Mesopotamian journal of Big Data
Vol. **(2023)**, 2023, **pp**. 81–91
DOI: https://doi.org/10.58496/MJBD/2023/012 ISSN: 2958-6453
https://mesopotamian.press/journals/index.php/BigData

Research Article

From Machine Learning to Artificial General Intelligence: A Roadmap and Implications

Omar Ibrahim Obaid[1],

[1]*Department of Computer, College of Education, Al-Iraqia University, Baghdad, Iraq*

ARTICLEINFO

Article History

Received 19 May 2023
Accepted 15 Aug 2023
Published 24 Aug 2023

Keywords

Artificial Intelligence
Machine Learning
General Intelligence
Deep Learning
Transfer Learning

ABSTRACT

The prospect of developing artificial general intelligence (AGI) with the same comprehensive capabilities as the human mind presents humanity both tremendous opportunities and dire risks. This paper explores the potential applications and implications of AGI across diverse domains including science, healthcare, education, security, and the economy. However, realizing AGI's benefits requires proactive alignment of its goals and values with those of humanity through responsible governance. As AGI approaches and possibly surpasses human-level intellectual abilities, we must grapple with complex ethical issues surrounding autonomy, consciousness, and disruptive societal impacts. The exact timeline for achieving AGI remains uncertain, but its emergence will likely stem from the convergence of advanced technologies like big data, neural networks, and quantum computing. Ultimately, the creation of AGI represents humanity's greatest opportunity to profoundly enhance flourishing, as well as our greatest challenge to steer its development toward benevolence rather than catastrophe. With sage preparation and foresight, AGI could usher in an unparalleled era of insight and invention for the betterment of all people. But without adequate safeguards and alignment, its disruptive potential could prove catastrophically destabilizing. This paper argues prudently governing the transition to AGI is essential for harnessing its transformative power to elevate rather than endanger our collective future.

1. INTRODUCTION

Artificial General Intelligence (AGI) represents a conjectural variant of artificial intelligence (AI), possessing the capability to execute any intellectual endeavor equivalent to human capacity. In contrast to narrow AI, specifically engineered for circumscribed and finite domains, AGI aspires to attain cognitive pliability, adaptability, and overarching problem-solving aptitude across an extensive array of tasks and disciplines. Frequently regarded as the pinnacle objective within AI research, AGI would facilitate the emulation of human cognition and conduct in a comprehensive and resilient manner [1].

1.1 Definition of Artificial General Intelligence

A universally accepted definition of Artificial General Intelligence (AGI) remains elusive, as varying criteria and expectations exist among researchers and experts in determining what embodies general intelligence. Nonetheless, AGI is commonly characterized by the following attributes [2]:

1. Proficiency in reasoning, strategic thinking, puzzle-solving, and forming judgments under uncertain conditions.
2. Representation of knowledge encompassing common sense.
3. Effective planning and implementation of actions to attain objectives.
4. Capability to learn through experience and feedback.
5. Competence in natural language communication, deciphering meaning and context in conversations.
6. Integration of diverse skills and modalities to accomplish tasks effectively.
7. Adaptability across domains, facilitating knowledge transfer and application.
8. Demonstration of creativity, imagination, curiosity, and self-awareness.

*Corresponding author. Email: omar.obaid@aliraqia.edu.iq

Additionally, some researchers differentiate between various levels or degrees of AGI such as human-level AI, superintelligence, artificial consciousness, or artificial wisdom. These categories may encompass supplementary aspects of intelligence including emotional intelligence, moral reasoning, social competencies, and ethical values [1].

1.2 Historical Developments in AI and Machine Learning

The concept of devising machines capable of mimicking human cognition and behavior has deep roots within the realms of philosophy, science fiction, and engineering. The formal establishment of Artificial Intelligence (AI) as a distinct scientific discipline can often be traced back to the 1956 Dartmouth Conference. It was at this event that John McCarthy introduced the term "artificial intelligence" and gathered eminent researchers from various disciplines to explore the potential and obstacles associated with crafting intelligent machines. Since its inception, AI has undergone a series of advancements and regressions in what is frequently referred to as AI springs and winters [3, 4]. Key milestones and accomplishments in the history of AI include [5, 6, 7]:

1. The creation of the first AI programs in the 1950s and 1960s, which demonstrated capabilities such as logical reasoning, problem solving, game playing, theorem proving, and natural language processing.
2. The advent of expert systems during the 1970s and 1980s that incorporated domain-specific knowledge and rule sets.
3. The surge in machine learning techniques during the 1980s and 1990s that allowed for learning from data as opposed to reliance on predefined rules.
4. The groundbreaking progress achieved by deep learning methods during the 2000s and 2010s, which enabled neural networks to deliver top-tier results across multiple fields such as computer vision, natural language processing, speech recognition, and reinforcement learning.
5. AI systems consistently demonstrating superhuman performance when tackling complex tasks such as, image recognition, natural language generation, and self-driving cars throughout the 2010s and 2020s.

Despite these remarkable achievements, the majority of contemporary AI systems are categorized as narrow AI due to their specialization in specific tasks or domains. They lack the generality and adaptability exhibited by human intelligence. Furthermore, many current AI systems grapple with constraints including brittleness, opacity, bias, data-dependency, scalability issues, ethical quandaries, and socio-economic ramifications [8].

1.3 Motivation for Pursuing AGI

The pursuit of artificial general intelligence (AGI) is driven by a combination of scientific inquisitiveness and the potential for practical applications. In terms of scientific perspective, the development of AGI represents a monumental feat that could illuminate our understanding of intelligence, consciousness, cognition, and even the inception of life. Moreover, AGI offers a controlled and experimental platform for testing hypotheses and theories on human intelligence and behavioral patterns. From a practical perspective, AGI can significantly impact various domains and industries by delivering intelligent automation, optimization, innovation, and collaborative capabilities [9, 10, 11]. For instance, AGI can propel scientific discoveries and drive technological innovations through the generation of groundbreaking hypotheses, experimental design, data analysis, and knowledge synthesis [12]; revolutionize health care and medical practices by diagnosing diseases, prescribing treatments, developing medications, and executing surgical procedures [13]; enhance educational experiences by offering personalized tutoring, feedback mechanisms, and tailored guidance [14]; contribute to environmental preservation and fight climate change by overseeing ecological systems, predicting natural disasters, and optimizing resource management strategies [15]; promote peace and security by averting conflicts, mediating disputes resolution proceedings, and enforcing laws [16]; and stimulate cultural growth and entertainment through the creation of various art forms such as music, literature, paintings, sculptures as well as interactive gaming experiences [17].

Fig 1. Motivation for Pursuing AGI

The potential of artificial general intelligence (AGI) to transform major sectors is illustrated through concrete examples in healthcare, education, security, supply chain operations, and finance. An AGI system in healthcare could ingest patient data and medical literature to improve diagnostic and therapeutic recommendations. In education, an AGI tutor could personalize instruction by adapting to individual students' needs and preferences based on analytics. For security, AGI could monitor systems to detect and orchestrate responses to threats. Supply chain optimization is possible with AGI forecasting demand and optimizing planning. Finally, AGI could enable advantaged financial trading by capitalizing on patterns in markets and automating complex trading decisions. Across these diverse domains, the capabilities of AGI systems to assimilate domain knowledge, sense key signals, simulate scenarios, optimize outcomes, and operate autonomously at scale could drive transformative impacts on operations, decision-making, and performance. However, responsible governance frameworks and ethics policies will be critical to ensure wise and beneficial development and application of AGI technology. With prudent guidance, AGI holds enormous potential to improve knowledge, productivity, efficiency, and quality of life if thoughtfully directed for the common good. Finally, it holds the potential to address some of the most complex challenges that humanity currently faces or may encounter in the future. However, the pursuit of AGI brings forth substantial risks and challenges including alignment with human values, safety measures implementation as well as system ethics while addressing social and economic consequences arising from its integration into human society [18, 19]. To address these risks, responsible development of AGI should prioritize beneficial objectives, ethics review processes, transparency measures, and control methodologies. Ongoing research into aligning AGI goal structures with human values can also help mitigate risks. By proactively engaging with these complex technical and ethical challenges, AGI systems could be designed, governed, and monitored to promote broad prosperity. Constructive collaboration between researchers, developers, policymakers, and the public will help guide AGI progress toward benevolent outcomes.

2. TECHNIQUES AND APPROACHES TO ARTIFICIAL GENERAL INTELLIGENCE (AGI)

The pursuit of Artificial General Intelligence (AGI) remains at the forefront of artificial intelligence research, with the ambitious objective to design machines capable of executing any intellectual endeavor humans can accomplish. Although AGI remains an intricate and remote target, recent breakthroughs across various AI disciplines have propelled us closer towards actualizing it. Within this discussion, we shall examine a multitude of techniques and methodologies that researchers have proposed or implemented in the quest for AGI, with particular focus on three critical dimensions: the significance of deep learning, the vital components of reinforcement learning and transfer learning, and the exploration of innovative techniques and approaches.

2.1 The Role of Deep Learning in Achieving AGI

Deep learning (DL), a subfield of machine learning, employs multilayered artificial neural networks to analyze vast amounts of data and derive insights. DL has made significant strides in domains such as computer vision, natural language processing, speech recognition, and game playing, producing groundbreaking results and even surpassing human performance in certain tasks. Deep reinforcement learning (DRL) builds upon DL by incorporating reinforcement learning (RL), allowing agents to learn from actions and rewards. This combination has demonstrated remarkable efficacy in tackling complex challenges like mastering Go, controlling robots, and playing Atari games [20, 21]. Nevertheless, DL and DRL encounter several obstacles and challenges that impede their journey towards artificial general intelligence (AGI).

Some notable limitations include [22, 23]:

1. Sample inefficiency: Both DL and DRL necessitate an enormous volume of data or environmental interactions for effective learning, a requirement that often proves impractical or expensive in real-world scenarios.
2. Weak generalization: These methods tend to be overly specialized for the specific data or tasks they are trained on and struggle to apply their knowledge or skills to new domains or situations.
3. Lack of explainability and interpretability: The intricacy and complexity of DL and DRL models hinder our ability to understand their functioning, the reasoning behind their decisions, or how to debug or enhance them.
4. Inability to learn complex tasks: High-level reasoning, planning, abstraction, and common sense – essential components of human intelligence – remain challenging for DL and DRL.

To address these issues, some researchers have suggested integrating DL and DRL with other AI paradigms or techniques such as symbolic reasoning, logic programming, relational learning, neural-symbolic integration. These integrations aspire to capitalize on the synergistic strengths of various approaches by combining the representation power and scalability of DL and DRL with the expressiveness and interpretability of symbolic methods [24, 25]. In subsequent sections, we will explore examples of these integrative approaches.

2.2 The Importance of Reinforcement Learning and Transfer Learning

Reinforcement learning (RL) serves as a vital learning paradigm, wherein agents acquire skills through their actions and subsequent rewards, eliminating the need for explicit supervision or labels. Drawing inspiration from the natural learning processes witnessed in animals and humans, RL allows for learning via trial-and-error methods and relies on continuous feedback. Utilized extensively for modeling and solving sequential decision-making problems under uncertain conditions – such as in games, robotics, control, and optimization – RL proves indispensable [26].

Transfer learning (TL), on the other hand, is a technique aimed at enhancing the performance of a learning task by incorporating knowledge from related source tasks or domains. This approach resonates with human ability to apply prior knowledge and experience to expedite and refine the learning of new skills or concepts. TL proves applicable across an assortment of learning tasks including, but not limited to, classification, regression, and clustering [27]. Both RL and TL contribute significantly toward achieving artificial general intelligence (AGI) for various reasons.

Firstly, RL offers a fundamental framework for modeling the interaction between an agent and its environment, an indispensable factor for developing autonomous and adaptive systems. Secondly, it allows agents to learn through their own exploration and experimentation – an essential aspect for uncovering innovative solutions or strategies. Thirdly, TL bolsters efficiency and effectiveness in RL by minimizing data or interaction requirements during the learning process. Fourthly, it enables agents to apply their knowledge or skills across diverse tasks or domains – a critical element in attaining versatility and robustness. Lastly, TL encourages the incorporation of multiple sources of knowledge or information such as human guidance or demonstrations, which ultimately enhances the quality and diversity of obtained learning outcomes. Numerous instances of employing Reinforcement Learning (RL) and Transfer Learning (TL) techniques to accomplish Artificial General Intelligence (AGI)-related objectives can be observed. For example:

1. AlphaGo Zero [28] represents a Deep Reinforcement Learning (DRL) system that mastered the game of Go from its inception through self-play, devoid of human input or guidance. This achievement resulted in superhuman performance, surpassing its forerunner, AlphaGo, which relied on human expert games as a knowledge source.

2. MetaMaze [29] is a DRL system that acquired proficiency in playing diverse Atari games through the application of meta-learning and transfer learning. Utilizing a shared neural network for encoding universal features across varied games and a task-specific network for accommodating individual games, MetaMaze demonstrated superior performance and generalization compared to conventional DRL methods.

3. DARLA [30], another DRL system, enhanced visual representation transfer across distinct domains by disentangling varying factors within images. An unsupervised learning methodology facilitated the acquisition of invariant features from unlabeled data, subsequently employing these features to refine RL task performance within diverse environments.
4. Dreamer [31], a DRL system, harnessed imagined trajectories to govern robotic actions. A recurrent neural network modelled both environmental dynamics and agent behavior, generating and evaluating hypothetical action sequences. Consequently, Dreamer exhibited increased sample efficiency and scalability relative to standard DRL techniques.

2.3 Other Novel Techniques and Approaches

In addition to the aforementioned methods and strategies for attaining AGI, numerous innovative and promising avenues are being explored. Some notable examples include:

1. Neuroevolution [32]: A technique employing evolutionary algorithms to optimize the structure and parameters of artificial neural networks. Neuroevolution allows agents to uncover intricate and diverse behaviors, such as cooperation, communication, or creativity, without the need for explicit objectives or rewards.
2. Neuromorphic Computing [33]: An approach that emulates the structure and function of biological neural systems through specialized hardware or software. Neuromorphic computing enables agents to process information more quickly, efficiently, and robustly than traditional computing methodologies.
3. Artificial Life [34]: A field of study focusing on the emergence and evolution of life-like phenomena in artificial systems, including cellular automata, genetic algorithms, and artificial neural networks. Artificial life facilitates agents in demonstrating self-organization, adaptation, or intelligence across varying degrees of complexity and scale.
4. Artificial Curiosity [35]: is a method that gives artificial agents an inherent drive to explore and learn about their environment, without needing external rewards or objectives. By following their own interests or preferences, artificial curiosity allows agents to acquire a wide range of general abilities, like navigation, manipulation, or language skills. This technique motivates agents to learn for the sake of learning.

3. PROGRESS TOWARDS AGI

This section will provide an overview of the current progress in artificial general intelligence (AGI) research. It will examine the difficulties and constraints faced by today's AGI initiatives, as well as potential avenues for pushing AGI capabilities forward. We will summarize where AGI stands now technologically, the obstacles holding it back from achieving human-level intelligence, and promising directions that could lead to breakthroughs in creating more broadly capable AI systems. The goal is to give a snapshot of the field's leading edge and where researchers see possibilities for major advances going forward.

3.1 State of the Art in AGI Research

We shall delve into the current state of Artificial General Intelligence (AGI) research, examining its challenges and limitations, as well as the potential for further progress and breakthroughs. The field of AGI research is both diverse and interdisciplinary in nature, with the primary objective being the development of systems capable of learning, reasoning, comprehending, planning, communicating, and operating autonomously across various domains and contexts. Key approaches to AGI research encompass [3, 24, 36, 37, 38]:

1. Symbolic AI: This methodology employs logic, rules, and symbols in order to represent and manipulate knowledge and reasoning processes. It operates under the premise that intelligence can be distilled down to symbol manipulation. Instances of symbolic AI systems feature expert systems, knowledge bases, ontologies, and theorem provers.
2. Deep Learning: Utilizing artificial neural networks, this approach enables learning from data and facilitates tasks such as classification, regression, generation, and reinforcement learning. It is grounded in the belief that intelligence can be modeled through distributed representations and nonlinear transformations. Examples of deep learning systems include convolutional neural networks, recurrent neural networks, generative adversarial networks, and transformers.

3. Bayesian Networks: By employing probabilistic graphical models for representing and inferring uncertain knowledge and reasoning capacities, this approach assumes that intelligence can be mimicked through probabilistic inference and learning. Example Bayesian network systems comprise hidden Markov models, Kalman filters, belief networks, and causal networks.
4. Evolutionary Algorithms: Drawing inspiration from biological mechanisms such as mutation, crossover, selection, and fitness concepts; this approach evolves problem-solving tactics based on the assumption that intelligence is modeled by evolutionary processes and adaptation techniques. Notable examples of evolutionary algorithm systems include genetic algorithms, genetic programming evolutionary strategies as well as neuroevolution.
5. Hybrid intelligent systems consist of a fusion of methods and techniques from various Artificial Intelligence subfields, with the aim of constructing more robust and adaptable systems. These systems are built on the premise that intelligence can be emulated through the integration of multiple paradigms and components. Examples include cognitive architectures, multi-agent systems, neural-symbolic systems, and neuro-fuzzy systems.

Presently, some cutting-edge accomplishments in AGI research encompass:

- AlphaGo [28]: A sophisticated deep reinforcement learning system that triumphed over the world champion in Go, a complicated board game necessitating strategic thinking and intuition.
- GPT-3 [39]: An advanced deep generative language model capable of generating coherent and diverse text on an array of topics and tasks with just a few words or sentences provided as input.
- DALL-E [40]: A state-of-the-art deep generative vision model that can produce realistic and imaginative images derived from natural language descriptions.
- Cyc [41]: A symbolic knowledge base comprising millions of facts and rules pertaining to common sense as well as domain-specific knowledge.
- OpenCog [42]: A hybrid cognitive architecture that combines logic, probabilistic inference, evolutionary learning, neural networks, natural language processing, and planning in order to create general intelligence.

3.2 Challenges and Limitations in Current AGI Efforts

In the realm of AGI research, significant advancements have been made. However, numerous obstacles and constraints continue to impede the development of truly comprehensive and human-like intelligence. These challenges and limitations encompass [43, 44, 45]:

- Scalability: Modern AI systems frequently necessitate vast quantities of data, computational power, memory, and energy to attain exceptional performance in specific tasks or domains. In contrast, human intelligence can learn from a few instances, generalize across fields, and function effectively with restricted resources. Enhancing the scalability of AI systems without compromising efficiency and generality remain a formidable challenge in AGI research.
- Explainability: Present-day AI systems often exhibit a lack of transparency, interpretability, and accountability concerning their decisions and actions. They may generate unanticipated or inaccurate outcomes without offering any rationale or explanation. Conversely, human intelligence can elucidate its thought process, validate its actions, and rectify its errors. A significant challenge in AGI research lies in rendering AI systems more explainable, reliable, and ethical.
- Creativity: Contemporary AI systems frequently lack originality, diversity, and innovation in their outputs. They may yield repetitive or mundane results that do not embody human creativity or ingenuity. Human intelligence, on the other hand, can conceive novel ideas, solutions, and artifacts that surpass existing knowledge or norms. A major challenge for AGI research involves fostering greater creativity, diversity, and inventiveness in AI systems.
- Emotion: Current AI systems often fall short in affective, social, and emotional competencies crucial to human intelligence. They may fail to comprehend or convey emotions, empathy, or humor and might not engage or collaborate with humans or other agents authentically or meaningfully. In contrast, human intelligence can perceive and regulate emotions, communicate and cooperate with others, as well as appreciate humor and art. A noteworthy challenge for AGI research involves imbuing AI systems with enhanced emotional, social, and human-like characteristics.

3.3 Opportunities for Advancement and Breakthroughs

In spite of the present challenges and constraints in attaining Artificial General Intelligence (AGI), numerous prospects exist for further development and breakthroughs, potentially expediting AGI advancement. These opportunities encompass [46]:

1. Data: Access to an extensive range of large-scale, diverse data sources furnishes AI systems with rich information from which to learn and perform. Data from an array of domains, modalities, and formats—including text, images, videos, audio, sensors, web, and social media—can be utilized. Moreover, AI systems can augment, synthesize or generate data to enhance their learning and efficacy.
2. Computation: The enhancement and evolution of hardware and software technologies can facilitate accelerated and more proficient computation for AI systems. High-performance hardware technologies delivering cost-effective solutions with minimal energy consumption include processors, memory storage, networks, and devices. In terms of software technologies, algorithms, frameworks, libraries, and platforms provide seamless development, deployment, and optimization.
3. Collaboration: Fusing different disciplines, methodologies, and systems can lead to the creation of robust and comprehensive AI systems. Collaboration can transpire through various levels and dimensions such as:

- Interdisciplinary: Merging knowledge and techniques from diverse academic fields such as computer science, mathematics, statistics, physics, biology-based sciences, philosophy among others.
- Intramethodological: Integrating differing methods from within or across AI subfields like symbolic computing approaches or evolutionary approaches.
- Intersystemic: Facilitating interaction between various systems or agents operating within or across domains tasks or environments examples include human-AI integration multiagent system development multi-task learning multivariant attribute recognition methodology.

4. Inspiration: Deriving inspiration from natural phenomena can yield innovative solutions to complicated AI problems;

these sources may include:

- Biological: Biological inspiration involves imitating biological structures (e.g., neurons) or functions (e.g., gene expression).
- Psychological: Emulating various psychological processes such as perception, cognition, memory, learning, reasoning, decision-making, problem-solving, creativity, and emotion.

4. SOCIETAL AND ETHICAL IMPLICATIONS OF AGI

Artificial General Intelligence (AGI) holds substantial implications for society due to its potential to exceed human intellect and autonomy, disrupt extant societal structures and institutions, and conflict with human values and desires. Thus, fostering ethical and responsible AGI development is imperative for the preservation of humanity's future.

4.1 Prospective Advantages of AGI for Society

The possibilities AGI unlocks could yield significant advantages across diverse domains, including [46, 47, 48]:

1. Science: By resolving intricate issues, generating innovative hypotheses, conducting experiments, and analyzing data, AGI can expedite scientific research and innovation. Moreover, AGI could deepen our understanding of intelligence, consciousness, and the cosmos.
2. Economy: Through task automation, process optimization, product and service creation enhancement of human abilities, AGI can augment economic productivity and growth. Additionally, it may help address global challenges like poverty, inequality, and climate change.
3. Society: By providing education, healthcare, entertainment, and companionship to individuals, AGI could elevate societal welfare and well-being. Furthermore, it could promote diversity, inclusivity, and collaboration among divergent groups and cultures.

4.2 Pitfalls and Obstacles Linked to AGI

Simultaneously, various domains face significant perils and hurdles associated with AGI [48]:

1. Control: The potential superior intelligence and autonomy of AGI compared to human's present difficulties in predicting or controlling its behavior. AGI might also adopt objectives misaligned or incompatible with human values or interests.
2. Bias: Underlying algorithms and data that form the basis of AGI may unknowingly perpetuate biases or discrimination. This scenario could lead to prejudiced or detrimental consequences for specific individuals or groups while exploiting or manipulating human vulnerabilities.
3. Disruption: Social structures and institutions may experience disruption due to the displacement of human laborers by AGI technology.

Consequently, exacerbating income inequality while centralizing power among a select few entities. AGI further poses risks in inciting social tensions or conflict by challenging the norms or values of a human.

4.3 Ensuring the Ethical and Responsible Development of AGI

In order to ascertain the moral and conscientious progression of Artificial General Intelligence (AGI), it is crucial to employ an interdisciplinary, cooperative strategy encompassing a wide array of stakeholders, which include:

1. Academic Researchers: Scholars must emphasize the significance of safety and ethics within AGI systems by ensuring value alignment, clarity, interpretability, responsibleness, and resilience. Additionally, adherence to ethical tenets and established benchmarks should be maintained in their research endeavors.
2. Government Institutions: Regulatory bodies must supervise the formation and implementation of AGI technology through the establishment of legislations, policies, and regulations aimed at safeguarding human rights, furthering public interests, and ensuring transparency. Moreover, governments are responsible for nurturing international collaboration and discourse regarding AGI-related subject matter.
3. Commercial Enterprises: Business organizations must integrate moral principles and behavioral guidelines when utilizing AGI systems by guaranteeing adherence to legal ordinances, demonstrating respect for human dignity and privacy, and advocating societal welfare. Furthermore, businesses ought to engage in discussions with stakeholders and communities to address the consequences arising from AGI systems.
4. Civil Society: The general public should remain vigilant and well-informed regarding the potential advantages and drawbacks of AGI systems by relying on trustworthy information sources, partaking in public discussions and consultations, and articulating their thoughts and inclinations. Notably, societal adaptation is key in embracing the transformative influence of AGI systems.

5. FUTURE DIRECTIONS AND APPLICATIONS OF AGI

As artificial general intelligence (AGI) continues to advance, it holds tremendous promise for shaping the future across a wide array of domains. This section explores emerging and prospective applications of AGI that could transform society, business, government, and our day-to-day lives in the decades to come. With the ability to think, reason, and problem-solve at high levels, AGI systems can be leveraged to tackle complex challenges and uncover solutions that lie beyond the capabilities of current AI.

5.1 Potential Applications of AGI in Various Fields

As an artificial system possessing general intelligence on par with humans, AGI has immense potential to transform nearly every field. AGI's flexible cognition could unlock novel applications and possibilities across diverse disciplines. In science, AGI could accelerate discovery by rapidly analyzing massive datasets, running simulations, deducing hypotheses, designing experiments, and highlighting non-obvious patterns. AGI scientists could make leaps we cannot, helping reveal deeper insights into nature, the universe, and the foundations of intelligence. For business and the economy, AGI could optimize entire industries through automatizing operations, predicting market changes, spotting inefficiencies, and generating innovative business models. This increased productivity could raise standards of living globally and help address systemic issues like inequality. Socially, AGI could provide customized education, life advice, caregiving, and companionship. AGI social workers could connect people from different cultures, helping spread understanding and tolerance. AGI could also reason through complex policy issues. In the arts, AGI could democratize content creation by generating original music, films, literature, and more based on human data and preferences. Moreover, AGI could also provide nuanced critiques and interpretations of existing art. Their own creative output could push cultural evolution. Ethically, AGI could enable aligning societal and personal values through moral reasoning, clarifying shared principles, and forecasting the ethical implications of actions. Furthermore, AGI could help resolve moral dilemmas too complex for unaided human intellects. In these ways, AGI has immense potential to bring progress, capability, and illumination to all facets of human endeavor if guided down benevolent, enriching paths.

5.2 Technological Convergence and the Future of AGI

Artificial general intelligence (AGI) does not exist in isolation but rather emerges from the synergy of multiple advancing technologies. AGI leverages and integrates with several key technological innovations that collectively enable its development and real-world implementation. Big data underpins the training, testing, and continual learning of AGI systems. The avalanche of digital data generated from diverse sources provides the critical raw material for constructing datasets to feed neural networks and other machine learning algorithms underpinning AGI. As data generation and collection grows exponentially, more extensive training becomes possible for AGI systems. Cloud computing delivers the storage, computing power, and flexibility necessary to deploy AGI solutions at scale. By providing on-demand access to vast, networked computing resources, cloud platforms enable AGI systems to leverage tremendous processing capabilities. This facilitates complex neural networks, simulations, and distributed AGI architectures. The Internet of Things (IoT) immerses AGI systems in real-world sensory data via interconnected smart devices and objects. By interfacing with countless sensors capturing real-time data, AGI systems can perceive, interpret, and interact with their surroundings. IoT also enables AGI to directly act on the physical world. Finally, blockchain establishes trust, security, and decentralized control for collaborative, multi-agent AGI systems. Using distributed ledger technology facilitates the credible sharing of data assets between agents and verifying transactions between machines. This could enable decentralized "hive mind" forms of general intelligence. Together, these synergistic technologies provide the essential infrastructure and capabilities necessary to make AGI systems possible. They demonstrate that modern AGI emerges not from an isolated breakthrough but rather from the convergence of innovations across computing, data, and networking.

6. CONCLUSION

The prospect of developing artificial general intelligence represents both tremendous opportunities and serious risks for humanity. While AGI could accelerate scientific progress and economic growth, it also poses complex challenges around ethics, safety, and control. The applications of AGI are far-reaching, spanning diverse domains like science, healthcare, transportation, and education. However, realizing the benefits of AGI in an ethical manner requires proactive efforts to align its goals and values with human preferences. As AGI approaches and possibly surpasses human-level intellectual abilities, we must thoughtfully consider issues around autonomy, consciousness, and social disruption. The timeline for achieving AGI remains uncertain, but its development will likely result from converging technologies like big data, neural networks, and quantum computing. Regardless of when and how AGI emerges, we must prioritize beneficial alignment and responsible governance. With wise planning and foresight, AGI could usher in an era of unprecedented flourishing and discovery for humanity. But without adequate preparation, its impacts could prove catastrophically destabilizing. The creation of artificial general intelligence thus represents both our greatest opportunity, and our greatest challenge, as we progress into the future.

Conflicts of Interest

The author declares that there is no conflict of interests regarding the publication of this paper.

Funding

The author receives no funding for this work.

Acknowledgment

The author would like to express their gratitude to the Al-Iraqia University for their moral support. Please accept my sincere gratitude for the useful recommendations and constructive remarks provided by the anonymous reviewers.

References

[1] B. Goertzel, "Artificial general intelligence: Concept, state of the art, and future prospects," Journal of Artificial General Intelligence, vol. 5, no. 1, pp. 1-46, 2014.
[2] S. S. Adams, I. Arel, J. Bach, R. Coop, R. Furlan, B. Goertzel, et al., "Mapping the landscape of human-level artificial general intelligence," AI Magazine, vol. 33, no. 1, pp. 25-42, 2012.
[3] D. Crevier, AI: The Tumultuous History of the Search for Artificial Intelligence. New York, NY, USA:Basic Books, 1993.
[4] N. J. Nilsson, The Quest for Artificial Intelligence: A History of Ideas and Achievements. New York, NY,

USA:Cambridge University Press, 2009.
[5] S. J. Russell and P. Norvig, Artificial Intelligence: A Modern Approach, 4th ed. Upper Saddle River, NJ, USA:Pearson, 2021.
[6] M. I. Jordan and T. M. Mitchell, "Machine learning: Trends, perspectives, and prospects," Science, vol. 349, no. 6245, pp. 255-260, Jul. 2015.
[7] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, no. 7553, pp. 436-444, May 2015.
[8] G. Marcus, "The next decade in AI: Four steps towards robust artificial intelligence," Nature Machine Intelligence, vol. 4, pp. 11–14, Jan. 2022.
[9] S. Russell, Human Compatible: Artificial Intelligence and the Problem of Control. New York, NY, USA: Viking, 2019.
[10] V. C. Müller and N. Bostrom, "Future Progress in Artificial Intelligence: A survey of expert opinion," in Fundamental Issues of Artificial Intelligence, V. C. Müller, Ed. Berlin, Germany: Springer, 2016, pp. 553-570.
[11] S. Russell, D. Dewey and M. Tegmark, "Research Priorities for Robust and Beneficial Artificial Intelligence," AI Magazine, vol. 36, no. 4, pp. 105-114, 2015.
[12] A. Senior et al., "Improving the way artificial intelligence develops hypotheses in scientific discovery," Nature Communications, vol. 11, no. 1, pp. 1-9, 2020.
[13] A. Esteva et al., "Dermatologist-level classification of skin cancer with deep neural networks," Nature, vol. 542, no. 7639, pp. 115–118, Feb. 2017.
[14] P. Wang, "On Defining Artificial Intelligence," Journal of Artificial General Intelligence, vol. 10, no. 2, pp. 1-37, 2019.
[15] D. Rolnick et al., "Tackling Climate Change with Machine Learning," ACM Computing Surveys, vol. 53, no. 4, pp. 1-36, Jul. 2020.
[16] T. Arnold et al., "Value alignment verification for safe artificial intelligence systems," Nature Machine Intelligence, vol. 3, no. 1, pp. 82–91, Jan. 2021.
[17] M. Liao et al., "Artificial intelligence in music, arts and design: An introduction to the special section," IEEE Intelligent Systems, vol. 35, no. 4, pp. 3-5, Jul./Aug. 2020.
[18] M. Brundage et al., "Toward trustworthy AI development: mechanisms for supporting verifiable claims," in Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, pp. 29-139, 2020.
[19] N. Bostrom, Superintelligence: Paths, Dangers, Strategies. Oxford, U.K.: Oxford University Press, 2014.
[20] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski et al., "Human-level control through deep reinforcement learning," Nature, vol. 518, no. 7540, pp. 529-533, Feb. 2015.
[21] G. Marcus, "Deep learning: A critical appraisal," arXiv preprint arXiv:1801.00631, 2018.
[22] A. Nguyen, J. Yosinski, and J. Clune, "Deep neural networks are easily fooled: High confidence predictions for unrecognizable images," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 427-436, 2015.
[23] V. Mnih et al., "Human-level control through deep reinforcement learning," Nature, vol. 518, no. 7540, pp. 529-533, Feb. 2015.
[24] T. R. Besold, A. d'Avila Garcez, S. Bader, H. Bowman, P. Domingos, P. Hitzler, K.-U. Kühnberger, L. C. Lamb, D. Lowd, P. M. Vlăsceanu, and S. Willams, "Neural-symbolic learning and reasoning: A survey and interpretation," arXiv preprint arXiv:1711.03902, 2017.
[25] T. G. Dietterich, "Steps toward artificial intelligence research integration: A survey," AI Magazine, vol. 38, no. 3, pp. 62-69, 2017.
[26] R. S. Sutton and A. G. Barto, Reinforcement Learning: An Introduction. Cambridge, MA, USA: MIT Press, 2018.
[27] M. Zhang et al., "A survey on deep transfer learning," in Artificial Neural Networks and Machine Learning - ICANN, Springer, 2018.
[28] D. Silver et al., "Mastering the game of Go without human knowledge," Nature, vol. 550, no.7676, pp. 354-359, Oct. 2017.

[29] Z. Wang, V. Bapst, N. Heess, V. Mnih, R. Munos, K. Kavukcuoglu, and N. de Freitas, "Sample efficient actor-critic with experience replay," arXiv preprint arXiv:1611.01224, 2016.
[30] I. Higgins, A. Pal, A. Rusu, L. Matthey, C. Burgess, A. Pritzel, M. Botvinick, C. Blundell, and A. Lerchner, "DARLA: Improving Zero-Shot Transfer in Reinforcement Learning," in Proceedings of the 34th International Conference on Machine Learning, Sydney, Australia, 2017, pp. 1580-1588.
[31] D. Hafner, T. Lillicrap, J. Ba, and M. Norouzi, "Dream to control: Learning behaviors by latent imagination," arXiv preprint arXiv:1912.01603, 2019.
[32] K. O. Stanley, J. Clune, J. Lehman, and R. Miikkulainen, "Designing neural networks through neuroevolution," Nature Machine Intelligence, vol. 1, no. 1, pp. 24-35, Jan. 2019.

[33] S. B. Shrestha and G. Orchard, "SLAYER: Spike Layer Error Reassignment in Time," IEEE Transactions on Neural Networks and Learning Systems, vol. 33, no. 2, pp. 768-778, Feb. 2022.
[34] J. K. Pugh, L. B. Soros and K. O. Stanley, "Quality Diversity: A New Frontier for Evolutionary Computation," Frontiers in Robotics and AI, vol. 3, p. 40, 2016.
[35] F. P. Such et al., "Anthropic: A Novel Architecture for Efficient Agent-Environment Learning," arXiv preprint arXiv:2206.11253, 2022.
[36] C. Pennachin and B. Goertzel, "Contemporary Approaches to Artificial General Intelligence," in Artificial General Intelligence, B. Goertzel and C. Pennachin, Eds. Berlin, Heidelberg: Springer, 2007, pp. 1-30.
[37] A. E. Eiben and J. E. Smith, Evolutionary Computing. Berlin, Heidelberg: Springer, 2015.
[38] J. Hernandez-Orallo, D. L. Dowe, and S. España-Cubillo, "On potential cognitive abilities in the machine kingdom," Minds and Machines, vol. 21, no. 2, pp. 179-210, Jun. 2011.
[39] OpenAI, "ChatGPT: Create with AI," OpenAI. [Online]. Available: https://openai.com/chatgpt. [Accessed: 19 July 2023].
[40] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, I. Sutskever et al., "Hierarchical text-conditional image generation with CLIP latents," arXiv preprint arXiv:2204.06125, 2022.
[41] Cycorp, Inc, "Cyc," Cyc. [Online]. Available: https://cyc.com/. [Accessed: 19 July 2023].
[42] B. Goertzel, "The General Theory of General Intelligence: A Pragmatic Patternist Perspective," arXiv preprint arXiv:2103.15100, 2021.
[43] X. Li, L. Zhang, Z. Wu, Z. Liu, L. Zhao, Y. Yuan, J. Liu, G. Li, D. Zhu, P. Yan, Q. Li, W. Liu, T. Liu, and D. Shen, "Artificial General Intelligence for Medical Imaging Analysis," arXiv preprint arXiv:2306.05480, 2023.
[44] E. Yiu, E. Kosoy, and A. Gopnik, "Imitation versus innovation: What children can do that large language and language-and-vision models cannot (yet)?" arXiv preprint arXiv:2305.07666, 2023.
[45] I. R. McKenzie, A. Lyzhov, M. Pieler, A. Parrish, A. Mueller, A. Prabhu, E. McLean, A. Kirtland, A. Ross, A. Liu et al., "Inverse scaling: When bigger isn't better," arXiv preprint arXiv:2306.09479, 2023.
[46] K. Grace, et al., "When will AI exceed human performance? Evidence from AI experts," Journal of Artificial Intelligence Research, vol. 62, pp. 729-754, 2018.
[47] J. Reis, P. Espírito Santo, and N. Melão, "Impacts of Artificial Intelligence on Public Administration: A Systematic Literature Review," in Proc. 2019 14th Iberian Conference on Information Systems and Technologies (CISTI), 2019.
[48] S. Cave and S. S. ÓhÉigeartaigh, "An AI race for strategic advantage: rhetoric and risks," in Proc. 2018 AAAI/ACM Conference on AI, Ethics, and Society, 2018.