaboration and cooperation, robot motion, robot perception, synthetic scenario generation, testing and evaluation, failure detection, and robot design. Deploying autonomous robots in complex applications currently requires significant time, effort, and expertise from humans, and we believe investigation into these opportunities will contribute significantly to the reduction or elimination of the current challenges and bottlenecks. Additionally, the pursuit of accelerating deployment using Generative AI will aid in the expansion of the HRI as a field. We argue that there are fundamental HRI challenges throughout the system design, development, testing, evaluation, and deployment stages that are not within current purview of the scientific community, but must be included and prioritized.

References

- Al-Hussaini, S.; Dhanaraj, N.; Gregory, J. M.; Jomy Joseph, R.; Thakar, S.; Shah, B. C.; Marvel, J. A.; and Gupta, S. K. 2022. Seeking Human Help to Manage Plan Failure Risks in Semi-Autonomous Mobile Manipulation. *Journal of Computing and Information Science in Engineering*, 22(5): 050906.
- Al-Hussaini, S.; Gregory, J. M.; Dhanaraj, N.; and Gupta, S. K. 2021. A simulation-based framework for generating alerts for human-supervised multi-robot teams in challenging environments. In *IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, 168–175. IEEE.
- Al-Hussaini, S.; Gregory, J. M.; Guan, Y.; and Gupta, S. K. 2020. Generating alerts to assist with task assignments in human-supervised multi-robot teams operating in challenging environments. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 11245–11252. IEEE.
- Aldausari, N.; Sowmya, A.; Marcus, N.; and Mohammadi, G. 2022. Video generative adversarial networks: a review. *ACM Computing Surveys (CSUR)*, 55(2): 1–25.
- Anjomshoe, S.; Najjar, A.; Calvaresi, D.; and Främling, K. 2019. Explainable agents and robots: Results from a systematic literature review. In *18th International Conference on Autonomous Agents and Multiagent Systems (AA-*

- MAS 2019), Montreal, Canada, May 13–17, 2019, 1078–1088. International Foundation for Autonomous Agents and Multiagent Systems.
- Driess, D.; Xia, F.; Sajjadi, M. S. M.; Lynch, C.; Chowdhery, A.; Ichter, B.; Wahid, A.; Tompson, J.; Vuong, Q.; Yu, T.; Huang, W.; Chebotar, Y.; Sermanet, P.; Duckworth, D.; Levine, S.; Vanhoucke, V.; Hausman, K.; Toussaint, M.; Greff, K.; Zeng, A.; Mordatch, I.; and Florence, P. 2023. PaLM-E: An Embodied Multimodal Language Model. *arXiv preprint arXiv:2303.03378*.
- Gregory, J.; Fink, J.; Stump, E.; Twigg, J.; Rogers, J.; Baran, D.; Fung, N.; and Young, S. 2016. Application of multi-robot systems to disaster-relief scenarios with limited communication. In *Field and Service Robotics: Results of the 10th International Conference*, 639–653. Springer.
- Gregory, J. M.; Al-Hussaini, S.; Agha-mohammadi, A.-a.; and Gupta, S. K. 2022a. Taxonomy of A Decision Support System for Adaptive Experimental Design in Field Robotics. *Artificial Intelligence for Human-Robot Interaction Symposium at AAAI Fall Symposium Series*.
- Gregory, J. M.; Reardon, C.; Lee, K.; White, G.; Ng, K.; and Sims, C. 2019. Enabling intuitive human-robot teaming using augmented reality and gesture control. *Artificial Intelligence for Human-Robot Interaction Symposium at AAAI Fall Symposium Series*.
- Gregory, J. M.; Sahu, D.; Lancaster, E.; Sanchez, F.; Rocks, T.; Kaukeinen, B.; Fink, J.; and Gupta, S. K. 2022b. Active Learning for Testing and Evaluation in Field Robotics: A Case Study in Autonomous, Off-Road Navigation. In *2022 International Conference on Robotics and Automation (ICRA)*, 8217–8223. IEEE.
- Ha, H.; Agrawal, S.; and Song, S. 2021. Fit2Form: 3D generative model for robot gripper form design. In *Conference on Robot Learning*, 176–187. PMLR.
- Jain, M.; Bengio, E.; Hernandez-Garcia, A.; Rector-Brooks, J.; Dossou, B. F.; Ekbote, C. A.; Fu, J.; Zhang, T.; Kilgour, M.; Zhang, D.; et al. 2022. Biological sequence design with gflownets. In *International Conference on Machine Learning*, 9786–9801. PMLR.
- Kim, S.; Lee, S.-g.; Song, J.; Kim, J.; and Yoon, S. 2018. FloWaveNet: A generative flow for raw audio. *International Conference on Machine Learning*.
- Kumra, S.; Joshi, S.; and Sahin, F. 2020. Antipodal robotic grasping using generative residual convolutional neural network. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 9626–9633. IEEE.
- Kusiak, A. 2020. Convolutional and generative adversarial neural networks in manufacturing. *International Journal of Production Research*, 58(5): 1594–1604.
- Li, Q.; Peng, Z.; Wu, H.; Feng, L.; and Zhou, B. 2022. Human-AI Shared Control via Policy Dissection. *Advances in Neural Information Processing Systems*, 35: 8853–8867.
- Liu, M.-Y.; Huang, X.; Yu, J.; Wang, T.-C.; and Mallya, A. 2021. Generative adversarial networks for image and video synthesis: Algorithms and applications. *Proceedings of the IEEE*, 109(5): 839–862.
- Lugmayr, A.; Danelljan, M.; Romero, A.; Yu, F.; Timofte, R.; and Van Gool, L. 2022. RePaint: Inpainting using Denoising Diffusion Probabilistic Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11461–11471.
- Mahowald, K.; Ivanova, A. A.; Blank, I. A.; Kanwisher, N.; Tenenbaum, J. B.; and Fedorenko, E. 2023. Dissociating language and thought in large language models: a cognitive perspective. *arXiv preprint arXiv:2301.06627*.
- Mihailova, M. 2021. To dally with Dalí: Deepfake (Inter) faces in the art museum. *Convergence*, 27(4): 882–898.
- Miljković, F.; Rodríguez-Pérez, R.; and Bajorath, J. 2021. Impact of artificial intelligence on compound discovery, design, and synthesis. *ACS omega*, 6(49): 33293–33299.
- OpenAI. 2023. GPT-4 Technical Report. *arXiv:2303.08774*.
- Pallagani, V.; Muppasani, B.; Murugesan, K.; Rossi, F.; Horesh, L.; Srivastava, B.; Fabiano, F.; and Loreggia, A. 2022. Plansformer: Generating Symbolic Plans using Transformers. *arXiv preprint arXiv:2212.08681*.
- Pataranutaporn, P.; Danry, V.; Blanchard, L.; Thakral, L.; Ohsugi, N.; Maes, P.; and Sra, M. 2023. Living Memories: AI-Generated Characters as Digital Mementos. In *Proceedings of the 28th International Conference on Intelligent User Interfaces*, 889–901.
- Pavlik, J. V. 2023. Collaborating with ChatGPT: Considering the implications of generative artificial intelligence for journalism and media education. *Journalism & Mass Communication Educator*, 78(1): 84–93.
- Raman, S. S.; Cohen, V.; Rosen, E.; Idrees, I.; Paulius, D.; and Tellex, S. 2022. Planning with Large Language Models via Corrective Re-prompting. *Foundation Models for Decision Making Workshop at Neural Information Processing Systems*.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*.
- Ren, H.; and Ben-Tzvi, P. 2020. Learning inverse kinematics and dynamics of a robotic manipulator using generative adversarial networks. *Robotics and Autonomous Systems*, 124: 103386.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Valmeekam, K.; Olmo, A.; Sreedharan, S.; and Kambhampati, S. 2022. Large Language Models Still Can’t Plan (A Benchmark for LLMs on Planning and Reasoning about Change). *arXiv preprint arXiv:2206.10498*.
- Xie, Y.; Yu, C.; Zhu, T.; Bai, J.; Gong, Z.; and Soh, H. 2023. Translating Natural Language to Planning Goals with Large-Language Models. *arXiv preprint arXiv:2302.05128*.
- Zhao, W. X.; Zhou, K.; Li, J.; Tang, T.; Wang, X.; Hou, Y.; Min, Y.; Zhang, B.; Zhang, J.; Dong, Z.; et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223*.