

Based on the research questions and other related questions, several key areas need further work for AI to achieve its objectives:

- **Reasoning:** Achieving human-level reasoning remains a significant challenge. There are many open questions on the best way to get there.
- **Factuality and Trustworthiness:** Ensuring AI systems are factual and trustworthy is crucial. The report indicates that the community does not believe these problems will be solved soon.
- **Agents:** Effectively integrating LLM-driven agents into multi-agent systems requires solving challenges related to collaboration, negotiation, scalability, transparency, and computational efficiency.
- **Evaluation:** There's a need for more rigorous evaluation methods for AI systems. This includes developing better ways to monitor deployed systems, evaluate the safety of agentic AI, increase transparency, and address human engagement.
- **Ethics and Safety:** Addressing ethical and safety concerns, including issues like misinformation, privacy, and responsibility, is essential.
- **Embodiment:** Developing intelligent situated agents with human-level performance across a broad range of tasks remains a key challenge.
- **Cognitive Science:** Further research is needed to develop human-like learning and reasoning methods, expand cognitive architectures, and improve social agents.
- **Hardware:** Overcoming hardware limitations, such as memory capacity, compute throughput, and power consumption, is crucial for advancing AI.
- **Social Good:** Designing contextually relevant AI systems, overcoming data limitations, ensuring sustainability, and

developing robust evaluation frameworks are critical for AI to effectively address social challenges.

- **Sustainability:** There is a need to minimize the energy and resource demands of AI, develop innovative applications for sustainability, and improve data collection and modeling to guide sustainable AI development.
- **Scientific Discovery:** Key challenges include improving communication between AI and human scientists, defining the hypothesis space, and addressing issues related to data accuracy, noise, and reproducibility.
- **Artificial General Intelligence (AGI):** Achieving AGI requires progress in areas such as developing architectures beyond transformers, improving long-term planning and reasoning, enhancing generalization, enabling continual learning, and ensuring alignment, interpretability, and safety.
- **Societal Influences:** More research is needed to understand and guide AI's impact on people and society.
- **Diversity of Approaches:** The report emphasizes the importance of supporting a diversity of research approaches, including non-neural methods, interdisciplinary collaboration, and new paradigms.
- **Interdisciplinary Research:** Expanding AI research to include diverse perspectives from other disciplines is crucial for addressing the societal, ethical, and cultural implications of AI.
- **Role of Academia:** Universities need to adapt to the era of "big AI" by finding their role in AI research, addressing challenges in hiring and retaining AI talent, and collaborating with industry.
- **Geopolitical Aspects:** Developing international governance frameworks, addressing geopolitical risks, and promoting ethical AI development are essential for navigating the global

implications of AI.